Structured Payment in Pawnshop Borrowing: Mandates vs. Choice*

Craig McIntosh Isaac Meza Joyce Sadka Enrique Seira
Francis J. DiTraglia †

This draft: October 11, 2024

Abstract

Pawn loans offer borrowers an unparalleled degree of repayment flexibility in exchange for a harsh penalty in case of default: forfeit of collateral worth more than the loan amount. Using a large RCT conducted in Mexico City, we document key stylized facts and explore the merits of replacing flexibility with structured repayment contracts in this important but understudied form of credit. A novel experimental design that includes a mandatory frequent-payments arm, a (status quo) flexible payments arm, and a choice between the two point-identifies not only the average treatment effect, but also the effects of treatment on the treated and the untreated along with the average selection on gains, allowing a rigorous study of mandates versus choice. Although the average treatment effect of assigning borrowers to structured payments is a 22% decrease in their financial cost and a 15% increase the probability that they will recover their pawn, only 11% of borrowers choose structured repayment contracts voluntarily. We show that structured repayment benefits nearly all borrowers, including those who would not freely choose it, and find no evidence of selection on gains.

Keywords: Mandates versus choice, paternalism treatment on the untreated, payment frequency, overconfidence.

JEL codes: G41, C93, O16, G21

*We thank Jose Maria Barrero, Camilo García-Jimeno, Andrei Gomberg, Emilio Gutierrez, Anett John, David Laibson, Aprajit Mahajan, Matt Rabin, Mauricio Romero, Charlie Sprenger, Séverine Toussaert, Chris Udry, Jonathan Zinman, and seminar participants at Dartmouth, ITAM, MSU, USC, UC Davis, UCSD, and USC who provided valuable feedback. Ricardo Olivares, Gerardo Melendez, and Alonso de Gortari provided excellent research assistance and Erick Molina helped with formatting. Research assistance was financed through faculty grants at ITAM. Our research partner had no say in the design or results of the experiment.

†Seira: MSU, enrique.seira@gmail.com (corresponding author); McIntosh: University of California San Diego, ctmcintosh@ucsd.edu; Meza: Harvard University, isaacmezalopez@g.harvard.edu; Sadka: ITAM, jsadka@itam.mx; DiTraglia: Oxford, francis.ditraglia@economics.ox.ac.uk

1 Introduction

Pawn lending is one of the oldest and most prevalent forms of credit in the world (Carter & Skiba, 2012). Pawn loans are typically small, and short-term. They involve no credit checks or proof of income, and are widely used by people without access to other forms of credit. Borrowers surrender highly liquid collateral as the pawn, typically gold, in return for a loan that is substantially smaller than the value of their pawn. These loans are very flexible from the borrowers' side: there are no scheduled payments, no prepayment penalties, and interest charged only on outstanding balance. But unlike mortgage lenders, pawn lenders retain the entire value of the collateral under default along with any payments already made, rather than simply recovering the outstanding value of the loan. There are more than 11,000 pawn shops across the US with 30 million clients and $14 billion yearly revenues; in China pawn lending is a 43 billion dollar industry.[1] Despite this, pawnshops have received little attention in the literature to date.

Pawn loans are different from conventional credit in important ways. First, pawn loans typically default at much higher rates than other forms of credit. For example, 44% of loans made by our partner lender are not repaid, and 29% of those who fail to repay make one payment prior to defaulting, money which is completely lost to them subsequently. Second, lender profits are higher under default than repayment. An industry standard contract in Mexico City lends 70% of the value of gold collateral at an interest rate of 7% for three months. Because gold is highly liquid, the profit from a borrower who defaults is 30% of the collateral value compared to 22.5% in interest if a loan is paid in full on the last day. Furthermore, pawn borrowers are typically economically vulnerable, and in our context are also naively optimistic; borrowers report a subjective probability of repayment that averages 92%. The unusual degree of flexibility in these markets would benefit borrowers if they were neo-classical agents, but the prevalence of costly default that hurts borrowers but benefits lenders, combined with apparently naive borrowers, calls this into question. Other lending modalities with dramatically lower equilibrium default, such as microfinance, typically rely on highly structured repayment regimes.

This paper asks whether structured repayments can help pawnshop borrowers avoid default. To do so, we present a unique three-armed experiment covering around 4,500 pawnshop clients across six of our partner lender's Mexico City branches. Borrowers in the Control arm received the standard status quo pawn contract (described above) which allows us to document payment trajectories and default in standard pawn markets. Borrowers in the 'Structured' arm were required to make three monthly payments, with each monthly payment including the accrued interest at that time, and faced a nominal

[1] `https://tinyurl.com/ybm56dpe`, `https://tinyurl.com/y9zdcgws`, `https://tinyurl.com/y59ptdam`. In Mexico the scale of pawn lending rivals that of microfinance: during the past three years a single large lender in the country made over 4 million loans to more than a million clients, compared to a total of 2.3 million microfinance clients in all of Mexico during 2009 (Pedroza, 2010).

fee of 2% of that month's payment if the payment was delinquent. Borrowers in the 'Choice' arm were allowed to choose between the Standard and Structured contracts. A design that features both assignment and choice permits a nuanced understanding of how treatment effects relate to the choices that individuals freely make: because our context featured no differential selection of clients across arms, we can recover a unique combination of economically-relevant causal estimands.[2]

Comparing the Structured and Control arms gives the Average Treatment Effect (ATE) of structured repayment in the pawnshop context. Comparing the Choice arm to the Control provides a standard Intention to Treat (ITT) effect of offering this option. The presence of all three arms in the same experiment then allows us to point identify treatment effects *separately* for those who would and would not choose structure, given the choice. Consider only the control and choice arms. No one in the control arm is treated, but some borrowers in the choice arm choose to be treated: choosers. We can view the comparison of these two arms as an experiment with one-sided non-compliance. Here assignment to the choice arm serves as an instrumental variable for treatment receipt. Under an exclusion restriction, a familiar argument shows that the IV estimator identifies the causal effect for people who are shifted by the instrument: choosers. This is the Treatment on the Treated (TOT). Now consider only the Choice and Structured arms. Everyone in the Structured arm is treated, but some borrowers in the Choice arm choose not to be treated: non-choosers. We can view the comparison of these two arms as a separate experiment, also with one-sided non-compliance. Here assignment to the Structured arm serves as an instrumental variable for treatment receipt. Under a symmetric exclusion restriction, an almost identical argument shows that the IV estimator identifies the causal effect for people who are shifted by the instrument: non-choosers. This is the Treatment on the Untreated (TUT). As far as we are aware, this is the only experimental paper that provides a direct comparison of the TUT and TOT in the same context. This comparison is key for juxtaposing voluntary versus mandatory treatment.

Our ATE results show that mandatory structured payments are strongly effective in lowering financial costs to borrowers and preventing default in pawnshop lending: the average borrower in the Structured arm pays financing costs inclusive of fees that are 22% lower than the control, and faces a probability of default that is 6.6 percentage points lower (15% of the mean). In terms of Annual Percentage Rates (APR), the financial cost of borrowing falls by 11 percentage points per loan (19% of the mean). In short, structured repayment saves borrowers money by charging them fees. The mechanism by which our monthly payment contract achieves these cost savings is intriguing: not only does it improve repayment and speed up the pace of payments on average, but it also decreases the amount of money paid by borrowers who will ultimately go on to

[2]See Section 3.4 for details.

default. While we find that borrowers who experience structured payment contracts are 6 percentage points more likely to borrow again from the same lender, a back-of-the-envelope calculation suggests that the structured contract nevertheless lowers lender profits by at least 14% per borrower. It is thus unsurprising that structured payment contracts were unavailable in Mexico city pawnshops outside of our experiment: lenders profit from default.[3]

Despite the large financial cost savings that it provides, there is low demand for structured repayment: only 11% of borrowers in the choice arm choose structure. The mystery is then why so few people choose a contract that is financially beneficial on average. Could it be that this small fraction of "choosers" simply represents the only individuals who benefit from structure? The strongly significant difference between the Structured arm and the Choice arm suggests that this is not the case. More rigorously, an exercise that non-parametrically bounds the distribution of individual treatment effects, following Fan & Park (2010), shows that structured repayment benefits more than twice the number of people who chose it, even under the most conservative assumptions. Imposing the two exclusion restrictions mentioned above allows us to go further, point identifying the average treatment effects for choosers (TOT) and non-choosers (TUT) along with the average selection on gains (ASG = TOT − TUT), a measure of the empirical relevance of Roy-type selection. Our estimated TOT implies that choosers actually see a very small and insignificant increase in the cost of their loans from structure. The TUT effect, in contrast, suggests that the financial cost savings to non-choosers are large: $192 pesos, equivalent to a 10.6 percentage point savings in APR. On average, the borrowers who would *not* choose commitment, would have faced substantially lower financial costs if they had. We find no evidence of selection on gains. Finally, we combine our experimental treatment and outcome data with survey responses collected for a subset of borrowers to estimate conditional average treatment effects, both TUTs and ATEs, using the Causal Random Forest algorithm of Athey *et al.* (2019). We estimate positive conditional average TUT effects for 93% of the borrowers who did *not* choose the structured contract. In short, it is extremely difficult to find identifiable groups of borrowers who incur higher financial costs under a structured contract, even when restricting attention to those who would not choose it voluntarily

So why do borrowers seem to be leaving money on the table? One possible explanation could be impatience: under a standard model of rational choice, a sufficiently impatient borrower might prefer the status quo contract in spite of the higher probability of default. We find, however, that discount rates of almost 4,000% would be required to rationalize such a preference, making impatience an unlikely explanation for the low take-up of structured repayment contracts. Another potential explanation frequently raised in

[3]Unlike most pawn lenders, our partner operates as a non-profit, donating all of their earnings to philanthropic causes.

the literature is time inconsistency: hyperbolic borrowers might choose a contract that maximizes short-term liquidity. However, we find that a standard survey measure of time inconsistency does not predict larger TUT effects. Our data suggest that overconfidence provides a better explanation for borrowers' failure to select an apparently superior contract. In particular, we find that among borrowers who choose not to commit to a structured repayment plan, those who overestimate their probability of recovery the most would also benefit the most from structure. Overconfident borrowers appear to be unwilling to use tools designed to prevent default because they incorrectly believe they are not at risk.

While the results discussed above suggest that structure is broadly beneficial, even for those who would not choose it voluntarily, our estimated conditional TUT effects are negative for 7% of the untreated. This raises the question of whether it is possible to assign structured contracts in a *targeted* way to only those borrowers who are likely to benefit from them. To answer this question, we revisit the causal forest results described above for the subset of borrowers who completed our baseline survey. Our exercise restricts attention to survey variables that would be hard to manipulate ("narrow" targeting). With this restriction, all of the targeting rules that we explore have relatively weak predictive power. Our best-performing targeting rule only lower the overall mis-targeting rate from 9.7%–the baseline when all borrowers are assigned to structure–to 9.5%. In short, it is difficult to beat a policy that imposes structured contracts for all borrowers.

We make four main contributions to the literature. First, as far as we are aware, ours is the first paper to simultaneously identify both TOT and TUT effects in a single experiment, without the need for auxiliary structural modeling assumptions. Our three-armed experimental design also identifies the average selection on gains (ASG = TOT - TUT), average selection bias (ASB)–the difference in untreated potential outcomes for choosers versus non-chooers–and the average selection on levels (ASL)–the analogous comparison for treated potential outcomes. Our design is new to the literature. Like us, Fowlie *et al.* (2021) employ a one-stage, three-armed experimental design. But because they identify two TOT effects for different treated groups, rather than a TUT and TOT effect, their design cannot speak to the question of selection-on-gains. The two-stage designs of Karlan & Zinman (2009) and Beaman *et al.* (2023) likewise do not identify TUT effects. Our design is potentially transferable to other settings in which researchers aim to understand the relationship between treatment effects and compliance. We provide a companion STATA package that implements regression-based estimators for all identified causal parameters, along with cluster-robust standard errors.

Second, the benefits of universal structure in this context allow us to contribute to the relatively small literature on paternalism and selection-on-gains. In contrast to Laibson (2018), who studies private paternalism from a theoretical perspective, we show empirically that pawn borrowers would benefit, and lenders would suffer, from a paternalistic policy that mandates commitment in the form of structured repayment contracts. Laibson has spoken of how principals may conceal commitment features that agents need but do not demand; our results suggest that this picture may be inverted in the upside-down world of pawn lending: principals provide unstructured contracts that *induce* default, despite the fact that borrowers would benefit from the imposition of more structure. A pioneering paper in the relationship between choice and treatment effects is Sadoff *et al.* (2019), who shows that individuals with the most time-inconsistent preferences are the least likely to demand commitment. Our papers are complementary; while they identify one particular mechanism generating low commitment take up (i.e. time inconsistency), we are able to identify a money-metric value for the amount forgone by not taking up commitment without the need to first elicit preferences before testing for negative selection. Relatedly, whereas Walters (2018) combines a distance-based instrument with structural modeling assumptions to show that students who select into more effective schools have smaller treatment effects (TOT $<$ TUT)., we identify TUT and TOT effects without the need for a structural model.

Third, our paper provides a detailed picture of pawn lending, a widespread and important financial service that, despite its prevalence and impact on economically vulnerable borrowers, has received little scholarly attention. While an older literature considers the exploitative potential of over-collateralization and underpriced collateral (Basu, 1984), the behavioral implications of such contracts remain largely unexplored. To our knowledge, ours is the first paper to focus specifically on pawn lending.

Finally, we contribute to a growing literature on the effects of payment frequency in lending. In the context of microfinance, (Field & Pande, 2008) find no effects from frequent repayment schemes. Other studies find that increasing repayment flexibility in microfinance (by giving grace periods to pay or options to delay payment) improves business performance (Field *et al.*, 2013; Barboni & Agarwal, 2023; Battaglia *et al.*, 2023). Our experiment in contrast, highlights some of the dangers of flexibility in the context of overcollateralized pawn lending. The setting of microfinance experiments differs from ours in an additional way: they are performed on top of already highly structured microfinance contracts, and they involve borrower pools who may have selected into that type of lending precisely because it provides structure (Bauer *et al.*, 2012). This difference may explain why Field & Pande (2008) find almost no default in their control group, compared to 44% in ours. In addition to studying the effects of payment frequency, we examine both take-up and the efficacy of voluntary commitment mechanisms. A number of papers have found low demand for commitment as we do–structured payment contracts with late payment fees are a form of commitment. While Ashraf *et al.* (2006), Giné *et al.* (2010), Bai *et al.* (2020), Royer *et al.* (2015), Sadoff *et al.* (2019) find low demand for commitment in many domains, others have found more robust demand (Kaur *et al.* (2015), Casaburi & Macchiavello (2019), Schilbach (2019), Tarozzi *et al.* (2009), Dupas & Robinson (2013)). Unlike all of these papers, however, we separately point-identify and estimate the effects of commitment for borrowers who would and would not choose it.

The remainder of the paper is structured as follows. Section 2 provides context and defines our main outcome variables. Section 3 describes the experiment and data sources, and shows pre-treatment balance across arms. Section 4 provides the standard ITT analysis of the experiment, while Section 5 shows how to identify, estimate and carry out inference for the TOT, TUT, ASG and related causal effects under our three-armed experimental design. Section 6 investigates why mandating structured repayment functions so well in this context and whether it can be more finely targeted. Section 7 concludes.

2 Context

2.1 Pawnshop borrowing

Pawn loans involve individuals leaving valuable liquid assets, typically jewelry, as collateral in exchange for an immediate cash loan. Collateral is typically more valuable than the loan amount, allowing lenders to give the loan immediately without checking a borrower's credit history. This makes pawn loans a popular way to get cash to pay for emergencies. In fact, they are one of the most prevalent forms of borrowing. There are more than 11,000 pawn shops across the US, with 30 million clients and $14 billion in yearly revenues.[4] Our partner pawn lender alone served more than 1 million clients in the last 3 years with more than 4 million contracts. For comparison there were 2.3 million micro-finance clients across all lenders in Mexico in 2009 (Pedroza, 2010).

Pawning is also one of the oldest forms of borrowing. Pawn lending existed in antiquity at least since the Roman Empire, and there are records of it in China about 1,500 years ago (Gregg, 2016). In spite of the high prevalence and long history, pawnshop borrowing has not received much attention in the economics literature. The closest widely studied product is payday lending. In developing countries, however, payday lending is likely small compared to pawnshop lending; the latter is faster and requires less documentation, making it more accessible to informal sector workers who receive their salaries in cash.

As with payday lending, pawnshop lending is controversial. Regulators have concerns with the sophistication of borrowers using it, speculating they may suffer from behavioral and cognitive deficiencies that lead to making sub-optimal choices, biases that are exacerbated by contract design.[5] There is some evidence in support of this view for payday

[4] See here, here, and here.

[5] The US congress has actually banned the payday lending industry from serving active military personnel, and some States in the US have imposed zoning restrictions, interest caps, and restrictions

borrowers[6] but none for the large pawn-lending industry. Our study reinforces the idea that a lack of sophistication may be an integral part of the way that standard pawn contracts are designed and structured by lenders.

2.2 Pawning Logistics and Contracts

To study this market, we partnered with one of the largest pawn shops in Mexico, an institution with more than one hundred branches spanning multiple states. This lender (whom we refer to as "Lender P") has a simple and typical business model.

Appraising and Lending Lender P takes gold jewelry as collateral in exchange for a fraction of the value of the piece, in cash. No other collateral and no credit history checks are needed. The transaction takes less than 10 minutes and is conducted at the branch in person between the client and the appraiser (i.e. a teller). The appraiser weighs the gold piece and runs tests on its purity. Based on these she assigns a gold value to the piece, stores it as collateral, and gives 70% of the gold value of the piece in cash to the client. The borrower signs a 2-page contract with the conditions of the loan and leaves with the cash. Appraisers were paid a flat monthly wage.

Contract Lender P had only one type of contract, henceforth the *status quo* contract. It stipulated that the interest rate was 7% *per month* compounded daily on the outstanding amount of the loan. The loan had a 90 days term with 15 days' grace period. The client could make payments at the branch at any time with no penalty for pre-payment. Under this status quo contract, there are no payment reminders or any other kind of interim contact between the lender and the borrower. Internal lender rules allocate payments first to interest and fees owed and up to the day of the payment, and to capital only after the first has been covered. If the client returns to pay the principal plus the accumulated interest within 105 days, she recovers her pawn, otherwise the pawnbroker keeps the piece *and* any payments already made. Before the contract expires, the client had the right to renew for another 3 months by going to the pawnshop, paying the accumulated interest, and signing a new contract with exactly the same terms and the same piece as the original contract (38% of borrowers renew at least once with a given pawn). This contract is standard in the industry. Pawnshops make money in three ways: by reselling

on serial borrowing as consumer protection measures against payday lending (Stegman, 2007).

[6] Bertrand & Morse (2011) write that "Under the view that the people borrowing from payday lenders are making an informed, utility-maximizing choice given the constraints that they face, one would not expect additional information disclosure about the payday product to alter their borrowing behavior", but to the contrary, they find that simply disclosing how financing costs add up reduced demand by 11%. Melzer (2011) finds that payday loan access leads to increased difficulty paying mortgage, rent and utilities bills.

the jewelry left as collateral on defaulted loans, by charging interest on non-defaulted loans, and by keeping the payments made on defaulted loans.

2.3 Measuring Borrowers' Financial Costs

Borrowers' financial costs are composed of two main categories: the cost of losing their collateral, and the interest and fees incurred during the life of the loan. For each given loan we observe if the client lost her pawn, $\mathbb{1}(\text{Default}_i)$. If a loan has been rolled over and is still outstanding, we consider it to be non-defaulted. This approach is conservative in our context (biases treatment effects towards zero), as we show in detail in the appendix. In our data 13% of experimental loans are ongoing (i.e. censored) when the data period ends. Regarding interest, our administrative data classifies payments made into three types according to their payment allocation rules: payments to principal P^C, payments on generated interests P^I, and payments on penalty fees P^F. We observe each and every payment made under each category, its amount and date.

We define a borrower's financial cost as the total monetary outflow—in cash or pawn value—from the borrower to the lender. This includes all payments the borrower made toward interest and fees, but also the net difference between the appraised value of the pawn and the loan amount (Value-Loan) in the event of default. When there is no default the borrower recovers her pawn and there is no loss of value for the borrower. Payments towards capital are considered a cost only when the borrower defaults, as she is not reimbursed for these payments. Note however that when she does not default payments to capital are not an actual outflow, as they sum up to the value of the loan the lender disbursed in the first place. The formula for financial cost for person i is thus as follows:

$$\text{Financial Cost}_i = \sum_t P_{it}^I + \sum_t P_{it}^F + \mathbb{1}(\text{Default}_i) \times \left(\text{Value}_i - \text{Loan}_i + \sum_t P_{it}^c \right)$$

where t indexes days, and $\mathbb{1}(\text{Default}_i)$ is an indicator function for defaulting. Because the period of the loan is only 90 days we do not apply discounting in calculating costs. In robustness checks reported below we show that our results are virtually unchanged when applying a wide range of time discounting factors.

We consider the above to be an accurate measure of financial cost in pesos. However, we also report results incorporating two non-financial costs: (i) using the subjective value of the pawn reported by the borrower in place of its appraised gold value, and (ii) adding a measure of travel expenses and the opportunity cost of time, as clients have to go to the branch in order to make payments.

As a second measure of cost we calculate the Annual Percentage Rate (APR) in order to express the cost as a percentage of the loan, per year, inclusive of default costs. The standard definition is given in the formula below. Note that in contrast with the standard calculation of APR, our measure of financial cost includes the expected cost from defaulting, and important distinction in this context.

$$(\text{APR})_i = \left(1 + \frac{\frac{\text{Financial Cost}_i}{\text{Loan}_i}}{\text{loan term}_i}\right)^{\text{loan term}_i} - 1$$

3 Experiment, Data and Stylized Facts

3.1 Treatment arms and randomization

The Structured contract For the purpose of the experiment we designed a new contract that is identical to the status quo contract except that, informed by the design of micro-lending contracts, it enhances the regularity and salience of payments as a way to encourage repayment (Morduch, 1999; Bauer *et al.*, 2012). It has the same interest rate (7% *per month*) which accumulates daily on outstanding debt, the same loan size/collateral ratio (70%), and the same loan term (90 days, and a grace period of 15 days). Borrowers' gold pawns are appraised in the same way by the same appraisers under both the new and status quo contracts. The Structured contract however requires the client to make regular monthly payments for the duration of the contract, with the principal and interest payments split evenly across the three months of the contract (day 30, 60 and 90 after loan disbursement). The importance of this monthly payment was made salient in the contract and payment receipts, and by the levying of a nominal fee (2% of minimum due) on individuals who fell behind in their payments. The fee was modest and intended to make the payment deadlines salient. As a benchmark, the transportation cost to visit the branch to make a payment is comparable to the fee, on average.

To elicit demand for the monthly payment contract, we include an arm that allows borrowers to opt into this contract if they choose. The existence of both a non-optional "structured" arm, and a choice arm in our design is key to estimating a battery of treatment effects under fairly mild assumptions. We next describe the three experimental arms in more detail.

Treatment Arms Treatments were randomized at the branch-day level. Each day a computer randomly assigned which types of contracts were on offer that day in the branch, and the IT system would only offer these. We have 3 different experimental arms[7]

[7]The experiment included other independent arms that involved no fee penalties and did not emphasize the structure of payments. These are being analyzed in a separate paper.

1. *Control* arm: consisted of branch-days offering the status quo contract described in Section 2, and only this contract.

2. *Mandatory structured payments* arm: consisted of branch-days requiring all borrowers to use the Structured contract described above.

3. *Choice* arm: consisted of branch-days offering the client *a choice* between the structured payments contract, and the status quo contract.

We did not allocate an equal number of days across arms, since we were interested in having more power in some of them. The number of branch days allocated to each were 84 to control, 80 to mandated structured payments, and 93 to choice. See Figure 1 for a CONSORT-style diagram of the study design and recruitment.

Figure 1: Experiment description

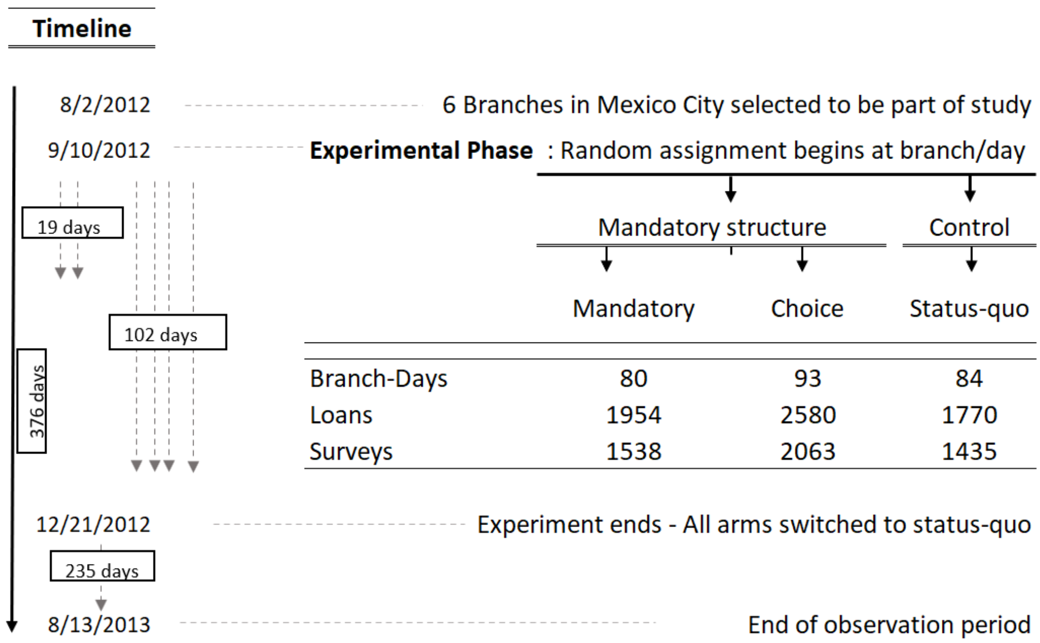

Randomization We implemented the experiment in six branches of Lender P beginning on September 9, 2012. The branches were selected by Lender P to be dispersed across Mexico City and have varying sizes. In four of them the experiment ran for 102 days, and in 2 of them we ran it for a shorter time to economize on data collection costs once we realized we would not be constrained by sample size. Branches are more than 5 km apart from each other, and there is no substitution among them; none of the borrowers appear in more than one of our branches.

Branch personnel did not know which treatment would be assigned to each day and were blind to the objective of the intervention. They were told that there were 3 different "types of contract-days", that the system chose randomly for any given date, and that it could happen for instance that two or more consecutive dates had the same contract. They were also told that this way of operating was in place in several of Lender P's branches (they did not know which ones), and that it would be in place for several months. Randomizing at the day level limits the problem of contamination arising from clients realizing that other clients get different contracts than theirs. It also limits potential manipulation by appraisers, who in the presence of individual-level randomization could potentially pick their preferred customer from the line or tell them to wait until their desired contract shows up on the screen. Intra-branch day correlation on the probability of default (ICC) is small, at 0.05, so we lose little power vis-a-vis individual-level randomization.

Some clients pawned more than one time during the duration of the experiment, with 14% pawning 2 times and 8% more than 2 times. To have a clean comparison we consider only the first pawn conducted during the experimental window. Moreover, 30% of first pawns involve more than 1 loan, as 2 or more pieces of gold were submitted. We treat each of them as separate loans. In the appendix we show that our results are robust to this analysis choice.

Timeline Figure 1 shows the experimental timeline along with the length of time for which we observe payments. For loans made in the first week of the experiment, we observe up to 338 subsequent days of loan information; for loans made in the last week we observe up to 235 days. Figure 1 also illustrates the number branch-days per arm, the number of loans, and the number of surveys.

Explaining the Contracts We made sure clients understood the contract terms. First, we had full-time enumerators explaining contract terms to clients. The explanation took about 3-5 minutes and continued until the client said she understood the contract terms. Enumerators then asked clients to explain the contract back to them before correcting any misunderstandings. Second, the appraiser gave clients the "Contract Terms Summary" and read it out loud to them before after their piece had been appraised but before they signed the contract. We are confident the overwhelming majority of clients understood the contracts and that those in the choice arm made informed choices.

3.2 Data

Administrative Data The study exploits two types of data: administrative data from the lender, and a short survey that we implemented. The administrative data contains a unique identifier for each client, an identifier for the piece she is pawning, and the transactions relating to that piece. These transactions include the value of the item as assessed by the appraiser, the amount of money loaned (70% of the item's value), the date of the pawn transaction, and the type of contract for that pawn: structured payments or status quo. Within the period of the loan, we followed each transaction related to that piece in the administrative data: when payments were made and for what amounts, whether there was default (i.e. the client lost her pawn), and whether any late-payment fees were imposed. After the experimental loan, we are able to track subsequent behavior and to see whether that borrower took a subsequent loan. We have this information for all the pawns that occurred in the experiment's 6 branches between August 2, 2012 and August 13, 2013. This includes all the pawns that took place during our experiment along with those that originated one month before and eight months after our experiment. Figure 1 shows the design and timing of the experiment, along with the sample sizes in each arm. The experiment comprises 6,304 pawns while our administrative data covers a total of 26,180 pawns.

Survey Data An additional team of enumerators was stationed at the entrance of each branch; they asked clients to complete a 5-minute survey before going to the teller window to appraise their piece and *before* they learned which contracts would be available on that day. The survey allows us to measure loan takeup since we know whether people entered the branch, completed the survey, and then, after learning the contract terms, decided not to pawn. As we will see, this is extremely infrequent; only about 3% of borrowers did that. This is not very surprising, borrowers come determined to borrow, many of them to solve an emergency, and are not very sensitive to contract terms (that may explain why they accept high interest rates and over-collateralization in the first place). The survey was intentionally brief. It measured demographics, proxies for income/wealth, education, present-biased preferences, experience pawning, if family or friends commonly asked for money, how time-consuming and costly it was to come to the branch, the subjective probability of recovering the piece that they intended to pawn, the subjective value of their piece in money terms (how much money they would sell it for), among others. The reported surveys are at the pawn level. One same borrower can have multiple surveys, one for each pawn, resulting in over 6,000 surveys.

3.3 Stylized facts

Given the scarcity of evidence about pawn loans in economics, documenting stylized facts is a useful contribution. This section provides some important ones.

Borrowers. Borrowers are not inexperienced: 87% of clients in our survey report they have pawned before. However they are economically vulnerable: 20% of them could not pay either water, electricity & gas or rent in the past 6 months. In terms of triggers and uses of the loan, 89% said they are pawning because of an emergency, and only 12% stated it was to be used in a "non-urgent expense". When asked why they are pawning this piece, 5% responded "lost a family member", "a medical emergency" (11%), or "an urgent expense" (72%). Only 32% of borrowers have completed college, 73% are women, with an average age of 43 years.

Many borrowers lose their pawn. Our context is also one with high borrower default: 44% of clients lose their pawn in a time span of 230 days from the date of pawning. One potential explanation for high default is that clients are really just knowingly selling their gold piece through a pawn contract on which they intend to default. This appears unlikely for several reasons: clients can easily sell the gold and obtain a higher amount of instant cash at gold-buying stores located close to almost all our pawnshop branches; the reported subjective value of the pawn is larger than the appraised value for 83% of clients (this also means it is inefficient that they lose their pawns). The default rates for the experimental period and branches is similar to that observed for all branches in subsequent periods (see Figure OA-2).

Gold buying shops. Gold purchasing shops locate close to our Lender's branches. Branches have an average of 2.8 of these shops in the vicinity. This fact is important given the high default rates which could indicate that borrowers are ex ante intending to "sell" their pawn. Given that they are receiving 70% of the value of the pawn next door to shops that would give them 100% of this value, we think this is prima facie unlikely.

Paying without recovering. Among those that lose their pawn, 29% paid a positive amount towards its recovery. On average this subset of borrowers paid 34% of the value of their loan (see Figure OA-1 in Appendix). This can only be rationalized if they expected to recover their pawn.

Overconfidence. Our baseline survey asked borrowers for their subjective probability of recovering the pawn. The average self-assessed recovery probability is 92%. This contrasts sharply with the actual recovery rate for these same borrowers, 56%. Overall 72% of borrowers report a 100% probability of repaying their loan.[8]

[8] A recent paper by Allcott *et al.* (2021) finds only the most inexperienced borrowers from a U.S. payday lender are overconfident, but their study examines priors about getting out of debt in the future rather than priors on repaying the current loan, as we do.

3.4 No differential selection or imbalance

We were aware that the major threat to experimental validity in this study was selection into borrowing based on contract terms. We now show in three different was that there was *no such selection*. First, among borrowers who entered the branch and completed the baseline survey (at which time they did not yet know their treatment status), 96.5% end up pawning. This rate is not only high but is virtually identical across treatment arms. Second, the number of loans per day did not differ across experimental arms, showing that borrowers did not stop pawning when the structured arm or choice arm was active. Third, borrower and loan characteristics are the same across arms, showing that there is no selection on observables.

An important implementation detail, mentioned above, is that appraisers were paid a flat wage, independent of how many contracts or which types of contracts borrowers signed. They therefore had no incentive to influence borrower choice. Consistent with this, adding or removing appraiser fixed effects leaves results unchanged.

Table 1 uses administrative data to show there are no differences in the number of loans, borrowers, or amount borrowed across arms. Each row corresponds to a regression of the dependent variable on indicators for treatment assignment (status quo/control, structured repayment, choice). In the first row ("Take-up") we use the sample of those individuals who answered the baseline survey (3710 subjects) and check which of them pawned based on the administrative data. The dependent variable indicates whether the person pawned. Because this survey was implemented at the entrance of the branch before they knew their treatment assignment, it is a good measure of take-up. The first result is that the overwhelming majority (96.5%) take the loan. Second, the same fraction of borrowers take up the loan across treatment arms. This is strong evidence that borrowers did not select into treatment arms, which were announced between survey recruitment and contract completion. The remaining rows use only administrative data. For each branch day, we calculate the number of borrowers pawning, the number of pawns used, pawns per borrower, and the total amount borrowed and regress each of those on treatment arms. We again cannot reject that any of those measures of demand are the same across arms. We also report medians to reduce the influence of outliers. Overall, this strongly suggests that the treatments did not induce differential selection.[9]

[9] Had we found differential selection, the cross-arm effects would still be relevant from the perspective of the lender, but would no longer represent intensive margin impacts on comparable groups of people.

Table 1: No selection across arms

	Control	Structure	Choice	p-value
Take-up	0.967	0.955	0.961	0.82
	(0.01)	(0.01)	(0.01)	
Number of borrowers	20.8	22.2	25.4	0.17
	(3.29)	(3.9)	(4.89)	
median	19	20	21	0.5
Number of pawns/borrower	1.4	1.4	1.4	0.43
	(0.08)	(0.04)	(0.05)	
median	1.4	1.3	1.3	0.6
Number of pawns	31	31.3	37.2	0.24
	(5.8)	(5.6)	(7.9)	
median	27	28	30	0.46
Amt borrowed/borrower	2266.8	2094	2115.2	0.18
	(101.8)	(83.7)	(99.9)	
median	2154.3	2041	2047.5	0.65
Total borrowed	47877	47813	54780	0.4
	(8005)	(9436)	(12587)	
median	37520	39420	40850	0.73
Obs	85	81	94	

Each row in this table corresponds to a regression at the branch-day level of each dependent variable against the experimental arms indicators (control, mandatory structure , choice). The table reports the coefficients on each of these indicators, as well as the p-value of an F-test of the null hypothesis of equality of the three coefficients. We also report results using median regression.

Table 2 presents summary statistics for the sample of actual borrowers across arms, showing that our randomization succeeded in achieving balance across the experimental arms. 73% of surveyed borrrowers are women, the average age is 43 years, 66% have completed at least a high school education, and 87% have pawned before. Finally, borrowers' subjective probability of recovering their pawn is close to 92%. The average subjective value they report for the items they pawn is 4084 MXN, much larger than the average appraised gold value of 3238 MXN. Importantly, all borrower's characteristics are balanced across arms, again consistent with no differential selection across arms. Tables 1 and Table 2 tests are the standard experimental integrity analysis. Together they constitute strong evidence of no borrower selection across arms. This was reinforced by reports from our enumerators and appraisers that those who entered the branch were determined to borrow and that there were no complaints about what loan contracts were available. Importantly, this does *not imply* that borrowers had no preference over alternative contracts. It simply means that the outside option of leaving the branch without a loan is unattractive.

Table 2: Borrower's characteristics are balanced

	Control	Structure	Choice	p-value
		Panel : Survey Data		
Subjective value	4084	3877	4173	0.51
	(186)	(193)	(172)	
Trouble paying bills	0.19	0.21	0.18	0.67
	(0.024)	(0.023)	(0.02)	
Present bias	0.14	0.13	0.13	0.89
	(0.02)	(0.01)	(0.01)	
Makes budget	0.62	0.59	0.65	0.29
	(0.028)	(0.036)	(0.021)	
Subj. pr. of recovery	91.89	91.65	93.61	0.09
	(0.721)	(1.031)	(0.582)	
Pawn before	0.87	0.89	0.9	0.25
	0.02	(0.013)	(0.011)	
Age	43.32	42.85	43.82	0.73
	(0.688)	(0.949)	(0.792)	
Female	0.73	0.72	0.71	0.88
	(0.023)	(0.019)	(0.02)	
+ High-school	0.66	0.67	0.65	0.84
	(0.027)	(0.022)	(0.018)	
Obs	1386	1469	1982	

The table uses survey data to show balance across arms. Survey data is not used for the main results of the paper. It is only used in Section 6. Each row in this table corresponds to a regression, where the level of observation is a borrower that pawned which was surveyed. The dependent variables are displayed in the first column. Each dependent variable is regressed against the experimental arms indicators (control, mandatory structure, choice). The table reports the coefficients on each of these indicators, as well as the p-value an F-test of the null hypothesis of equality of the three coefficients. "Subjective value" of the pawn refers to how much would the client be willing to sell the pawn for in pesos, "Trouble paying bills" is an indicator for the borrower reporting having problems paying electricity, water and services in the last 6-months, present bias is constructed as in Ashraf *et al.* (2006); "Makes buget" as an indicator for whether the household make expenses budget for the month ahead of time. The subjective probability of recovery was elicited a la Manski (from 0 to 100 what is the probability that you will recoup your pawn); pawned before is a dummy=1 if the client declares to have pawned before (although not necessarily with Lender P); age is age of the borrower in years, +High-school is a dummy that indicates if the client has completed high school.

4 Average Treatment Effects

We begin by estimating average treatment effects of assignment to the Mandatory structured payments and the Choice arms, relative to those assigned to the Control arm. As we explain below, only about 11% of those in the choice arm chose the monthly payment option (Figure OA-3 shows coefficient plots for the characteristics that determine choosing structured payments in the Choice arm).

Specification Table 3 presents estimates and standard errors from a standard pooled experimental regression

$$y_{ij} = \alpha + \beta^M T_i^M + \beta^C T_i^C + \gamma X_{ij} + \epsilon_{ij} \tag{1}$$

Table 3: Effects on Financial Cost

	FC	Components of FC					APR
		Interest pymnt	Fee pymnt	Def×Ppl pymnt	Lost pawn value	Default	
		$\sum_t P^I_{it}$	$\sum_t P^F_{it}$	$\mathbb{1}(\text{Def}_i) \times \sum_t P^C_{it}$	$\mathbb{1}(\text{Def}_i) \times \text{Value-Loan}_i$	$\mathbb{1}(\text{Def}_i)$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Mandatory structured	-204.0***	-157.3***	32.1***	-1.27	-78.8**	-0.066***	-0.11***
	(48.1)	(34.9)	(1.43)	(3.10)	(31.6)	(0.023)	(0.019)
Choice	-38.9	-24.9	1.34**	-0.93	-15.4	-0.023	-0.0086
	(49.8)	(38.4)	(0.54)	(3.02)	(33.1)	(0.021)	(0.019)
Observations	6304	6304	6304	6304	6304	6304	6304
R-squared	0.013	0.022	0.151	0.003	0.007	0.013	0.031
Control Mean	942.4	545.9	0	5.96	396.5	0.44	0.57

This table shows the treatment effects for our core pecuniary outcomes. Each column is a different regression for different outcomes on an indicator for the mandated and choice arms, following specification in equation 1. Columns (1) & (7) analyze our core financial cost measures, while the rest of the columns decompose these into finer components. A few borrowers take more than one loan on the first day they appear in an experimental branch. These are treated as different observations. Additional results, available upon request, show that our results are robust to different ways of handling multiple loans for each borrower. Each regression includes branch and day-of-week FE. Standard errors are clustered at the branch-day level.

where i indexes the borrower, j indexes branch, T_i^M and T_i^C are indicator variables for receiving the Mandatory or Choice arms, X_{ij} are branch and day-of-week fixed effects. Standard errors are clustered at the branch-day level, the unit of treatment assignment[10]. Given full compliance in the Mandatory arm, the coefficient β^M is the ATE of mandatory structure while β^C is the ITT of structure in the Choice arm on the outcome variable y_{ij}. Our two primary outcome variables are financial cost in pesos and Annual Percentage Rate (APR), as defined in Section 2.3.

Results for these outcomes appear in columns (1) and (7) of Table 3. The remaining columns decompose the financial cost and APR outcomes into their components: interest payments (col 2), payment towards any fees incurred (col 3), payments toward the principal (col 4). Column 5 shows the value of lost pawn conditional on losing it. In column 6 the dependent variable is a dummy indicating default. Finally, column 7 rescales financial cost as a function of loan size to estimate causal effects on incurred APR.[11]

Results The results are stark. The mandatory structure arm yields large and significant decreases in the cost of loans to clients, as measured either by financial cost or APR. Despite causing an increase in fees, the Mandated arm leads to a decrease of 204 pesos in

[10] A minority of borrowers pawned on more than one day during the experiment: 14% pawned on two distinct days, and 8% on three or more days. To avoid contamination from earlier treatments to which these individuals were exposed, we restrict our sample to each borrower's *first visit.* Note that a borrower may pawn multiple items her first visit. We include these as separate observations. Because our standard errors are clustered at the branch-day level, they automatically account for any dependence in error terms arising from multiple pawns by the same borrower on her first visit.

[11] As we explained above, loans can be extended for an additional 3 months by paying the interest owed and restarting the loan under the same treatment conditions. This means that some loans extend for more than 3 months. We consider the entire flow of cost for the duration of our sample.

the costs of borrowing (out of a group mean of 942 in the status quo), equivalent to 22% reduction as a fraction of mean cost. These cost savings arise from a 6.6 percentage point (pp) decrease in the probability of default (out of a baseline mean of 44pp, implying cost savings of 79 pesos), and also from lower interest payments since, as we will document below, the commitment-style contract speeds up payments so the interest rate applies to a smaller principal. This translates into a large reduction in APR. A credit product that has an effective average APR of 57% in the status quo arm (inclusive of default) is reduced to a cost of 46% through the imposition of a more regularized payment structure. This is in stark contrast to Field & Pande (2008) for instance.

In contrast, to the impressive effectiveness of the mandatory structure arm, the Choice arm fails to deliver significant changes in any measure, with the exception of an increase in fees, for which we are highly powered, since this outcome is zero for every control observation. Giving borrowers the choice of contract did not decrease financial cost, whereas forcing them into a structured payment contract dramatically reduced it. As we explore later in the paper, the null effect of the Choice arm arises because few borrowers demanded structure (consistent with Ashraf *et al.* (2006), Giné *et al.* (2010), Bai *et al.* (2020), Royer *et al.* (2015), Sadoff *et al.* (2019)), with 89% choosing the less effective status-quo contract.

Payment speed To shed light on the mechanisms behind the ATEs discussed above, Table OA-2 shows how structure affects a number of intermediate outcomes. One can group the types of intermediate outcomes into two categories: measures of the speed of pawn recovery, and measures of the decision of when to default. While the first payment for borrowers in the status-quo contract occurs on average only on day 82 (on a 90 days contract), borrowers in the mandatory structure arm start paying 13.8 days earlier on average (col 1). Not only do they start paying earlier, the first payment is also 9.7% larger (col 2), and a larger fraction (7.9%) actually pay in full and recover their pawn in the first payment, compared to 30% in the status quo contract (col 3). The resolution of the loan (either by payment or default) is shortened by 27.9 days (col 4), and conditional on recovery (an endogenous control) by 17.9 days (col 5).

Payment Bifurcation A very undesirable outcome from the borrower's perspective is to pay towards the loan without paying in full, i.e. to default on the loan despite making some payments. In this case, borrowers lose both the collateral and any payments made toward recovery. One could be concerned that by encouraging them to pay monthly, more borrowers might end up in this dire scenario in the mandatory structure contract. Column 6 shows this is not the case. On the contrary, 7 percentage points fewer borrowers end up in this situation, compared to a status quo contract mean of 12 percent. Conditional on defaulting those assigned to the mandatory structure contract have paid 4.1% less of their loan (col 7), and a 14 percent higher proportion of borrowers pay zero conditional on defaulting, an outcome analogous to "selling their pawn" (col 8). One interpretation of this bifurcation is that the mandatory structure contract forces borrowers to think earlier about whether they will indeed be able to eventually recover their pawn, and separates borrowers into those "selling" their pawn and those recovering it, reducing the share of undecided borrowers that end up paying interest and end up losing the pawn anyway. This mechanism may also help to explain why the mandatory structure contract does not increase the number of visits to the branch to pay (col 10): those recovering their pawn visit more, but those defaulting have 0.20 fewer visits (col 11). Finally, Column 9 shows that treatment effects are concentrated in the intensive margin as treatment does not affect the fraction of clients who pay a positive amount towards pawn recovery.

Other Costs Although the paper focuses on financial costs, we consider three additional costs here. First, we include the cost of visiting the branch to make a payment. This includes the self-reported transport cost (most use public transport), as well as the opportunity cost of time. To err on the conservative side, we subtract a whole day's minimum wage for the day of the visit. Second, we consider a rough proxy of the value of liquidity that borrowers lose by paying sooner. To do this we add the interest costs on to the payments in the mandatory structure arm and recompute treatment effects with these payments compounding daily (as if they had to borrow in order to make the more rapid payments). Third we consider the *subjective* value of the pawn. All results presented thus far value the collateral as appraised by the lender, but the piece may be worth more to the borrower than its gold value. For many of them the pawned jewelry has sentimental value. This is reflected in the subjective valuation they reported in the survey, which is almost twice as large as the appraised gold value of the pawn on average. Our third extra cost considers this higher subjective value. Table OA-4 shows that our results are robust to all these changes.

It is worth emphasizing that our results do not directly measure borrower welfare. Financial cost savings and welfare could diverge in a number of ways. For example, requiring monthly payments could cause an increase in borrowers' levels of stress, an effect that our measures do not take into account. Moreover, given the short time frame of pawn loans, we do not explicitly account for discounting. As shown in Section 6.1, however, one would need to assume unreasonably high discount rates to cast doubt on our findings. Similarly, we cannot measure the financial cost of other loans that borrowers may potentially take out to make interim payments on their initial pawn loan. Given that the "penalty" from a missed interim payment is small–less than the transport cost to visit the branch–we consider it unlikely that borrowers will have taken out even more expensive loans to avoid the penalty. Having said this, a large literature focuses on the cost of servicing loans. We follow in the steps of this literature.

Censoring of Loan Completion The window of time during which we were able to observe borrower behavior was limited in each branch, meaning that there were loans that we do not see completed (particularly those pawns that were rolled over for one or two further 90-day spells). Overall, 13% of all experimental loans are censored, meaning that they neither default nor repay within the observation window. In the prior analysis we handle this issue in a conservative way by using outcomes such as "did not default" which are well-defined even when we do not observe the completion of the loan, resulting in an estimate biased toward zero by the more rapid loan repayment observed in the Structured arm. In the Online Appendix we consider this issue in more detail. Most importantly, Table OA-3 conducts a bounding exercise that examines how large the effects of this censoring could possibly be by making bracketing assumptions about repayment on censored loans in the treatment and control, respectively. Comfortingly, Panel B of this table shows that even the most muted possible effect in the bounding exercise still recovers impacts of Structure on financial costs and APR that are negative and significant at the 99% level. Using a lasso-logit model to predict the outcome of censored loans, Panel E shows the APR impact of Structure increases from a 11 pp reduction (headline results) to a 17 pp reduction. Hence there appears to be no scope for this censoring issue to overturn our results, and our core results (implicitly assuming censored loans are not defaulted) is almost certainly an under-estimate of the true impacts.

Lender profit One of the questions we are interested in is whether lender profits decrease with structured payment. If that is the case, lenders will have little incentive to include this feature in their contracts. For a single loan, lenders get what borrowers lose in a zero-sum way so that a borrower's Financial Cost_i for loan i is the profit from that loan for the lender. Table 3 showed that borrowers pay less to lenders *per loan* under the structured payments contract. However, in Table OA-5 in the appendix, we show that borrowers assigned to the mandatory structure arm are 6.7% more likely than in the status quo arm to be repeat customers during our sample period. This brings some extra benefit to lenders that we take into account in a simple back-of-the-envelope calculation. From column 1 of Table 3 we know {Financial Cost_i | Contract=X} for $X=$ {status quo, structure} contracts are respectively \$942MXN and \$738MXN. From Table OA-5 we know Pr(Repeat|Contract $= X$) for $X=$ {status quo, structure} are 32% and 38.7%, respectively. We can use these four numbers to calculate a discounted sum of profits for client j comparing a scenario where only status quo loans are offered with a second scenario where only structured payment loans are offered.[12] Assuming independence across repawning events within a given contract type, one can write: $\text{Profit}_{j,X} = \sum_{t=0}^{T} \delta^t \{\text{Financial Cost}_j | \text{Contract} = X\} \times \Pr(\text{Repeat}|\text{Contract} = X)^t$, where

[12] We do this for simplicity. Otherwise, we would have to assume something about how borrowers switch dynamically across types loans, and this is something we do not observe in the data.

δ is the discount rate. We find that for every $\delta \in [0,1]$ and for every T the status quo results in higher profit for the lender. That is, the extra repeat probability for the structured payment contract is not enough to compensate for the loss in each of those loans.[13] Figure OA-4 in the appendix shows the examples for different values of δ, T.

5 Choice and Heterogeneous Treatment Effects

The results from Section 4 show that commitment *works*: clients who were assigned to the mandatory structure arm experienced substantially lower financial costs on average. In spite of this, given the opportunity, only 11% of borrowers chose this structure freely. If the effect of structure were homogeneous, this would be enough to conclude that the 89% who did not choose it would have been financially better-off if they had. In a world of heterogeneous treatment effects, however, low demand for commitment could still be consistent with borrowers adhering to a standard model of rational choice. The borrowers who did not choose structure could simply be those who don't need it. Indeed, we find strong evidence that the effect of mandatory structure varies substantially across individuals in our experiment: we test and reject the null hypothesis of homogeneous treatment effects using the methodology of Chernozhukov *et al.* (2018) (details available upon request). So the question remains: can we explain the 89% of borrowers not choosing commitment as being those that don't benefit financially from it? In this section, importantly, we show that structure would lower average financial costs even for the subset of borrowers who choose *not* to commit voluntarily. To simplify the exposition in this and all sections that follow, we re-define all outcome variables so that *beneficial* treatment effects are *positive*. Using this convention, a positive treatment effect of structure on financial cost, for example, reflects the average cost *savings* caused by structure.

5.1 Bounding the Distribution of Individual Treatment Effects

Calculating the number of individuals who benefit from mandatory structure requires the distribution of individual treatment effects. While this distribution cannot be point identified, it can be bounded. Let (Y_{i0}, Y_{i1}) be i's potential outcomes under the control and mandatory structure conditions and define $\Delta_i \equiv Y_{i0} - Y_{i1}$. Let (F_0, F_1, F_Δ) be the marginal distributions of (Y_{i0}, Y_{i1}, Δ) and define

$$\underline{F}(\delta) \equiv \max\left\{0, \sup_y F_1(y) - F_0(y-\delta)\right\}, \quad \overline{F}(\delta) \equiv 1 + \min\left\{0, \inf_y F_1(y) - F_0(y-\delta)\right\}.$$

[13] An alternative calculation is as follows: $\text{Profit}_{j,X} = \{\text{Financial Cost}_j | \text{Contract} = X\}\{1 + n \times \Pr(\text{Repeat}_j)\}$ where n is are exogenously arising need to borrower, which are lead to repeat pawning. In this case, the status quo contract dominates (profit-wise) the structured payment contract for all n.

Since F_0 and F_1 are point identified under random assignment, so are $\underline{F}$ and $\overline{F}$, and the pointwise sharp bounds for F_Δ are $\underline{F}(\delta) \leq F_\Delta(\delta) \leq \overline{F}(\delta)$ (Fan & Park, 2010). Figure OA-9 plots these bounds for the APR outcome in our experiment. To bound the share of borrowers who benefit from mandatory structure, we merely substitute $\delta = 0$ into the preceding, since $\mathbb{P}(\Delta_i > 0) = 1 - F_\Delta(0)$. Our point estimates of $\underline{F}(0)$ and $\overline{F}(0)$ are 0.03 and 0.77 respectively, with 95% confidence intervals of [0.025, 0.050] and [0.75, 0.80]. Since we are interested in $\mathbb{P}(\Delta_i > 0) = 1 - F_\Delta(0)$, it follows that *at least 23%*, and at most 97%, of individuals borrowers benefit from mandatory structure.[14] That is, even under the most conservative possible assumptions, we find evidence of depressed demand: more than 23% would benefit from structured repayment but only 11% demand it. These bounds are wide because they concern the distribution of *individual treatment effects* and make no assumptions beyond our experimental randomization. By adding assumptions it is possible to say more. We turn to this next.

5.2 Potential Outcomes and Exclusion

This section shows that, under an additional mild exclusion restriction assumption, our three-armed experimental design point-identifies a range of additional treatment effects of interest. Foremost among them is the effect of treatment on the *untreated*–those who would not demand structured repayment contracts–which is needed to assess the benefits mandating a contractual feature. We begin by providing a full definition of the potential outcomes in our empirical setting. Let $Z_i \in \{0, 1, 2\}$ denote the treatment arm to which to participant i was assigned: $Z_i = 0$ denotes the mandatory no-structure arm, $Z_i = 1$ denotes the mandatory structured arm, and $Z_i = 2$ denotes the choice arm. Now let D_i be the treatment that participant i actually *received*, where $D_i = 0$ denotes no-structure and $D_i = 1$ denotes structure. We assume perfect compliance in the $Z_i = 0$ and $Z_i = 1$ arms. It is only in the $Z_i = 2$ arm that participants are free to choose between alternative contracts. Let $C_i \in \{0, 1\}$ denote a participant's "choice type." If $C_i = 1$ then participant i *would choose structure*, given the option; if $C_i = 0$ she would not. As shorthand, we call borrowers with $C_i = 1$ "choosers" and those with $C_i = 0$ "non-choosers." Whereas a participant's choice type C_i is only observed if she is allocated to the choice arm ($Z_i = 2$), her treatment D_i and experimental arm Z_i are always observed. Given the design of our experiment, these quantities are related by

$$D_i = \mathbb{1}(Z_i \neq 2)Z_i + \mathbb{1}(Z_i = 2)C_i. \tag{2}$$

We maintain the stable unit treatment value assumption (SUTVA) throughout. This

[14] Confidence intervals are constructed using the asymptotic distribution of $(\underline{F}, \overline{F})$. See Fan & Park (2010).

means that borrower i's outcomes depend only on her *own* values of Z_i and D_i, not those of any other person in the experiment. Under this assumption, a fully general model for the potential outcomes in our experiment would take the form $Y_i(d, z)$ for $d \in \{0, 1\}$ and $z \in \{0, 1, 2\}$, allowing participant i's potential outcome to depend *both* on the treatment she actually receives, D_i, and the experimental arm to which she is assigned, Z_i. This model is too general, however, to point identify meaningful causal effects. For this reason, we consider two mild exclusion restrictions.

To state these restrictions we first define some additional notation. Because our experimental design implies that any borrower with with $Z_i = 0$ has $D_i = 0$, we abbreviate the potential outcome $Y_i(d = 0, z = 0)$ as Y_{i0}. Similarly, since any borrower with $Z_i = 1$ has $D_i = 1$, we abbreviate $Y_i(d = 1, z = 1)$ as Y_{i1}. Using this notation, we can write the pair of exclusion restrictions which we refer to as "No Agency Effects" (NAE). In words, these restrictions say that having agency–being given a choice–has no *direct* effect on outcomes; it only affects them *through the chosen treatment.* Mathematically, the first NAE restriction is

$$Y_i(d = 0, z = 2) = Y_i(d = 0, z = 0) \equiv Y_{i0}. \tag{3}$$

Equation 3 only restricts the potential outcomes of non-choosers, individuals with $C_i = 0$, because they are the only borrowers for whom $D_i = 0$ when $Z_i = 2$. In words, this condition assumes that every non-chooser experiences the same potential outcome regardless of whether she is assigned to the choice arm or the control arm. Similarly, we assume a second NAE exclusion restriction given by

$$Y_i(d = 1, z = 2) = Y_i(d = 1, z = 1) \equiv Y_{i1}. \tag{4}$$

Equation 4 only restricts the potential outcomes of choosers, individuals with $C_i = 1$, because they are the only borrowers for whom $D_i = 1$ when $Z_i = 2$. In words, this condition assumes that every chooser experiences the same potential outcome regardless of whether she is assigned to the treatment arm or the choice arm.

Mathematically (3) and (4) have the same structure as the standard LATE exclusion restriction that $Y_i(d, z)$ depends only on d, not on z. Substantively, however, they are slightly different, given that there is no explicit reference to the "chosen" versus "mandatory" treatment distinction in the usual LATE setup. In essence, (3) and (4) assume that being assigned a particular treatment has the same result as choosing it for yourself, provided that you are assigned the same treatment that you *would have chosen.* If the mere fact of having been given a choice has a direct effect on outcomes, one or both of our exclusion restrictions will be violated. One can imagine situations in which this might be the case. For example, even someone who would have voluntarily chosen to undergo drug rehabilitation, given the choice, might respond differently when coerced. In our empirical setting, however, both (3) and (4) are plausible.[15] Moreover, each has testable implications that we fail to reject: see Appendix E. Under (3) and (4), the observed outcome Y_i is related to (Y_{i0}, Y_{i1}) by

$$Y_i = \mathbb{1}(Z_i = 0)Y_{i0} + \mathbb{1}(Z_i = 1)Y_{i1} + \mathbb{1}(Z_i = 2)\left[(1 - C_i)Y_{i0} + C_iY_{i1}\right]. \tag{5}$$

Equation 5 is the key to understanding the results that follows. Random assignment of $Z_i = 0$ and $Z_i = 1$ identifies the marginal distributions of Y_{i0} and Y_{i1} for the population as a whole. Random assignment of $Z_i = 2$ likewise point identifies the share of choosers ($C_i = 1$), the distribution of Y_{i1} for choosers, and the distribution of Y_{i0} for non-choosers ($C_i = 0$). Because Z_i is assigned independently of pre-treatment covariates X_i, we also identify the *conditional* distributions of Y_{i0} and Y_{i1} given X_i.

5.3 The "Mandates vs. Choice" Design

As a direct consequence of (5), our experimental design—henceforth the "Universal vs. Choice design"—point identifies a number of interesting and economically-relevant causal quantities. First it identifies the treatment on the treated (TOT) and untreated (TUT) effects:

$$\text{TOT} \equiv \mathbb{E}(Y_{i1} - Y_{i0}|C_i = 1), \quad \text{TUT} \equiv \mathbb{E}(Y_{i1} - Y_{i0}|C_i = 0).$$

The TOT is the causal effect of structure on borrowers who would voluntarily choose it, while the TUT is the causal effect on borrowers who would not. If the TUT is positive, then borrowers who did not choose structure would have experienced better outcomes, *on average*, if they had. Selection on gains implies TOT>TUT>ATE, as in a canonical Roy model. Because our design identifies all three quantities, it allows us to test this implication directly and to calculate the average selection on gains (ASG), namely the difference between the TOT and TUT effects:

$$\text{ASG} \equiv \mathbb{E}[Y_{i1} - Y_{i0}|C_i = 1] - \mathbb{E}[Y_{i1} - Y_{i0}|C_i = 0] = \text{TOT} - \text{TUT}.$$

A number of recent papers compare estimates of the TOT and TUT to better understand who selects into treatment and why, e.g. Cornelissen *et al.* (2018) and Walters (2018). This line of work relies, at least to some extent, upon structural modeling assumptions to extrapolate from the reduced-form quantities that are identified by the data

[15] Closely related assumptions are common, if often tacit. Chamberlain (2011) explicitly assumes that choosing and being assigned a treatment have the same effect. A significant literature using compulsory schooling laws to estimate the returns to education tacitly assumes that schooling *in general* has the same returns, regardless of whether it is mandated or chosen freely. Similarly, fertilizer yields are tacitly assumed to generate the same returns regardless of whether the fertilizer was chosen by farmers or provided by the government, as long as it is the same fertilizer.

Figure 2: Graphical Intuition for the Mandates vs. Choice Design.

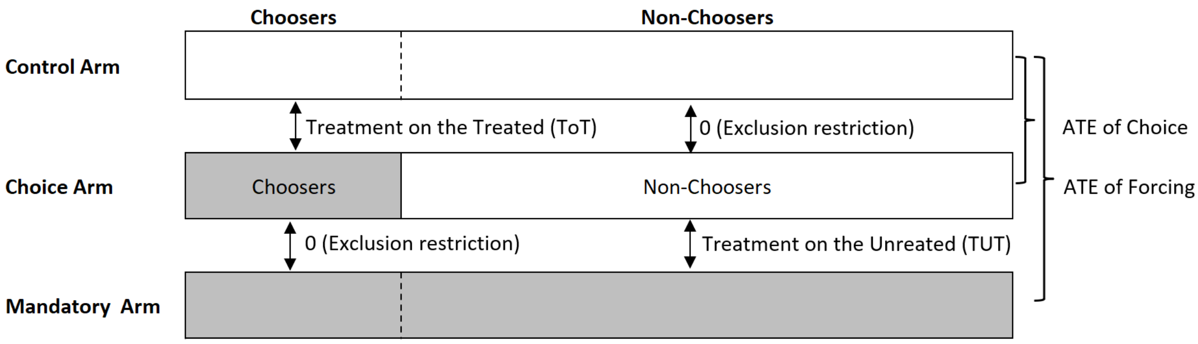

Gray rectangles denote borrowers with a structured contract; white rectangles denote borrowers with a status quo contract. A comparison of means across control and Structured arms identifies the ATE of Structure. The difference of mean outcomes across the choice and control "nets out" the non-choosers, and hence equals the TOT multiplied by the share of choosers. Similarly, the difference of means across the structured and choice arms "nets out" the choosers and equals the TUT multiplied by the share of non-choosers. The share of choosers, illustrated using dashed vertical lines, is equal on average across arms under random assignment.

alone to more interesting, and economically relevant, causal parameters.[16] An alternative approach aims to avoid structural assumptions by calculating conditional local average treatment effects (LATE) given observed covariates X and re-weighting them according to the distribution of covariates in some population of interest to yield, for example, an average treatment effect (Aronow & Carnegie, 2013; Angrist & Fernandez-Val, 2013). But there is no free lunch: this "LATE-and-reweight" approach relies upon assumptions of its own, most crucially the assumption that there is *no selection-on-gains* conditional on X, i.e. that the conditional LATE equals the conditional ATE. In contrast to both approaches, the Mandates vs. Choice design uses exogenous experimental variation to point identify the ATE, TOT, and TUT without ruling out unobserved selection-on-gains or relying on additional structural modeling assumptions.

The key insight can be read directly from (2) and (5); Figure 2 provides graphical intuition. Viewing Z_i as an instrumental variable, the "Mandates vs. Choice" design can be interpreted as a *pair* of RCTs, each subject to one-sided non-compliance. The first of these compares $Z_i = 0$ to $Z_i = 2$. For each individual with $Z_i = 0$ we have $D_i = 0$ and observe Y_{i0}. For those with $Z_i = 2$ we have $D_i = C_i$ and observe $(1 - C_i)Y_{i0} + C_iY_{i1}$. This is identical to a "randomized encouragement" design in which treatment is only available to those who are encouraged: $Z_i = 2$. Under this interpretation, those with $C_i = 1$ are

[16]While the marginal treatment effects (MTE) approach (Heckman & Vytlacil, 2007) can in principle be used to identify the TOT and TUT without parametric restrictions, doing so requires an instrumental variable Z with sufficiently rich support that the probability of treatment take-up given Z varies continuously between zero and one. In practice, instrumental variables are usually discrete and, even when continuous, typically have a more modest effect on take-up.

"the compliers" and it follows by a standard argument (see Section 8) that

$$\frac{\mathbb{E}(Y_i|Z_i=2)-\mathbb{E}(Y_i|Z_i=0)}{\mathbb{E}(D_i|Z_i=2)-\mathbb{E}(D_i|Z_i=0)}=\frac{\mathbb{E}(Y_i|Z_i=2)-\mathbb{E}(Y_i|Z_i=0)}{\mathbb{E}(D_i|Z_i=2)}=\text{TOT} \tag{6}$$

since $\mathbb{E}(D_i|Z_i=0)=0$ by (2). A closely related argument can be used to construct a Wald estimand that identifies the TUT. Here we consider $Z_i=1$ to be the "encouragement" and compare the outcomes for these individuals to those with $Z_i=2$. If $Z_i=1$ then $D_i=1$ and we observe Y_{i1}. If instead $Z_i=2$ then $D_i=C_i$ and we observe $(1-C_i)Y_{i0}+C_iY_{i1}$. Again, we can view this as an experiment with one-sided non-compliance, but now the situation is reversed. Everyone with $Z_i=1$ is treated, but some people with $Z_i=2$ are "always-takers" who obtain the treatment ($D_i=1$) despite having been allocated to the "control" arm $Z_i=2$. Under this interpretation, the "compliers" are those with $C_i=0$: when $Z_i=1$ they take the treatment, and when $Z_i=2$, they do not. Thus,

$$\frac{\mathbb{E}(Y_i|Z_i=1)-\mathbb{E}(Y_i|Z_i=2)}{\mathbb{E}(D_i|Z_i=1)-\mathbb{E}(D_i|Z_i=2)}=\frac{\mathbb{E}(Y_i|Z_i=1)-\mathbb{E}(Y_i|Z_i=2)}{1-\mathbb{E}(D_i|Z_i=2)}=\text{TUT} \tag{7}$$

since $\mathbb{E}(D_i|Z_i=1)=1$ by (2). Because they identify both the TOT and TUT, (6) and (7) also identify the average selection on gains: ASG = TOT − TUT.

The Mandates vs. Choice design also identifies the average selection bias (ASB) and average selection on levels (ASL). In particular,

$$\text{ASB}\equiv\mathbb{E}(Y_{i0}|C_i=1)-\mathbb{E}(Y_{i0}|C_i=0)=\frac{\mathbb{E}(Y|Z=0)-\mathbb{E}(Y|Z=2,D=0)}{\mathbb{E}(D|Z=2)} \tag{8}$$

$$\text{ASL}\equiv\mathbb{E}(Y_{i1}|C_i=1)-\mathbb{E}(Y_{i1}|C_i=0)=\frac{\mathbb{E}(Y|Z=2,D=1)-\mathbb{E}(Y|Z=1)}{1-\mathbb{E}(D|Z=2)} \tag{9}$$

as shown in Section 8. The ASB tells us whether borrowers who voluntarily choose structure are those who are worse off, on average, under the status quo. Similarly, the ASL tells us whether borrowers who voluntarily choose structure are those who are better off, on average, under the structured contract.

Equations (6)–(9) are useful for understanding why our experimental design identifies the TOT, TUT, ASG, ASB, and ASL but they are less convenient for estimation and inference. Appendix F explains how to compute each of these quantities from a small number of just-identified, linear instrumental variables regressions, along with appropriate cluster-robust standard errors. These estimators and standard errors are implemented in our companion STATA package.

Table 4 calculates the causal quantities described above—TOT, TUT, ASG, ASB, and ASL—for our experimental data, along with robust standard errors for each. For purposes of comparison, the table also presents the ATE results from Section 4 above (row 1),[17] along with the corresponding average potential outcomes $\mathbb{E}[Y_0]$ and $\mathbb{E}[Y_1]$ (rows 4–5). The columns of the table correspond to different outcome variables defined above. We are now in a position to report one of the important results in the paper. For all four outcome definitions, the TUT effect is positive, statistically and economically significant, and comparable in magnitude to the ATE. In other words: the structured payments contract is on average *beneficial* to the people who *would not choose it*, and these benefits are large. That is, mandating structured payment contracts would decrease the financial cost. Borrowers are leaving money in the table by not choosing it.

Some comments are in order. First, this result holds on average. Some non-choosers could lose from a mandate. Section 6.2 assesses personalized targeting strategies. Second, financial cost is not welfare, and further work with more comprehensive measures (stress, consumption, performance in other loans) needs to be done. We provide proof of concept that at least financial cost may not be minimized by free choice and provide a method to rigorously compare take-up mandates against consumer choice. Third, given that providing such contracts may not be in the lender's interest, borrowers may not even be presented with this choice.

Table 4 also estimates the TOT effect. Due to the low take-up rate in the choice arm, TOT effects are imprecisely estimated. Only one of them, %(1-Default) from column (3), is statistically significant. This carries over into our estimates of the average selection on gains. Our point estimates are *negative* for all but the (1 - Default) outcome, but none is statistically distinguishable from zero. For the (1 - Default) outcome, we have sufficient precision to conclude that the average selection bias (ASB) is large and negative. This means that borrowers who choose structure would have faced a higher probability of default under the status quo contract than borrowers who do not choose structure.

Overall, Table 4 suggests that structure works but that *not enough* people choose it: the structured contract saves money on average even to those who would *not* choose it voluntarily. We believe this result is new in the household finance literature. It also illustrates how the "Mandates vs. Choice" design can be used more generally to understand the relationship between choice and treatment effect heterogeneity, in particular the causal effect among non-choosers.[18]

[17] Coefficients are not exactly the same since Table 4 includes branch and day-of-week FE.

[18] For instance, in the debate on financial commitment take-up, some papers argue that this effect is low (Ashraf *et al.* (2006), Giné *et al.* (2010), Bai *et al.* (2020), Royer *et al.* (2015) while others argue it is high (Kaur *et al.* (2015), Casaburi & Macchiavello (2019), Schilbach (2019), Tarozzi *et al.* (2009), Dupas & Robinson (2013)), but none estimates the benefits for non-takers. In a different domain, doctors claim that medical treatment abandonment is too high in a broad range of diseases (McDonald *et al.*, 2002), without knowing the treatment benefits for those that abandon.

Table 4: Five Treatment Effects Estimates: TOT, TUT, ASG, ASB, ASL

	APR % benefit	FC benefit	% (1-Default)	% (1-Refinance)
	(1)	(2)	(3)	(4)
ATE	9.41***	183.0***	7.74***	6.34**
	(2.06)	(50.8)	(2.50)	(2.90)
ToT	-0.59	111.9	37.4*	-25.9
	(21.4)	(528.3)	(21.6)	(29.1)
TuT	10.6***	191.5***	4.20*	10.2***
	(2.47)	(50.8)	(2.41)	(2.90)
$\mathbb{E}[Y_1]$	-47.4***	-759.4***	64.2***	67.2***
	(1.42)	(27.3)	(1.69)	(1.70)
$\mathbb{E}[Y_0]$	-56.8***	-942.4***	56.4***	60.9***
	(1.49)	(42.9)	(1.84)	(2.35)
ASG	-11.2	-79.6	33.2	-36.1
	(22.9)	(556.2)	(22.6)	(30.6)
ASB	15.8	291.5	-39.1*	22.7
	(22.3)	(551.2)	(22.3)	(30.1)
ASL	4.58	211.9***	-5.90	-13.4***
	(3.55)	(59.5)	(4.29)	(4.20)
Observations	6304	6304	6304	6304
H_0 : ATE-TuT=0	0.62	0.89	0.14	0.23
H_0 : ATE-ToT=0	0.63	0.89	0.14	0.24
H_0 : ASG=0	0.63	0.89	0.14	0.24
H_0 : ASG$\geq$ 0	0.69	0.56	0.071	0.88

This table presents results computed using the derivations from Section 5.3. The APR and financial cost outcomes have been multiplied by -1 so that a positive causal effect *benefits* the borrower in each of the four columns. The bottom panel presents p-values for a number of null hypothesis tests of treatment effect heterogeneity.

6 The Case for Mandating Structure

6.1 Why does mandating frequent payments work in this context?

A behavioral literature has highlighted voluntary commitment as an attractive way of allowing the "right" people to self-select. The argument for compulsory treatment, centers on the surprising result that TUT > 0. We now investigate four potential explanations for the positive TUT: the need to learn, time discounting, present bias, and overconfidence.

Our experiment introduced a new contract into an environment that had not previously featured structured repayment; perhaps clients required experience to understand its benefits. Are clients who experienced the commitment contract more likely to choose structure subsequently than those assigned to the status-quo contract? Appendix Table OA-6 shows *no difference* in choosing structure for those assigned to mandatory structure versus those assigned to status quo control that came back to pawn again (22%

did). While these exercises cannot completely exclude the possibility that learning would play a role with more exposure, they provide no indication that the lack of voluntary compliance is simply a matter of inexperience with structure in the short run. Note also that there will be no learning if lenders don't provide a structured payment contract in the first place.

Can high borrower time discount rates explain low demand for structured payments? After all, monthly payments are front-loaded while pawn recovery is back-loaded, even if by only a few days. Figure OA-6 calculates TUT effects for different borrower discount rates given the actual timing of repayment and pawn recovery. The solid line gives the TUT effect adjusted for a specified annual discount rate, while the shaded regions give the associated 95% & 90% confidence interval. We see that non-choosers continue to experience significant decreases in the net present value of financial costs up to annual discount rates of 1,000%, and the NPV remains positive, although not significant only at 5,000% discount rates. These are huge discount rates.[19] As such, discounting is unlikely to explain why those who benefit, on average, from structure fail to choose it.

If the benefits of structure among non-choosers cannot be explained by standard models of rational choice, the canonical behavioral story would center on time inconsistency. While commitment is useful to anyone with hyperbolic time preferences, only those who are sophisticated—i.e. aware that they are hyperbolic discounters—will demand it. A large share of "naïve" hyperbolics in the population–borrowers who are unaware that they are hyperbolic discounters–could therefore drive a large and positive TUT. Our baseline survey included standard questions about discount rates between tomorrow and a month in the future versus discount rates between three and four months in the future. This allows us to classify borrowers who display more impatience over immediate delays as present biased. This measure of financial hyperbolicity is widely used in survey research, although it is not without problems.[20]

If we could perfectly measure present bias and sophistication, we could divide the sample into three groups: sophisticated hyperbolics (who chose structure), time-consistent non-choosers (for whom structure will not be effective), and naïve hyperbolic non-choosers (who will benefit from structure even if not chosen by them). If present bias fully explains the low take-up rate of voluntary commitment, we should find that the TUT for present-biased borrowers is positive. This is because among the group of non-takers, a

[19] In a similar exercise for subprime borrowers in the US, Adams *et al.* (2009) argue that even discount rates smaller than these would be too large to be reasonable and can be safely discarded as explanations.

[20] We also asked a simple self-reported question about temptation, and obtain similar answers with this alternate measure. Recent empirical work has shown the superiority of more elaborate measures such as "convex time budgets" (Andreoni *et al.*, 2015) while questioning the interpretation of measures of hyperbolicity that are not based on consumption (Andreoni & Sprenger, 2012; Cohen *et al.*, 2020), suggesting that real effort tasks provide a better measure (Augenblick *et al.*, 2015). Given that we had only a few minutes to interview real pawnshop clients prior to a commercial transaction, our simple measure was a necessary compromise.

comparison of present-biased borrowers against everyone else is a comparison of naïve hyperbolics against time-consistent non-choosers. The left panel of Figure OA-7 in the Appendix carries out a feasible version of this exercise using our survey measure of present bias. We find no indication that present-bias explains our positive estimated TUT.[21]

Finally, motivated by the high subjective probabilities of recovery we found in the baseline survey, and by the fact that a large fraction of borrowers pay towards recovery and then lose the pawn and these payments, we explore overconfidence in recovery as a potential explanation of low demand for structured payments. After all, recall that while 72% of survey respondents believe they have a 100% chance of recovering their pawn, in reality only 43% will go on to do so. Borrowers who are certain to recover their pawn perceive no benefit of committing to repay under the threat of fees and, therefore, no benefit of choosing the structured payment contract. To explore this, as in the preceding paragraph, we compare the overall TUT estimate to estimates computed for two sub-groups, defined by a binary variable that we call "sure-confidence". This measure equals one for any individual who says at the time of borrowing that they have a 100% probability of recovering their pawn, and zero otherwise. In contrast to our results for present bias, we find that sure-confidence *is* a strong predictor of TUT. Figure OA-7 shows that the TUT is almost entirely confined to the sure confident individuals, while those that say they have some chance of defaulting have TUT close to zero. It seems that the effectiveness of paternalism in our experiment may be driven by *overconfident* borrowers who, heedless of the risk of default, fail to choose structure despite benefiting substantially when it is mandated.[22]

6.2 Targeting Paternalism using Causal Forests

On average, the structured contract benefits both those who would choose it and those who would not. In Section 5.1, we briefly went *beyond* average effects by presenting bounds on the distribution of individual treatment effects. Because they made no assumptions beyond random assignment, these bounds were relatively wide. We then applied NAE assumptions to estimate average treatment effects for choosers and non-choosers.

In this section, we again impose NAE exclusions, but this time conditional on observed covariates, not just on average. This enables us to use causal forests to estimate *conditional* average treatment effects (CATE) as well as CTOT and CTUT effects. This

[21]Similarly, we find that experience pawning does not predict TUT.

[22]In Figure OA-8 we plot the coefficient estimates from a regression that predicts sure confidence with a battery of individual-level characteristics. Older males are more likely to be sure-confident, as are those with more education. Taken at face value, the sure-confident also report less financial stress, less trouble paying bills, and to be more frequently relied upon financially by family members. Viewed through a behavioral lens, however, it is also possible that the type of person who is over-confident in their ability to repay a loan also exaggerate their degree of economic security in their response to survey questions. In any case, it appears that sure confidence may be difficult to predict with easily-observed and objective demographic criteria, a point to which we return below.

stronger assumption allows us to say more about personalized treatment effects and, therefore, to assess different ways to target structured payments. Among other things, it will potentially allow us to identify groups of borrowers who are on average *harmed* by structure. To estimate the conditional average treatment effects we use the "generalized random forest" methods of Athey *et al.* (2019) (see Appendix D for details).[23] We implement these estimates using the covariates from our survey.

Figure 3: Heterogeneous Treatment Effects.

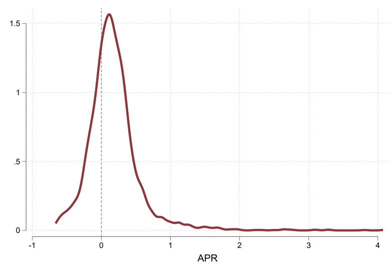

(a) ToT

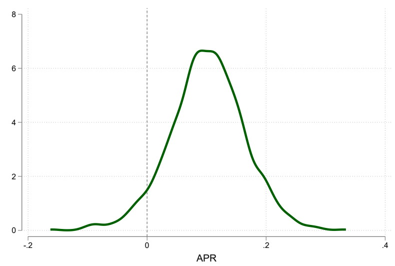

(b) TuT

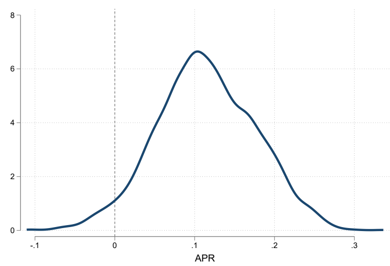

(c) ATE

To estimate conditional average treatment effects given administrative and survey data and the function `causal_forest()` of the `grf` R package; to estimate conditional TOT and TUT effects we use the `instrumental_forest()` function from the same package. For more information see Appendix D.

Figure 3 plots densities of the estimated CATE, CTOT, and CTUT effects from the generalized random forest models described above. In each case, the outcome variable is APR benefit, i.e. the reduction in APR from a structured payments contract. As we see from the figure, the conditional average effects are highly heterogeneous but overwhelmingly positive. The TUT density is particularly interesting for the question of paternalism since, as emphasized above, it presents conditional average effects for borrowers who would not voluntarily choose structure. Only 7% of our estimated conditional TUTs are

[23]The generalized random forest approach of Athey *et al.* (2019) produces conditional average effect estimators that are asymptotically normal. Our inferences in this section are carried out by simulating from the normal limit distributions using the estimated standard errors. As covariates, we use 22 questions in our survey (what we call the *wide* set of covariates in Appendix D. The survey questionnaire (with the highlighted questions being the ones for the wide set of covariates) can be found in OA-1 in the Appendix.

negative (95% confidence interval of [4%, 9%]).[24] Figure 3 strengthens the argument that many borrowers will save on financial cost by choosing the structured payment contract.

We now use the CTUT estimates to quantify how much extra financial cost would borrowers incur under different scenarios. To simplify language, we define a "mistake" for a non-chooser to be a conditional TUT estimate that significantly exceeds some large and positive APR threshold, e.g., 10 percentage points. At any such threshold, we can calculate the percentage of non-choosers in the choice arm who have benefited by more than that threshold from having chosen structure. Again, we do not measure welfare, only financial cost.

The results of this exercise can be read off from the green curve in Figure 4. Defining $F_{\text{TUT}}(\delta)$ to be the CDF corresponding to the density of CTUT estimates from Figure 3, the green curve in Figure 4 is merely $[1 - F_{\text{TUT}}(\delta)] \times 100\%$. In other words, the vertical axis for the green curve plots the percentage of non-choosers who made a "mistake" of a certain threshold from the horizontal axis by not choosing. Note that we use a *positive* APR threshold to denote a mistake for both choosers (red) and non-choosers (green). This ensures that bigger mistakes are always to the *right* of smaller mistakes. The shaded region gives associated 95% pointwise confidence bounds.

Figure 4: "Mistakes" in the choice arm.

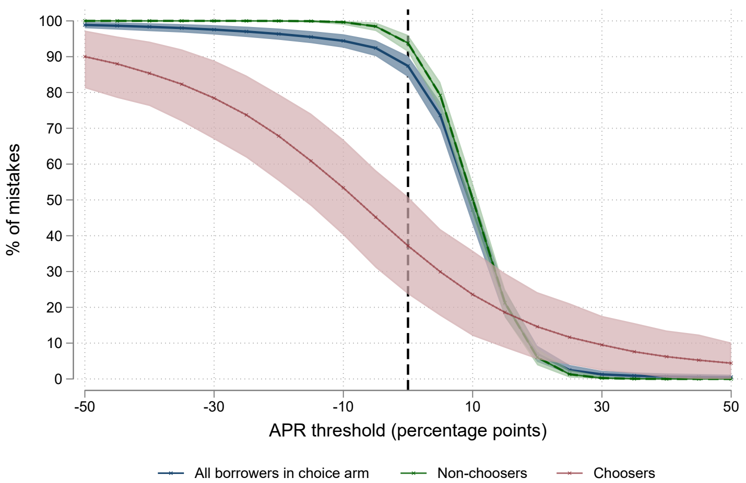

This figure presents conditional average TUT and TOT effects for the APR outcome from Figure 3 in an alternative manner, to consider the fraction of borrowers in the choice arm who made "mistakes" in their decision to accept or refuse the structured contract. A "mistake" for a non-chooser is defined as a positive conditional TUT effect that significantly exceeds a specified threshold APR value.

[24]To be clear, this is a probability statement about conditional average effects over the distribution of covariates. In particular, we estimate that $\int \mathbb{1}\{\mathbb{E}[Y_1 - Y_0|X = \mathbf{x}, C = 0] < 0\}\, f(\mathbf{x}|C = 0)\, \mathrm{d}\mathbf{x}$ is approximately 0.07. The share of non-choosers with negative conditional *average* treatment effects need not equal the share with a negative *individual* effects, i.e. $\mathbb{P}(Y_1 < Y_0|C = 0)$. But the more treatment effect heterogeneity that X_i explains, the closer these two values become.

For choosers we follow an analogous approach, defining a "mistake" as a *negative* CTOT exceeding a particular APR threshold. The results for choosers can be read from the red curve in Figure 4. If $F_{\text{TOT}}(\delta)$ denotes the CDF corresponding to the density of conditional TOT estimates from Figure 3, then the red curve is merely $F_{\text{TOT}}(-\delta) \times 100\%$. The red shaded region gives associated 95% pointwise confidence bounds. The blue curve in Figure 4 gives the *overall* percentage of borrowers in the choice arm who made a "mistake" at a particular APR threshold. This total is computed by taking a weighted average of the green (non-choosers) and red (choosers) curves, with weights equal to their shares in the choice arm.

The results in Figure 4 suggest that a large fraction of non-choosers would have experienced lower financial costs had they chosen the structured payment contract, even at an APR threshold as large as 10 percentage points, we estimate that more than half of them should have chosen the structured payment contract to lower financial costs. In contrast, relatively few choosers appear to have made mistakes by choosing it. This now allows us to make a stronger statement in favor of paternalism in our context; not only does structure generate large benefits *on average*, but it also benefits the *vast majority* of borrowers who would be coerced under a policy of mandatory structure.

6.3 Can we target paternalism accurately?

Even though the overwhelming majority of borrowers would save on financing costs, not all of them would, according to our estimates. Is it possible to target paternalism only towards those who would benefit from it?

The answer depends on how predictive the covariate information is of treatment effects. Note that to implement a targeting protocol, the covariate information used in the model would need to be either verifiable or hard to manipulate. In our survey and administrative data, this leaves us with only a few covariates: age, gender, high school education or above, desired loan size, and whether that individual has ever pawned before. We call these the "narrow" covariate set, to contrast with the full set of survey variables, which we call the "wide" covariate set. We take the estimated $\widehat{\text{CATE}}(X_i)$ using the wide covariates as our ground truth.[25]

To investigate targeting we ask how well we can predict this ground truth using the narrow covariate set, which is the set of less-manipulable variables that the pawnlender has. We use both a logistic regression and a random forest classification model using the narrow covariate set to predict positive benefits at the individual level. We compare the in-sample performance of targeting rules against a policy of universal mandated structure.

[25] We assign a dummy variable equal to one for the 90% of borrowers who have a positive CATE using the wide covariate set. Figure OA-11 shows the relationship between these two different CATEs. We generate substantially less heterogeneity when using the narrow covariate set, although we still reject the null hypothesis of no treatment effect heterogeneity with the test from Chernozhukov *et al.* (2018).

Table 5: Type I & II errors using targeting narrow rules

Rule	% incorrectly assigned to control	% incorrectly assigned to treatment	Overall Error Rate
All to Status-quo	90.22	0	90.22
All to Structure	0	9.78	9.78
Optimal	0	0	0
Narrow rule (RF)	4.38	5.21	9.59
Narrow rule (Logit)	6.9	7.76	14.66
Allow choice	93.81	37.18	87.4

This table reports error rates for six different rules for allocating individuals to structured repayment. Row 1 assigns all borrowers to control, Row 2 all to structured payments. Row 3 uses 'optimal targeting according to the CATE from the wide covariate set. Row 4 uses a random forest classification, and Row 5 a logit model, both with only the narrow covariate set. Row 6 uses the choices made by borrowers.

Table 5 shows error rates for six possible assignment rules: assigning all borrowers to Status quo, all to structured payments, the optimal (wide-fores), the random forest with narrow covariates, the Logistic model with narrow covariates, and the actual choice made by borrowers. All models take the wide RF as the ground truth. While the narrow RF correctly assigns to control roughly half of those who do not benefit from the structures payment contract, it also incorrectly allocates to control 4.38% that would gain from structured payments. As such, it only improves the overall correct targeting rate by about half of a percentage point relative to universal structure. The Logit assignment rule is even less successful at predicting benefits and harms, with a higher share of borrowers incorrectly assigned to both treatment and control, meaning that the overall correct targeting rate for the Logit is 5 pp *lower* than universal structure to structured payments. Self-targeting through choice proves to be little better than assigning everyone to the Status quo contract, given the low take-up rate and the presence of both Type I and Type II errors in the choice arm.

The takeaway is that given (a) low take-up, (b) the large fraction of the sample benefiting from structure, (c) and the weak predictive power of the narrow covariates, universal paternalism assigning all borrowers to the strutured contract appears to be an attractive targeting method in our context.

7 Conclusion

This paper makes several contributions. First, it analyzes the important but understudied industry of pawn loans. Likely, hundreds of millions of people use pawn loans, and yet economists have not studied them. Overcolateralization and sunk payments make pawn loans a product different than microfinance and payday loans in the incentives they create. In fact, while recent literature shows that more flexibility in payments may be beneficial in microfinance, this does not seem to be the case in our pawn lending context. We show that a simple change to contract terms results in substantial financial savings for pawn borrowers: mandatory structure lowers the APR from 57% to 46%, and reduces the fraction of borrowers who default by 6.6 pp, or 15%. That this new contract generated large benefits for borrowers and yet was not offered, and that a contract that generated default was the industry standard instead, is related to the idea of "veiled paternalism" Laibson (2018), put on its head, embedding features that lead to high borrower costs.

Second, a novel "Mandates-Choice" design paired with NAE assumptions allows us to go beyond ATE results and draw an important set of conclusions about the relationship between take-up and heterogeneous treatment effects, providing a complement to marginal treatment effects methods. Estimating treatment effects for non-takers is critical in thinking about paternalism, and we show one rigorous way to do it that is new in the literature and distinct from Karlan-Zinman designs.

Third, in terms of our empirical results, we find substantial benefits of treatment for *non-choosers* and no evidence of selection on gains by borrowers who choose structure. This is the first such result in the household finance literature and a first step in thinking about paternalistically imposing frequent payment contracts. Our results show that over-optimism of pawn recovery is the characteristic most strongly associated with benefiting from the commitment despite not having chosen it, and it maybe an explanation for the low demand for structured repayment. Given that the rate of voluntary commitment in our sample is only 11% and given the limited benefits from personalized targeting with our data, in order to achieve widespread benefits in this context compulsory structure appears to be an attractive policy. With a now well-established toolkit of regular small payments and incentives delivering small default rates in microfinance, regulators may fruitfully investigate the possibility of requiring pawnbrokers to embed features of commitment and regularity into their repayment structures in more consistent ways.

There are several questions left for future research. Would results hold using more comprehensive measures of welfare? How fast would borrowers learn about the benefits of structured payments if regulators forced lenders to provide this contractual feature? How would lender competition affect the incentives to provide such a contract? Much needs to be done to improve the lives of the millions of borrowers using panw loans.

References

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

Allcott, Hunt, Kim, Joshua, Taubinsky, Dmitry, & Zinman, Jonathan. 2021. Are High-

Interest Loans Predatory? Theory and Evidence from Payday Lending. *The Review of Economic Studies*, **89**(3), 1041–1084.

Andreoni, James, & Sprenger, Charles. 2012. Estimating time preferences from convex budgets. *American Economic Review*, **102**(7), 3333–3356.

Andreoni, James, Kuhn, Michael A, & Sprenger, Charles. 2015. Measuring time preferences: A comparison of experimental methods. *Journal of Economic Behavior & Organization*, **116**, 451–464.

Angrist, Joshua D., & Fernandez-Val, Ivan. 2013. ExtrapoLATE-ing: External Validity and Overidentification in the LATE Framework. *Page 401 of: Advances in Economics and Econometrics: Tenth World Congress*, vol. 3. Cambridge University Press.

Aronow, Peter M, & Carnegie, Allison. 2013. Beyond LATE: Estimation of the average treatment effect with an instrumental variable. *Political Analysis*, **21**(4), 492–506.

Ashraf, Nava, Karlan, Dean, & Yin, Wesley. 2006. Tying Odysseus to the Mast: Evidence From a Commitment Savings Product in the Philippines*. *The Quarterly Journal of Economics*, **121**(2), 635–672.

Athey, Susan, Tibshirani, Julie, & Wager, Stefan. 2019. Generalized random forests. *Ann. Statist.*, **47**(2), 1148–1178.

Augenblick, Ned, Niederle, Muriel, & Sprenger, Charles. 2015. Working over time: Dynamic inconsistency in real effort tasks. *The Quarterly Journal of Economics*, **130**(3), 1067–1115.

Bai, Liang, Handel, Benjamin, Miguel, Edward, & Rao, Gautam. 2020. Self-Control and Demand for Preventive Health: Evidence from Hypertension in India. *Review of Economics and Statistics, Forthcoming.*

Barboni, Giorgia, & Agarwal, Parul. 2023. How do flexible microfinance contracts improve repayment rates and business outcomes? experimental evidence from india. *Experimental Evidence from India (February 14, 2023).*

Basu, Kaushik. 1984. Implicit interest rates, usury and isolation in backward agriculture. *Cambridge Journal of Economics*, **8**(2), 145–159.

Battaglia, Marianna, Gulesci, Selim, & Madestam, Andreas. 2023. Repayment Flexibility and Risk Taking: Experimental Evidence from Credit Contracts. *The Review of Economic Studies*, **91**(5), 2635–2675.

Bauer, Michal, Chytilová, Julie, & Morduch, Jonathan. 2012. Behavioral foundations of microcredit: Experimental and survey evidence from rural India. *American Economic Review*, **102**(2), 1118–39.

Beaman, Lori, Karlan, Dean, Thuysbaert, Bram, & Udry, Christopher. 2023. Selection Into Credit Markets: Evidence From Agriculture in Mali. *Econometrica*, **91**(5), 1595–1627.

Bertrand, Marianne, & Morse, Adair. 2011. Information Disclosure, Cognitive Biases, and Payday Borrowing. *The Journal of Finance*, **66**(6), 1865–1893.

Carter, Susan Payne, & Skiba, Paige Marta. 2012. Pawnshops, behavioral economics, and self-regulation. *Rev. Banking & Fin. L.*, **32**, 193.

Casaburi, Lorenzo, & Macchiavello, Rocco. 2019. Demand and Supply of Infrequent Payments as a Commitment Device: Evidence from Kenya. *American Economic Review*, **109**(2), 523–55.

Chamberlain, Gary. 2011. 1011 Bayesian Aspects of Treatment Choice. *In: The Oxford Handbook of Bayesian Econometrics*. Oxford University Press.

Chernozhukov, Victor, Demirer, Mert, Duflo, Esther, & Fernandez-Val, Ivan. 2018. *Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in India*. Tech. rept. National Bureau of Economic Research.

Cohen, Jonathan, Ericson, Keith Marzilli, Laibson, David, & White, John Myles. 2020. Measuring time preferences. *Journal of Economic Literature*, **58**(2), 299–347.

Cornelissen, Thomas, Dustmann, Christian, Raute, Anna, & Schönberg, Uta. 2018. Who benefits from universal child care? Estimating marginal returns to early child care attendance. *Journal of Political Economy*, **126**(6), 2356–2409.

DiTraglia, Francis J. & Garcia-Jimeno, Camilo. 2019. Identifying the Effect of a Misclassified, Binary, Endogenous Regressor. *Journal of Econometrics*, **209**(2), 376–390.

Dupas, Pascaline, & Robinson, Jonathan. 2013. Why Don't the Poor Save More? Evidence from Health Savings Experiments. *American Economic Review*, **103**(4), 1138–71.

Fan, Yanqin, & Park, Sang Soo. 2010. Sharp bounds on the distribution of treatment effects and their statistical inference. *Econometric Theory*, **26**(3), 931–951.

Field, Erica, & Pande, Rohini. 2008. Repayment Frequency and Default in Microfinance: Evidence from India. *Journal of the European Economic Association*, **6**(2/3), 501–509.

Field, Erica, Pande, Rohini, Papp, John, & Rigol, Natalia. 2013. Does the Classic Microfinance Model Discourage Entrepreneurship among the Poor? Experimental Evidence from India. *American Economic Review*, **103**(6), 2196–2226.

Fowlie, Meredith, Wolfram, Catherine, Baylis, Patrick, Spurlock, C Anna, Todd-Blick, Annika, & Cappers, Peter. 2021. Default effects and follow-on behaviour: Evidence from an electricity pricing program. *The Review of Economic Studies*, **88**(6), 2886–2934.

Giné, Xavier, Karlan, Dean, & Zinman, Jonathan. 2010. Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation. *American Economic Journal: Applied Economics*, **2**(4), 213–35.

Gregg, Samuel. 2016. How Medieval Monks Changed the Face of Banking. *American Banker*, **1**(88).

Heckman, James J, & Vytlacil, Edward J. 2007. Econometric evaluation of social programs, part II: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments. *Handbook of econometrics*, **6**, 4875–5143.

Huber, Martin, & Mellace, Giovanni. 2015. Testing instrument validity for late identification based on inequality moment constraints. *The Review of Economics and Statistics*, **97**(2), 398–411.

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

Kaur, Supreet, Kremer, Michael, & Mullainathan, Sendhil. 2015. Self-Control at Work. *Journal of Political Economy*, **123**(6), 1227–1277.

Laibson, David. 2018. Private Paternalism, the Commitment Puzzle, and Model-Free Equilibrium. *AEA Papers and Proceedings*, **108**(May), 1–21.

McDonald, Heather P., Garg, Amit X., & Haynes, R. Brian. 2002. Interventions to Enhance Patient Adherence to Medication PrescriptionsScientific Review. *JAMA*, **288**(22), 2868–2879.

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

Morduch, Jonathan. 1999. The microfinance promise. *Journal of economic literature*, **37**(4), 1569–1614.

Pedroza, Paola. 2010. *Microfinanzas en América Latina y el Caribe: El sector en Cifras*. Tech. rept. Interamerican Development Bank Report.

Royer, Heather, Stehr, Mark, & Sydnor, Justin. 2015. Incentives, Commitments, and Habit Formation in Exercise: Evidence from a Field Experiment with Workers at a Fortune-500 Company. *American Economic Journal: Applied Economics*, **7**(3), 51–84.

Sadoff, Sally, Samek, Anya, & Sprenger, Charles. 2019. Dynamic Inconsistency in Food Choice: Experimental Evidence from Two Food Deserts. *The Review of Economic Studies*, **87**(4), 1954–1988.

Schilbach, Frank. 2019. Alcohol and Self-Control: A Field Experiment in India. *American Economic Review*, **109**(4), 1290–1322.

Stegman, Michael A. 2007. Payday Lending. *Journal of Economic Perspectives*, **21**(1), 169–190.

Tarozzi, Alessandro, Mahajan, Aprajit, Yoong, Joanne, & Blackburn, Brian. 2009. Commitment Mechanisms and Compliance with Health-Protecting Behavior: Preliminary Evidence from Orissa, India. *American Economic Review*, **99**(2), 231–35.

Walters, Christopher R. 2018. The Demand for Effective Charter Schools. *Journal of Political Economy*, **126**(6), 2179–2223.

8 Proofs

This section gives a formal derivation of the identification results presented in Equations (6)–(9) of subsection 5.3. To simplify the presentation, we omit i subscripts throughout.

Assumption 1 (Randomized Choice Design and Exclusion Restriction)**.**

(i) Z *is independent of* (Y_0, Y_1, C)

(ii) $D = \mathbb{1}(Z \neq 2)Z + \mathbb{1}(Z = 2)C$

(iii) $Y = \mathbb{1}(Z = 0)Y_0 + \mathbb{1}(Z = 1)Y_1 + \mathbb{1}(Z = 2)[(1 - C)Y_0 + CY_1]$

Lemma 1. *Under Assumption 1,*

(i) $\mathbb{E}(D|Z = 2) = \mathbb{P}(C = 1)$

(ii) $\mathbb{E}(Y|Z = 0) = \mathbb{E}(Y_0)$

(iii) $\mathbb{E}(Y|Z = 1) = \mathbb{E}(Y_1)$

(iv) $\mathbb{E}(Y|D = 0, Z = 2) = \mathbb{E}(Y_0|C = 0)$

(v) $\mathbb{E}(Y|D = 1, Z = 2) = \mathbb{E}(Y_1|C = 1)$*.*

Proof. Part (i) follows because $Z = 2$ implies $D = C$ and Z is independent of C. Parts (ii) and (iii) follow similarly: given $Z = 0$ we have $Y = Y_0$, given $Z = 1$ we have $Y = Y_1$, and Z is independent of (Y_0, Y_1). For parts (iv) and (v), first note that Assumption 1 (iii) implies that Z is conditionally independent of (Y_0, Y_1) given C. Now, $Z = 2$ implies that $D = 0$ if and only if $C = 0$. Hence, $\mathbb{E}(Y|D = 0, Z = 2) = \mathbb{E}(Y_0|C = 0)$ establishing part (iv). For part (v) $Z = 2$ implies that $D = 1$ if and only if $C = 1$ from which it follows that $\mathbb{E}(Y|D = 1, Z = 2) = \mathbb{E}(Y_1|C = 1)$. □

Proposition 1. *Under Assumption 1,*

(i) $TOT \equiv \mathbb{E}(Y_1 - Y_0|C = 1) = \dfrac{\mathbb{E}(Y|Z = 2) - \mathbb{E}(Y|Z = 0)}{\mathbb{E}(D|Z = 2)}$

(ii) $TUT \equiv \mathbb{E}(Y_1 - Y_0|C = 0) = \dfrac{\mathbb{E}(Y|Z = 1) - \mathbb{E}(Y|Z = 2)}{1 - \mathbb{E}(D|Z = 2)}$

(iii) $ASB \equiv \mathbb{E}(Y_0|C=1) - \mathbb{E}(Y_0|C=0) = \dfrac{\mathbb{E}(Y|Z=0) - \mathbb{E}(Y|Z=2, D=0)}{\mathbb{E}(D|Z=2)}$

(iv) $ASL \equiv \mathbb{E}(Y_1|C=1) - \mathbb{E}(Y_1|C=0) = \dfrac{\mathbb{E}(Y|Z=2, D=1) - \mathbb{E}(Y|Z=1)}{1 - \mathbb{E}(D|Z=2)}$.

Proof. Parts (i) and (iii) we require an expression for $\mathbb{E}(Y_0|C=1)$ in terms of (Y, D, Z). By Lemma 1(ii) and iterated expectations

$$\mathbb{E}(Y|Z=0) = \mathbb{E}(Y_0) = \mathbb{E}(Y_0|C=0)\mathbb{P}(C=0) + \mathbb{E}(Y_0|C=1)\mathbb{P}(C=1).$$

Re-arranging and substituting Lemma 1(i) and (iv),

$$\mathbb{E}(Y_0|C=1) = \frac{\mathbb{E}(Y|Z=0) - \mathbb{E}(Y|Z=2, D=0)\mathbb{E}(1-D|Z=2)}{\mathbb{E}(D|Z=2)}. \tag{10}$$

Part (i) follows by combining (10) with Lemma 1(v) and simplifying; part (iii) follows by combining (10) with Lemma 1(iv) and simplifying. Similarly, for parts (ii) and (iv) we require an expression for $\mathbb{E}(Y_1|C=0)$ in terms of observables. By Lemma 1(iii) and iterated expectations,

$$\mathbb{E}(Y|Z=1) = \mathbb{E}(Y_1) = \mathbb{E}(Y_1|C=0)\mathbb{P}(C=0) + \mathbb{E}(Y_1|C=1)\mathbb{P}(C=1).$$

Re-arranging and substituting Lemma 1(i) and (v),

$$\mathbb{E}(Y_1|C=0) = \frac{\mathbb{E}(Y|Z=1) - \mathbb{E}(Y|Z=2, D=1)\mathbb{E}(D|Z=2)}{\mathbb{E}(1-D|Z=2)}. \tag{11}$$

Part (ii) follows by combining (11) with Lemma 1(iv) and simplifying; part (iv) follows by combining (11) with Lemma 1(v) and simplifying. □

ONLINE APPENDIX: The forcing-choice design and paternalism in pawnshop borrowing

Craig McIntosh Isaac Meza Joyce Sadka Enrique Seira Francis J. DiTraglia

A Additional materials

Figure OA-1: Behavior of borrowers who lost their pawn.

(a) Elapsed days to first payment (b) Elapsed days to last payment

(c) Payments as % of loan (d) Number of payments

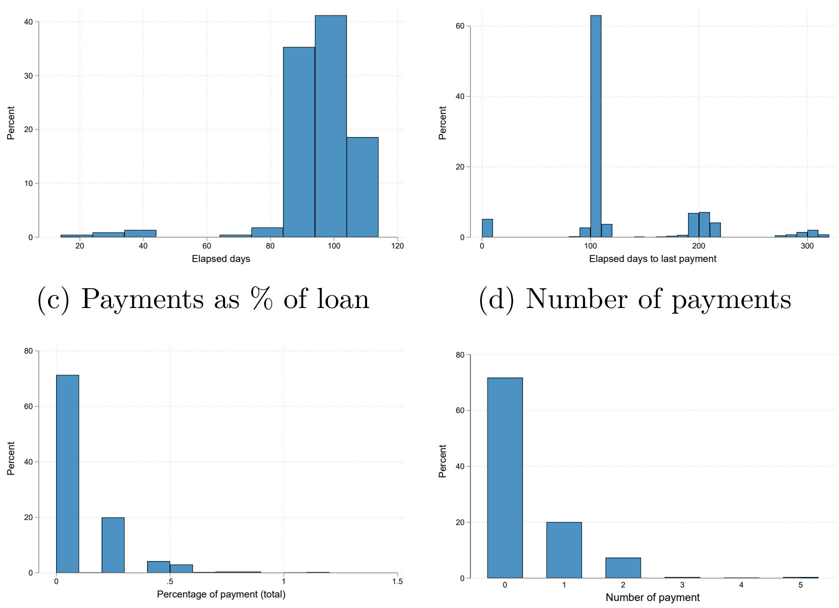

This figure provides more details on the behavior of clients who were assigned to the control group and did not recover their pawn. Panel (a) shows a histogram of days elapsed from the pawn to the first payment, while panel (b) displays a histogram of days elapsed until the last payment. Some borrowers make payments after day 105, the end of the grace period: if they pay all interest owed, they can “restart” the loan. This amounts to starting a new loan with the same conditions and same pawn. Panel (c) shows a histogram of the fraction of the loan paid, while panel (d) presents a barplot of the number of times that borrowers went to the branch to make payments.

Figure OA-2: Weekly default rates experimental branches and all branches

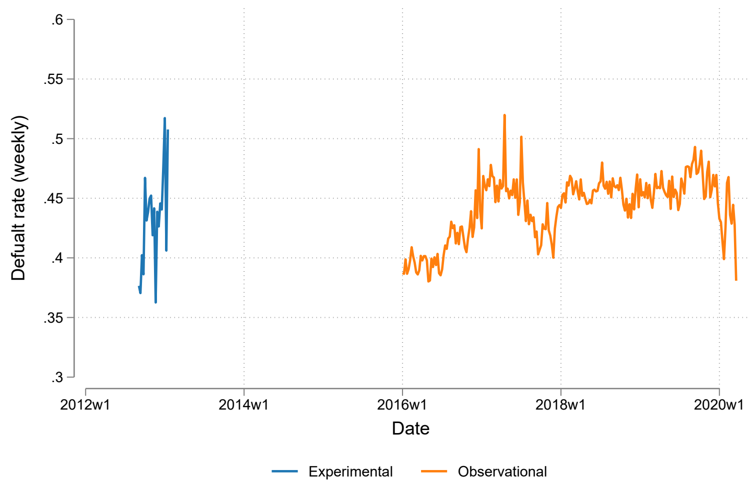

The above figure shows the weekly default rates for our experimental sample, and for 4 years after our experiment. We show that the default rates in the experimental sample are not atypical.

Figure OA-3: Determinants of choice.

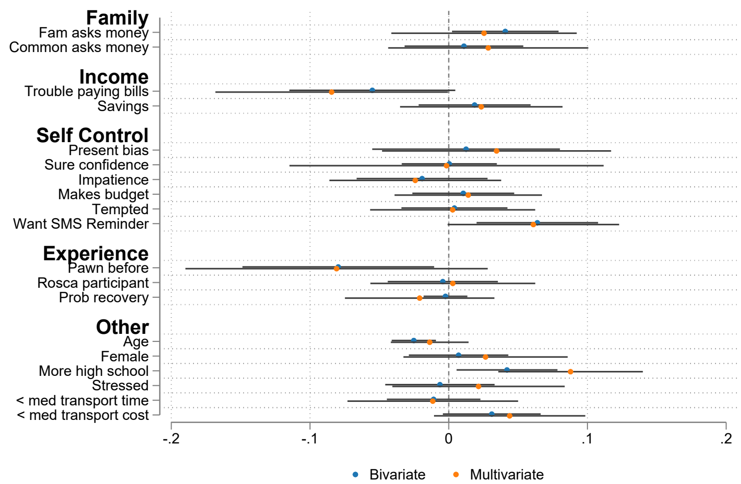

The above figure shows the determinants in a bivariate and multivariate OLS regression of choosing commitment. Choice commitment is a binary variable equal to one, whenever subjects choose the mandatory commitment contract in the choice arm.

Figure OA-4: Profit difference Mandatory vs Control arm.

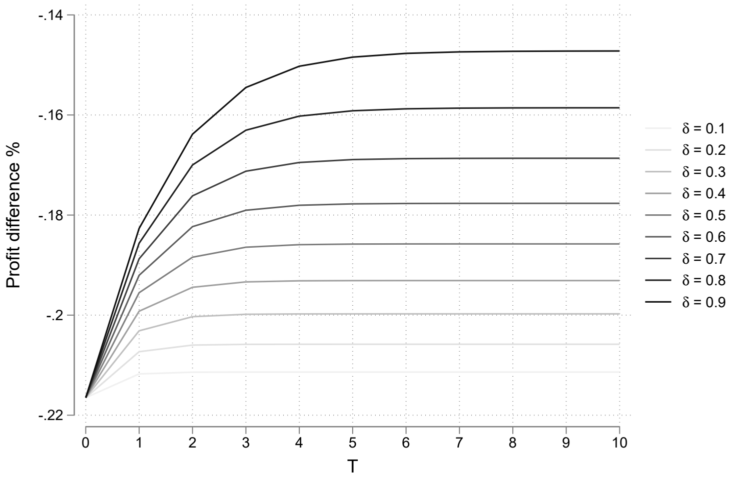

The above figure shows the difference in profit $\text{Profit}_{j,X} = \sum_{t=0}^{T} \delta^t \{\text{Financial Cost}_j | \text{Contract} = X\} \times \Pr(\text{Repeat}|\text{Contract} = X)^t$ for the Mandatory vs Control contracts for different values of the discount rate δ and time horizon T.

Table OA-1: Baseline survey questions (translated to English)

	Baseline Survey
1	**Your pawn was:**
	(a) Inherite, (b) a gift, (c) bought by me, (d) lend to me, (e) other ____________
2	**Mark with an "X" in the line below how likely is that you recover your pawn.**
	Where 0 is impossible and 100 is completely certain
3	**How much do you think the item you plan to pawn is worth? ________________ pesos**
4	**Gender**
5	**Age**
6	**Civil Status**
	(a) married, (b) single, (c) divorced, (d) widowed
7	**Work status**
	(a) employed, (b) own business, (c) houseshores, (d) don't work, (e) retired, (f) study
8	**Education**
	(a) no formal education, (b) primary, (c) middle school, (d) highschool, (e) more than highschool
9	**In the last month, did a friend or family member asked you for money?**
	(a) yes (b) no
10	**What would you like to have: 100 pesos tomorrow or 150 pesos in one month?**
11	**How often do you feel stressed by your economic situation?**
	(a) always, (b) very often, (c) sometimes, (d) never
12	**What is the main reason you want to pawn?**
	(a) Need the money because somebody in my family lost his/her job
	(b) Need the money to pay for a sickness in the family
	(c) Need the money for an urgent expense
	(d) Need the money for some non urgent expense.
13	**How stressed do you feel from the situation that led to pawn?**
	(a) very stressed, (b) somwhat stressed, (c) a little stressed, (d) not stressed
14	**In 3 months, I expect to have a __________________ situation**
	(a) better, (b) similar, (c) worse
15	**Have you panwned before?**
	(a) yes (b) no
16	**How many times have you pawned on a Lender P branch?**
	(a) NO___ (b) 1-2 times ___ (c) 3-5 times____ (d) More than 5____
17	**If you are saving money and a family member wants to use it for something**
	(a) I would only give him the money for an urgent expenze
	(b) I would give him the money even if it was not an urgent expense
	(c) I would not give him/her the money regardless
	(d) No one would ask me for my money
18	**Do you make an expenses budget for the month ahead of time?**
	(a) always, (b) very often, (c) sometimes, (d) never
19	**Do you have other items you could pawn?**
	(a) yes (b) no
20	**Do you have savings?**
	(a) yes (b) no
21	**Do you participate in a ROSCA?**
	(a) yes (b) no
22	**Is it common that family or friends ask for money?**
	(a) yes (b) no
23	**How much did you spend to come to the branch today? $________________ pesos**
24	**How much time does it usually take to come to this branch? ____________**
25	**How much does your family spend in a normal week? $________________ pesos**
26	**How much do you manage to save in a normal week? $________________ pesos**
27	**Does it happen to you that you spend more than you wanted because you fall into temptation?**
	(a) never, (b) almost never, (c) sometimes, (d) very often
28	**In the last 6 months, has it happened that at some point you lacked money to pay**
	(a) rent? (b) food (c)food (d) medicine (e) electricity (f) heating (g) telephone (i) water
29	**What would you like to have: 100 pesos in 3 months or 150 pesos in four months?**
30	**Would you like to receive (free) reminders for upcomming payments?**
	(a) yes (b) no

Translation of the baseline questionnaire. Highlighted questions were used in the wide forest.

A.1 Intermediate Outcomes

Table OA-2: Effects on intermediate outcomes

	Panel A : Speed of payment				
	Days to 1st payment	% of payment in 1st visit	Pr(Recovery in 1st visit)	Loan duration (days)	Loan duration | recovery
	(1)	(2)	(3)	(4)	(5)
Mandatory structured	-13.8***	9.76***	0.079***	-27.9***	-17.9***
	(1.61)	(2.74)	(0.026)	(4.35)	(3.88)
Choice	-3.51**	-0.58	-0.010	-0.18	-1.35
	(1.57)	(2.23)	(0.022)	(4.33)	(4.19)
Observations	4412	6304	6304	6304	3031
R-squared	0.055	0.017	0.016	0.054	0.041
Control Mean	82.8	45.8	0.30	136.6	103.9

	Panel B : Variables related to default				Panel C : Visit variables	
	Pr(+ payment & default)	% of pay | def	Pr(Selling pawn | def)	Pr(Selling pawn)	# of visits	# of visits | def
	(6)	(7)	(8)	(9)	(10)	(11)
Mandatory structured	-0.071***	-4.12***	0.14***	0.0050	-0.031	-0.20***
	(0.015)	(1.28)	(0.034)	(0.021)	(0.049)	(0.050)
Choice	-0.027**	-1.82*	0.050*	0.0035	0.085	-0.079*
	(0.014)	(1.05)	(0.029)	(0.019)	(0.053)	(0.043)
Observations	6304	2492	2492	6304	6304	2492
R-squared	0.011	0.024	0.034	0.016	0.022	0.028
Control Mean	0.12	9.68	0.71	0.31	1.14	0.39

This table explores treatment effects in "intermediate variables". Each column represents regression output for different dependent variables following equation (1). Panel A focuses on variables related to the speed of payment. While Panel B focuses on variables related to default, and Panel C related to visits. Each regression includes branch and day-of-week FE. Standard errors are clustered at the branch-day level.

A.2 Censoring

Some loans in our sample are "censored" in that they continue beyond our observation period. For these loans, we do not know whether the borrower ultimately defaulted or recovered her pawn. We have also shown that one effect of the forcing arm is to accelerate repayment, meaning that it is less likely for loans in this arm to be censored. This issue is illustrated in Figure OA-5, which shows the CDF of loan completion (either default or recovery in Panel (a)) and loan recovery (Panel (b)) by the number of days since first pawn. Two features of these graphs are salient for our analysis. The first is the extent to which loan outcomes are observed more quickly in the mandatory commitment arm. This is primarily due to the substantially higher rate of repayment of Mandatory Commitment loans at 120 days (15 pp higher than the other arms). The second is the very low rate at which loans are recovered in any arm after 120 days. In the 180-320 day window loans are largely dormant, suggesting that many of the censored loans will in fact end in default.

The confluence of censoring and a treatment effect on censoring is potentially problematic from an experimental point of view. The approach taken in the headline results is a conservative one in that it inherently assumes that all of the loans outstanding at the end of the observation window will be repaid, making it so that the acceleration of payment observed in the Mandatory arm does not translate mechanically into the that treatment decreasing default. Nonetheless, to be certain that this issue is not driving our results we conduct a bounding exercise to understand how large the effects of this problem can possibly be.

One way of considering the effect that this issue could have on our results is to make extreme assumptions about the outcome of these loans in the treatment and control so as to bound the possible influence of censoring. In Table OA-3 we compare the Mandatory and Control arms, bounding the censoring issue by reversing assumptions about the outcome of censored loans in the treatment versus the control. Panel B provides the lower bound for the treatment effect (closest to zero) by assuming censored control loans are always repaid and treatment loans never are; even in this lower-bound case the treatment effect is cost-reducing and significant at the 1% significance level and indeed the magnitude of this lower bound estimate is only 6% closer to zero than our headline result. Panel C estimates the upper bound by making the reverse assumption. Comfortingly, even with these extreme assumptions the significance on the main treatment effects never flips and treatment effects on financial cost and interests payments remain negative and significant in all scenarios. So there appears to be no scope for the censoring issue to overturn our main results.

Finally, Panel E of this table conducts a logit prediction model that uses all of the available information on loans that were completed to predict the outcome of loans that were not. This is a "best guess" of the outcome on censored loans. Using this prediction, we replicate the main experimental results and find that the treatment effect on financial cost increases from -204 (main results) to -264 (censored loans predicted), and the APR from -11% to -17%. Hence, while the censoring issue does have an effect on the magnitude of our estimated treatment effects, these checks confirm that (a) the core results are fully robust to censoring, and (b) the headline approach that we take to the issue is conservative and likely understates the true magnitude of impacts.

Table OA-3: Bounding censoring

	FC	Interest pymnt	Principal pymnt	Lost pawn value	Default	APR
		Panel A :	Control = 0	Mandatory structured = 0		
	(1)	(2)	(3)	(4)	(5)	(6)
Mandatory structured	-236.0***	-191.7***	-0.63	-75.9**	-0.064***	-0.14***
	(48.1)	(37.6)	(3.01)	(30.5)	(0.023)	(0.022)
Observations	3724	3724	3724	3724	3724	3724
R-sq	0.016	0.025	0.004	0.012	0.019	0.043
Control Mean	989.9	593.4	5.96	396.5	0.44	0.61
		Panel B :	Control = 0	Mandatory structured = 1		
	(7)	(8)	(9)	(10)	(11)	(12)
Mandatory structured	-191.2***	-207.7***	1.17	-15.1	0.0083	-0.076***
	(49.7)	(37.4)	(3.45)	(31.2)	(0.024)	(0.026)
Observations	3724	3724	3724	3724	3724	3724
R-sq	0.013	0.026	0.004	0.009	0.014	0.023
Control Mean	989.9	593.4	5.96	396.5	0.44	0.61
		Panel C :	Control = 1	Mandatory structured = 0		
	(13)	(14)	(15)	(16)	(17)	(18)
Mandatory structured	-319.0***	-140.4***	-2.33	-210.3***	-0.21***	-0.24***
	(50.9)	(34.1)	(3.16)	(30.3)	(0.023)	(0.027)
Observations	3724	3724	3724	3724	3724	3724
R-sq	0.021	0.020	0.004	0.021	0.053	0.061
Control Mean	1069.2	545.9	7.69	523.3	0.57	0.70
		Panel D :	Control = 1	Mandatory structured = 1		
	(19)	(20)	(21)	(22)	(23)	(24)
Mandatory structured	-274.2***	-156.3***	-0.53	-149.6***	-0.13***	-0.17***
	(52.5)	(33.8)	(3.58)	(31.1)	(0.024)	(0.030)
Observations	3724	3724	3724	3724	3724	3724
R-sq	0.017	0.021	0.003	0.013	0.028	0.032
Control Mean	1069.2	545.9	7.69	523.3	0.57	0.70
		Panel E :	Prediction with lasso-logit model			
	(25)	(26)	(27)	(28)	(29)	(30)
Mandatory structured	-264.9***	-169.6***	-1.43	-127.4***	-0.12***	-0.17***
	(53.8)	(37.2)	(3.52)	(33.1)	(0.025)	(0.028)
Choice	-42.4	-29.1	-2.66	-14.6	-0.017	0.0026
	(56.9)	(41.8)	(3.24)	(34.9)	(0.024)	(0.029)
Observations	6304	6304	6304	6304	6304	6304
R-sq	0.018	0.022	0.002	0.010	0.016	0.042
Control Mean	1034.5	563.4	7.69	471.2	0.52	0.66

Given the censored loans, i.e. loans that have not finished by the end of the observation period, we estimate 'a la Manski' bounds for these loans, meaning that we impute all loans to either *default*= 1 or *recovery*= 0 depending on the treatment arm. Different panels perform different imputations for the censored loans for all possible combinations for the imputation, and computes the ATE for the same outcomes of Table 3. Panel A, for instance, assumes that all outstanding loans are fully payed. Panel B is the most conservative imputation since it assumes all outstanding loans in the control arm are payed, while all the outstanding loans in the mandatory commitment arm default. Panel C, on the other hand, is the most optimistic scenario opposite to that of Panel B. Panel D assumes all remaining loans default. The last panel makes the imputation to the censored loans according to the best prediction using a piecewise lasso logit model for default. In concrete, we build two logit models with lasso regularization, depending whether the loan duration is less than 220 days (two cycles) or more than 220 days. For prediction we use the former whenever the last recorded payment was done within 220 days, and the latter otherwise. Both models includes loan characteristics (loan size, branch), and payment behavior (loan duration so far, days to first payment, % of first payment, % of payments at 30, 60, 90, and 105 days, and % of interest payed at 105 days), but the latter model also includes % of payments at 150, 180, and 210 days. This predictive model achieves an accuracy rate of 92% both in-sample and out-of-sample. Note that in all panels we maintain significant results for Financial Cost as dependent variable, while only in the most conservative scenario (Panel B) we lose significance for the APR outcome.

Figure OA-5: Survival graph.

(a) Ended contract

(b) Recovery

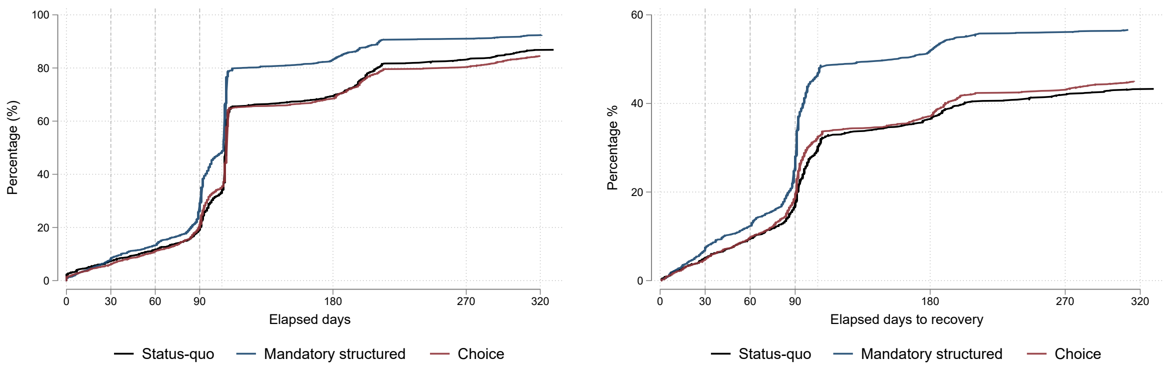

This Figure shows the CDF of loan completion either default or recovery in Panel (a), or loan recovery in Panel (b), by the number of days since first pawn.

A.3 Robustness accounting for other costs

Table OA-4: Effects on more comprehensive cost measures

	FC	FC (subj.value)	FC + tc	FC - interest	FC (subj.value) + tc - int
	(1)	(2)	(3)	(4)	(5)
Mandatory structured	-204.0***	-299.9***	-207.7***	-98.5***	-146.3**
	(48.1)	(83.3)	(49.0)	(36.7)	(72.8)
Choice	-38.9	-56.4	-32.6	-30.7	-25.3
	(49.8)	(83.5)	(50.9)	(39.2)	(74.4)
Observations	6304	6304	6304	6304	6304
R-squared	0.013	0.009	0.014	0.005	0.006
Control Mean	942.4	1389.9	1026.1	480.7	927.7

	APR	APR (subj.value)	APR + tc	APR - interest	APR (subj.value) + tc - int
	(6)	(7)	(8)	(9)	(10)
Mandatory structured	-0.11***	-0.22***	-0.13***	-0.062***	-0.097**
	(0.019)	(0.051)	(0.028)	(0.019)	(0.044)
Choice	-0.0086	-0.053	-0.0035	-0.031*	-0.043
	(0.019)	(0.045)	(0.028)	(0.018)	(0.040)
Observations	6304	6304	6304	6304	6304
R-squared	0.031	0.011	0.027	0.008	0.007
Control Mean	0.57	1.12	0.72	0.31	0.84

This table augments the measure of financial cost presented in Table 3 with measures of transaction costs, subjective costs, and adjustments for liquidity costs. Panel A reports financial cost in pesos, while Panel B shows APR. Columns (1) and (6) replicate our previous results for comparability. Columns (2) and (7) of Table OA-4 use the subjective value of the pawn reported by the borrower rather than its appraised value. Columns (3) and (8) adjust for self-reported transport costs per visit plus an entire day's wage, both multiplied by the number of visits that each individual made.[26] Columns (4) and (9) adjust to consider the liquidity cost. Finally, columns (5) and (10) include all three changes together. The main takeaway from the table is that results are quite robust to including a much expanded measure of costs. Each regression includes branch and day-of-week FE. Standard errors are clustered at the branch-day level.

A.4 Repawning

We now estimate equation 1 with Different dependent variables. In column 1, the dependent variable indicates, for each borrower in the experiment, whether he/she pawned again after the first loan in the experiment (up to the end period of our data set 338 days after the experiment began). The result is that the likelihood of repeat business increases by 6.7%. While this appears to be *prima facie* evidence of greater satisfaction among borrowers in the mandatory arm, the interpretation is complicated by the fact that monthly payments may themselves trigger more borrowing to pay them. Note that from the lender's perspective, why repeat pawning happens may not be as important as the fact that it happens. Illiquidity does not seem to drive this result, given that the effect on re-pawning comes after 90 days (during the period of contract demanded payments) and not before (see columns 2 and 3). Column 4 only considers new loans that use different collateral from that of the initial one. We do this to foreclose the explanation that those in the Mandatory arm, being more liquidity-constrained, return to pawn a second pawn to be able to pay the monthly payments of their first loan. However, we cannot reject a zero effect on pawning a different collateral. Column 5 focuses on the (endogenous) subsample of those recovering their pawn in both arms of the experiment. This means that *both* arms have recovered their pawn and could re-pawn if they so wish, and also that the liquidity demands from the monthly contract are *no longer* there as the contract has been closed. We find that the difference between the Forcing contract and the status quo is even larger in this subsample, with the former having 11pp higher likelihood of being a repeat client during our sample period.

Table OA-5: Effects on Repeat Pawning

	Ever pawns again (ITT)				
		After 90 days	Within 90 days	Different collateral	Cond. on rec
	(1)	(2)	(3)	(4)	(5)
Mandatory structured	0.067*	0.037***	0.032	0.044	0.11**
	(0.035)	(0.013)	(0.027)	(0.030)	(0.047)
Choice	0.040	0.0098	0.030	0.038	0.057
	(0.031)	(0.0087)	(0.026)	(0.028)	(0.042)
Observations	4441	4441	4441	4441	2170
R-squared	0.003	0.006	0.001	0.002	0.008
Control Mean	0.32	0.020	0.30	0.30	0.35

This table estimates the specification of equation 1 but at the level of the borrower (not the loan). Each column represents a regression with a different outcome variable. In column 1, the dependent variable indicates, for each borrower in the experiment, whether he/she pawned again after the first loan in the experiment (up to the end period of our data set 338 days after the experiment began). Column 2 is analogous, but only pawning after 90 days of the first loan is considered. Column 3 instead considers pawning before 90 days. Column 4 is analogous to column 1 but focuses on the pawning of a gold piece that is different from the one in the first experimental loan. Column 5 is analogous to column 1, but conditioning on the sample that recovered the first loan. Standard errors are clustered at the branch-day level.

B Alternative explanations

B.1 Learning

The experiment was not designed to study learning, and the short time frame and sample size limit what we can say. Table OA-6 presents information about borrowers' *future* pawning behavior as a function of treatment assignment. Column (1) considers the 228 clients who returned only a second time to pawn again at a day/branch that was randomly assigned to the choice arm. Each of the two rows in this column presents a difference of mean commitment take-up rates, and associated standard error. The first row compares those who were *initially* assigned to mandatory structure against those where were assigned to control; the second row compares those who were initially assigned to the choice arm to those who were assigned to the other two arms. In each case, there is no statistically discernible difference in the rates of commitment take-up. Granted, this is a selected sample because the decision to pawn again is potentially endogenous to the initial treatment allocation. For this reason, Column (2) considers the full sample of 4441 borrowers by re-defining the outcome variable to be an indicator for returning to pawn again at a branch/day when commitment was offered *and* choosing commitment. This composite outcome variable is not subject to the sample selection problem (although it is directly driven by the decision to repeat borrow). The comparison in the two rows remains the same: mandatory structure versus control in row one and choice versus mandatory arms in row two. Again, there is no statistically discernible difference in commitment take-up rates in either row. While these exercises cannot completely exclude the possibility that learning plays a role, they provide no indication that the lack of voluntary compliance is simply a matter of inexperience with commitment.

Table OA-6: Effect of Prior Assignment on Subsequent Choice

	Choose structure in $t+1$	Ever choose structure in $t+1$
t	(1)	(2)
Mandatory structured (ATE)	-0.0047	0.00014
	(0.048)	(0.0027)
Choice (ITT)	0.034	0.0015
	(0.057)	(0.0030)
Observations	228	4441
R-sq	0.004	0.000
DepVarMean	0.092	0.0047

Column (1) reports results for the 228 borrowers who returned to pawn again at a day/branch that was randomly assigned to the choice arm, enabling us to observe whether they chose commitment or the status quo contract. Each row presents a difference in mean commitment take-up rates and associated standard errors. The first row (ATE) compares borrowers who were initially assigned to mandatory structure against those were assigned to the control condition. The second row (ITT) compares borrowers who were initially assigned to the choice commitment condition to those who were not. Whereas column (1) conditions on the (endogenously) selected sample of borrowers who return to pawn again, column (2) considers the full sample by re-defining the "outcome" to be an indicator for whether a borrower pawned again on a day when choice was offered *and* chose commitment.

B.2 Discount rates

Figure OA-6: Financial benefit TUT effect for different discount rates.

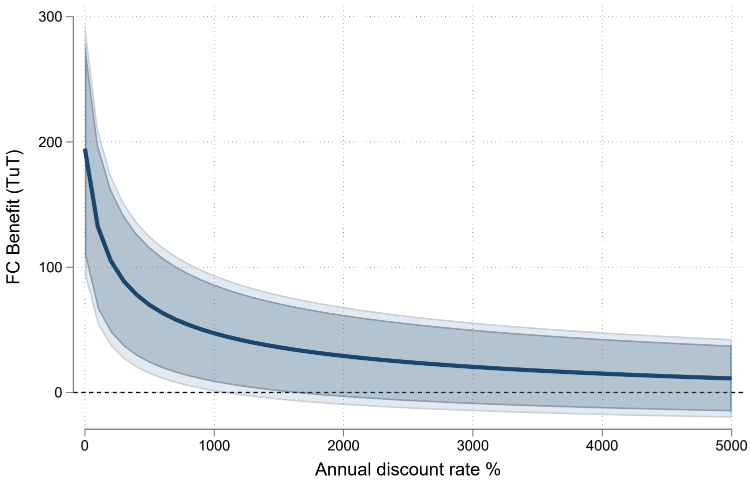

This Figure re-estimates the treatment on the untreated (TUT) effect from Table 4, introducing a daily discount factor in the definition of financial benefit. At a given annual discount rate in percentage points (x-axis) the solid line gives the adjusted TUT and the shaded regions 90% & 95% confidence bands. A discount factor of one corresponds to the estimate from Table 4. As seen from the figure, borrowers would need to face unrealistically large discount rates to reverse our headline result of a large, positive, and statistically significant TUT effect.

B.3 Present Bias

Present bias. If the benefits of commitment among non-choosers cannot be explained by standard models of rational choice, the canonical behavioral story would center on time inconsistency. While commitment is useful to anyone with hyperbolic time preferences, only those who are sophisticated–i.e. aware that they are hyperbolic discounters–will demand it. A large share of "naïve" hyperbolics in the population–borrowers who are unaware that they are hyperbolic discounters–could therefore drive a large and positive TUT. Our baseline survey included standard questions about discount rates between today and a month in the future versus discount rates between three and four months in the future. This allows us to classify borrowers who display more impatience over immediate delays as present biased. This measure of financial hyperbolicity is widely used in survey research, although it is not without problems.[27]

[27] Our measure is dichotomous, and it is not incentivized. Recent empirical work has shown the superiority of more elaborate measures such as "convex time budgets" (Andreoni *et al.*, 2015) while questioning the interpretation of measures of hyperbolicity that are not based on consumption (Andreoni & Sprenger, 2012; Cohen *et al.*, 2020), suggesting that real effort tasks provide a better measure (Augenblick *et al.*, 2015). Given that we had only a few minutes to interview real pawnshop clients prior to a commercial transaction, our simple measure was a necessary compromise.

Figure OA-7: Heterogeneity of the TUT by behavioral variables.

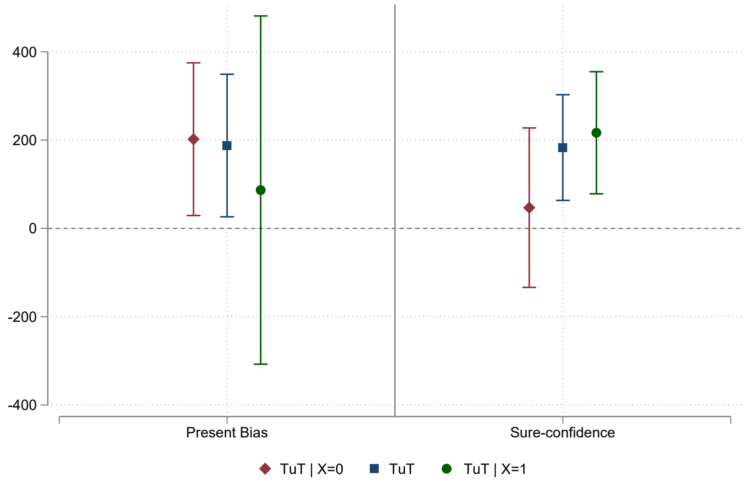

Each panel in this figure shows how the estimated treatment on the untreated (TUT) effect varies with a binary survey variable X_i. In the left panel (P.B.), $X_i = 1$ if borrower i is "present-biased" based on her responses to the time preference questions from our survey. In the right panel (Sure-confidence) $X_i = 1$ if borrower i reported that she was certain to recover her pawn, zero otherwise.

If we could perfectly measure present bias and sophistication, we could divide the sample into three groups: sophisticated hyperbolics (who chose commitment), time-consistent non-choosers (for whom forcing will not be effective), and naïve hyperbolic non-choosers (who will benefit from mandatory structure). If present bias fully explains the low take-up rate of voluntary commitment, we should find that the TUT for present-biased borrowers is positive. This is because among the group of non-takers, a comparison of present-biased borrowers against everyone else is a comparison of naïve hyperbolics against time-consistent non-choosers.

The left panel of Figure OA-7 carries out a feasible version of this exercise using our survey measure of present bias. The overall TUT estimate along with a 95% confidence interval is given in blue.[28] The corresponding TUT estimate and confidence interval for present-biased borrowers identified with the survey question is given in green; results for all other borrowers are shown in red. The overall TUT is a weighted average of the impact in these two sub-groups. The TUT among the present biased is insignificant and less than half the size of the strongly significant TUT among those who are *not* present biased. Therefore, taking our survey measure of hyperbolicity at face value, we find no indication that present-bias explains our positive estimated TUT.

[28]For all borrowers who answered our present-bias survey questions.

B.4 Sure Confidence

Figure OA-8: Determinants sure confidence.

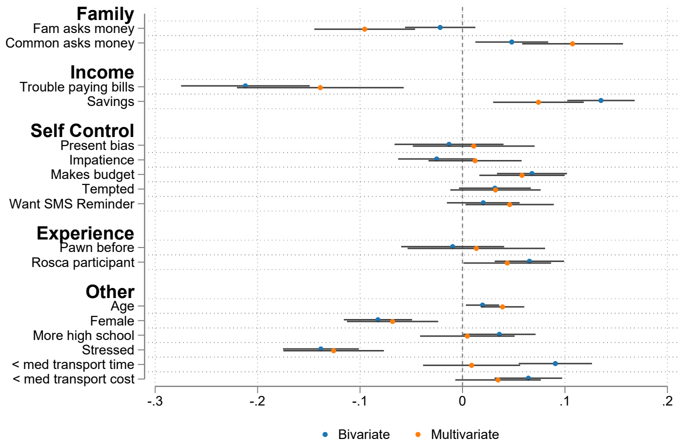

The above figure shows the determinants in a bivariate and multivariate OLS regression of sure confidence among the non-choosers. Sure confidence is a binary variable defined to be one when people report a 100% probability of recovery.

C Bounds, FOSD and Rank Invariance

Figure OA-9: Fan & Park bounds for benefit in APR%.

(a) APR Bounds

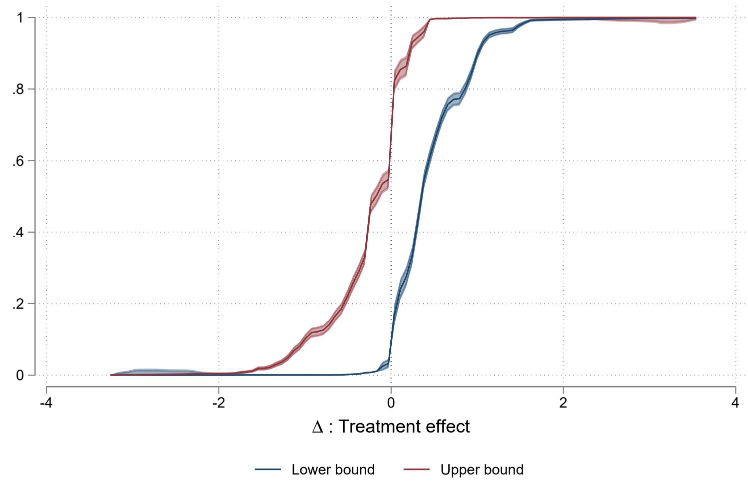

This figure depicts the Fan & Park (2010) bounds on the distribution F_Δ of individual treatment effects $\Delta \equiv (Y_1 - Y_0)$, described in Section 5.1, for the APR outcome. The dark red curve and light red shaded region give the estimated upper bound function $\overline{F}$ for F_Δ and associated (pointwise) 95% confidence interval. The dark blue curve and light blue shaded region give the estimated lower bound function $\underline{F}$ for F_Δ and associated (pointwise) 95% confidence interval. Confidence intervals are computed using the asymptotic distribution for the bounds. Evaluating the bounds at $\delta = 0$, we see that between 23% and 97% of borrowers have a positive individual treatment effect.

Figure OA-10: Distribution of treatment effects under rank invariance.

(a) APR benefit (b) Financial benefit

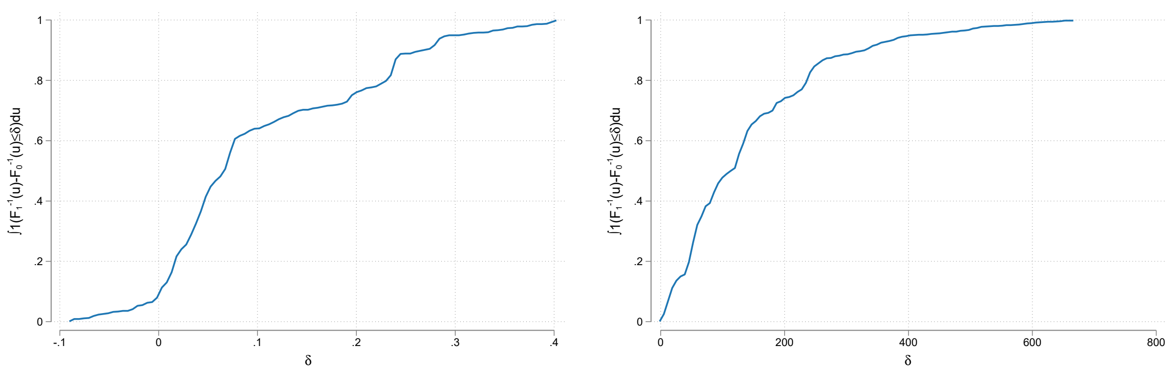

This figure shows the CDF of individual treatment effects under the assumption of rank invariance, computed from $F_\Delta(\delta) = \int_0^1 \mathbb{1}\{F_1^{-1}(u) - F_0^{-1}(u) \le \delta\}\, du$ where F_1^{-1} and F_0^{-1} are the quantile functions of Y_1 and Y_0.

D Causal Random Forest, CATE, and 'mistakes'

To estimate the conditional average treatment effects shown in Figure 3, we use administrative and survey data, and the function `causal_forest()` of the `grf` R package; to estimate conditional TOT and TUT effects we use the `instrumental_forest()` function from the same package. In each case, we use the default parameter values from the `grf` package with one exception: we increase the number of trees from the default value of 2000 to 5000. The functions `causal_forest()` and `instrumental_forest()` implement special cases of the "generalized random forest" methods of Athey *et al.* (2019). In broad strokes, these functions combine a large number of regression trees that recursively partition the covariate space to estimate conditional average effects. The trees are "honest" in that observations used to determine the optimal partition are not used to estimate effects, and vice-versa. While closely related to more familiar "regression-tree" random forests, the generalized random forest approach explicitly targets the parameter of interest–a conditional ATE or IV estimand–when choosing the optimal covariate partition. [29]

[29] For more details, see Athey *et al.* (2019) and the `grf` documentation: `https://grf-labs.github.io/grf/`. When constructing our random forest estimates of heterogeneous treatment effects, we use observations for all borrowers who answered at least *part* of the intake survey. We impute the median response for the missing values, while also including an indicator whether the variable originally had a missing value. Results are similar if we manually include interactions between the original/imputed variable and an indicator for missingness. This is as expected, given that tree-based methods by their nature "automatically" consider interactions of arbitrary orders.

Figure OA-11: Conditional ATEs from “wide” and “narrow” covariate sets.

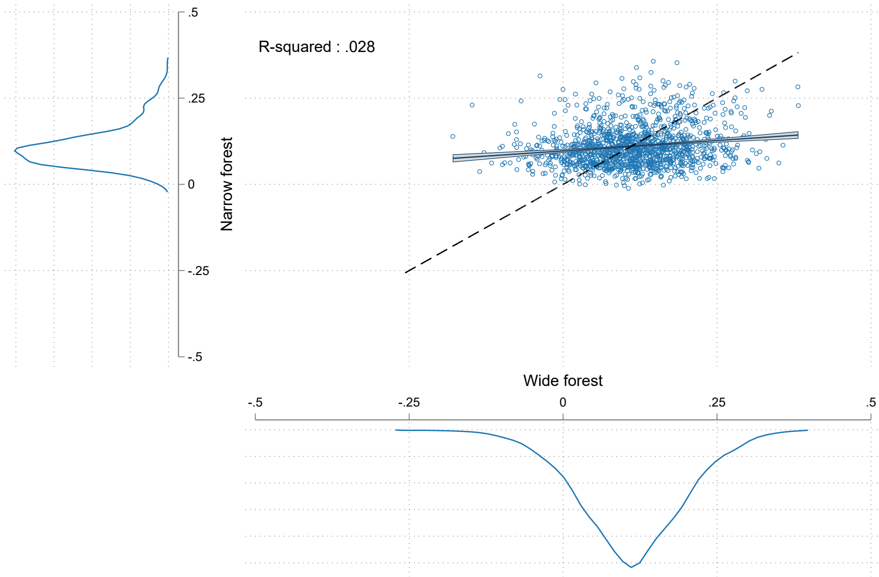

This figure plots the relationship between the causal forest conditional ATE estimates from Section 6.2 that use the “wide” set of covariates (all intake survey responses) and those based on a restricted “narrow” set of covariates (age, gender, HS education, and previous borrowing). The scatterplot graphs one estimate versus the other, with the “wide” covariate set on the horizontal axis and the “narrow” set on the vertical axis. The density plots on each axis show the estimated marginal distribution of conditional ATEs under each covariate set. The density for the “wide” covariate set is considerably more dispersed, as the causal forest based on this set of covariates captures considerably more treatment effect heterogeneity.

E Testable Implications of the Exclusion Restriction

As above, let $Y_0 \equiv Y(d=0, z=0)$ and $Y_1 \equiv Y(d=1, z=1)$ denote the potential outcomes under *mandatory treatment*: Y_0 is the potential outcome when assigned to the status quo contract and Y_1 when mandated to the structured contract. Further let $Y_{0,2} \equiv Y(d=0, z=2)$ and $Y_{1,2} \equiv Y(d=1, z=2)$ denote the potential outcomes under *free choice of treatment*: $Y_{0,2}$ is the potential outcome when choosing the status quo contract and $Y_{1,2}$ when choosing the commitment contract. Using this notation, (3) becomes $Y_0 = Y_{0,2}$ and while (4) becomes $Y_1 = Y_{1,2}$. Without imposing these, Assumption 1(iii) becomes

$$Y = \mathbb{1}(Z=0)Y_0 + \mathbb{1}(Z=1)Y_1 + \mathbb{1}(Z=2)\left[(1-C)Y_{0,2} + CY_{1,2}\right]$$

but parts (i) and (ii) continue to hold. Accordingly, parts (i)–(iii) of Lemma 1 are unchanged, while parts (iv) and (v) become

$$\mathbb{E}(Y|D=0, Z=2) = \mathbb{E}(Y_{0,2}|C=0), \quad \mathbb{E}(Y|D=1, Z=2) = \mathbb{E}(Y_{0,1}|C=1).$$

Using these expressions, the testable restrictions we consider here are as follows:

$$\mathbb{E}(Y_0|C=0) = \mathbb{E}(Y_{0,2}|C=0) \tag{12}$$

$$\mathbb{E}(Y_1|C=1) = \mathbb{E}(Y_{1,2}|C=1). \tag{13}$$

Because they refer to different groups of people–choosers versus non-choosers–either of (12) and (13) could hold when the other is violated. For this reason we consider each in turn. Our approach is closely related to arguments from Huber & Mellace (2015) and DiTraglia & Garcia-Jimeno (2019), among others.

Consider first (12). Let $p \equiv \mathbb{P}(C=1) = \mathbb{P}(D=1|Z=2)$ denote the share of choosers in the population. This value is point identified regardless of whether the exclusion restriction holds. Because Z was randomly assigned, a fraction p of borrowers with $Z=0$ are choosers while the remaining $(1-p)$ are non-choosers. It follows that, regardless of whether the exclusion restriction holds, the observed distribution of $Y|Z=0$ is a mixture of $Y_0|C=0$ and $Y_0|C=1$ with mixing weights $(1-p)$ and p. This allows us to construct a pair of bounds for $\mathbb{E}(Y_0|C=0)$ as follows. The non-choosers must lie *somewhere* in the distribution of $Y|Z=0$. Consider the two most extreme possibilities: they could occupy the bottom $(1-p)\times 100\%$ of the distribution or the top $(1-p)\times 100\%$ of the distribution. For this reason, computing the average of the *truncated* distribution of $Y|Z=0$, cutting out the top $p\times 100\%$, provides a lower bound for the average of Y_0 among non-choosers. Similarly, cutting out the bottom $p\times 100\%$ provides an upper bound. Let y^0_{1-p} denote the $(1-p)$ quantile of $Y|Z=0$ and y^0_p denote the p quantile of the same distribution.

Using this notation, the bounds are given by

$$\mathbb{E}\left(Y|Z=0,Y\leq y_{1-p}^{0}\right)\leq\mathbb{E}(Y_0|C=0)\leq\mathbb{E}\left(Y|Z=0,Y\geq y_{p}^{0}\right)$$

These bounds do not rely on the exclusion restriction. Under Equation 12, however, we know that $\mathbb{E}(Y_0|C=0)=\mathbb{E}(Y|D=0,Z=2)$. Therefore, if the exclusion restriction for non-choosers holds, we must have

$$\mathbb{E}\left(Y|Z=0,Y\leq y_{1-p}^{0}\right)\leq\mathbb{E}(Y|D=0,Z=2)\leq\mathbb{E}\left(Y|Z=0,Y\geq y_{p}^{0}\right). \tag{14}$$

Equation 14 provides a pair of testable implications of (12). If either inequality is violated, then the exclusion restriction for non-choosers fails. In our experiment, $\widehat{p}=\widehat{\mathbb{P}}(D=1|Z=2)=0.11$. For the APR outcome we estimate

$$\widehat{\mathbb{E}}(Y_{\text{APR}}|Z=0,Y_{\text{APR}}\leq y_{0.89}^{0})=0.48,\quad \widehat{\mathbb{E}}(Y_{\text{APR}}|Z=0,Y_{\text{APR}}\geq y_{0.11}^{0})=0.62.$$

Since $\widehat{\mathbb{E}}(Y_{\text{APR}}|D=0,Z=2)=0.58$ falls between these bounds, we find no evidence against the exclusion restriction for non-choosers. The same result holds for the financial cost outcome: results available upon request.

We can use an analogous approach to construct testable implications for 13, yielding

$$\mathbb{E}\left(Y|Z=1,Y\leq y_{p}^{1}\right)\leq\mathbb{E}(Y|D=1,Z=2)\leq\mathbb{E}\left(Y|Z=1,Y\geq y_{1-p}^{1}\right). \tag{15}$$

where y_p^1 and y_{1-p}^1 are the p and $1-p$ quantiles of the distribution of $Y|Z=1$. If either inequality is violated, then the exclusion restriction from Equation 13 fails. Again, in our experiment $\widehat{p}=0.11$. For the APR outcome we estimate

$$\widehat{\mathbb{E}}(Y|Z=1,Y\leq y_{0.11}^{1})=0.06,\quad \widehat{\mathbb{E}}(Y|Z=1,Y\geq y_{0.89}^{1})=1.28$$

Since $\widehat{\mathbb{E}}(Y_{\text{APR}}|D=1,Z=2)=0.43$ falls between these bounds, we find no evidence against the exclusion restriction for the choosers. The same holds for the financial cost outcome: results available upon request.

F Estimation and Inference

F.1 Regression-based Estimation of TOT, TUT, ASG, ASL, and ASB

Let $Z_0 \equiv \mathbb{1}\{Z = 0\}$, $Z_1 \equiv \mathbb{1}\{Z = 1\}$, and $Z_2 \equiv \mathbb{1}\{Z = 2\}$. Under standard regularity conditions, the following proposition shows that an IV regression of Y on an intercept, Z_1 and $Z_2 D$ with instruments $(1, Z_0, Z_1)$ provides consistent estimates the ATE and TOT, while an IV regression of Y on an intercept, $-Z_0$ and $-Z_2(1-D)$ with the same instrument set consistently estimates the ATE and TUT effects.

Proposition 2. *Under Assumption 1,*

(i) $Y = \mathbb{E}(Y_0) + ATE \times Z_1 + TOT \times Z_2 D + U$

(ii) $Y = \mathbb{E}(Y_1) + ATE \times -Z_0 + TUT \times -Z_2(1-D) + V$

where $\mathbb{E}(U|Z) = \mathbb{E}(V|Z) = 0$.

Proof of Proposition 2. For part (i), since $Z_2 D = Z_2 C$ and $(Z_0+Z_1+Z_2) = 1$, Assumption 1 (iii) implies $Y = Y_0 + Z_1(Y_1 - Y_0) + Z_2 D(Y_1 - Y_0)$. Now define

$$U \equiv [Y_0 - \mathbb{E}(Y_0)] + Z_1[(Y_1 - Y_0) - \text{ATE}] + Z_2 D[(Y_1 - Y_0) - \text{TOT}].$$

Since $Z_2 D = Z_2 C$ and Z is independent of (Y_1, Y_0) by Assumption 1 (i), it follows that $\mathbb{E}(U|Z) = Z_2 \mathbb{E}\left[C\left\{(Y_1 - Y_0) - \text{TOT}\right\}|Z\right]$. Thus, by iterated expectations,

$$\mathbb{E}\left[C\left\{(Y_1 - Y_0) - \text{TOT}\right\}|Z\right] = \mathbb{P}(C = 1)\left[\mathbb{E}(Y_1 - Y_0|C = 1) - \text{TOT}\right] = 0$$

since Z is independent of (Y_0, Y_1) given C, an implication of Assumption 1 (i).

For part (ii), since $Z_2(1-C) = Z_2(1-D)$ and $(Z_1 + Z_2) = 1 - Z_0$, Assumption 1 (iii) implies $Y = Y_1 - Z_0(Y_1 - Y_0) - Z_2(1-D)(Y_1 - Y_0)$. Define

$$V \equiv [Y_1 - \mathbb{E}(Y_1)] - Z_0[(Y_1 - Y_0) - \text{ATE}] - Z_2(1-D)[(Y_1 - Y_0) - \text{TUT}].$$

Since $Z_2(1-D) = Z_2(1-C)$ and Z is independent of (Y_0, Y_1) by Assumption 1 (i), $\mathbb{E}(V|Z) = -Z_2\mathbb{E}[(1-C)\left\{(Y_1 - Y_0) - \text{TUT}\right\}|Z]$. Thus, by iterated expectations,

$$\mathbb{E}[(1-C)\left\{(Y_1 - Y_0) - \text{TUT}\right\}|Z] = \mathbb{P}(C = 0|Z)\left[\mathbb{E}(Y_1 - Y_0|C = 0) - \text{TUT}\right] = 0$$

since Z is independent of (Y_0, Y_1) given C, an implication of Assumption 1 (i). $\square$

Since ASG = TOT − TUT, the preceding proposition provides a consistent estimate of the ASG effect. The ASB effect, $\mathbb{E}(Y_0|C = 1) - \mathbb{E}(Y_0|C = 0)$, can likewise be estimated by taking the difference of coefficients across two linear IV regressions with *no intercept* and instrument sets (Z_0, Z_2), as shown in the following proposition.

Proposition 3. *Under Assumption 1*

(i) $(1-D)Y = \mathbb{E}(Y_0) \times Z_0 + \mathbb{E}(Y_0|C=0) \times (1-D)Z_2 + U_0$

(ii) $(1-D)Y = \mathbb{E}(Y_0) \times (Z_0 + Z_2) + \mathbb{E}(Y_0|C=1) \times -DZ_2 + U_1$

where $\mathbb{E}(U_0|Z) = \mathbb{E}(U_1|Z) = 0$.

Proof. Assumption 1 (ii) implies $(1-D) = Z_0 + Z_2(1-C)$. Hence,

$$(1-D)Y = Z_0Y_0 + Z_2(1-C)[(1-C)Y_0 + CY_1] = Z_0Y_0 + Z_2(1-C)Y_0$$

by Assumption 1 (iii), since $Z_j^2 = Z_j$ for any j and $Z_jZ_k = 0$ for any $j \neq k$ and, similarly, $(1-C)^2 = (1-C)$ and $C(1-C) = 0$. Therefore, since $Z_2(1-C) = Z_2(1-D)$,

$$(1-D)Y = Z_0Y_0 + Z_2(1-D)Y_0, \quad (1-D)Y = (Z_0+Z_2)Y_0 + (-DZ_2)Y_0.$$

Now, define

$$\begin{aligned} U_0 &\equiv Z_0[Y_0 - \mathbb{E}(Y_0)] + Z_2(1-D)[Y_0 - \mathbb{E}(Y_0|C=0)] \\ U_1 &\equiv (Z_0+Z_2)[Y_0 - \mathbb{E}(Y_0)] + (-Z_2D)[Y_0 - \mathbb{E}(Y_0|C=1)]. \end{aligned}$$

Since $Z_2(1-D) = Z_2(1-C)$, and Z is independent of Y_0,

$$\mathbb{E}(U_0|Z) = Z_2\mathbb{E}[Y_0 - \mathbb{E}(Y_0|C=0)|C=0, Z] = 0$$

by iterated expectations and the fact that Z is conditionally independent of Y_0 given C. Since $Z_2D = Z_2C$, a nearly identical argument gives

$$\mathbb{E}(U_1|Z) = -Z_2\mathbb{E}[Y_0 - \mathbb{E}(Y_0|C=0)|C=1, Z] = 0. \qquad \square$$

The final result in this section implies that the ASL effect, $\mathbb{E}(Y_1|C=1) - \mathbb{E}(Y_1|C=0)$, can be estimated as the difference of coefficients across two linear IV regressions with *no intercept* and instrument set (Z_1, Z_2).

Proposition 4. *Under Assumption 1,*

(i) $DY = \mathbb{E}(Y_1) \times (Z_1 + Z_2) + \mathbb{E}(Y_1|C=0) \times (D-1)Z_2 + V_0$

(ii) $DY = \mathbb{E}(Y_1) \times Z_1 + \mathbb{E}(Y_1|C=1) \times DZ_2 + V_1$

where $\mathbb{E}(V_0|Z) = \mathbb{E}(V_1|Z) = 0$.

Proof. By Assumption 1, $D = Z_1 + Z_2C$. Hence, by Assumption 1 (iii),

$$DY = Z_1Y_1 + Z_2C[(1 - C)Y_0 + CY_1] = Z_1Y_1 + Z_2CY_1$$

because $Z_j^2 = Z_j$ for any j and $Z_jZ_k = 0$ for any $j \neq k$ and, similarly, $(1-C)^2 = (1-C)$ and $C(1 - C) = 0$. Therefore, since $Z_2(1 - C) = Z_2(1 - D)$,

$$DY = (Z_1 + Z_2)Y_1 + Z_2(D - 1)Y_1, \quad DY = Z_1Y_1 + Z_2DY_1.$$

Now, define

$$\begin{aligned} V_0 &= (Z_1 + Z_2)[Y_1 - \mathbb{E}(Y_1)] + Z_2(D - 1)[Y_1 - \mathbb{E}(Y_1|C = 0)] \\ V_1 &= Z_1[Y_1 - \mathbb{E}(Y_1)] + Z_2D[Y_1 - \mathbb{E}(Y_1|C = 1)]. \end{aligned}$$

Since $Z_2(1 - D) = Z_2(1 - C)$ and Z is independent of Y_1,

$$\mathbb{E}(V_0|Z) = -Z_2\mathbb{E}[Y_1 - \mathbb{E}(Y_1|C = 0)|C = 0, Z] = 0$$

by iterated expectations and the fact that Z is conditionally independent of Y_1 given C. Since $Z_2D = Z_2C$, a similar argument gives

$$\mathbb{E}(V_1|Z) = Z_2\mathbb{E}[Y_1 - \mathbb{E}(Y_1|C = 1)|C = 1, Z] = 0. \qquad \square$$

F.2 Inference for ASG, ASB, and ASL

We now explain how to carry out cluster-robust inference for the ASG, ASB, and ASL effects, as implemented in our companion STATA package. Each of these effects can be expressed as a difference of coefficients from two just-identified linear IV regressions. The ASG effect is the difference of the TOT and TUT effects from Proposition 2. Similarly, the ASB effect is the difference of $\mathbb{E}(Y_0|C = 1)$ and $\mathbb{E}(Y_0|C = 0)$ from Proposition 3 while the ASL effect is the difference of $\mathbb{E}(Y_1|C = 1)$ and $\mathbb{E}(Y_1|C = 0)$ from Proposition 4. Within each pair of IV regressions the outcome variable and instrument set is identical; only the regressors differ. Since our estimators of all three effects share the same structure, our discussion abstracts from the specific regressors and instruments used in each case.

Let $g = 1, ..., G$ index clusters and $i = 1, ..., N_g$ index individuals within a particular cluster g. In our experiment, a cluster is a branch-day combination and the experimentally-assigned treatment (control, mandatory, or choice arm) is assigned at the cluster level. We assume that observations are iid across clusters but potentially correlated within cluster. Now consider a pair of just-identified linear IV regressions given by $Y_{ig} = \boldsymbol{X}'_{1,ig}\boldsymbol{\theta}_0 + U_{ig}$ and $Y_{ig} = \boldsymbol{X}'_{0,ig}\boldsymbol{\theta}_1 + V_{ig}$ with common instrument vector $\boldsymbol{W}_{ig}$. Stacking observations in the usual manner, e.g. $\mathbf{W}_g' \equiv \begin{bmatrix}\boldsymbol{W}_{1g} & \cdots & \boldsymbol{W}_{N_g g}\end{bmatrix}$ and $\mathbf{W}' = \begin{bmatrix}\mathbf{W}_1' & \cdots & \mathbf{W}_G'\end{bmatrix}$ we can write the preceding equations in matrix form as $\mathbf{Y} = \mathbf{X}_1\boldsymbol{\theta}_\mathbf{1} + \mathbf{U}$ and $\mathbf{Y} = \mathbf{X}_0\boldsymbol{\theta}_\mathbf{0} + \mathbf{V}$ with instrument matrix $\mathbf{W}$. Now, the IV estimators for $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_0$ can be expressed as

$$
\begin{aligned}
\widehat{\boldsymbol{\theta}}_1 &= \left(\mathbf{W}'\mathbf{X}_1\right)^{-1}\mathbf{W}'\mathbf{Y} = \boldsymbol{\theta}_1 + \left(\mathbf{W}'\mathbf{X}_1\right)^{-1}\mathbf{W}'\mathbf{U}\\
\widehat{\boldsymbol{\theta}}_0 &= \left(\mathbf{W}'\mathbf{X}_0\right)^{-1}\mathbf{W}'\mathbf{Y} = \boldsymbol{\theta}_0 + \left(\mathbf{W}'\mathbf{X}_0\right)^{-1}\mathbf{W}'\mathbf{V}.
\end{aligned}
$$

By our experimental design and exclusion restriction, $\mathbf{W}_{ig}$ is independent of U_{ig} both unconditionally and conditional on cluster size. Hence, by a standard argument and under mild regularity conditions, the following expression provides a consistent, cluster robust estimator of $\widehat{\text{Avar}}(\widehat{\boldsymbol{\theta}}_1 - \widehat{\boldsymbol{\theta}}_0)$

$$
\widehat{\text{Avar}}(\widehat{\boldsymbol{\theta}}_1 - \widehat{\boldsymbol{\theta}}_0) = \begin{bmatrix}\left(\mathbf{W}'\mathbf{X}_1\right)^{-1} & -\left(\mathbf{W}'\mathbf{X}_0\right)^{-1}\end{bmatrix}\begin{bmatrix}\mathbf{S}_{UU} & \mathbf{S}_{UV}\\ \mathbf{S}_{UV}' & \mathbf{S}_{VV}\end{bmatrix}\begin{bmatrix}\left(\mathbf{X}_1'\mathbf{W}\right)^{-1}\\ -\left(\mathbf{X}_0'\mathbf{W}\right)^{-1}\end{bmatrix}
$$

where we define the IV residuals $\widehat{\mathbf{U}}_g \equiv \mathbf{Y}_g - \mathbf{X}_{1,g}\widehat{\boldsymbol{\theta}}_1$ and $\widehat{\mathbf{V}}_g \equiv \mathbf{Y}_g - \mathbf{X}_{0,g}\widehat{\boldsymbol{\theta}}_0$ along with the matrices $\mathbf{S}_{UU} \equiv \sum_{g=1}^G \mathbf{W}_g'\widehat{\mathbf{U}}_g\widehat{\mathbf{U}}_g'\mathbf{W}_g$, $\mathbf{S}_{UV} \equiv \sum_{g=1}^G \mathbf{W}_g'\widehat{\mathbf{U}}_g\widehat{\mathbf{V}}_g'\mathbf{W}_g$, and finally $\mathbf{S}_{VV} \equiv \sum_{g=1}^G \mathbf{W}_g'\widehat{\mathbf{V}}_g\widehat{\mathbf{V}}_g'\mathbf{W}_g$. In our application the number of clusters, G, is large. If desired, an *ad hoc* degrees of freedom correction can be applied by multiplying the associated standard errors by $\sqrt{G/(G-1)}$.