Working it out: Randomized modification and entrepreneurial effort in a collateralized debt market*

Christopher Eaglin† Apoorv Gupta‡ Filippo Mezzanotti§ Jonathan Zinman¶

October 2025

Abstract

We enrich a standard debt overhang model with liquidity constraints to guide the design and interpretation of a collateralized debt modification experiment on a publicly traded lender's delinquent vehicle loans to minibus entrepreneurs. Liquidity constraints add another borrower incentive compatibility constraint that interacts with debt overhang to shape repayment and effort. Consistent with model predictions, we find: debt reduction does not affect liquidity constrained borrowers; payment reduction improves both repayment and effort for borrowers with sufficient vehicle equity; payment reduction induces repayment without effort increases for low-equity borrowers. These results suggest a pecking order strategy for modification practice and policy.

*For helpful comments, we thank Deniz Aydin, Victor Lyonnet, Pascal Noel, Avantika Pal, Janis Skrastins, Jack Willis, Yufeng Wu; seminar participants at Brown, MIT, Northwestern, Northwestern (Kellogg), NYU, Rice (Jones), The College of Mexico, The JP Morgan Chase Institute, USC (Marshall), Vanderbilt; and conference participants at NBER Summer Institute, Chicago Household Finance Conference, London Business School Summer Symposium, Chicago Entrepreneurship Workshop, RAPS/RCFS Europe Conference, Holden Conference in Finance and Real Estate, Managing in Emerging Markets Conference, Columbia Management, Analytics and Data Conference, and IPA-GPRL 2024 Research Gathering. We also thank Jialu Sun for the excellent research assistance. We gratefully acknowledge financial support from the Dartmouth Rockefeller Center, the Financial Institutions and Markets Research Center at Northwestern Kellogg School of Management, Fuqua School of Business Dean's Fund and the Crown Fund at Duke University. Special thanks to the management and staff of the cooperating lender for implementing the experimental protocols and sharing data. We pre-registered the field experiment component of this study as AEARCTR-0013052. The pre-registry and pre-analysis plan are available at: https://www.socialscienceregistry.org/trials/13052.

† Duke University, Fuqua School of Business: chris.eaglin@duke.edu

‡ Dartmouth College: apoorv.gupta@dartmouth.edu

§ Kellogg School of Management, Northwestern University and NBER: filippo.mezzanotti@northwestern.edu

¶ Dartmouth College, IPA, J-PAL, and NBER: jzinman@dartmouth.edu

1 Introduction

Loan modification has long been important for lenders, policymakers, and courts — and is increasingly so as technological advances facilitate riskier lending. Effective modification depends on the nature and extent of liquidity constraints, moral hazard, and externalities, yet identification challenges and data limitations have constrained attempts to generate pertinent empirical evidence.

Most existing studies concern households, leaving the effects of business loan modifications largely unexplored despite decades of corporate finance research on debt overhang (e.g., Myers (1977); Kalemli-Özcan, Laeven and Moreno (2022); Jordà, Kornejew, Schularick and Taylor (2022)). Existing studies largely concern policy-driven modifications in response to large aggregate shocks (e.g., Kanz (2016); Ganong and Noel (2020); Gyongyosi and Verner (2024)), despite most modifications occurring in the normal course of lending business and with remarkably high frequency.[1] Those studies tend to rely on quasi-experimental variation from policies that exclude borrowers deemed to have the strongest incentives for strategic default (Ganong and Noel 2022, p.1060), even though lenders must reckon with such borrowers. The three existing randomized control trials (RCT) on loan modifications concern unsecured consumer lending (Dobbie and Song 2020; Aydin 2024; Burlando, Kuhn, Prina, and Wilson 2025), but collateralized lending is contractually distinct (e.g., Bester (1985); Berger and Udell (1990); Gertler, Green and Wolfram (2024); Collier, Ellis and Keys (2025)) and economically important in most settings, including ours.[2]

We address these gaps by combining a collateralized debt modification RCT with rich administrative data on entrepreneurial effort and borrower incentives. Our RCT design and analysis are guided by a theoretical model that enriches the standard debt overhang framework with a borrower liquidity constraint and its interaction with the threat of repossession (specifically, with the value lost in the state of the world where the borrower defaults and the pledged asset is repossessed).[3]

The RCT is implemented by a publicly-traded lender in South Africa on a near-universe of its over 3,000 delinquent loans with \$76 million in debt outstanding as of November 2023, and compares three common approaches to loan modification. Our control arm implements the lender's standard modification: capitalizing arrears by extending maturity, keeping monthly payment unchanged. Our treatment arms start by administering the lender's standard modification and then further implement *payment reduction* (via additional maturity extension that lowers the monthly installment amount) or *debt reduction* (an interest write-down that leaves monthly payment unchanged). Per standard practice in our setting, these modifications are offered on an opt-out basis and yield take-up rates close to 100%, effectively shutting down selection channels.

The loans finance activity in an economically vital product market— minibus taxi services, the primary form of transit in many low- and middle-income countries—where small business

[1] For example, Bidder, Crouzet, Jacobson and Siemer (2024) finds that 37% of single-lender corporate loans from large U.S. banks are modified at some point.

[2] See Section 2 for details on our setting. In the U.S., more than \$14.7 trillion (out of \$18 trillion) of household debt (Federal Reserve Bank of New York 2024), and more than \$1 trillion (out of \$1.3 trillion) of small business debt, was secured by a physical asset in 2023 (Federal Deposit Insurance Corporation 2023).

[3] Our focus on the role of liquidity constraints is similar in spirit to He and Xiong (2012), as we discuss below.

borrowers pledge their business' primary productive asset as collateral and face substantial liquidity constraints and repossession risk (as detailed in Section 2.2).[4] GPS data on driving activity, from devices embedded in the financed vehicles per loan covenants, provides unusually granular and accurate measures of entrepreneurial effort.[5] Together with data on repayments to our partner lender and outside lenders, our study paints an unusually complete picture of how contract terms shape borrower behavior over a 12-month horizon.

Our model examines how liquidity constraints and strategic incentives jointly influence borrower behavior under collateralized debt. Given our access to effort data, we model strategic incentives as a debt overhang problem, since that class of models focuses on when and how high debt burden prevents borrowers from fully capturing the long-run returns from exerting or investing effort (Myers 1977). Standard debt overhang models predict that a substantial reduction in the overall debt burden is sufficient to restore repayment incentives and effort, but we show this prediction no longer holds if liquidity constraints are binding. The possibility that liquidity constraints affect *repayment* behavior has long been a consideration for empirical work on collateralized borrowing (e.g., Adams, Einav and Levin (2009); Ganong and Noel (2022); Low (Forthcoming)), yet that work does not consider effort responses and standard debt overhang models do not allow for liquidity constraints.

We show that for liquidity-constrained borrowers, reducing total debt may still fail to prevent repossession of their pledged asset, leading to inefficient business termination ex-post even when the borrower is incentivized to repay, and thereby discouraging effort and repayment ex-ante. In contrast, modifications that ease short-run cash flow pressures while leaving the total debt burden unchanged—such as our payment reduction treatment—can have a more immediate effect on effort and repayment, if debt burden is not too high. We also show that, when facing additional default costs beyond the future loss of the collateralized asset or business itself, borrowers with high debt burden may increase repayment but not effort when their available liquidity improves sufficiently. In short, our model highlights the likely importance of liquidity constraints and the threat of collateral repossession as tandem drivers of borrower behavior, and the importance of accounting for these factors when designing and targeting loan modifications (and ex-ante contracts, as we discuss below).

The crux of our model is a simple insight: if *either* strategic incentives or liquidity constraints can distort borrower behavior, then lenders (and policymakers) interested in changing behavior must seek to ensure that *both* distortions are not binding. Functionally, liquidity issues introduce a second borrower incentive compatibility constraint. This generates four sets of qualitative, testable predictions that are distinct from those produced by a standard debt overhang model. These predictions motivate our experimental design, and our focus on effort as well as repayment responses.

[4]Spillovers from the product market to the macroeconomy are plausibly substantial here, as studied extensively in mortgage and housing markets, and documented in various minibus taxi markets across the world when driver strikes impede economic activity (Eaglin (Forthcoming)).

[5]Other studies have found substantial measurement error in business self-reports of input utilization or effort and instead rely on enumerator observation (e.g., Walker, Shah, Miguel, Egger, Soliman and Graff (2024)).

Our empirical results are broadly consistent with each of them.[6]

First, we find no evidence that debt reduction raises effort or repayment on average, contrary to the standard model's prediction. Although our confidence intervals for average treatment effects (ATEs) do not exclude economically meaningful gains, the null results should not be attributed to a small treatment size: the debt reduction we study is larger than in prior work and, in percentage terms, comparable to our payment reduction treatment (12% vs. 10.8%) – which does yield significant positive ATEs. Moreover, some borrowers do respond to debt reduction, as reflected in heterogeneous treatment effects (HTEs) previewed in our third testable prediction below.

Second, we do find evidence that payment reduction increases effort and repayment, although ATE estimates for the former are noisy. Payment reduction also reduces use of outside credit lines, as measured using credit bureau data, further corroborating the importance of liquidity *per se*.

Third, we find evidence that liquidity constraints and debt overhang interact in the two ways predicted by the model. One is that payment reduction is relatively effective at changing behavior for borrowers with relatively high equity in their vehicle at baseline. The second way is that debt reduction actually is effective for borrowers with greater predicted liquidity during the experiment (using predictive approaches to assess treatment heterogeneity per Kent (2020)).[7] These two HTEs further validate our model, and the second one also pushes against interpreting our lack of ATEs on debt reduction as stemming from a low-powered treatment.

Fourth, we find evidence that repayment and effort responses diverge as predicted by our model; namely, when liquidity constraints are relaxed but the debt overhang constraint still binds. This combination is most likely to hold for two subgroups in our experiment: those who get payment reduction and have low baseline vehicle equity, and those who get payment reduction or have higher predicted liquidity and do not get debt reduction. We find strong evidence that repayment increases substantially relative to effort in these two subgroups, but not in the rest of sample, just as predicted by our theory. These results also provide an empirical signature of the importance of additional default costs beyond loss of the collateralized asset (see Section 2.2 for details), since the model shows that repayment and effort diverge only if such costs are large enough (for evidence on the importance of additional default costs in the residential mortgage market see, e.g., Ganong and Noel (2022)).

Our results also have implications for lender modification strategies. At a coarse level, they validate lenders' apparent strong revealed preference for maturity extensions (payment reduction),[8] which are not only more effective than debt write-downs but also much cheaper in NPV terms. More finely, they suggest a targeting pecking order of sorts: first target payment reduction to borrowers

[6] We pre-registered most of the empirical tests motivated by our testable predictions. Section 5 details which tests are and are not pre-registered and why.

[7] Measuring baseline equity and predicting liquidity are non-trivial undertakings, and we devote considerable attention to measurement, and related inference issues, in Sections 5.2.1 and 5.2.2.

[8] About 75% of residential mortgage modifications in the U.S. in recent years have been maturity extensions only (Federal Housing Finance Agency 2025). Our reading of Bidder, Crouzet, Jacobson and Siemer (2024) suggests that the proportion is even higher in their sample of commercial loans, if one considers only modifications that are favorable to the borrower.

with sufficient equity, then consider giving some debt reduction to any remaining borrowers with sufficient liquidity, then consider giving both payment and debt reduction to the residual group (i.e., to those with both low equity and low liquidity).[9]

Our work contributes to several literatures cited above, some of which focus on collateralized debt due to its economic importance and distinct features, and others that consider debt contracting and small business financing more broadly.

One set of contributions is to work on the effects of collateralized debt modifications. Our novel sample, data, identification, and model are each key here. We consider voluntary debt modification by a lender, not the policy-driven modifications that have been the focus of related literatures thus far. Our sample includes borrowers who are *a priori* most likely to default strategically due to having very negative equity positions, in contrast to most prior work on mortgage modifications.[10] Our data on entrepreneurial effort permits inferences about efficiency along a key new margin. Our field experiment is, to our knowledge, the first one on modifications in a collateralized debt market, and the first one on modifications in a SME lending market.[11] Together with testable predictions generated by our theoretical model, our experiment generates the suggestive guidance on modification design and targeting that we previewed above, including the new insight that modifications should take both strategic incentives and liquidity constraints into account even if *default* has only a single trigger (e.g., even if liquidity is the primary driver of default as found in Ganong and Noel (2020, 2022), and Low (Forthcoming)). We also generate empirical evidence on potential externalities for policy consideration, finding no evidence of credit risk spillovers to other lenders or of changes in risky driving behavior that could affect other road users and insurers.

We also contribute to empirical work on the drivers of collateralized loan default, by identifying effects of changes in contract terms. We connect to prior work in several ways, starting with a setup that makes our model applicable to any stage of loan contracting. Empirically, about half of the loans originated by our lender enter delinquent status at some point, and everyone in our experiment is brought current at the time of randomization. As such our empirical results identify effects on a sample of borrowers close to the margin of default. This complements prior work using randomized variation in collateral requirements on agricultural loans in Kenya (Jack, Kremer, de Laat and Suri 2023) and school fee loans in Uganda (Gertler, Green and Wolfram 2024), by randomizing debt amount and maturity for SMEs. We also complement work on consumer vehicle and mortgage markets, which has used non-randomized sources of variation in contract terms and

[9] Modifying both maturity and total debt on the same loan seems to be exceedingly rare in commercial lending practice, both in Bidder et al.'s data (only 3% of mods) and anecdotally. Indeed, our partner lender deemed it infeasible due to various operational, funding, and accounting constraints. It seems to be more common in residential mortgage lending, with about 26% of Fannie Mae and Freddie Mac loan modifications doing both over the period from 2023:Q4 through 2025:Q1 (Federal Housing Finance Agency 2025).

[10] Per Ganong and Noel (2022, p.1060): "... by construction, the prior literature does not study borrowers excluded from mortgage modifications —which often have stringent eligibility criteria designed to exclude strategic defaulters — and borrowers who are deeply underwater."

[11] Prior work estimating effects of vehicle loan modifications uses non-randomized sources of variation and focuses on consumer credit and bankruptcy provisions therein (Chakrabarti and Pattison 2019), with some incidental coverage in work on COVID-era forbearance (Cherry, Jiang, Matvos, Piskorski and Seru 2021).

also examined the role of shocks to income or collateral value (e.g., Einav, Jenkins and Levin (2012); Foote and Willen (2018); Garmaise, Jansen and Winegar (2025)). We help connect work on drivers of default to the abovementioned work on modifications and other dimensions of collateralized debt contracting discussed below, with our novel insights on the contracting implications that arise when liquidity creates an additional borrower incentive compatibility constraint and effort and repayment responses can diverge.

We also contribute to the voluminous literature on incentive problems and debt overhang (e.g., Myers (1977); Lamont (1995); Hennessy (2004); Diamond and He (2014)), by adding liquidity constraints to the standard framework while accounting for the threat of collateral repossession. Our focus on liquidity constraints is similar in spirit to He and Xiong (2012), which theoretically models the debt market for large corporations and highlights how market-level liquidity shocks can generate credit risk and exacerbate debt overhang problems. Our focus on the role of repossession builds on theory and empirics identifying incentive effects of collateral in various settings (Bester 1985; Chan and Thakor 1987; Tirole 2006; O'Malley 2021; Jack, Kremer, de Laat and Suri 2023; Gertler, Green and Wolfram 2024; Collier, Ellis and Keys 2025). Our framework challenges the standard intuition that more stringent collateral requirements necessarily improve borrower incentives, by showing that repossession, when induced by liquidity constraints, can reduce effort in a manner similar to traditional debt overhang. The intuition is straightforward: if liquidity constraints can trigger the loss of a critical business asset and the accompanying inefficient termination of the firm, entrepreneurs exert less effort ex-ante.

We also contribute to work on how financial frictions shape entrepreneurs' decision making (e.g., Kerr and Nanda (2011); Buera, Kaboski and Shin (2015)), including equipment financing (e.g., Ma, Murfin and Pratt (2022)), by focusing on how liquidity constraints mediate effort responses to contract incentives. Our linked data on financing and effort, unusual in the study of closely-held businesses, is key here. Many studies have of course examined whether and how liquidity constraints constrain small firm expansion, in developing countries and elsewhere. That work includes examination of the marginal returns to additional units of capital and labor (e.g., de Mel, McKenzie and Woodruff (2008, 2019); Banerjee and Duflo (2014)).[12] Our contribution is to show that liquidity constraints can depress utilization of inframarginal units of capital and labor, thereby suggesting that liquidity constraints contribute to the increasingly well-documented slack in input utilization (Walker, Shah, Miguel, Egger, Soliman and Graff 2024; McMillan and Kebede 2025).

In sum, our study is novel in several respects. We enrich a standard debt overhang model with a liquidity constraint and its interaction with the standard debt overhang constraint. This enriched theoretical framework generates four distinct predictions that we test empirically. We consider voluntary debt modification by a lender, not the policy-driven modifications that have been the focus of related literatures thus far. We focus on small businesses, including borrowers with the strongest incentives for strategic default, unlike most papers on mortgage modifications. We have

[12] See also Aydin and Kim (2025), which implements a RCT to examine how changes in debt capacity affect firms' borrowing and investments.

and use rich measures of borrower effort, which is typically observed relatively coarsely if at all, especially in work on SME lending thus far. We have a field experiment on collateralized debt modifications. As such we contribute to and help connect the various abovementioned literatures.

2 Setting and data overview

This section provides background information about our setting, including the product market (minibus taxi mass transit), our partner lender and the financing market (loans collateralized by the vehicles), and an overview of borrower characteristics (we defer details on our experiment sample until Section 4.2). We then provide an overview of our various data sources and key variables.

2.1 Minibus taxi market and firms

As in many developing countries, a private minibus transport market sprang up decades ago to meet excess demand for mass transit and has grown to become the modal mode of mass vehicle transport in South Africa (Statistics South Africa 2020). A typical minibus in this market is a 16-seater manufactured by Toyota (Panel (a), Figure A.1). Approximately 40% of the nation's population reports taking a minibus on a daily basis, with 80% riding at least once per year (Kerr 2017). There are an estimated 250,000 minibus taxis spanning all of the populated areas in South Africa, generating about R100 billion in revenue annually in 2021 and thereby accounting for approximately 3% of the annual GDP (Competition Commission of South Africa 2021).[13]

Minibus taxi service is indeed a hybrid between bus and taxi services. Like a bus, it runs along a defined route. Routes are defined as a path between two points in space. The start and end points on the route are taxi ranks (akin to a bus station, e.g. Panel (b), Figure A.1) and/or bus stops. Like a taxi, route service is unscheduled and there are no formal stops: passengers hail a minibus using hand signals, and the driver picks up and drops off passengers anywhere along the route, at his discretion.

Services on a given route are governed by one of about 1,200 informal taxi associations. Associations are membership-based, with membership comprised of firms licensed to operate on a given route. Through strictly enforced rules, associations control entry through permitting, and further limit competition by limiting each driver to a single primary route. Associations also completely control pricing and do not allow price competition on a given route, while changing prices only infrequently. Each association controls one or more routes, and each route is controlled by one association. Casual empiricism suggests that associations optimize as monopsonists; in any case, minibus taxi operators typically face stiff competition for customers. For example, in the City of Cape Town, one of the most populous metropolitan areas in South Africa with approximately 1,000 minibus routes, 200 routes have an operator financed by our partner lender alone. Among those, the median route has 4 operators with a standard deviation of 8 operators. Taxi operators

[13] R denotes South African Rand, with 1 US Dollar (USD) worth about 18.7 South African Rand (R) in December 2024, at the end of our experiment.

compensate associations through a large upfront fee upon acceptance into the association, and then smaller recurring fees (one can think of these as membership, licensing, and/or permitting fees). These fees cover the association's costs of managing the route (including maintaining and staffing the taxi ranks, resolving disputes, and enforcing licensing and pricing), regulatory compliance and political lobbying, and likely also some rent transfer to association leadership (Bähre 2014; Kerr 2018).

Minibus taxi businesses look much like most businesses across the world: they are small, closely-held, informal, and owner-operated.[14] Many and perhaps most businesses are comprised of a single owner-operator, although it is not uncommon for an owner to hire an additional driver.[15] Anecdotally, most owners have ambitions to expand, and some owners do end up managing a small fleet of minibuses and drivers. The vehicle (or vehicles) is the firm's primary productive asset, and to a first approximation its only one. New vehicles cost about R500,000 at time of purchase. As detailed below, vehicles are almost always financed and pledged as collateral. Despite having a relatively simple balance sheet (see Section 2.2), and the associations' control over entry and pricing, the firm's optimization problem is non-trivial. Most to the point for our purposes, owners and operators do have control over several other key aspects of the business, including when and how aggressively to drive on their primary route, whether to petition the association for permission to do off-route trips (which are usually long-haul and outside of commuting hours), vehicle improvements and maintenance, borrowing and other aspects of cash flow management, and loan repayment. We discuss implications for measuring entrepreneurial effort below in Section 2.3.

Another key characteristic that minibus taxi firms share with small businesses across the world is that they are quite liquidity constrained. Given their informality, they borrow through personal loans taken by the owners. Most owners have limited or checkered credit histories, and this is reflected in low credit scores (mean= 622 and standard deviation (SD)= 25 at origination in our lender's portfolio).[16] Unsurprisingly then, minibus taxi credit is difficult to access, and expensive and collateralized for approved applicants. The firms who do make it into our lender's portfolio have limited access to working capital financing, with many lacking a credit line at all (per credit bureau data), and most firms who do have one at high utilization (see Section 5.2.2).

As such collateralized debt is key, and the minibus-secured loan (or loans) comprises the bulk of liabilities for minibus-operating firms: an estimated 82% of total loan balance outstanding in October 2023 for the lender's portfolio. This seems to be fairly consistent with the evidence for

[14]This is similar to the US, where the modal small business is owner-operated (see Goetz, Hyatt, Kroff, Sandusky and Stinson (2025) and `https://www.fedsmallbusiness.org/reports/survey/2023/2023-report-on-nonemployer-firms`), and to the minibus sector in other countries (Kelley, Lane and Schönholzer 2024; Björkegren, Duhaut, Nagpal and Tsivanidis 2025).

[15]We abstract from owner-driver contracting due in part to data limitations and in part to conventions in debt overhang models, which typically make no distinction between firm owners and employees. Conversations with owners, drivers, and associations support the implicit assumption that owner and driver incentives are fairly well-aligned, due to relational contracting, owner and association monitoring and information sharing, and limited outside options for drivers. Any improvements in owner-driver contracting (Kelley, Lane and Schönholzer 2024) likely would bring our model even closer to reality and strengthen our empirical results.

[16]As in the U.S., credit scores in South Africa range from 300 to 850 and rank borrowers by their estimated probability of staying current.

small businesses across the world (Beck, Demirgüç-Kunt and Maksimovic 2008), and in South Africa, where an estimated 75 to 80% of household and small business debt is secured by a physical asset (World Bank 2022).

2.2 Financing: Our lender, the market, and loan terms

Our cooperating lender is one of the five largest minibus financiers in South Africa, with about 32,000 loans and R11.6 billion in principal outstanding in its portfolio as of October 2023. The lender has operated since 2006, has been publicly traded since 2012, and has a market share of about 15%. The lender finances both new and used minibuses, with new comprising about 70% of the portfolio. Alongside this core financing business, the lender also insures, sells, and repairs vehicles.

Loan proceeds are disbursed only after verifying adequate vehicle insurance coverage and installation of a global positioning satellite (GPS) telemetric device in the minibus. Thus far, the lender has primarily used GPS to locate vehicles in the event that repossession is warranted. As detailed in the next sub-section, we use this data to measure borrowers' entrepreneurial effort.

Figure A.2 shows the distribution of origination loan terms in the lender's portfolio at baseline. Panel (a) shows that most loan amounts fall in the R350,000 to R600,000 range, with a median loan size of approximately R485,000 (mean= R487,000 and SD= R67,000). Panel (b) shows that about two-thirds of loans have a contracted maturity of 72 months, with 60, 66 and 84 months making up most of the remaining sample. There are no prepayment penalties, yet prepayment in full is uncommon; only 3.5% of the 13,300 loans originated in 2016 and 2017 and scheduled to mature before our experiment were prepaid in full. Downpayments are modest, usually in the 0 to 5% inter-quartile range (overall mean= 1.4% and SD= 5.4%), as borrowers typically have few liquid assets to offer and lenders thus choose to deal with credit risk through screening, high interest rates, monitoring, and repossession of the pledged, financed vehicle (as detailed in the rest of this sub-section). Figure A.2 Panel (c) shows that interest rates are high relative to prime collateralized loans, with a mean of 21.6% (SD= 3.5%).[17]

Lending in this market is indeed quite risky. Even after rejecting most applications following rigorous underwriting,[18] our lender is left with a borrower pool characterized by the low credit scores, limited access to working capital financing, and few liquid assets as described above. The low downpayments (high loan-to-value) and minibus taxi product market conditions described above further imply that debt service is challenging for borrowers, and indeed the lender's underwriting usually projects its borrowers to have high debt-to-income ratios. Low downpayments, together with discrete depreciation from the "drive-off-the-lot" effect, create endemic concerns about strategic default from deeply underwater equity positions in the collateralized asset (see Sections 4.2 and

[17] These rates are high in real terms in as well, as the average annual national inflation during 2019-2023 was 5%.

[18] As in the subprime consumer auto lending market in the U.S., lenders screen and underwrite applications with risk-based pricing models that consider credit history, vehicle condition, and driving records (Einav, Jenkins and Levin 2012; Jansen, Pierce, Snyder and Nguyen 2024). Our lender also requires a detailed business plan that includes the proposed route and affiliated taxi association.

5.2.1 for more details).

Default rates are high, as one would expect given the borrower and loan characteristics documented above. The share of loans 90+ days delinquent stood at 18% at the start of 2021, gradually increasing to about 22% by mid-2023. In the universe of 13,300 loans originated by the lender in 2016 and 2017 and originally scheduled to mature before the start of our experiment in November 2023, 57% became 90+ days delinquent at some point and 47% had their vehicle repossessed. Repossession is costly and rarely leaves the lender whole (due to the negative vehicle equity positions noted above, lengthy proceedings, legal and operating costs, and lost goodwill with associations). This motivates lenders, including ours, to do modifications (recall our discussion in Section 1 on the frequency of loan modification in various debt markets). Our lender's dissatisfaction with the performance of its standard modification (see Section 4.1 for details) made it receptive to experimenting with the more aggressive and theory-driven modifications we analyze here.

Default rates are consistently high despite strong incentives for the borrower to repay, starting with the threat of losing their productive asset in an economy with limited outside employment options (e.g., a 33% official unemployment rate in 2025:Q2). Borrowers face additional costs of defaulting beyond loss of the vehicle, including dealing with vigorous collection efforts by the lender, credit reporting, loss of eligibility for future loans from the lender, court costs, and potential judgments.

Costly and frequent collateralized defaults, in a setting with many signatures of prevalent liquidity constraints, motivate some of our enrichments to standard models of debt contracting. We will detail those in Section 3, after describing how we measure repayment and effort next.

2.3 Data sources and variable measurement

For both entrepreneurial effort and loan measures, we take similar approaches to summarizing multiple measures and to considering various horizons for outcome measurement. For summary purposes, we construct pre-registered, standardized indices of multiple correlated component measures that provide informative signals about the underlying construct of interest. This approach also reduces the number of statistical hypotheses tested. For measurement horizons, we pre-registered monthly and 12-month versions of each outcome. For stock variables like loan performance, we pre-register the 12-month version as the snapshot of the measure at the 12th-month. For flow variables like effort, we pre-register an aggregation rule based on the underlying construct of interest, as discussed below. Appendix Section C.1 provides details on construction of these measures and choices regarding data cleaning based on our pre-registered rules. Appendix Table A.1 documents the pairwise correlations between index components.

Minibus loan performance. The lender shared loan performance data on its entire portfolio, in monthly snapshots pulled from January 2021 through November 2024. As such we have data for loans originated as far back as January 2016.

In principle, a summary measure of loan performance should capture risk-adjusted profits. In practice, most lenders lack such a summary statistic at the loan level because they do not track all variable costs, or allocate fixed costs, accordingly. As such, after consulting with the lender, we pre-registered a standardized summary index based on normalized values of three equally-weighted component measures: (1) an indicator of being current on the loan, defined as having R100 or less past due after the payment due date; (2) arrears amount $\times$ -1; and (3) (arrears amount scaled by the required monthly payment) $\times$ -1. This prioritization of current repayment status (i.e., of avoiding delinquent states) as the key metric of success reflects common practice among lenders, their funders, and regulators.

One measurement issue is how to deal with loans that leave the lender's books during our 12-month analysis period. This is unproblematic, practically and conceptually. In practice, this event is rare (7.8% of loans in our sample),[19] accounts for only 1.8% of the potential loan-month observations, and we do not find evidence that it is affected by either modification treatment. Conceptually, since we know the repayment status at the time of leaving the book, we can appropriately fill in the blanks. Specifically, we consider a loan repaid in full as current, and delinquent if not repaid in full, per each of our component measures, in each of its months after the leaving the books (Appendix C.1 provides details). As such, we do not think of leaving the books as presenting an attrition issue per se.

Outside borrowing. We use data from Experian, one of the major credit bureaus in South Africa,[20] to measure outside borrowing, repayment on outside debt sources, and overall credit access. We use these for sample description (Section 4.2), as outcome measures (Sections 5.1.2 and 5.4), and as inputs to predicting liquidity constraints during the experiment for heterogeneous treatment effect estimation (Section 5.2.2).

Entrepreneurial effort. We use driving data to measure effort, the key borrower choice variable in debt overhang models. Such models, including ours in Section 3, conceptualize effort broadly as an activity to produce income that could in turn be used to service debt.

As noted above, the lender requires that each financed vehicle have an operating GPS device, for the purposes of tracking the location of the collateral. The device reports the vehicle's location throughout the day, allowing us to extract several signals about driving behavior (and, implicitly, the resulting income generation).

Specifically, we construct a standardized summary index of "entrepreneurial effort" by averaging

[19] A loan can leave the books for one of three reasons in our sample. One is simply reaching maturity with full repayment. This is rare in our sample by construction, because we excluded loans within three months of maturity from our experiment (in total 1.4% of our borrowers in our experiment sample have maturity of less than 12 months under the modification). A second is prepayment in full, which is typically low in the full book (Section 2.2) and also during our experiment (0.7%). A third is repossession, which is common overall (Section 2.2), as is *initiating* the repossession process in our sample. But repossession itself is rare over our 12-month experiment horizon (5.7%), because per standard modification practice every loan in our sample is brought current before random assignment (Section 4.1), the lender initiates proceedings only after a loan is 90 days or more past due, and then repossession actually occurs several months later due to legal and logistical constraints.

[20] Consumer credit bureau reporting practices and market structure in South Africa are similar to the U.S..

the following equally-weighted component measures, each aggregated to the monthly level after excluding data points based on pre-registered rules for identifying likely data recording errors: (1) distance driven; (2) time driven; (3) number of days worked in the month (i.e., number of days with non-zero trips); and (4) the number of hours spent on the job (the total duration between start of vehicle's first trip and the end of vehicle's last trip during the day).

Observational data shows a very strong association between effort reductions and sharp declines in repayment performance. Borrowers who stay current maintain stable levels of driving per our effort index, while borrowers who fall behind do so in tandem with sizable reductions in driving (about 0.3 standard deviations on average, see Figure A.3).

Conceptually, the association between driving and income generation is likely even stronger than between driving and repayment (see Sections 3.5 and 5.3 for theory and empirics on when effort and repayment diverge). Driving is minibus taxi firms' only income-generating activity. It requires both labor (the driver) and capital (the vehicle). That is, the key inputs are strong complements (perhaps even in the Leontief sense): labor utilization requires vehicle utilization, and vice versa. As such, our effort measure spans utilization of the borrower's key inputs, which aligns with how debt overhang models conceptualize income generation, at least in the short-run. One also expects that more productive firms (e.g., firms with a more reliable vehicle) will drive more at any point in time, and that those higher marginal products of input utilization will translate into more income. As such our effort measures likely reflect ongoing capital investment to at least some extent (e.g., maintenance required to keep the vehicle running), although the possibility of deferred maintenance does complicate the mapping from our data on input utilization to income over longer horizons.

Another attractive feature of our effort measure is that it captures margins of business activity over which borrowers have the most control, and likely also the most potential elasticity with respect to changes in debt contract terms. Recall from Section 2.1 that individual firms have no control over prices, and little control over where they drive day-to-day. Consequently, when and how much they drive are key margins of adjustment (including securing association approval to do long-haul trips). And there are ample reasons to think there is room for adjustment on the margin: fixed costs of capital and labor acquisition, institutional features of our product market (e.g., queuing up for passengers), and empirical evidence of slack in input utilization (Walker, Shah, Miguel, Egger, Soliman and Graff 2024) spanning decades and countries (Taubman and Gottschalk 1971).

Once a borrower leaves the loan book we can no longer measure their effort. In practice, as detailed above for repayment measures, this affects only 1.8% of potential borrower-month observations in our sample. Following our pre-analysis plan, we do not impute effort values for these observations. There are two reasons for this. First, our focus is on effort in the context of the relationship between lender and borrower. A loan leaves the books if and only if it is repaid or otherwise discharged, and at that point it is "game over" with respect to inducing effort to repay that particular loan. As such and as with repayment, we do not think of this as presenting an attrition issue per se. Second, even if we were conceptually interested in imputing missing effort, we have no evident way of doing so accurately (in contrast to repayment).

In terms of analysis, although we also present month-by-month estimates, our main focus is on the borrower's aggregate effort during the experimental period. This is motivated by the idea that the lender (and social planner) presumably cares most about the *total* amount of effort exerted by the borrower. We operationalize this for our 12-month measures by simply summing our monthly measures (and then dividing by 12 to facilitate the comparison with the monthly outcomes). We provide more details in Appendix C.1.

In sum, our effort measure maps well into the theoretical object of interest, both empirically and conceptually.

3 Debt Overhang Models Without and With Liquidity Constraints

This section outlines a simple framework that guides our experimental design, empirical tests, and interpretation. The goal is to isolate, in the most tractable way, how debt obligations shape repayment incentives and entrepreneurial effort in a setting where a collateralized asset is used to generate value. Although highly stylized and qualitative, the framework captures the hypothesized mechanisms we aim to test empirically. Appendix B works through some generalizations and finds that they leave the main qualitative insights unchanged.

3.1 Preliminaries: Effort choice without debt

To start, we consider a simple two-period setting in which an unlevered entrepreneur chooses the level of effort e to exert in running a business in the first period. As discussed in Section 2.3, e summarizes all business inputs, in keeping with, for example, corporate finance's focus on investment broadly defined. The choice of e determines first-period profits, through a standard strictly concave production function $f(e)$, with $f' > 0$ and $f'' < 0$. In the second period, the business is sold at a price that depends on the first-period profits: $V(f(e))$, where $V' > 0$ and $V'' \leq 0$.[21] Our assumptions imply that more effort today is tied to higher business equity – or, more generally, higher continuation value – in the future. This assumption is consistent with our measurement of entrepreneurial effort in the data: as discussed in Section 2.3, our measure can be interpreted as a sufficient statistic for business effort, encompassing dimensions such as when and how much to drive, vehicle maintenance, upgrades, and reputation building that are crucial for determining the continuation value of the business.[22]

This two-period setup reflects the importance of effort broadly defined: an entrepreneur's current actions affect both the immediate payoff from operating the business and its continuation value in the future. The entrepreneur discounts future payoffs with a factor $\beta < 1$ and faces a linear cost of effort, and thus chooses e to maximize the total discounted value of current and future returns

[21]For instance, an interesting special case (which is empirically relevant in many corporate valuation applications) is where V is the multiple of recent earnings at which a business can be sold (i.e., V is just a linear function).

[22]Importantly, a model characterized by the opposite assumption (i.e., $V' < 0$) would be inconsistent with the empirical results provided later.

minus the cost of effort:

$$\max_{e} \quad f(e) + \beta V(f(e)) - e$$

The first-order condition for optimal effort e^{FB} is then:

$$f'(e^{FB})(1 + \beta V'(f(e^{FB}))) = 1$$

This condition highlights that the entrepreneur of course chooses effort to equate its marginal cost to its marginal benefit, with the latter including its effects on both immediate profits and future valuation.

3.2 Standard debt overhang: no liquidity constraints

We now introduce a debt obligation $D = D_1 + \beta D_2$, consisting of payments D_1 and D_2 due in periods 1 and 2 respectively.[23] In our context and many others, D is collateralized by and finances a productive or otherwise valuable asset (a firm's minibus taxi in our setting).[24]

The entrepreneur retains the option to default in either period (i.e., to not pay D_t), at the cost of losing V. For now we make two simplifying assumptions about the cost of default. One is that the loss of V entails the loss of all equity in the business, including any equity in its vehicle. Another simplifying assumption is that there are no additional costs of default. We relax the latter assumption starting in Section 3.5, and Appendix B shows how our results hold under a general formulation of default costs.

This loss of V is directly tied to the importance of collateral repossession in our setting and many others: borrowers cannot default, thereby losing their financed and pledged asset, and still operate the business. In other words, the continuation value $V(f(e))$ can be realized if and only if the borrower does not default on the loan.[25]

The entrepreneur now trades off the value of repayment against the value of walking away from the debt obligation when deciding whether to repay and how much effort to exert. We assume the borrower first decides on the effort level e, and then sequentially makes the payment decisions P_1 and P_2. This simple framework delivers the classical debt overhang effect, as in Myers (1977): if D is too high, incentive compatibility is lost because the entrepreneur no longer captures sufficient returns to their effort. Specifically, the entrepreneur decides to exert first-best effort and repay in

[23] For ease of exposition, the model discussed in this section also assumes that $D \geq D_2$. Appendix B, relaxes this assumption and shows that the key message of the model remains qualitatively unchanged.

[24] To be clear, this is not a model of optimal lending, and we place no structural restrictions on D_1 and D_2. One way this maps to loan modification situations is if the loan was originated under different conditions (e.g., different expectations about future cash flows) which no longer hold due to a subsequent shock. Another way this maps to loan modification situations, including ours, is if a nonperforming loan is reset to current status by capitalizing arrears, and the lender is unsure about the optimal modification contract. More generally, this setup captures any stage of the contracting game where the lender is unsure about the optimum and interested in how changes to the contract or the borrower's business affect the borrower's behavior.

[25] In contrast, default on an unsecured loan is less likely to preclude continued operation of the business, given unsecured debtor protections (e.g., bankruptcy options for borrowers and limited recourse for lenders) and substantial costs of trying to enforce contracts through the legal system.

full if and only if the discounted continuation value weakly exceeds the debt burden:

$$\beta V(f(e^{FB})) - D \geq 0 \tag{1}$$

When this condition fails to hold, effort falls below the first-best level to e^{SB} (which is implicitly defined by $f'(e^{SB}) = 1$), since the entrepreneur no longer fully internalizes the continuation value.

In this standard debt overhang framework, a reduction in debt obligations D can increase both repayment and effort. Importantly, what matters is the overall size of the reduction to D, not how it is distributed between short- and long-term obligations. For instance, a lender could relax the debt overhang constraint by reducing the long-term obligations D_2 while keeping the short-run payment D_1 unchanged.[26] This equivalence will no longer hold once we introduce a liquidity constraint, as we do next.

3.3 Incorporating liquidity constraints

We now add a liquidity constraint, which is absent from standard debt overhang models despite its evident empirical importance in our setting (Section 2) and many others (e.g., Adams, Einav and Levin (2009); Ganong and Noel (2022); Low (Forthcoming)). The liquidity constraint introduces an additional driver of default, and the entrepreneur's decisions are now shaped by the combined effects of the standard model's strategic concerns and short-term cash flow limitations. This mechanism amplifies the repossession threat, as the borrower may be effectively forced to shut down the business even when its continuation value exceeds the outstanding liabilities.

Specifically, we assume that the entrepreneur is constrained to use only the cash generated by their business net of C before making any repayment in the first period, and that outside financing is unavailable.[27] The entrepreneur now exerts low effort e^{SB} if *either* the overall debt burden is too large (as in equation (1)) *or* the liquidity constraint is binding:

$$\underbrace{\beta V(f(e^{FB})) - D < 0}_{\text{Debt overhang constraint}} \quad \vee \quad \underbrace{f(e^{FB}) - C - D_1 < 0}_{\text{Liquidity constraint}} \tag{2}$$

The intuition is that liquidity constraints effectively reduce the returns to effort, by creating a state-of-the-world where there is inefficient liquidation of the business: the entrepreneur no longer reaps the future benefit of today's effort e increasing tomorrow's valuation V, if C precludes making the payment D_2 required to realize V.[28]

[26] The more general version of the model discussed in Appendix B offers an even starker result: in that setting, the borrower's behavior depends directly on both D and D_2, implying that a lender can lower long-term debt at the cost of higher immediate payments to satisfy incentive compatibility.

[27] For example, C could represent an unexpected shock to expenses or revenue. Although the cash-on-hand assumption may be overly restrictive in some settings, it holds to a first approximation in many settings, including ours (as documented in Section 2) and even in large swaths of wealthier countries (e.g., consider macroeconomics' recent focus on the high prevalence of "hand-to-mouth" households, and the use of a hard borrowing constraint assumption in many literatures). Moreover, our qualitative findings likely hold even if some marginal financing is available, as long as it is sufficiently expensive.

[28] This idea is similar to He and Xiong (2012), which studies the debt market for large corporations and highlights

This simple enrichment of the standard debt overhang problem immediately produces two striking new implications. One is that the threat of collateral repossession no longer unambiguously improves incentive alignment between lender and borrower, as in most prior work on the incentivizing effects of collateral.[29] Now the risk of repossession due to liquidity constraints– the risk of a state where the borrower, from a strategic perspective, is incentivized to repay and continue operating the business, but cannot repay– can dampen effort and repayment.

The second new implication is that the lender now has two borrower incentive compatibility constraints to consider, and this changes contracting strategy. Generally, the lender should consider whether *both* constraints are satisfied, since *either* one binding can be sufficient to trigger default. More specifically, the lender now has less assurance, relative to the standard debt overhang framework, that reducing D_2 will improve repayment (and effort). Conversely, reducing D_1 may be necessary. In fact reducing D_1 may be sufficient to achieve incentive compatibility for a liquidity constrained borrower who is not bound by debt overhang, even if D remains unchanged.

3.4 Distinct testable predictions

The two simple models with debt presented above imply contrasting predictions about the impacts of different loan modifications on borrower behavior. These predictions motivate our experimental design and related empirical tests, and guide interpretation of test results, and so we summarize them here for subsequent reference in Sections 4 and 5.

Prediction 1 (Debt reduction). *In the standard debt overhang model, substantial debt reduction will increase effort and repayments. Conversely, in a model that also incorporates liquidity constraints, debt reduction will not necessarily increase effort and repayments.*

As discussed above, in the standard model it is only high debt burden D that potentially reduces borrower incentives to exert effort and repay. A sufficiently large debt reduction thus improves both margins. Once liquidity constraints are incorporated, however, reducing D may prove ineffective: a borrower with limited cash-on-hand may still be unable to make payments, and the prospect of that liquidity-induced default reduces the borrower's incentives to exert effort and repay.

Prediction 2 (Payment reduction). *In the standard debt overhang model, reducing short-run payment obligations without affecting the present value of debt has no effect on effort or repayment. In contrast, when liquidity constraints are present, reducing short-run payments may stimulate higher repayment and greater effort.*

As discussed above, in the standard debt overhang model the only relevant quantity affecting borrower decisions is the present value of the debt D.[30] Hence, lowering D_1 affects behavior only

a different mechanism through which liquidity risk can generate credit risk and exacerbate strategic incentives in a classical debt overhang setting.

[29] As discussed in Appendix B, our framework does allow for the possibility that liquidity constraints can lead to effort in excess of first-best. We do not focus on that case because it is inconsistent with our empirical results.

[30] Here we gloss over different NPVs for the borrower and lender, due to different discount rates, because in our setting both discount rates are high enough to make the difference between them immaterial for our qualitative predictions.

if it changes D. But if liquidity constraints bind, then reducing the short-term payment (e.g., by extending maturity to lower the monthly payment as is commonly done in modifications) can increase repayment and effort, even if the present value of the debt D remains unchanged. This response reflects that, for some borrowers, default is driven purely by liquidity.

Prediction 3 (Debt Overhang $\times$ Liquidity). *Liquidity constraints interact with the standard debt overhang problem to shape borrower responses to contract terms. For example, payment reductions will be more effective for borrowers with substantial ex-ante equity (i.e., with higher V relative to D), while debt reductions will be more effective for those whose liquidity constraints are less severe.*

This prediction follows directly from the presence of two potentially binding incentive compatibility constraints in equation (2). Lowering payments relaxes liquidity constraints and increases effort for borrowers who do not face a binding debt overhang constraint. Similarly, reducing the debt burden D is more likely to be effective for borrowers who do not face a binding liquidity constraint. Conversely, if both constraints bind, modifications that target only liquidity constraints or D will not be effective, with the latter prediction also present, implicitly, in the second half of Prediction 1.

3.5 When effort and repayment diverge, and an additional prediction thereon

As we detail in Appendix B, the entrepreneur's effort and repayment decisions endogenously move together in the two models with debt presented above: they are qualitatively interchangeable. But it turns out that this tight link depends on the assumption of no additional default costs beyond losing the continuation value V.

We illustrate this in Appendix B by adding a common specification of additional default costs, specifically an immediate fixed default cost ϕ.[31] This extension leaves Predictions 1-3 qualitatively unchanged but does introduce a new dynamic: in some cases, a borrower may decide to repay while staying at a second-best effort level. The intuition is that the benefit of postponing the immediate fixed cost of default, $(1-\beta)\phi$, may exceed the cost of continuing to make the regular payment D_1, thereby inducing short-term repayment even if the borrower plans to strategically default in the long run.[32] This mechanism operates only when borrowers are not liquidity constrained; if the constraint binds, then they effectively have no choice over repayment. As such we now have:

Prediction 4 (Effort vs. Repayment). *If there are significant immediate costs of default in addition to lost continuation value, repayment and effort responses may diverge. This situation can arise only when borrowers are not liquidity constrained.*

Taking this prediction to the data helps address two important questions and further illustrates the value of effort data. First, it can help identify whether additional default costs are important

[31] This specification is meant to capture stigma, stress, and/or adjustment costs (e.g., moving in the mortgage setting, finding different work in our setting). One can think of it as also capturing the NPV of a flow of future costs from default; e.g., lost relationship value with the lender and/or reduced access to credit more generally.

[32] As we discuss in Appendix B, this case arises when we relax the regularity condition $D \geq D_2$.

driver of borrower behavior, which is an important question for quantitative modeling, policy, and mechanism design (e.g. for credit reporting). Second, it is essential for addressing the question of whether repayment behavior is a summary statistic for effort. If not, improved repayment does not necessarily indicate the absence of debt overhang. More broadly, when the two actions diverge, effort (which affects real output) is likely a better social welfare proxy than repayment (which is a transfer).

3.6 Empirical research design implications

The four predictions developed above motivate our research design and empirical tests. In design terms, each of the predictions motivates comparing effects of a modification that reduces debt to one that alleviates liquidity constraints while leaving the total debt obligation unchanged. They also highlight the value of doing so in a setting where one can measure effort as well as repayment. In terms of specific empirical tests, Predictions 1 and 2 suggest estimating average treatment effects (ATEs) of the two different approaches to loan modification, while Predictions 3 and 4 suggest three key tests for heterogeneous treatment effects (HTEs) that we detail in Sections 5.2.1, 5.2.2, and 5.3.

4 Experimental design and implementation

We worked with our partner lender to design an experiment that would yield practical guidance on modification design while helping us test the model predictions detailed in the previous section. At a high-level, we have an experiment that randomly assigns a nonperforming loan to one of three arms, with equal probability:

a. Control: the lender's standard modification, which simply capitalizes arrears while extending maturity to hold monthly payment constant. In terms of our model, this modification does not change either total debt burden D or short-term payment obligation D_1.

b. Debt reduction: standard modification and then a substantial reduction in total debt burden, leaving monthly payment unchanged. In terms of our model, this treatment is designed to reduce D while leaving D_1 unchanged. This is designed to alleviate any binding debt overhang constraint in a standard model, but will not necessarily change borrower behavior in our model if liquidity constraints bind.

c. Payment reduction: standard modification and then a substantial reduction in the required monthly installment payment, leaving debt burden unchanged. In terms of our model, this treatment is designed to leave D unchanged while reducing D_1. A standard debt overhang model predicts that this treatment will be ineffective at changing borrower behavior, while our model predicts that it could, by alleviating liquidity constraints.

The rest of this section details this design and its implementation, including our experiment sample and first-stage estimates of treatment intensity. Section 5 will then detail how we use the design to help implement tests of the four Predictions developed in Section 3.

4.1 Design, constrained randomization, and modification implementation

We test our treatments relative to the lender's standard practice for modifying nonperforming loans. This "baseline modification" capitalizes arrears into principal while extending maturity to keep the monthly payment and other contract terms unchanged. Table 1 illustrates the mechanics.

Table 1: Illustrative example of baseline modification

		Loan characteristics					
	Loan status:	Principal Outstanding (in R) (1)	Accumulated Arrears (in R) (2)	D = NPV of Debt (in R) (3) = (1) + (2)	Interest Rate (4)	Monthly Payment (in R) (5)	Remaining Maturity (in months) (6)
pre-modification	in arrears	388,570	54,040	442,610	24%	12,824	47.0
post-modification	current	442,610	—	442,610	24%	12,824	59.2

Notes: NPV calculation in Column (3) uses the loan interest rate for discounting, per standard practice.

Anecdotally, other minibus taxi lenders typically take a similar approach to modifications and other forms of debt restructuring are less common in this market. Since 2020, Fannie Mae and Freddie Mac have been promoting a mortgage modification in the U.S. that is effectively very similar to our baseline one and now represents 44% of modifications to their loans: capitalizing up to 6 months of arrears into an interest-free balloon payment due at maturity.[33]

Despite being standard practice, the baseline modification does not seem particularly attractive. Theoretically, it neither alleviates liquidity constraints nor reduces debt burden. Empirically, the lender was motivated to experiment by prior observational data showing that repayment performance tends to improve only modestly following a baseline modification.

In our experiment, we use this baseline modification as the starting point to conduct the other loan modifications. This approach has two benefits beyond serving as a natural benchmark for practice and policy. First, bringing everyone in the experiment current by capitalizing arrears sharpens the focus to the effects of changing contract terms *per se*.[34] Second, bringing everyone current potentially expands external validity: one can think of our sample as representing not just delinquent borrowers, but also borrowers who are current and close to the margin of (re-)default.

Our treatments assign a 20% reduction in debt or monthly payment relative to baseline. Sizable

[33] The interest-free component to Fannie and Freddie's new modification does technically provide some debt reduction, but its amount is almost always quite small at typical mortgage interest rates and remaining maturities.

[34] A modification brings the borrower current. Consequently, a design where the control group gets no modification, and a treatment group gets modified, faces an identification challenge of disentangling effects of changing contract terms *per se* from effects of the initial change to repayment status. For example, the latter could have effects by reducing repossession risk (a delinquent borrower is closer to repossession status), and/or through credit reporting that changes the status of the loan from (seriously) delinquent to modified.

modifications like these are common in other markets, as documented above (Bidder, Crouzet, Jacobson and Siemer 2024; Federal Housing Finance Agency 2025). Our implemented reductions are subject to three constraints imposed by the lender. First, the maturity of any loan on the lender's book could not exceed 10 years (due to the lender's funding covenants). Second, debt reduction can only be implemented by reducing the interest rate (per standard modification practice).[35] Third, the resulting new interest rate must be at least 14% (the lender's cost of capital was about 13%).

The mechanics are as follows: in November 2023, we worked with the lender to identify an experimental sample that would include nearly all of its poorly performing loans, as detailed in Section 4.2. We next performed the baseline modification on each loan in the sample before randomly assigning each loan, with equal probability, to one of the three arms: baseline modification only (control group); baseline modification + interest rate reduction (debt reduction); baseline modification + maturity extension (monthly payment reduction). The latter is a free option to pay less monthly, given the lack of a prepayment penalty (Section 5.1.2 shows that some payment reduction borrowers do in fact pay more than their new installment amount). The randomization conditions on eight strata pertinent for the constrained randomization.[36] Table 2 provides an example of a typical loan to illustrate how loan characteristics change across various arms of the experiment, and our next two sub-sections check randomization balance and the first-stage (including the lender's adherence to the randomization, and the effective treatment magnitudes given the randomization constraints).

Table 2: Illustrative Example of Randomized Treatments

	Loan characteristics			
Experimental arms:	Interest Rate (1)	Monthly Payment (in R) (2)	Remaining Maturity (in months) (3)	D = NPV of Debt (in R) (4)
Baseline modification	24.0%	12,824	59.2	442,610
Debt reduction	17.4%	12,824	48.2	394,392
Payment reduction	24.0%	11,170	79.4	442,610

Notes: Column (4) shows the NPV of the total debt owed, using the baseline loan interest rate.

After modifying the contract, the lender contacted each borrower in our experiment sample through SMS (Figure A.4) and phone calls (Figure A.5), per its standard practices. Each message linked to borrower-specific information on the modified terms. Borrowers had five business days to opt-out of the modification by repaying their outstanding arrears with the lender, and did so at rates of 0.9% in the control arm, 1.3% in the payment reduction arm, and 1.2% in the debt burden reduction arm. The lender subsequently called each borrower in the sample who did not opt-out

[35] Given our focus on debt overhang models, we value the total amount owed, D, from the borrower's perspective, using the borrower's baseline interest rate on their minibus loan as the discount rate.

[36] Specifically, we stratify on all combinations of indicators for: whether the loan would receive an above-median interest write-down amount if assigned to that arm × whether the loan would receive an above-median reduction in monthly payment if assigned to that arm × above-median baseline estimated debt-to-income ratio.

to further highlight the modified contract terms.

4.2 Experimental Sample Characteristics and Balance Tests

We worked with the lender to create an experiment sample frame of all 3,848 loans eligible for baseline loan modification per the lender's standard criteria, which are designed to target borrowers who are not so deeply in default as to be on the brink of repossession proceedings.[37] Effectively, this implies borrowers who are 30 to 270 days delinquent at baseline. For each of these 3,848 loans we then estimated what the actual modification would be under each arm of our experiment, per the constrained randomization, and limited the experiment sample to the 3,186 loans that would be eligible to receive any of the three modifications regardless of their eventual random assignment.

Total debt outstanding at baseline for the loans in our experiment was R1.4 billion, and Table A.2 presents summary statistics and balance tests for our key variables just prior to loan modifications. Column 1 (and Column 2) shows control group characteristics, with key summary statistics including means (SDs) of: interest rate 24% (3%), principal outstanding R388,570 (R130,540), arrears outstanding R54,040 (R26,060), remaining maturity 47 months (18 months), monthly loan installment R12,996 (R1,730), and loan-to-value (LTV) of 0.88 (0.22).

The heterogeneity in LTV is key in two respects: it provides the basis for our pre-registered heterogeneous treatment effect test of the first part of Prediction 3, and it shows that we have many borrowers who are (deeply) underwater on their collateralized asset (see also Figure A.6). The latter point suggests that our sample includes better representation of borrowers with the strongest incentives to strategically default– those with a binding debt overhang constraint, in our models' terms– than studies of mortgage modifications (please recall our footnote 10). We also note that mean vehicle age is 3.4 years (SD= 2.55 years), loans were originated 2 years ago on average (SD= 1 year), 76% of borrowers are men, mean borrower age is 51 years old (SD= 11 years), and 30% of borrowers have more than one loan with our lender.[38]

Columns 3 and 4 of Table A.2 report estimates from two separate multivariate OLS regressions of assignment to debt reduction or payment reduction treatment on all of the row variables, conditional on randomization strata fixed effects. We find only one coefficient with a p-value of ≤ 0.05, which is about what one would expect to find by chance. The p-value from an F-test of the joint significance of the regressors is 0.56 for the debt reduction arm (column 2) and 0.75 for payment reduction (column 3). All told, these balance checks confirm the prior of a successful randomization.[39]

[37] The standard eligibility requirements exclude loans that: (i) had arrears less than 1x or greater than 9x their required monthly payment; (ii) had received a baseline modification in the past; (iii) had outstanding maturity of $\leq$ 3 months or $\geq$ 118 months; (iv) were currently undergoing repossession proceedings; or (v) had a vehicle with GPS device that was no longer reporting data.

[38] Because our model abstracts from the case where the borrower has more than one loan with the lender and optimizes jointly across all loans, and our experiment does not prevent such borrowers in our sample from getting different treatments, we check robustness to dropping multiple-loan borrowers or just different-treatment borrowers. We find that our key inferences regarding randomization balance and treatment effects are largely unchanged after dropping each of these sub-samples.

[39] The relatively large estimates on the baseline interest rate in Table A.2 are likely due to its collinearity with one of the randomization stratum. Indeed, univariate comparisons produce the expected results: relative to the control

4.3 Empirical strategy and first-stage

Our main treatment effect estimation specification estimates intent-to-treat (ITT) effects of debt reduction and monthly payment reduction as follows:

$$y_{it} = \alpha + \beta^{DR}\ 1(\text{debt reduction})_i + \beta^{PR}\ 1(\text{payment reduction})_i + \Gamma_t'\ \alpha_{s(i)} + \epsilon_{it} \tag{3}$$

where y_{it} is an outcome or first-stage measure for loan i in time t, ϵ_{it} is the error term, and $\alpha_{s(i)}$ is a vector of fixed effects for our eight randomization strata. We pre-registered estimating treatment effects over the entire 12-month experiment period (in which case t={month 12, or months [1,12]}, depending on the outcome, as detailed in Sections 2.3 and Appendix C.1), and monthly.[40] The omitted category is the control group, which is assigned to get the baseline modification only.

This specification for estimating ATEs provides the framework for implementing our tests of Predictions 1 and 2 in Section 5.1. We subsequently add additional variables to test for heterogeneous treatment effects per Prediction 3 and Prediction 4 in Sections 5.2 and 5.3.

We start here with our pre-registered tests of treatment intensity, which can be interpreted as the first-stages in our analysis. These estimates capture the net effects of the randomization constraints and few opt-outs detailed above, and of any non-compliance by the lender (we have not detected any), on the magnitude of our treatments.

Table 3 reports results from estimating equation 3 on contract terms during our 12-month experiment period (with Figure A.7 illustrating the distribution of treatment intensities in two arms). Column 1 shows that the interest rate is 6.6 percentage points (s.e.= 1.0 pp) lower in the debt reduction group compared to the baseline modification control group. This represents a 16.7% reduction in total debt (compared to our target of 20%), representing a 12% reduction in NPV. Conversely, the interest rate is unchanged in the payment reduction group, as intended. Column 2 shows that the monthly installment is lower by R1,392 (s.e.= R76.5) in the payment reduction group (control mean= R12,784). This 10.8% reduction indicates that our payment reduction treatment was more diluted by the randomization constraints than debt reduction. Monthly installment is unchanged in the debt reduction group, as intended. Column 3 shows how maturity adjusts to engineer the results in the previous columns – specifically, the debt reduction holding monthly payment constant in row 1, and the monthly repayment reduction holding debt constant in row 2. Maturity decreases by 11 months (s.e.= 0.67 months) in the debt reduction arm (row 1), and increases by 19.25 months (s.e.= 0.90 months) in the payment reduction arm (row 2).

The lender committed to leave the randomly assigned modifications in place for at least 12 months (hence our pre-registration of a 12-month horizon for estimating treatment effects), and

group, the interest rate is only 0.1 p.p. (s.e.=0.1 p.p.) lower in the debt reduction arm and 0.1 p.p. (s.e.=0.1 p.p.) lower in the payment reduction arm, with a p-value of 0.49 on the difference between debt reduction and payment reduction arm.

[40] We estimate heteroskedasticity-consistent (Huber-White robust) standard errors and never use more than one observation per unit of randomization (the loan) when estimating treatment effects. Because we do have some borrowers with multiple loans in our sample, we have also checked whether inferences change if we cluster at the borrower level. They do not.

month-by-month first-stage estimates confirm that effects on the key contract terms are largely unchanged over this horizon (Figure A.8).

These loan modifications are economically significant. On average, the payment reduction corresponds to about R20,000 in lower payments annually, or about 6% of the average annual national income.[41] Borrowers in the debt reduction arm see their total debt reduced by an average of about R53,000 in today's dollars, which, as noted above is 12% of the total owed after baseline modification.[42] This reduction is substantial compared to typical incomes (see above) and median wealth in South Africa (R96,000, per UBS (2023)), and comparable to modifications studied in prior work. Dobbie and Song (2020) report that the credit card RCT they examine generates a *maximum* interest reduction of about 11.8% of the lender's NPV (presumably this would be smaller in terms of the borrower's NPV and average NPVs). Aydin (2024) does not report treatment size estimates in NPV terms, but we infer a mean interest rate reduction of about one-fourth (24%) of the baseline APR. The comparable estimate in our RCT is 27%. In the quasi-experimental literature on consumer mortgage modifications, Ganong and Noel (2020) studies a policy intervention that generates a nominal reduction in loan balances of about 20% (no present value adjustment reported).

Tying this back to the debt overhang models in Section 3, our first-stage results confirm that the payment reduction arm gets a substantial reduction in D_1 that leaves D more or less unchanged, while the debt reduction arm gets a substantial reduction in D that leaves D_1 more or less unchanged. Our next section details how we use this variation to help test the four predictions detailed in Section 3.

5 Results

This section analyzes intent-to-treat (ITT) effects of loan modifications using specifications based on equation 3, with a focus on implementing tests of the model predictions laid out in Section 3.4 and 3.5. We begin by presenting average treatment effects (ATEs) on repayment and effort, and then test for heterogeneous treatment effects (HTEs) by proxies for strengths of debt overhang and liquidity constraints. We conclude with estimates of effects on some potentially key externalities.

5.1 Average Treatment Effects

5.1.1 Minibus Loan Performance

ATE estimates of effects on minibus loan repayment provide a first test of our model's Predictions 1 and 2. Per our pre-registered specifications, Table 4 presents estimates on loan performance snapshots taken 12-months after treatment assignment, and Figure 1 presents month-by-month estimates. Recall from Section 2 that we selected the loan performance measures in consultation with lender to capture its key metrics, and normalize each of them such that TE estimates are in standard deviation (SD) units (see Appendix Section C.1 for details on variable construction).

[41] https://www.wearedevelopers.com/en/magazine/311/south-africa-average-salary

[42] The lender's NPV cost is higher, at around R83,000, when discounted at its cost of funds.

Starting with Prediction 1, the first row of Table 4 shows no evidence that debt reduction changes the average repayment behavior. The point estimates are each small (0.02 SD or less), and the maximum effect size contained in the confidence intervals is about 0.1 SD. This pattern is inconsistent with the standard debt overhang model's prediction that substantial debt reduction should induce improved repayment. It is thus a first piece of evidence from TE estimates suggesting that the standard debt overhang model is missing something.

Turning to Prediction 2, the second row of Table 4 shows evidence of substantial improvement in repayment performance from payment reduction. For example, Column 1 shows that our summary index improves by 0.13 SD (s.e.= 0.04). Columns 2-4 show that effect sizes across the three index components range from 0.07 SD to 0.17, further suggesting that payment reduction increases the likelihood of being current on the loan and lowers arrears accumulation. This is another piece of evidence suggesting that our model with liquidity constraints better captures the reality of our setting.

Figure 1 plots month-by-month versions of Table 4 for the minibus loan repayment summary index, and it yields similar inferences. There is no evidence that debt reduction improves repayment in any month, while payment reduction improves performance in every month except for perhaps the first month of the intervention.

5.1.2 Additional results on minibus loan repayment, and on outside liquidity

Before turning to our estimates of ATEs on effort, we present some additional ATEs on borrower repayment behavior that help characterize the role of liquidity constraints. The additional tables and figures discussed here report estimates for both modifications (as in our main tables and figures), but we focus our discussion here on payment reduction, since the additional estimates on debt reduction here straightforwardly corroborate our inference of its ineffectiveness.

Appendix Table A.3 considers ATEs on additional measures of minibus loan repayment. We did not pre-register these and have added them in response to presentation comments. The main objective of this table is to unpack our main result that payment reduction improves the summary index of the lender's key repayment metrics, which emphasize avoiding delinquent states. Columns 1-3 suggest that the improvement is driven by decreases in late payments and increases in overpayments (relative to the new, lowered installment amount) during the 12-month experimental period, not by any increase in paying the exact new installment amount.[43] This is consistent with our payment reduction design intent to give borrowers flexibility to pay somewhat less than their original installment amount, and the accompanying prior that not every borrower would exercise that option. Consistent with the nature of the treatment (i.e., a reduction in payments due to the lender), Column 4 confirms that borrowers in this group paid less than the control group over the experimental period.

Taken together, these results suggest that payment reduction generates a tradeoff for the lender:

[43] We define "exact" as falling within a 1% bandwidth around the required installment amount x 12, and define the underpayment and overpayment indicators in Columns 1 and 3 relative to that benchmark.

it improves the likelihood that a borrower remains current, but reduces to the total amount repaid relative to the control group over our 12-month horizon. Although the lender was of course hoping there would be no tradeoff, it was neither surprised nor concerned about the finding there is one. There are several reasons for this view, starting with the ancillary benefits of prioritizing current repayment status: it aligns with funder, regulator, and analyst preferences, subject to the constraints on maturity length that our payment reduction adheres to (as detailed Section 4) and are designed to ensure sufficient repayment amounts at any point in time. One underpinning of this view is the strong prior that current status bodes well for future full repayment. This prior aligns with our model's insight that maintaining current status (or, conversely, lowering the probability of default) reduces the risk of inefficient liquidation from repossession, thereby relaxing the incentive problem and inducing repayment.

Besides making it more feasible for the borrower to stay current (and thereby better aligning its incentives with the lender's), payment reduction might also free up cash for other productive uses. One such use, and one we can observe using the credit bureau data described in Section 2.3, is reducing expensive outside borrowing for working capital (most outside borrowing takes the form of unsecured credit lines, and as such is almost surely even more expensive than the minibus taxi loan). Accordingly we pre-registered estimating ATEs on two summary indices of average outside borrowing over the 12-month period. Table A.4 Columns 1 and 2 present these results, which do in fact indicate modest (0.05 SD) but potentially meaningful reductions. Column 3 suggests that this reduction in outside borrowing does not come at the expense of reduced access at the end of the 12-month period, as measured by our pre-registered summary index. That is, we do not find evidence that payment reduction changes outside options to borrow (although the point estimates do not rule out economically meaningful effects in either direction). Following our pre-registration, we also present month-by-month estimates for each of the outcomes in this table (Figure A.9), in the interest of completeness.

All told, these additional results on minibus loan repayment and outside borrowing sharpen and strengthen the inference that liquidity constraints strongly shape how borrowers respond to debt contract terms.

5.1.3 Entrepreneurial Effort

ATE estimates of effects on entrepreneurial effort provide a second test of our model's Predictions 1 and 2, and some initial hints regarding Predictions 3 and 4. Following our pre-analysis plan, we measure effort as detailed in Section 2, considering both month-by-month totals (Figure 2) and the 12-month total divided by 12 (Table 5).

As with repayment, we find no evidence of debt reduction on effort. The point estimates for both the 12-month total and each month are close to zero. The confidence intervals do include economically meaningful effects in either direction (of up to $|0.1|$), but ultimately the full picture of our theoretical and empirical results will lead us to weigh the point estimates heavily when making inferences. For now, we note simply that the lack of empirical evidence for effort responding to debt reduction is consistent with the second part of Prediction 1, and provides another piece of evidence suggesting that the standard debt overhang model is missing something.

Table 5 and Figure 2 also report ATE estimates of the payment reduction on effort, showing imprecisely estimated modest increases (e.g., about 0.04 SD on the 12-month measures). As such, whether these results shed light on Prediction 2 (that payment reduction can increase effort, in contrast to the standard debt overhang model) depends greatly on the strength of one's priors. The lack of a sharp inference here makes sense in light of the first part of Prediction 3, which suggests that inconclusive ATEs could mask a strong HTE for those with a nonbinding debt overhang constraint (combined with a null effect for those with a binding one). We will implement that test in Section 5.2.1.

Before turning to HTE estimation, note that comparing payment reduction ATEs on effort vs. repayment provides our first piece of suggestive evidence supporting Prediction 4. For example, the 12-month effect on the effort index ATE is 0.09 SD (s.e.=0.05) lower than the repayment index ATE in Table 4 (0.04 vs. 0.13), which is consistent with our prior that effort and repayment responses do in fact diverge when liquidity constraints are less binding (here due to payment reduction). Section 5.3 will provide a sharper test, based on HTEs.

5.2 Heterogeneous treatment effects

Prediction 3 spells out two tests for heterogeneous treatment effects (HTEs). These are the two sides of the coin of the model's more general prediction that borrower behavior is driven by the interaction of liquidity constraints and debt overhang (i.e., by both incentive compatibility constraints). These tests are also substitutes for an experimental arm that provides both payment and debt reduction. This is practically useful because our lender, like many others, exhibits a strong revealed preference for doing just one or the other (see footnote 9, page 4 for details).

5.2.1 Heterogeneous treatment effects by baseline vehicle equity

We now turn to implementing our test of the first part of Prediction 3: payment reduction will be more effective for borrowers with high-enough equity (more specifically, high-enough continuation value from not defaulting on the debt). This is one side of the coin of the model's more general prediction that both liquidity and standard debt overhang constraints must be non-binding for a change in contract terms to induce repayment and effort.

This test requires a measure of the borrower's continuation value V, which is of course nontrivial to estimate. Any measurement error likely attenuates estimates of the model's predicted HTE here; our tests of the model's other predicted HTEs will instead rely on randomized variation in debt overhang, through our debt reduction treatment (see Sections 5.2.2 and 5.3).

We estimate V here based on vehicle equity. The idea is that, given our model setup, V captures the (continuation) value of the borrower's business, which in turn is largely determined by the net value of the vehicle.[44] Specifically, we use baseline loan-to-value (LTV), per common practice in

[44]A more valuable vehicle not only has higher resale value but also reflects higher continuation value, in the sense

research, business, and policy. The L in LTV is simply D at baseline, as we observe directly from the lender's data. We detail our approach to measuring vehicle value at baseline, using market value data for vehicle models, in Appendix C.2. We pre-registered a high vs. low equity cutoff at median baseline LTV in our experiment sample (Figure A.6), based on convention and hopes of maximizing statistical power.[45] For brevity, we focus our discussion below on the 12-month repayment and effort index outcome measures in Table 6. For completeness, we also present the month-by-month results for these indices in Figures 3 and A.10, and the 12-month index component results in Table A.5.

Consistent with Prediction 3, Table 6 Column 1 shows a larger point estimate on the repayment index for higher- than lower-equity borrowers (0.16 vs. 0.10), although the p-value on the 0.06 SD difference is too high at 0.38 to merit strong inference without strong priors. Column 2 shows stronger evidence of an HTE on the effort index, with point estimates of -0.05 for lower-equity and 0.12 for higher-equity and the 0.17 SD difference having a p-value of 0.06. We see two alternative explanations for the stronger HTE on equity for effort than repayment. One is that, statistically speaking, we simply fail to detect a larger true HTE on repayment. Another is that the repayment HTE is in fact weaker or non-existent, and that this is consistent with Prediction 4's focus on how repayment and effort responses can diverge. It could be that, in the presence of additional default costs, liquidity drives repayment decisions while effort responds only when both the debt overhang constraint is nonbinding due to high-enough equity and liquidity constraints are relaxed due to payment reduction. We will consider this explanation, and other evidence consistent with Prediction 4, in Section 5.3.

One issue for inference is whether the HTEs in Table 6 are causal and capture the effect of variation in equity *per se*. Even if not, they are still useful in two respects. First, they still speak to our model's descriptive and predictive value, so long as our measure of equity is correlated with whatever causally drives borrower behavior. Second and similarly, they are still useful for practice and policy in identifying an observable characteristic (equity position in the collateralized asset) that can be used for targeting payment reduction to borrowers who are most likely to respond to it (Adelino, Ferreira and Oliveira 2025).

If the HTEs do capture causal mediating effects of equity, they have an additional application beyond testing model predictions and providing guidance on how to target payment reduction. Specifically, causality would imply that one can also use our theory and empirics to make inferences about how borrowers respond to shocks to equity (Bernstein 2021; Ganong and Noel 2022). We now provide some evidence consistent with this interpretation, while emphasizing that the test in next subsection (Section 5.2.2) offers complementary evidence based on randomized increases to equity through debt reduction.

We do find some additional, circumstantial evidence that the HTEs presented in Table 6 capture a causal effect of equity *per se*. To get a sense of potential alternative interpretations and how to

that the higher value should reflect a higher functionality for the business purpose.

[45] Our model is qualitative and hence does not provide clear guidance on how much equity is "high enough".

address them empirically, we start by using an R-squared decomposition model (Huettner and Sunder 2012) of baseline LTV. This exercise estimates the contributions of a rich set of vehicle, borrower, and loan characteristics, at origination and/or baseline, to variation in baseline LTV (see Appendix C.2 for details).

The R-squared decomposition estimates that remaining maturity at baseline explains 51% of the variation in baseline LTV. None of the other nine variables (which are parameterized categorically where appropriate) explains more than 13%, presumably due to the lack of major variation in other vehicle characteristics, contract terms (including downpayment), and borrower characteristics (please refer back to Section 2 for more details). Remaining maturity could indeed proxy for unobserved business characteristics that drive our HTE results, rather than or in addition to equity *per se*. Borrowers with shorter remaining maturity in our experiment are those who defaulted later in the life of the loan (conditioning, as we do on maturity at origination), and so they might for instance be higher-quality entrepreneurs.

As such we take the approach of controlling for remaining maturity, under various functional form assumptions (see Appendix C.2 for details), to test the robustness of our main, pre-registered specification– and specifically of our key HTEs in Table 6.[46] Tables A.6 and A.7 report these robustness checks.[47] Their HTEs are similar qualitatively (p-values on effort HTEs actually fall a bit relative to our main specification) and larger quantitatively (e.g., our new estimates imply TEs on effort for high-equity borrowers ranging from 0.14 to 0.32, vs. 0.12 in our main specification).

Summing up, Table 6 indicates support for the first part of Prediction 3 – payment reduction is more effective for higher-equity borrowers – and for the model's broader prediction that borrowers will exert more effort only if both liquidity and debt overhang constraints are not too binding. We also find support for a causal interpretation of higher equity's role in relaxing the debt overhang constraint. The next sub-section develops additional evidence that a causal interpretation is warranted.

5.2.2 Heterogeneous treatment effects by predicted liquidity (repayment probability)

We now detail how we implement our sharpest test of the second part of Prediction 3: debt reduction will be more effective for borrowers with enough liquidity. This is the other side of the coin of the model's more general prediction that, for a change in contract terms to induce repayment and effort, both liquidity and standard debt overhang constraints must be non-binding. Note that this test relies on experimental variation in the debt overhang constraint, through variation in D

[46] We did not pre-register these robustness checks because we did not know ex-ante whether we would find an HTE on baseline equity. Subsequent to finding one, we received useful comments in presentations that led us to the decomposition and controlling-for-maturity exercises implemented here.

[47] We conduct three robustness checks. First, we replicate the heterogeneity table controlling for remaining maturity (standardized to mean 0, SD 1), including interactions with treatment dummies (Table A.6, Panel A). Second, we relax linearity by treating maturity as quartiles (Table A.6, Panel B). Because this saturated specification prevents separate estimation for high- and low-equity borrowers, we instead focus on their relative difference. Third, we re-estimate the baseline model excluding borrowers in the tails of the maturity distribution (Table A.7), where omitted-variable concerns are greatest.

(equation 2), instead of the baseline variation in V used in the previous sub-section. Accordingly, our estimates here both test our model and provide further evidence on the causal role of borrower equity position in the collateralized asset in creating or relaxing debt overhang constraints on borrower behavior.

In principle, one way to measure liquidity for the purposes of implementing this test would be to use baseline variation in access to outside credit lines. In practice, we lack sufficient variation for such a test to be powered.[48] Instead, we proxy for liquidity by predicting repayment likelihood absent any modification. This proxy makes sense in light of the evidence presented thus far on the importance of liquidity for repayment. The idea is to estimate, for each borrower i, a parameter $\mathcal{L}_i$ capturing the probability that any liquidity shock that drove i into default status, and hence into our experimental sample, is short-lived. Borrowers with higher $\mathcal{L}_i$ are more likely to resume full payments during the experimental period, even without payment reduction or debt reduction. Through the lens of our model, these are borrowers who may have been liquidity constrained in the recent past but are likely to become less constrained in the near future.

We thus expect borrowers with higher estimated $\mathcal{L}_i$ ($\hat{\mathcal{L}}_i$) to be closer to the standard debt overhang model and hence more responsive to debt reduction: the second part of Prediction 3 implies that the effect of debt reduction is increasing in $\hat{\mathcal{L}}_i$. Conversely, it is unclear if the response to payment reduction should be affected by $\hat{\mathcal{L}}_i$.[49]

To generate $\hat{\mathcal{L}}_i$, we first predict minibus loan repayment status six months into the experiment, inputting a rich set of 28 pre-treatment borrower, vehicle, and loan characteristics into a machine learning algorithm trained on the *control group* (Appendix Section D details our predictors, prediction function estimation, and prediction procedures).[50] Our approach follows recent work on predictive approaches to assessing treatment heterogeneity (PATH), which advocates using the control group to model counterfactual outcomes for treated units (Kent 2020; Chernozhukov, Demirer, Duflo and Fernández-Val 2025).[51]

We then use our estimate of $\hat{\mathcal{L}}_i$ for each borrower to estimate HTEs. Lacking a strong prior on HTE functional form, we follow our equity HTE implementation, with 1=(at or above median $\hat{\mathcal{L}}_i$) indicating our proxy for less binding liquidity constraints. Table 7 validates our predictive approach, showing exactly the pattern we expect to find on the high $\hat{\mathcal{L}}_i$ indicator main effects: a large and strong positive correlation with repayment.

The debt reduction rows in Table 7 (and Table A.8) provide the test motivated by the second part of Prediction 3, and show that debt reduction indeed is relatively effective at inducing behavior change among the group of borrowers with less binding liquidity constraints. The repayment and

[48] Per standard benchmarks at least 85% of our sample appears to be highly credit constrained at baseline, due to lacking a line at all (36%) or to having 50% or greater utilization (49%). The latter criterion probably errs on the side of a higher-than-standard utilization threshold; for example, credit bureaus consider 30% utilization to be high.

[49] The prediction would be sharper if were confident that high $\hat{\mathcal{L}}_i$ indicated completely nonbinding liquidity constraints.

[50] Recall that every borrower in the experiment, including the control group, is made current at the beginning of the experiment by capitalizing their arrears.

[51] We did not think of using this approach prior to launching the experiment and consequently did not pre-register anything along these lines.

effort indices increase by 0.17 (s.e.=0.07) and 0.18 (s.e.=0.09) in response to debt reduction for this group, relative to the more liquidity constrained group which has much lower estimates of -0.07 (s.e.=0.05) and -0.10 (s.e.=0.06). The implied heterogeneity is economically large and statistically distinguishable from zero (with p-values of 0.01 and 0.04). We do not find similar heterogeneity for payment reduction (although the confidence intervals do not rule it out); as noted above, our model does not deliver a sharp prediction on these results.

Overall, the findings here show that the lack of debt reduction ATEs in Tables 4 and 5 masks the exact sort of heterogeneity predicted by our model: borrowers who face less binding liquidity constraints actually do behave as the standard debt overhang model predicts. They also point to a causal role for debt reduction in relaxing the debt overhang constraint, in tandem with the results presented in the previous sub-section supporting the first part of Prediction 3.

5.3 Do effort and repayment responses diverge, and for whom?

Section 5.1 documents positive average treatment effects of payment reduction on repayment but null (or at least statistically weaker) effects on effort. This divergence between repayment and effort is consistent with Prediction 4. That prediction has a sharper implication that we explore here: we expect to see divergence specifically among borrowers who are less liquidity constrained but do face the debt overhang constraint.

We test this hypothesis by examining whether repayment and effort responses do in fact diverge in our two experimental subgroups most likely to face a binding debt overhang constraint and nonbinding liquidity constraint. One such subgroup consists of borrowers with low baseline equity (i.e., for whom debt overhang constraint is likely to bind) receiving a payment reduction (i.e., for whom liquidity constraints are relaxed). Returning to Table 6, one sees strong evidence of the predicted divergence: we observe a 0.10 SD (s.e.=0.05) increase in minibus repayment index (column 1), with no evident effect on effort (column 2). The estimated 0.15 SD difference in TEs between the repayment and effort indices has a p-value of 0.03, suggesting that we should be confident in inferring that repayment and effort responses diverge as predicted by our model.

The second subgroup is comprised of borrowers who are less bound by liquidity constraints and did not receive debt reduction: this includes those receiving payment reduction per random assignment (a treatment that reduces liquidity constraints), and those with higher $\hat{\mathcal{L}}$ in the control group (who faced less binding liquidity constrains to begin with). And since these borrowers did not receive any debt reduction, the debt overhang constraint is more likely to bind for them. Table A.9 estimates TEs for this subgroup and again finds strong evidence of our model's predicted divergence.[52] The repayment index increases by 0.25 SD (s.e.=0.04) while the effect on the effort index is 0.01 (s.e.=0.05).[53]

[52]We only recent conceived of this test and consequently did not pre-register it.

[53]Table A.9 also, as expected, almost exactly reproduces the finding in Table 7 that debt reduction is relatively effective for borrowers with less binding liquidity constraints. Although the results on the debt reduction TE across the two tables may appear different at first glance due to our use of different specifications across the two tables, they are in fact almost equivalent. For example, note that the implied debt reduction TE on the repayment index

Overall, we find strong empirical support for Prediction 4. This further validates the importance of liquidity constraints and its interaction with collateral repossession in shaping borrower responses to contract terms. It also provides an empirical signature that additional default costs (i.e., costs above-and-beyond repossession of the collateralized asset) are present, since our model only generates divergence between repayment and effort responses when such costs are present (Section 3.5).

5.4 Key externalities

Although a full welfare analysis lies beyond the scope of this paper, we do consider the possibility that loan modifications generate important externalities. We test for two such effects: on repayment to other lenders, and on risky driving (which potentially affects other road users, and insurers). For each of these two outcome families we construct a pre-registered standardized 12-month summary index that averages across monthly observations (please see Appendix Section C.1 for details on component variable definitions and index construction).

The corresponding results are presented in Appendix Table A.10. Columns 1 and 2 show no strong evidence of payment reduction ATEs on repayment to outside lenders, although there is some suggestion that debt reduction has a negative effect. Columns 3 and 4 show no evidence of strong effects on risky driving.

All told, we do not find strong evidence that loan modifications impose material costs on non-contracting parties. This supports what we understand to be the standard modeling assumption that such costs are zero, at least as a first pass.

6 Discussion

Our theory and empirics suggest several implications for modeling, practice, policy, and future research opportunities. We also discuss external validity and key caveats.

6.1 How and why liquidity constraints are key

Our theory and empirics point to various mechanisms through which liquidity constraints (LCs) affect collateralized borrower behavior. Underlying each of these mechanisms is the fundamental result that a binding LC introduces another borrower incentive compatibility constraint, in addition to the standard debt overhang (DO) constraint, that principals should take into account.

In particular, our results highlight that LCs do more than simply depress repayment by reducing ability-to-pay: they can, and in our setting do, tend to dampen the standard DO effect (Prediction 1), drive behavior even if the standard DO constraint is not binding (Prediction 2), and interact with the DO constraint when it does (Prediction 3). LCs can also depress effort, even in the short-run, when failing to repay leads to repossession of the collateralized asset in the longer-run.

for the high predicted repayment probability group in Table 7 is -0.072+0.171+0.265 =0.364, while in Table A.9 it is -0.072+0.411 =0.339. The two estimates are not exactly the same due to the inclusion of strata fixed effects.

This reduces incentives to exert short-run effort that produces subsequent payoffs (i.e., reduces incentives to invest, broadly defined). Our results related to Prediction 4 moreover highlight that the incentivizing effects of higher default costs apart from repossession are limited by LCs: there will be cases where a strategic borrower would repay in the short-run, effectively exercising the option to delay incurring the extra default costs, *if* they had or could get the liquidity to do so without exerting additional effort. Further investigation of these dynamics is a promising line of future research.

6.2 A shadow of repossession threat?

Our theoretical and empirical results add nuance to the long-held and empirically-validated prediction that repossession risk, and the incentive effects of collateral more broadly, reduce default. We show that there can be a countervailing force when both LCs and standard DO bind. The first half of our Prediction 3 speaks to this indirectly, and future work could test a sharper prediction with the help of exogenous variation in repossession risk: an increase in said risk actually could reduce effort and repayment for borrowers with binding LCs or low equity in the collateralized asset.

6.3 Standard debt overhang remains key

Our theoretical and empirical results related to Predictions 1 and 3 show that failures of debt reduction to "move the needle" on effort and repayment (as is the case with our average treatment effects) need not imply that the standard DO effect is absent. Rather, they may be signatures that both LCs and the standard DO constraint are binding, as our heterogeneous treatment effect results indicate.

6.4 Implications for loan restructuring practice and policy

Although our work here does not speak to optimal contracting (see our caveats discussion below), our results do suggest a pecking order of sorts for lender (and policymaker) modification strategy, in light of lenders' understandable revealed preference for NPV-neutral maturity extensions (payment reduction, leaving debt owed unchanged) over NPV-reducing debt reduction.[54] First, lenders generally should err on the side of offering payment reduction in settings where liquidity constraints likely bind. Second, lenders should consider trying to target payment reduction to borrowers with sufficiently high continuation value, as proxied for example by sufficiently high equity positions in their collateralized asset (or in their business more broadly). Third, debt reduction alone may be sufficient to induce effort and repayment from borrowers with non-binding LCs, and debt reduction and payment relief in tandem may be required to induce the desired changes in borrowers with both low equity and binding LCs (i.e., borrowers for whom both LCs and the standard DO constraint bind).

[54] Another reason lenders and policymakers may want to avoid debt forgiveness is that it seems to create greater dynamic, reputational concerns than do NPV-neutral modifications (Kanz 2016).

Our empirical results also could serve as estimates of several parameters one might include in a quantitative model used for welfare and policy analysis: elasticities of borrower behavior with respect to contract terms (or to changes in liquidity and equity positions more generally), and any externalities these responses create for other lenders, insurers, and road users.

6.5 External validity

We consider three dimensions of external validity here, and a fourth in the caveats discussion further below. First, we posit that several features of our setting provide external validity for loan originations and well-performing loans, even though our randomized variation is on loan modification terms. As noted above, one feature of our design is that everyone in the experiment gets a baseline modification that brings their loan current. As such, one can interpret our results as applying to non-delinquent borrowers who are close to the margin of default. Most borrowers likely are close to that margin in our setting — recall that 57% of our lender's loans originated in 2016-2017 entered delinquent status at some point — and many other non-prime lending markets. Another feature is that our model and empirics each speak to comparative statics that are important for any stage of loan (re)contracting: how borrower behavior responds to changes in liquidity, or to their equity position in their collateralized asset.

Second, we posit external validity to many other collateralized debt markets. Binding liquidity constraints seem to be more the rule than the exception for small businesses and consumers worldwide. More specifically, our market has much in common with the U.S. subprime consumer vehicle market that has attracted much research and policy scrutiny (Adams, Einav and Levin (2009); Jansen, Kruger and Maturana (Forthcoming)). Our setting also shares important characteristics with household mortgage markets. Our borrowers' minibus loan is almost always their primary debt obligation, as is the case for most home mortgage borrowers. Default costs are plausibly quite high— here due to lost income generation in addition to the usual costs incurred in a well-functioning credit market; in the mortgage case due to various factors (see Ganong and Noel (2022, p.1057) and references therein). Spillovers from the product market to the macroeconomy are also plausibly quite substantial, as studied extensively in mortgage and housing markets, and documented in various minibus taxi markets across the world when driver strikes impede economic activity (Eaglin (Forthcoming)).

Third, we speculate that our model will be useful for studying unsecured debt markets as well. Our intuition is that unsecured contracts create lower repossession risk for the borrower, making the specter of liquidity-induced business shutdown less relevant. If this intuition is correct, then the standard DO problem is more likely to bind in unsecured markets. This could help explain why loan modification RCTs on unsecured debt do find ATEs on debt reduction (Dobbie and Song 2020; Aydin 2024), while we do not. It also seems plausible that one-time default costs (in our model, additional default costs outside of V) are lower in unsecured markets, which could generate the testable prediction that borrower effort and repayment decisions are more tightly aligned in unsecured than secured markets. Properly investigating these intuitions is a promising line of future research.

6.6 Caveats

One potentially key caveat is that we lack empirical estimates beyond a 12-month horizon. But we posit that at least two features of our study mitigate this concern and conversely argue in favor of external validity for longer horizons and contracting parties' decision making generally. First, the parties have high discount rates. This is evident in their financing costs of course: borrowers' costs for collateralized debt averaged 24% per year and our lender's cost of capital stood at 13% per year. It is further validated by our various results on how liquidity constraints shape borrower behavior, and by our lender's planning horizon– indeed the research team had to expend substantial internal political capital to convince the lender that keeping the randomization in place for at least 12 months was worthwhile. Second, our lack of an ATE on debt reduction mitigates any concern that any short-term benefits the lender (or market) gleans from that modification might be offset by longer-term costs from signaling lack of commitment to contract enforcement.

Another key caveat, as noted above, is that neither our theory nor empirics speak to optimality. They do however potentially provide some guidance on how to structure models that one could use for welfare analysis, and some estimates of key parameters that one might include in a quantitative model.

7 Conclusion

We enrich a standard debt overhang model with liquidity constraints to guide the design and interpretation of an experiment on the near-universe of a publicly traded lender's delinquent vehicle loans to minibus taxi small businesses, in a setting where we can measure entrepreneurial effort using GPS data on driving behavior. The experiment compares the two most common approaches to substantial loan modification—payment reduction via maturity extension vs. debt reduction via the interest rate— to each other and also identifies their effects above-and-beyond a standard, less substantial modification. The standard "mod" is unattractive in theory but allows us to harmonize the experimental sample in ways that isolate effects of contract terms and provide external validity for originations as well as modifications.

The model's key insight is that liquidity constraints create an additional borrower incentive compatibility constraint that interacts with debt overhang to shape repayment and effort. These interactions produce several novel implications in theory and practice. One is the proposition that lenders and policymakers should attend to both constraints if they seek to change borrower behavior at any stage of the contracting game, even though one mechanism or the other—i.e., either liquidity constraints or debt overhang—is sufficient to trigger default on its own. Another is that the threat of collateral repossession need not align the borrower's incentives with the lender's when liquidity constraints bind, and actually can depress effort, even when the pledged asset is quite valuable to the borrower.

We take our model to data with several testable predictions about treatment effects produced by our experiment. Consistent with these predictions, we do not find strong evidence that debt reduction changes borrower behavior on average (due to prevalent liquidity constraints), while payment reduction does. The model also predicts a particular pattern of heterogeneous treatment effects that is borne out in the experiment: payment reduction is relatively effective at changing the behavior of borrowers who are less likely to have debt overhang incentives to default, debt reduction actually is effective for borrowers with more liquidity, and repayment and effort responses diverge when liquidity constraints are relaxed but debt overhang still binds. The latter result—where payment reduction increases repayment but not effort—also validates the importance of additional default costs beyond loss of the collateralized asset, since our model shows that repayment and effort diverge only if such costs are large enough. Such cases also highlight the value of effort data– when repayment is no longer a sufficient statistic for effort, data on the latter provides both a better welfare proxy and additional traction for taking theory to data.

Our theory and empirics, combined with institutional and descriptive evidence, also have implications for lender modification strategies. At a coarse level, they validate lenders' apparent strong revealed preference for maturity extensions, which are not only more effective than debt write-downs in our theory and experiment but also much less costly in NPV terms, and likely less costly in reputational terms. More finely, they suggest a targeting pecking order of sorts. First, target payment relief to borrowers with sufficiently low debt burden; e.g., those with sufficient equity in their business and/or pledged asset. Second, consider giving some debt relief to any remaining borrowers with sufficient liquidity, particularly when debt overhang is likely to bind– as seems more likely to be the case for unsecured debt, due to lower default costs for the borrower. Finally, consider giving both payment and debt relief to the residual group (i.e., to those with both debt overhang and low liquidity). Aspects of this strategy might also be applied, experimentally, at the origination stage, for example by erring on the side of longer maturities.

All told, our theory and empirics help bridge work on household and entrepreneurial finance, on loan modifications and drivers of default, and on secured and unsecured debt contracting. It also provides a potential jumping off point for several lines of future research. One is using our theoretical and empirical frameworks to study optimal contracting, perhaps in a quantitative model. Key related considerations are choices and interactions between secured and unsecured debt, and the potential role of ex-ante commitment to loan modification instead of only relying on standard ex-post approaches. Another is explicitly accounting for investment and/or uncertainty, and the attendant dynamics, including for liquidity constraints and their potential endogeneity. Yet another would scrutinize our inference that collateral disincentivizes borrower effort when liquidity or debt overhang constraints bind. A promising approach is to create or find some exogenous variation in repossession risk while holding liquidity and debt burden constant.

References

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

Adelino, Manuel, Miguel Ferreira, and Miguel Oliveira. 2025. "The Heterogeneous Effects of Household Debt Relief." *ECB Working Paper.*

Aydin, Deniz. 2024. "Forbearance vs. Interest Rates: Experimental Tests of Liquidity and Strategic Default Triggers." *Working Paper.*

Aydin, Deniz, and Olivia Kim. 2025. "Precautionary Debt Capacity." *Working Paper.*

Banerjee, Abhijit V., and Esther Duflo. 2014. "Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program." *The Review of Economic Studies*, 81(2): 572–607.

Beck, Thorsten, Asli Demirgüç-Kunt, and Vojislav Maksimovic. 2008. "Financing Patterns Around the World: Are Small Firms Different?" *Journal of Financial Economics*, 89(3): 467–487.

Berger, Allen N., and Gregory F. Udell. 1990. "Collateral, Loan Quality and Bank risk." *Journal of Monetary Economics*, 25(1): 21–42.

Bernstein, Asaf. 2021. "Negative Home Equity and Household Labor Supply." *The Journal of Finance*, 76(6): 2963–2995.

Bester, Helmut. 1985. "Screening vs. Rationing in Credit Markets with Imperfect Information." *American Economic Review*, 75(4): 850–55.

Bidder, Rhys, Nicolas Crouzet, Margaret M. Jacobson, and Michael Siemer. 2024. "Debt Flexibility." *Working Paper.*

Björkegren, Daniel, Alice Duhaut, Geetika Nagpal, and Nick Tsivanidis. 2025. "Public and Private Transit: Evidence from Lagos." *Working Paper.*

Buera, Francisco J, Joseph P Kaboski, and Yongseok Shin. 2015. "Entrepreneurship and Financial Frictions: A Macrodevelopment Perspective." *Annual Review of Economics*, 7(1): 409–436.

Burlando, Alfredo, Michael A. Kuhn, Silvia Prina, , and Brock M. Wilson. 2025. "Rehabilitating Delinquent Digital Borrowers." *Working Paper.*

Bähre, Erik. 2014. "A Trickle-Up Economy: Mutuality, Freedom and Violence in Cape Town's Taxi Associations." *Africa: Journal of the International African Institute*, 84(4): 576–594.

Chakrabarti, Rajashri, and Nathaniel Pattison. 2019. "Auto Credit and the 2005 Bankruptcy Reform: The Impact of Eliminating Cramdowns." *The Review of Financial Studies*, 32(12): 4734–4766.

Chan, Yuk-Shee, and Anjan V. Thakor. 1987. "Collateral and Competitive Equilibria with Moral Hazard and Private Information." *The Journal of Finance*, 42(2): 345–363.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val. 2025. "Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, With an Application to Immunization in India." *Econometrica*, 93(4): 1121–1164.

Cherry, Susan, Erica Jiang, Gregor Matvos, Tomasz Piskorski, and Amit Seru. 2021. "Government and Private Household Debt Relief during COVID-19." *Brookings Papers on Economic Activity*, 141–199.

Collier, Benjamin L., Cameron M. Ellis, and Benjamin J. Keys. 2025. "The Cost of Consumer Collateral: Evidence From Bunching." *Econometrica*, 93(3): 779–819.

Competition Commission of South Africa. 2021. "Market Inquiry Into the Land Based Public Passenger Transport Sector." Competition Commission of South Africa.

de Mel, Suresh, David McKenzie, and Christopher Woodruff. 2008. "Returns to Capital in Microenterprises: Evidence from a Field Experiment." *The Quarterly Journal of Economics*, 123(4): 1329–1372.

de Mel, Suresh, David McKenzie, and Christopher Woodruff. 2019. "Labor Drops: Experimental Evidence on the Return to Additional Labor in Microenterprises." *American Economic Journal: Applied Economics*, 11(1): 202–35.

Diamond, Douglas W, and Zhiguo He. 2014. "A Theory of Debt Maturity: the Long and Short of Debt Overhang." *The Journal of Finance*, 69(2): 719–762.

Dobbie, Will, and Jae Song. 2020. "Targeted Debt Relief and the Origins of Financial Distress: Experimental Evidence from Distressed Credit Card Borrowers." *American Economic Review*, 110(4): 984–1018.

Eaglin, F. Christopher. Forthcoming. "The Need for Speed: The Impact of Capital Constraints on Strategic Misconduct." *Management Science*.

Einav, Liran, Mark Jenkins, and Jonathan Levin. 2012. "Contract Pricing in Consumer Credit Markets." *Econometrica*, 80(4): 1387–1432.

Federal Deposit Insurance Corporation. 2023. "FDIC Small Business Lending Survey."

Federal Housing Finance Agency. 2025. "Foreclosure Prevention and Refinance Report (First Quarter 2025)." Federal Housing Finance Agency.

Federal Reserve Bank of New York. 2024. "Household Debt and Credit Report." Federal Reserve Bank of New York.

Foote, Christopher L., and Paul S. Willen. 2018. "Mortgage-Default Research and the Recent Foreclosure Crisis." *Annual Review of Financial Economics*, 10(Volume 10, 2018): 59–100.

Ganong, Peter, and Pascal Noel. 2020. "Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession." *American Economic Review*, 110(10): 3100–3138.

Ganong, Peter, and Pascal Noel. 2022. "Why do Borrowers Default on Mortgages?" *The Quarterly Journal of Economics*, 138(2): 1001–1065.

Garmaise, Mark, Mark Jansen, and Adam Winegar. 2025. "Collateral Damage: Low-Income Borrowers Depend on Income-Based Lending." *Working Paper.*

Gertler, Paul, Brett Green, and Catherine Wolfram. 2024. "Digital Collateral." *The Quarterly Journal of Economics*, 139(3): 1713–1766.

Goetz, Christopher, Henry Hyatt, Zachary Kroff, Kristin Sandusky, and Martha Stinson. 2025. "Business Owners and the Self-employed: Thirty-three Million (and Counting!)." *Working Paper.*

Gyongyosi, Gyozo, and Emil Verner. 2024. "Household Debt Relief and the Debt Laffer Curve." *Working Paper.*

Hennessy, Christopher A. 2004. "Tobin's Q, Debt Overhang, and Investment." *The Journal of Finance*, 59(4): 1717–1742.

He, Zhiguo, and Wei Xiong. 2012. "Rollover Risk and Credit Risk." *The Journal of Finance*, 67(2): 391–430.

Huettner, Frank, and Marco Sunder. 2012. "Axiomatic Arguments for Decomposing Goodness of Fit According to Shapley and Owen Values." *Electronic Journal of Statistics*, 6: 1239–1250.

Jack, William, Michael Kremer, Joost de Laat, and Tavneet Suri. 2023. "Credit Access, Selection, and Incentives in a Market for Asset-Collateralized Loans: Evidence From Kenya." *The Review of Economic Studies*, 90(6): 3153–3185.

Jansen, Mark, Lamar Pierce, Jason Snyder, and Hieu Nguyen. 2024. "Product Sales Incentive Spillovers to the Lending Market: Evidence from Subprime Auto Loan Defaults." *Management Science*, 70(8): 5463–5480.

Jansen, Mark, Samuel Kruger, and Gonzalo Maturana. Forthcoming. "Dealer Financing in the Subprime Auto Market: Markups and Implicit Subsidies." *Management Science.*

Jordà, Òscar, Martin Kornejew, Moritz Schularick, and Alan M Taylor. 2022. "Zombies at Large? Corporate Debt Overhang and the Macroeconomy." *The Review of Financial Studies*, 35(10): 4561–4586.

Kalemli-Özcan, Şebnem, Luc Laeven, and David Moreno. 2022. "Debt Overhang, Rollover Risk, and Corporate Investment: Evidence from the European Crisis." *Journal of the European Economic Association*, 20(6): 2353–2395.

Kanz, Martin. 2016. "What Does Debt Relief Do for Development? Evidence from India's Bailout for Rural Households." *American Economic Journal: Applied Economics*, 8(4): 66–99.

Kelley, Erin M., Gregory Lane, and David Schönholzer. 2024. "Monitoring in Small Firms: Experimental Evidence from Kenyan Public Transit." *American Economic Review*, 114(10): 3119–60.

Kent, et al. 2020. "The Predictive Approaches to Treatment Effect Heterogeneity (PATH) Statement." *Annals of Internal Medicine*, 172(1): 35–45.

Kerr, Andrew. 2017. "Tax(i)ing the Poor? Commuting Costs in South African Cities." *South African Journal of Economics*, 85(3): 321–340.

Kerr, Andrew. 2018. *Minibus Taxis, Public Transport, and the Poor.* World Bank.

Kerr, William R, and Ramana Nanda. 2011. "Financing Constraints and Entrepreneurship." *Handbook of Research on Innovation and Entrepreneurship. Cheltenham: Elgar*, 88–103.

Lamont, Owen. 1995. "Corporate-Debt Overhang and Macroeconomic Expectations." *The American Economic Review*, 1106–1117.

Low, David. Forthcoming. "What Triggers Mortgage Default? New Evidence from Linked Administrative and Survey Data." *The Review of Economics and Statistics.*

Ma, Song, Justin Murfin, and Ryan Pratt. 2022. "Young Firms, Old Capital." *Journal of Financial Economics*, 146(1): 331–356.

McMillan, Margaret S., and Hundanol A. Kebede. 2025. "Firm Capacity Underutilization and the Measurement of Productivity." *NBER Working Paper 34279.*

Myers, Stewart C. 1977. "Determinants of Corporate Borrowing." *Journal of Financial Economics*, 5(2): 147–175.

O'Malley, Terry. 2021. "The Impact of Repossession Risk on Mortgage Default." *The Journal of Finance*, 76(2): 623–650.

Statistics South Africa. 2020. "National Household Travel Survey 2020." Statistics South Africa, Pretoria, South Africa.

Taubman, Paul, and Peter Gottschalk. 1971. "The Average Workweek of Capital in Manufacturing." *Journal of the American Statistical Association*, 66(335): 448–455.

Tirole, Jean. 2006. "The Theory of Corporate Finance." HAL Post-Print.

UBS. 2023. "Global Wealth Databook." UBS.

Walker, Michael W, Nachiket Shah, Edward Miguel, Dennis Egger, Felix Samy Soliman, and Tilman Graff. 2024. "Slack and Economic Development." *NBER Working Paper.*

World Bank. 2022. "South Africa: Financial Sector Assessment Technical Note (Financial Inclusion)." World Bank.

Figures and Tables

Figure 1: Month-by-month ATEs of loan modifications on minibus loan repayment summary index

(a) Debt Reduction (b) Payment Reduction

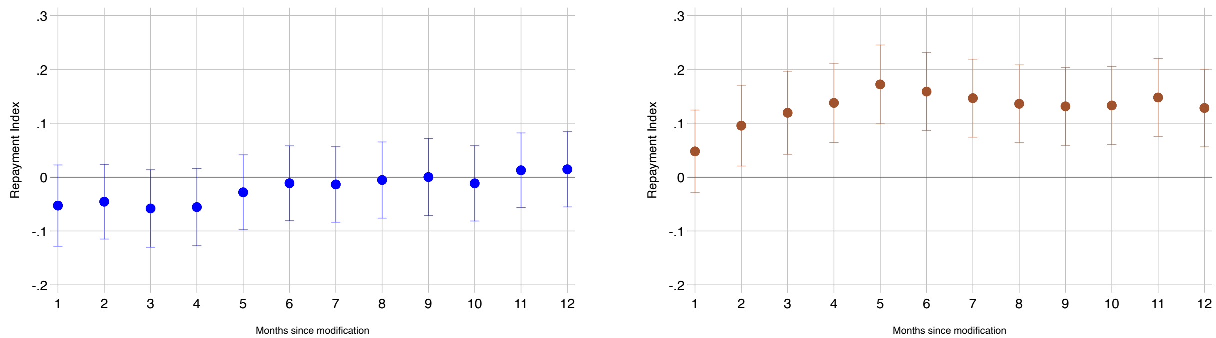

Notes: These panels report month-by-month intent-to-treat (ITT) average treatment effect (ATE) estimates per equation 3. LHS variable is standardized and thus has mean of approximately zero and TE estimates in standard deviation units; please see Section 2.3 and Appendix Section C.1 for additional details on repayment variable definitions. Vertical lines indicate 95% confidence intervals.

Figure 2: Month-by-month ATEs of loan modifications on entrepreneurial effort summary index

(a) Debt Reduction (b) Payment Reduction

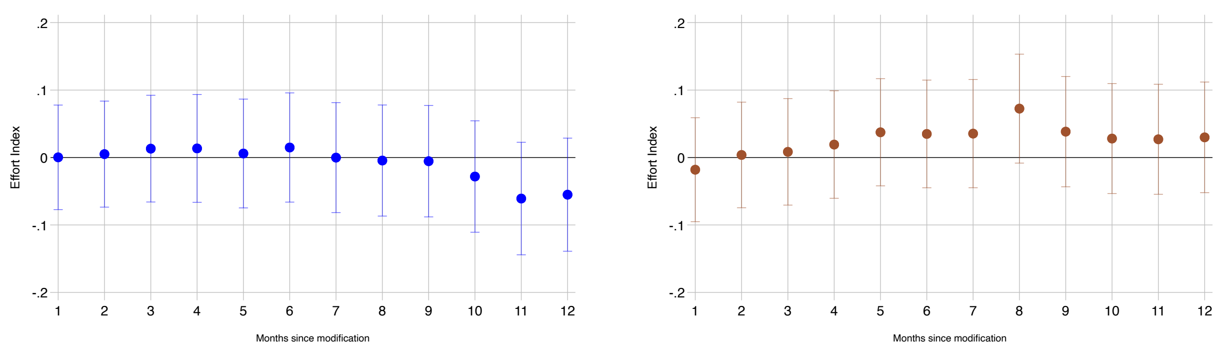

Notes: These panels report month-by-month intent-to-treat (ITT) average treatment effect (ATE) estimates per equation 3. LHS variable is standardized and thus has mean of approximately zero and TE estimates in standard deviation units; please see Section 2.3 and Appendix Section C.1 for additional details on effort variable definitions. Vertical lines indicate 95% confidence intervals.

Figure 3: Month-by-month estimates of payment reduction TEs, by baseline equity

Panel A. Minibus loan repayment summary

(a) borrowers with low baseline equity

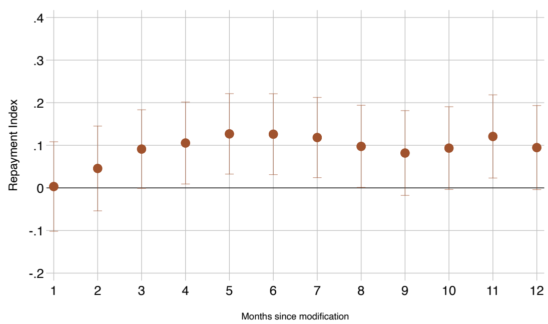

(b) borrowers with high baseline equity

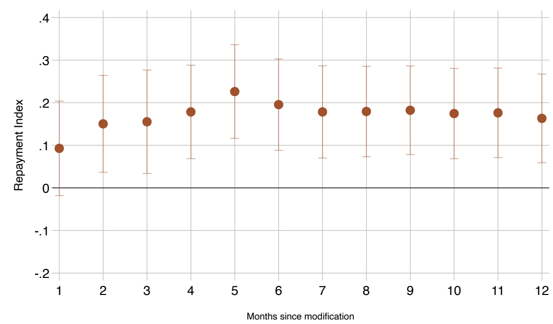

Panel B. Entrepreneurial effort summary index

(a) borrowers with low baseline equity

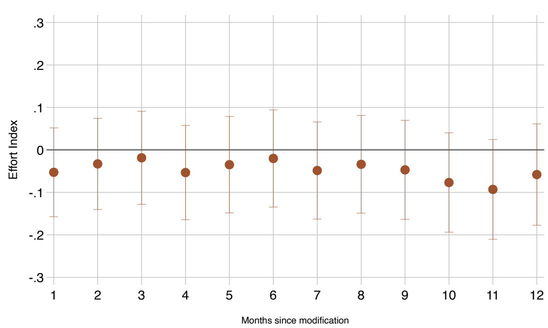

(b) borrowers with high baseline equity

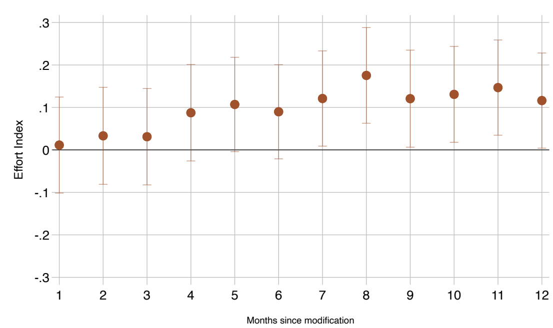

Notes: These panels report month-by-month intent-to-treat (ITT) average treatment effect (ATE) estimates per equation 3, separately for low and high baseline vehicle equity groups. The latter (former) group is defined as having measured baseline loan-to-value (LTV) less than or equal to (greater than) the median in our experimental sample. Please see Appendix C.2 for additional details on LTV measurement. LHS variables are standardized and thus have means of approximately zero and TE estimates in standard deviation units; please see Section 2.3 and Appendix Section C.1 for additional details on repayment and effort variable definitions. Vertical lines indicate 95% confidence intervals.

Table 3: First-stages: ATEs on contract terms

	Interest Rate (in p.p.) (1)	Monthly Installment (in R) (2)	Remaining Maturity (in months) (3)
β^{DR} : 1(debt reduction)	-6.58***	87.18	-11.11***
	(0.10)	(77.28)	(0.67)
β^{PR} : 1(payment reduction)	-0.12	-1,391.94***	19.24***
	(0.10)	(76.55)	(0.90)
Observations	3,186	3,186	3,186
R-squared	0.713	0.268	0.484
Control mean	23.90	12,784	54.85
$p[\beta^{DR} = \beta^{PR}]$	0.00	0.00	0.00
Strata f.e.	✓	✓	✓

Notes: Each column presents intent-to-treat (ITT) estimates of average treatment effects (ATEs) from a single OLS regression per equation 3. Dependent variable is described in the column heading and measured as an average over month-end snapshots over the 12-month experimental period (see Figure A.8 for monthly treatment effects (TE) estimates). Regressors are described in the rows. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 4: ATEs on minibus loan repayment

		Index components		
	Repayment Index (1)	1(current) (2)	Arrears Amount (3)	Arrears (Scaled) (4)
β^{DR} : 1(debt reduction)	0.014	0.021	-0.004	-0.018
	(0.036)	(0.040)	(0.042)	(0.042)
β^{PR} : 1(payment reduction)	0.128***	0.150***	-0.165***	-0.070
	(0.037)	(0.043)	(0.043)	(0.043)
Observations	3,186	3,186	3,186	3,186
R-squared	0.077	0.047	0.077	0.054
$p[\beta^{DR} = \beta^{PR}]$	0.00	0.00	0.00	0.21
Strata f.e.	✓	✓	✓	✓

Notes: Each column presents intent-to-treat (ITT) estimates of average treatment effects (ATEs) from a single OLS regression per equation 3. Dependent variable is described in the column heading and measured at end of our 12-month experiment period (see Figure 1 for monthly TE estimates). All LHS variables are standardized and thus have means of approximately zero and TE estimates in standard deviation units; please see Section 2.3 and Appendix Section C.1 for additional details on repayment variable definitions. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 5: ATEs on entrepreneurial effort

		Index components			
	Effort Index (1)	Distance driven (2)	Time driven (3)	Time Spent on job (4)	Total days driven (5)
β^{DR} : 1(debt reduction)	-0.011	-0.013	-0.038	-0.003	0.010
	(0.044)	(0.048)	(0.048)	(0.047)	(0.046)
β^{PR} : 1(payment reduction)	0.039	0.035	0.001	0.052	0.069
	(0.044)	(0.049)	(0.048)	(0.046)	(0.046)
Observations	3,178	3,178	3,178	3,178	3,178
R-squared	0.021	0.009	0.010	0.035	0.035
$p[\beta^{DR} = \beta^{PR}]$	0.24	0.30	0.40	0.22	0.19
Strata f.e.	✓	✓	✓	✓	✓

Notes: Each column presents intent-to-treat (ITT) estimates of average treatment effects (ATEs) from a single OLS regression per equation 3. Dependent variable is described in the column heading, measured as a sum over our 12-month experiment period, and divided by 12 for comparability with the monthly estimates in Figure 2. All LHS variables are standardized and thus have means of approximately zero and TE estimates in standard deviation units; please see Section 2.3 and Appendix Section C.1 for additional details on effort variable definitions. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 6: HTEs by Baseline Vehicle Equity

	Repayment Index (1)	Effort Index (2)
β^{DR+LE}: 1(debt reduction) $\times$ 1(low basline equity)	-0.033	-0.060
	(0.050)	(0.061)
β^{DR+HE}: 1(debt reduction) $\times$ 1(high baseline equity)	0.069	0.031
	(0.050)	(0.064)
β^{PR+LE}: 1(payment reduction) $\times$ 1(low basline equity)	0.099**	-0.046
	(0.050)	(0.061)
β^{PR+HE}: 1(payment reduction) $\times$ 1(high baseline equity)	0.164***	0.119*
	(0.053)	(0.063)
1(high baseline equity)	0.187***	-0.226***
	(0.055)	(0.068)
Observations	3,186	3,178
R-squared	0.093	0.026
$p[\beta^{DR+LE} = \beta^{PR+LE}]$	0.01	0.82
$p[\beta^{DR+HE} = \beta^{PR+HE}]$	0.07	0.15
$p[\beta^{DR+HE} = \beta^{DR+LE}]$	0.15	0.31
$p[\beta^{PR+HE} = \beta^{PR+LE}]$	0.38	0.06
Strata f.e.	✓	✓

Notes: Each column presents OLS estimates from a single regression. Dependent variable is described in the column heading and measured at 12-months for repayment and over 12-months for effort (please see Section 2.3 and Appendix Section C.1 for details). These variables are standardized and thus have means of approximately zero and TE estimates in standard deviation units. The effort index is rescaled for comparability to the monthly treatment effect estimates in Figure 3 (Figure A.10 has the monthly estimates for repayment). High (low) baseline equity takes the value of one if the measured baseline loan-to-value (LTV) is less than or equal to (greater than) the median value in our experimental sample, and zero otherwise. Baseline LTV is calculated as the total face value of loan outstanding for the borrowers (which is the sum of remaining loan principal plus arrears) divided by the estimated market value of the vehicle at baseline (please see Appendix C.2 for additional details). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 7: HTEs by Predicted Repayment Probability

	Repayment Index (1)	Effort Index (2)
1(debt reduction)	-0.072	-0.104
	(0.049)	(0.063)
1(debt reduction) × 1(high repayment prob.)	0.171**	0.181**
	(0.070)	(0.089)
1(payment reduction)	0.096*	0.020
	(0.050)	(0.062)
1(payment reduction) × 1(high repayment prob.)	0.082	0.036
	(0.073)	(0.088)
1(high repayment prob.)	0.265***	-0.058
	(0.053)	(0.066)
Observations	3,178	3,178
R-squared	0.112	0.022
Strata f.e.	✓	✓

Notes: Each column presents OLS estimates from a single regression. Dependent variable is described in the column heading and measured at 12-months for repayment and over 12-months for effort (please see Section 2.3 and Appendix Section C.1 for details). These variables are standardized and thus have means of approximately zero and TE estimates in standard deviation units. The effort index is rescaled for comparability to monthly treatment effect estimates. 1(high repayment prob.) takes the value of one if the predicted value $\hat{\mathcal{L}}_i$ (i.e., probability of being current on the loan six months into the experiment) is at or above the median value in our experimental sample, and zero otherwise (please see Appendix Section D for details on estimating $\hat{\mathcal{L}}_i$). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

A Appendix Figures and Tables (Online Appendix)

Figure A.1: A Toyota Quantum Minibux Taxi

Notes: The top picture shows the 16-seater minibus from Toyota that constitute 80% of the vehicles in our sample. The bottom picture shows a taxi rank.

Figure A.2: Distribution of loan terms at origination for the lender's entire portfolio

(a) Loan Amount (in R 1,000)

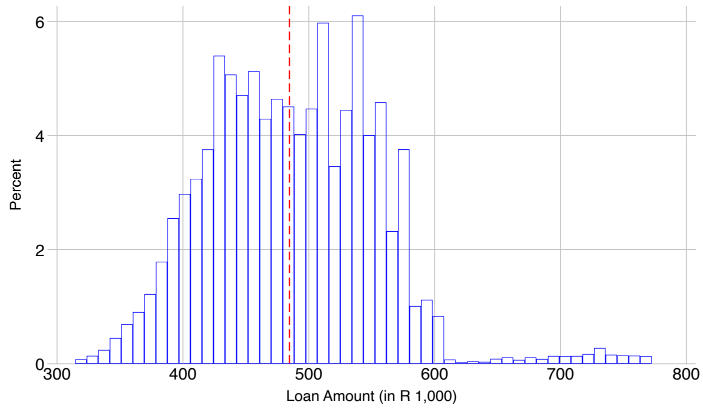

(b) Maturity in months (at Origination)

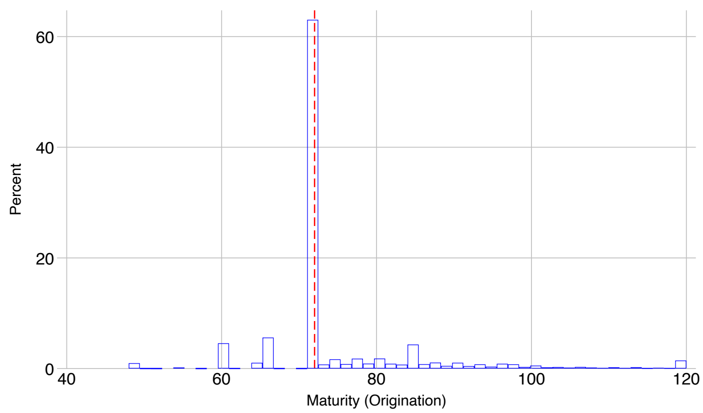

(c) Interest Rate (%)

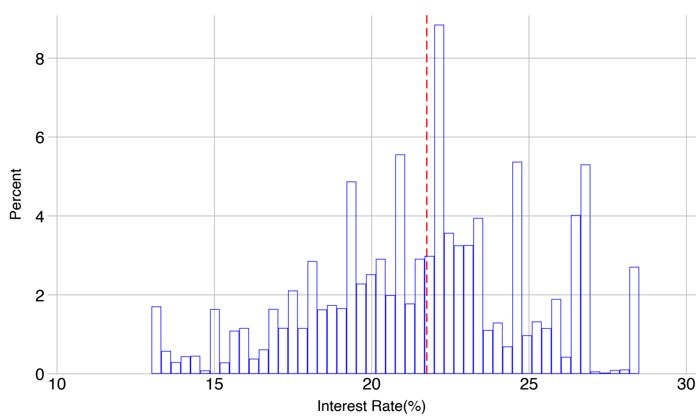

Notes: $N \approx 32{,}000$ and is comprised of all loan accounts active with the lender as of October 2023 (the month before the experiment rollout). The red vertical and dashed line on each subplot corresponds to the median value for that contract term.

Figure A.3: Co-evolution of Entrepreneurial Effort and 90+ Day Delinquency

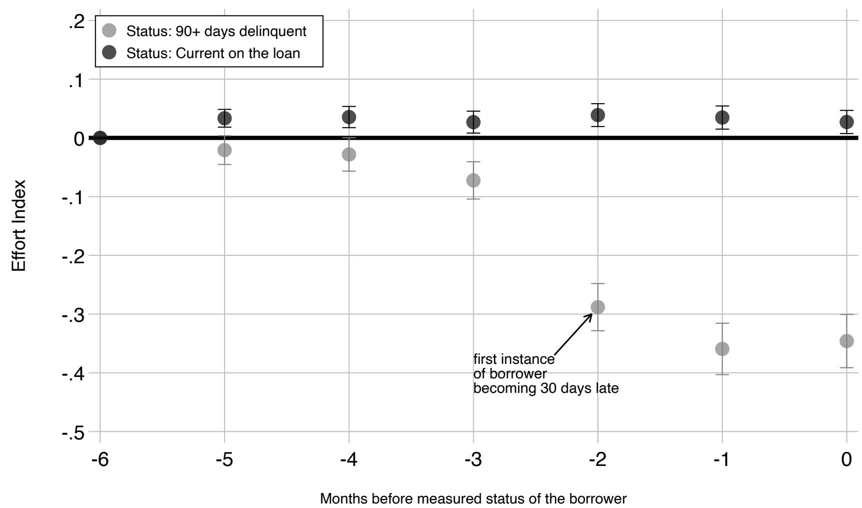

Notes: Black and gray dots plot the point estimates β^i from separate regressions for each of two repayment statuses i in month t: Effort Index$_{i,t} = \alpha_i + \alpha^i_{t_0} + \sum_{k=0}^{-5} \beta^i_k \,.\, \mathbf{1}(t = t^i_0 + k) + \epsilon_{it}$. Vertical bars show 95% confidence intervals on β^i. Repayment status is $\{1(\text{current})$ or $1(90+$ days delinquent$)\}$. The unit of observation is the measurement of i taken in each month $t_0 = [\text{June 2024}, \text{January 2025}]$, with the analysis corresponding to our experiment window because the June 2024 observations look back through December 2023. Min k=-5, because the omitted month is month 6 before t_0. For each account that is ever 90+ days delinquent in the June 2024-January 2025 window, we only include the observation from the month in which the account is first observed as 90+ days past due. For accounts current as of t_0, we include only those that have also been current in each of the preceding six months. For accounts that are current multiple times in the June 2024 to January 2025 window, we randomly select one of those monthly observations to include in the regression sample. Standard errors are clustered at the account level. We exclude accounts that are neither current nor seriously delinquent, and/or that are in our experimental sample, because here we are interested in estimating the typical correlation between effort and repayment status rather than anything particular to or inclusive of modifications.

Figure A.4: SMS messages sent by the lender per our experiment

Baseline modification arm

%Deal_Number%: is currently in Arrears. In an effort to assist you, your account has been reviewed and approved for a Term Extension. Follow the link to your document. Password is your ID number. To opt out call [redacted] *or email* [redacted]

Debt reduction arm

%Deal Number%: is currently in Arrears. In an effort to assist you, your account has been reviewed and approved for a Term Modification and Rate Reduction. Follow the link to your document. Password is your ID number. To opt out call [redacted] *or email* [redacted]

Payment reduction arm

%Deal_Number%: is currently in Arrears. In an effort to assist you, your account has been reviewed and approved for a Term Extension and Instalment Reduction. Follow the link to your document. Password is your ID number. To opt out call [redacted] *or email* [redacted]

Notes: Borrower and lender identifying information has been redacted.

Figure A.5: Scripts of calls made by the lender team per our experiment

Good day. My name is ______. I am calling from [lender]. [lender] is launching a program to assist clients who are behind on their loan payments. We see that you are struggling and are here to help. You are _____ in arrears, as of the [current date]. You have been selected to be part of our debt modification program. We would like to offer you the following offer [*EXPLAIN ASSIGNED OFFER, SEE BELOW*] and have adjusted your contract. Are you happy to proceed under these terms?

[*IF YES – please proceed to "We understand. . .", BELOW*]

[*IF NO*] - I understand you are skeptical of this offer. This offer will [*DESCRIBE ASSIGNED OFFER, SEE BELOW*] and rehabilitate your status with [lender]. Are you sure you are uninterested in taking this offer?

[*IF AGAIN NO*] - could you please let me know why you are uninterested in this offer? What sort of support from [lender] would be useful to assist you in making your monthly instalments in the short term and long term?

[*AFTER DISCUSSION – please proceed to "We understand. . .", BELOW*]

We understand that we are approaching the holiday season and this is a busy time for the industry. If you make additional cash, we encourage you to pay in more to your account to pay down your loan balance. Doing so now can you assist you later if you fall into trouble.

For the sake of clarity, this is the last offer of assistance we will extend to you. If you believe you will fall behind again in your account, please phone us immediately to inform us.

If you fall behind on your payments again, we will be forced to take immediate action to collect our arrears and repossess your vehicle.

Thank you for your time.

Term Extension [A.k.a. "baseline modification", in our paper's parlance]

We have capitalized your arrears into your account and extended your term. The new term of your loan will be _____. Your instalment will remain the same. This offer will also have a positive impact on your credit report by changing months in arrears to zero.

Interest Writedown [A.k.a. "debt reduction"]

We have capitalized your arrears into your account and reduced your interest rate to ______. That will lower your total loan interest payment by _____. The new term of your loan will be _____. Your instalment will remain the same. This offer will also have a positive impact on your credit report by changing months in arrears to zero.

Instalment Reduction and Term Extension [A.k.a. "payment reduction"]

We have capitalized your arrears into your account and lowered your monthly instalment by _____. The new term of your loan will be _____. Your new instalment will be ______. This offer will also have a positive impact on your credit report by changing months in arrears to zero.

Notes: Identifying information has been removed.

Figure A.6: Distribution of Baseline LTV in the experimental sample

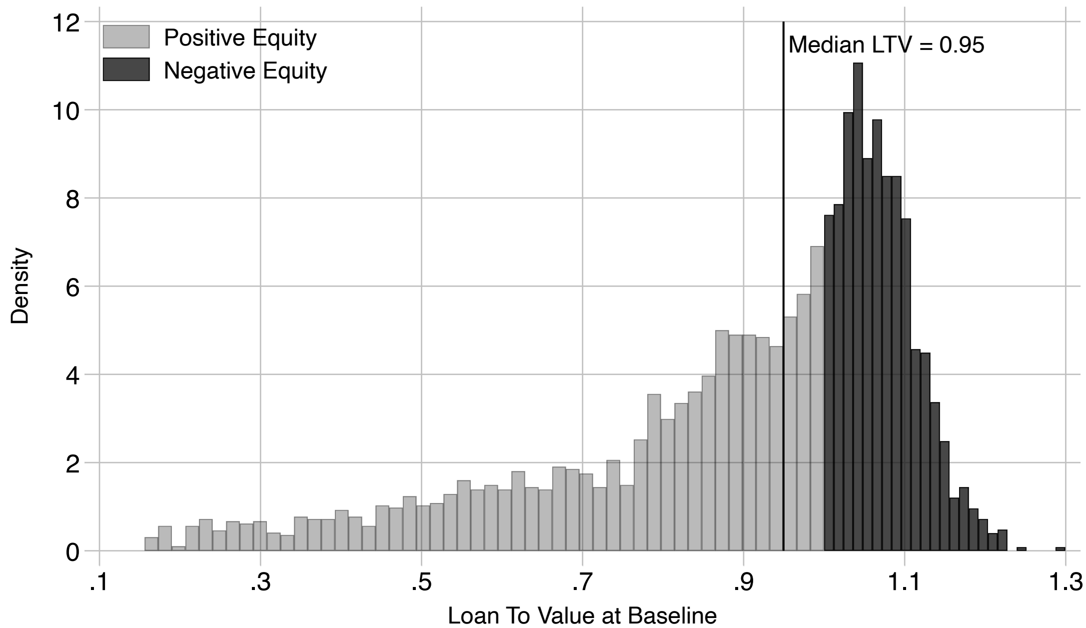

Notes: LTV (loan-to-value) is defined as the outstanding loan balance (principal + arrears) divided by our estimate of the market value of the collateral (i.e., the minibus). The vertical line represents the median value (LTV = 0.95) in our sample. Please see Appendix C.2 for additional details.

Figure A.7: Effective treatment intensity

(a) Debt Reduction

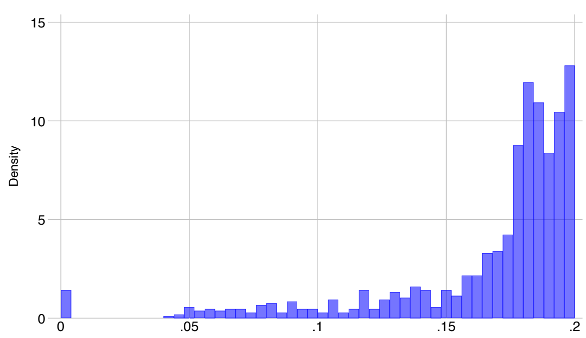

(b) Payment Reduction

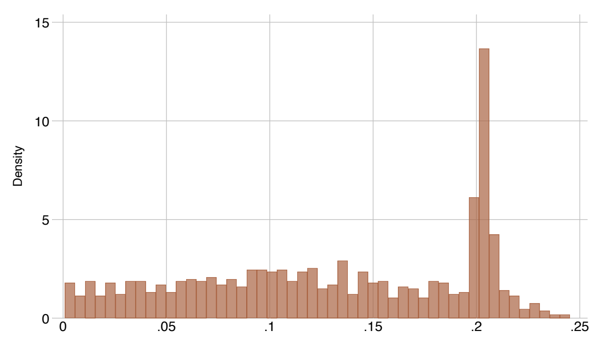

Notes: Recall that we targeted 20% reductions for each loan but randomized subject to several constraints detailed in Section 4.1.

Figure A.8: First-stages: Month-by-month ATEs of loan modifications on contract terms

Panel A. Interest Rate (%)

(a) Debt Reduction

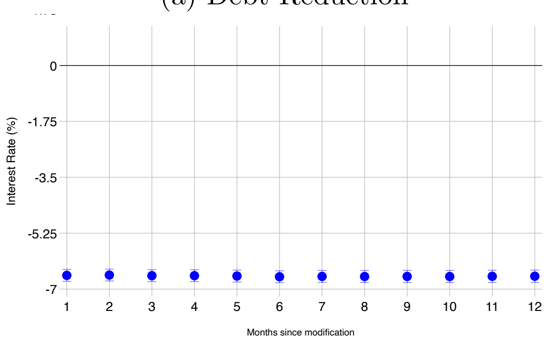

(b) Payment Reduction

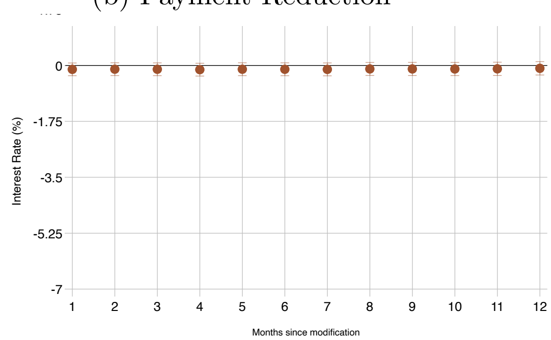

Panel B. Minimum Required Monthly Payment (Rand)

(a) Debt Reduction

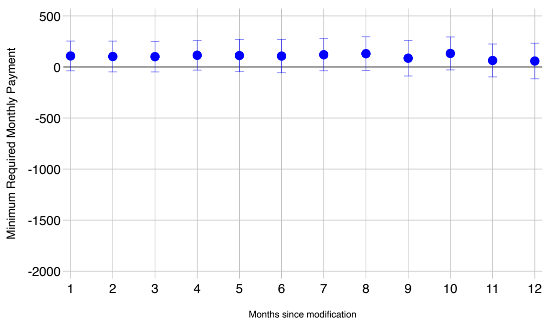

(b) Payment Reduction

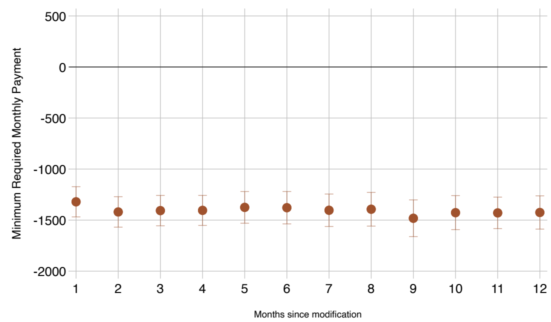

Panel C. Remaining Maturity (Months)

(a) Debt Reduction

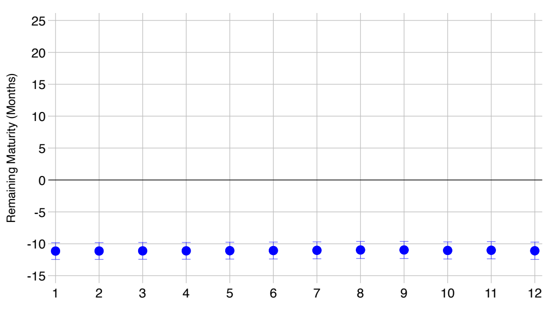

(b) Payment Reduction

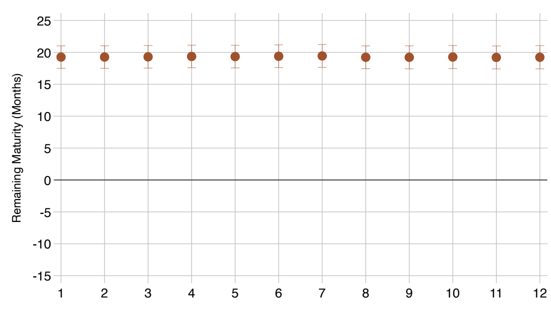

Notes: Month-by-month intent-to-treat (ITT) average treatment effect (ATE) estimates per equation 3. Please see Section 2.3 and Appendix Section C.1 for additional details on variable definitions. Vertical lines indicate 95% confidence intervals.

Figure A.9: Month-by-month ATEs of loan modifications on outside debt

Panel A. Outside Borrowing Index

(a) Debt Reduction

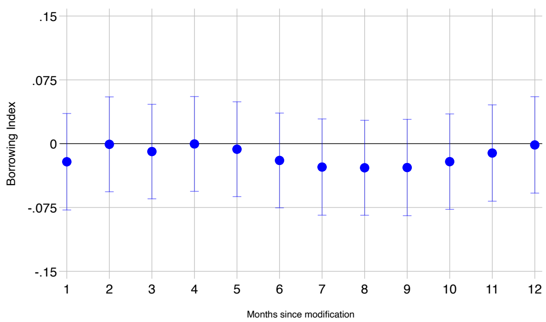

(b) Payment Reduction

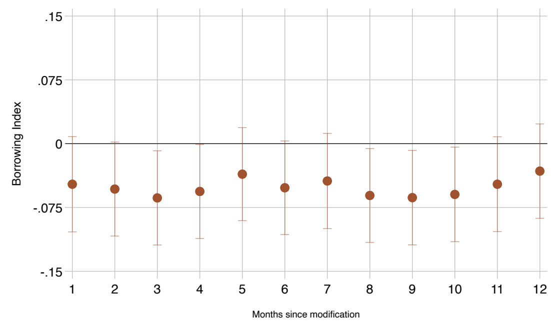

Panel B. Outside Borrowing Index (winsorized)

(a) Debt Reduction

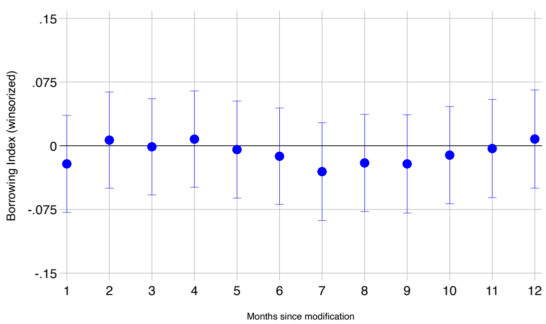

(b) Payment Reduction

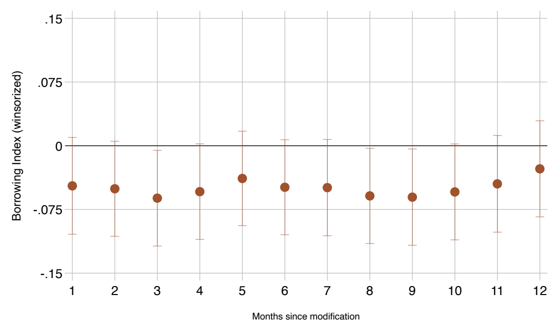

Panel C. Outside Credit Access Index

(a) Debt Reduction

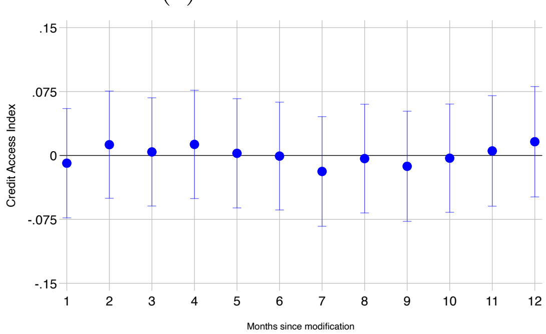

(b) Payment Reduction

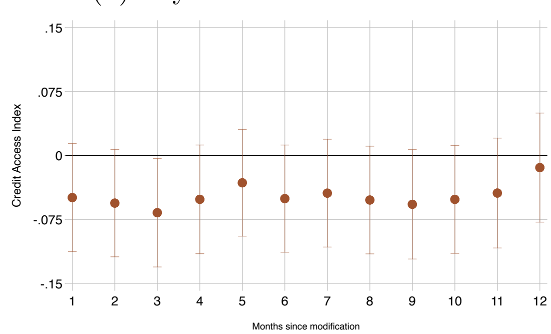

Notes: Month-by-month intent-to-treat (ITT) average treatment effect (ATE) estimates per equation 3. LHS variables are standardized and thus have means of approximately zero and TE estimates in standard deviation units; please see Section 2.3 and Appendix Section C.1 for additional details on variable definitions. Vertical lines indicate 95% confidence intervals.

Figure A.10: Month-by-month estimates of debt reduction TEs, by baseline equity

Panel A. Minibus loan repayment index

(a) borrowers with low baseline equity

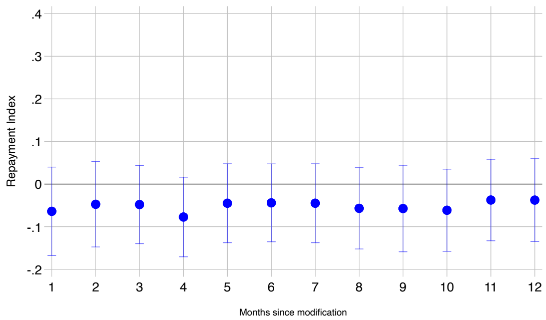

(b) borrowers with high baseline equity

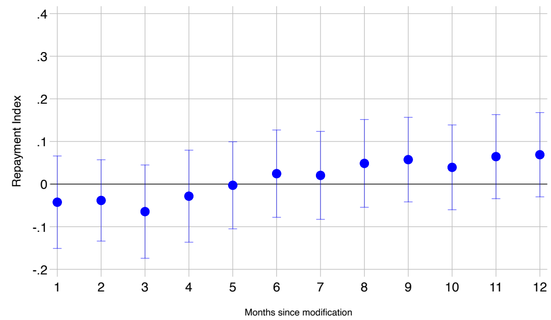

Panel B. Entrepreneurial effort index

(a) borrowers with low baseline equity

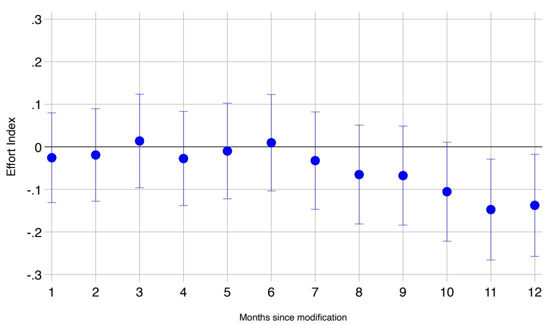

(b) borrowers with high baseline equity

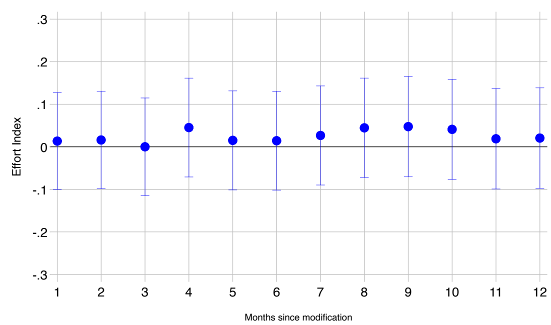

Notes: These panels report month-by-month intent-to-treat (ITT) average treatment effect (ATE) estimates per equation 3, separately for low and high baseline vehicle equity groups. The latter (former) group is defined as having measured baseline loan-to-value (LTV) less than or equal to (greater than) the median in our experimental sample. Please see Appendix C.2 for additional details on LTV measurement. LHS variables are standardized and thus have means of approximately zero and TE estimates in standard deviation units; please see Section 2.3 and Appendix Section C.1 for additional details on repayment and effort variable definitions. Vertical lines indicate 95% confidence intervals.

Figure A.11: Variance Decomposition of Baseline Vehicle LTV

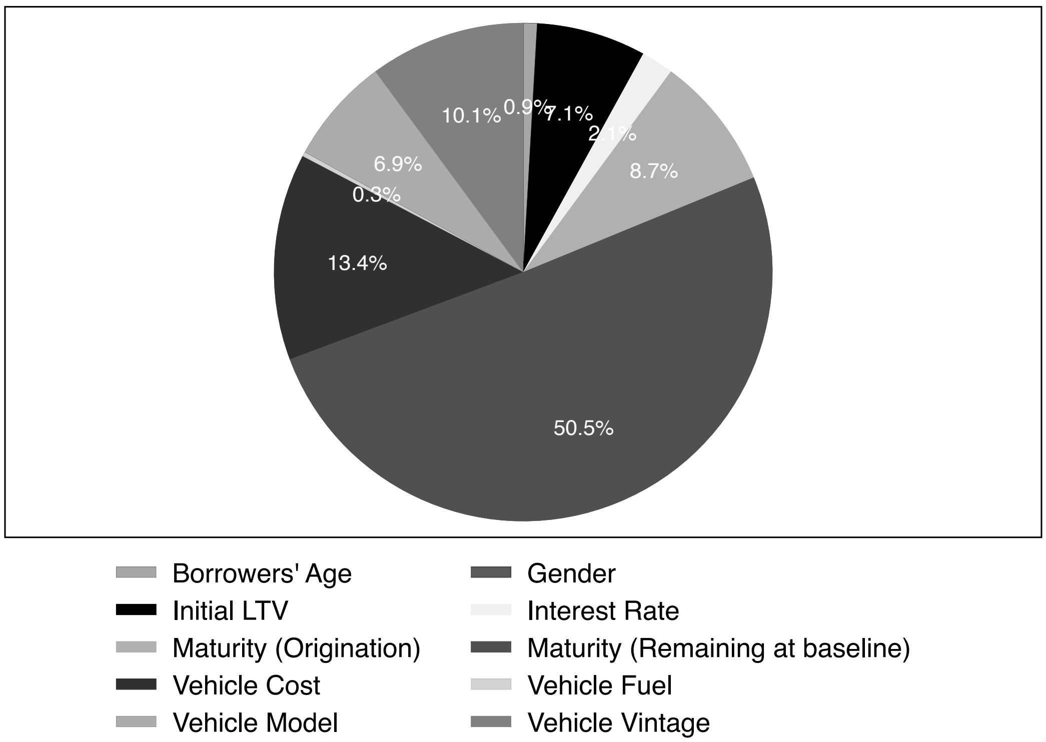

Notes: Experimental sample (N=3,186). Results from an R-squared decomposition model (Huettner and Sunder 2012), using the Shapley value framework to estimate how much of the variance in vehicle loan-to-value (LTV) is explained by different groups of covariates measured at baseline and/or origination. Please see Appendix C.2 for details on the covariates.

Table A.1: Correlations across index components for each key outcome summary index

Panel A. Minibus loan performance

	Repayment Index	1(current)	Arrears Amount × -1	Arrears (Scaled) × -1
Repayment Index	1			
1(current)	0.756	1		
Arrears Amount × -1	0.920	0.483	1	
Arrears (Scaled) × -1	0.918	0.478	0.904	1

Panel B. Outside Borrowing Index

	Outside Borrowing Index	1(non -zero debt)	Credit Utilization	Balance (Installment loans)	Borrowings (Credit Lines)
Outside Borrowing Index	1				
1(non-zero debt)	0.763	1			
Credit Utilization	0.733	0.603	1		
Balance (Installment Loans)	0.547	0.211	0.068	1	
Borrowings (Credit Lines)	0.640	0.233	0.296	0.188	1

Panel C. Outside Credit Access Index

	Outside Credit Access Index	1(has card)	1(has credit line)	1(has installment loan)	Number of credit lines
Outside Credit Access Index	1				
1(has card)	0.842	1			
1(has credit line)	0.860	0.765	1		
1(has installment loan)	0.525	0.166	0.208	1	
Number of credit lines	0.825	0.640	0.651	0.227	1

Panel D. Effort Index

	Effort Index	Distance Driven	Time Driven	Time in job (hours)	Total days driven
Effort Index	1				
Distance Driven	0.896	1			
Time Driven	0.944	0.846	1		
Time in job	0.952	0.763	0.861	1	
Total days driven	0.926	0.723	0.804	0.916	1

Notes: Experimental sample (N=3,186). Please see Section 2.3 and Appendix Section C.1 for additional details on variable definitions and construction.

Table A.2: Experimental sample: Summary statistics and balance

Sample: *All Baseline: 3,186 loans*	**Control mean** (1)	**Control SD** (2)	**Debt Reduction** (3)	**Payment Reduction** (4)
Panel A. Borrower characteristics:				
Credit Score	592.19	22.15	0.00	–0.00
			(0.00)	(0.00)
1(male)	0.76	0.43	–0.01	0.01
			(0.02)	(0.02)
Borrower's Age (in years)	51.10	10.97	–0.00	0.00
			(0.00)	(0.00)
No. of outstanding loans with lender	1.48	1.06	–0.02**	0.01
			(0.01)	(0.01)
Panel B. Loan characteristics:				
1(vehicle is new)	0.69	0.46	0.01	–0.02
			(0.03)	(0.03)
Age of the Vehicle	3.41	2.55	0.00	–0.00
			(0.01)	(0.01)
Loan Principal (R1,000)	388.57	130.54	–0.00	0.00
			(0.00)	(0.00)
Arrears (R1,000)	54.04	26.06	0.00	–0.00
			(0.00)	(0.00)
Maturity (Origination)	73.95	6.95	–0.00	0.00
			(0.00)	(0.00)
Loan to Value (Origination)	1.00	0.05	0.25	–0.13
			(0.23)	(0.24)
Loan to Value (Baseline)	0.88	0.22	0.02	0.10
			(0.18)	(0.18)
Panel C. Baseline contract terms:				
Interest Rate	0.24	0.03	–1.00	–0.15
			(0.64)	(0.62)
Installments (R1,000)	13.00	1.73	0.02	–0.00
			(0.02)	(0.02)
Remaining Maturity	47.49	18.38	0.00	–0.00
			(0.00)	(0.00)
Panel D. Baseline outcomes (standardized):				
Repayment Index	0.01	0.65	0.01	–0.01
			(0.02)	(0.02)
Effort Index	–0.01	0.91	0.00	0.00
			(0.01)	(0.01)
Outside Credit Access Index	0.01	0.76	0.02	–0.02
			(0.02)	(0.02)
Outside Borrowing Index	0.01	0.68	–0.00	–0.00
			(0.02)	(0.02)
p-value (joint F-test)			[0.560]	[0.754]
Observations			3,186	3,186
Observations per study arm	1,063		1,062	1,061

Notes: Columns (1) and (2) report baseline summary statistics from the control group only. Columns (3) and (4) each report results from a separate multivariate OLS regression of an indicator for that treatment assignment on the variables described in the rows. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.3: ATEs on other minibus loan repayment behaviors

	Repayment metrics ...			
	1(Under Payment) (1)	1(Payment in Full) (2)	1(Over Payment) (3)	Total Payment (in R) (4)
β^{DR} : 1(debt reduction)	-0.002	0.009	-0.007	-3,186.9
	(0.012)	(0.008)	(0.010)	(2,843.0)
β^{PR} : 1(payment reduction)	-0.039***	0.003	0.036***	-6,617.0**
	(0.013)	(0.008)	(0.011)	(2,841.1)
Observations	3,186	3,186	3,186	3,186
R-squared	0.032	0.006	0.029	0.026
Control mean	0.91	0.03	0.06	123,654
$p[\beta^{DR} = \beta^{PR}]$	0.01	0.49	0.00	0.21
Strata f.e.	✓	✓	✓	✓

Notes: Each column presents intent-to-treat (ITT) estimates of average treatment effects (ATEs) from a single OLS regression per equation 3. Dependent variable is described in the column heading and measured over the 12-month experimental period. Payment in full is defined as falling within |1%| of total minimum amounts due over the 12 months. Under- and over- payment are defined relative to payment in full. $^{***}p <$ 0.01, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.4: ATEs on outside borrowing and outside credit access

	Outside Borrowing Index (1)	Outside Borrowing Index (Winsorized at 5%) (2)	Outside Credit Access Index (3)
β^{DR} : 1(debt reduction)	-0.015	-0.009	0.016
	(0.026)	(0.026)	(0.033)
β^{PR} : 1(payment reduction)	-0.051**	-0.050*	-0.014
	(0.025)	(0.026)	(0.033)
Observations	3,186	3,186	3,186
R-squared	0.070	0.073	0.027
$p[\beta^{DR} = \beta^{PR}]$	0.14	0.11	0.36
Strata f.e.	✓	✓	✓

Notes: Each column presents intent-to-treat (ITT) estimates of average treatment effects (ATEs) from a single OLS regression per equation 3. Dependent variable is described in the column heading and measured over 12-months for outside borrowing and at end of our 12-month experiment period for credit access (see Figure A.9 for monthly TE estimates). All LHS variables are standardized and thus have means of approximately zero and TE estimates in standard deviation units; please see Section 2.3 and Appendix Section C.1 for additional details on variable definitions. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.5: HTEs by baseline vehicle equity on index components

	Repayment Index Components			Effort Index Components			
	1(current)	Arrears Amount	Arrears (Scaled)	Distance driven	Time driven	Time Spent on job	Total days driven
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
β^{DR+LE}: 1(debt reduction)	-0.035	0.042	0.022	-0.056	-0.080	-0.074	-0.028
× 1(low baseline equity)	(0.046)	(0.065)	(0.062)	(0.068)	(0.067)	(0.064)	(0.063)
β^{DR+HE}: 1(debt reduction)	0.083	-0.060	-0.065	0.023	-0.002	0.060	0.042
× 1(high baseline equity)	(0.066)	(0.053)	(0.057)	(0.069)	(0.069)	(0.067)	(0.067)
β^{PR+LE}: 1(payment reduction)	0.078	-0.172***	-0.048	-0.039	-0.067	-0.063	-0.014
× 1(low baseline equity)	(0.050)	(0.066)	(0.061)	(0.069)	(0.066)	(0.064)	(0.063)
β^{PR+HE}: 1(payment reduction)	0.227***	-0.166***	-0.100	0.104	0.064	0.161**	0.147**
× 1(high baseline equity)	(0.070)	(0.054)	(0.061)	(0.069)	(0.068)	(0.066)	(0.066)
1(high baseline equity)	0.111*	-0.244***	-0.205***	-0.221***	-0.224***	-0.274***	-0.183***
	(0.059)	(0.068)	(0.064)	(0.075)	(0.073)	(0.071)	(0.071)
Observations	3,186	3,186	3,186	3,178	3,178	3,178	3,178
R-squared	0.056	0.093	0.067	0.013	0.014	0.041	0.037
$p[\beta^{DR+LE} = \beta^{PR+LE}]$	0.02	0.00	0.23	0.80	0.85	0.85	0.83
$p[\beta^{DR+HE} = \beta^{PR+HE}]$	0.04	0.04	0.56	0.22	0.33	0.12	0.10
$p[\beta^{DR+HE} = \beta^{DR+LE}]$	0.15	0.22	0.30	0.41	0.42	0.15	0.45
$p[\beta^{PR+HE} = \beta^{PR+LE}]$	0.08	0.94	0.55	0.14	0.17	0.01	0.08
Strata f.e.	✓	✓	✓	✓	✓	✓	✓

Notes: Same as Table 6, except that here each dependent variable is a summary index component. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.6: HTEs by baseline vehicle equity: Robustness to controlling for remaining maturity

	Repayment Index		Effort Index	
	(1)	(2)	(3)	(4)
Panel A. Z_i **= Normalized baseline remaining maturity**				
β^{DR+LE}: 1(debt reduction) × 1(low baseline equity)	-0.033	-0.000	-0.060	-0.087
	(0.050)	(0.063)	(0.061)	(0.077)
β^{DR+HE}: 1(debt reduction) × 1(high baseline equity)	0.069	0.031	0.031	0.058
	(0.050)	(0.061)	(0.064)	(0.081)
β^{PR+LE}: 1(payment reduction) × 1(low baseline equity)	0.099**	0.104	-0.046	-0.116
	(0.050)	(0.064)	(0.061)	(0.077)
β^{PR+HE}: 1(payment reduction) × 1(high baseline equity)	0.164***	0.160**	0.119*	0.187**
	(0.053)	(0.064)	(0.063)	(0.078)
1(high baseline equity)	0.187***	0.096	-0.226***	-0.252***
	(0.055)	(0.076)	(0.068)	(0.096)
Observations	3,186	3,186	3,178	3,178
R-squared	0.093	0.099	0.026	0.027
$p[\beta^{DR+LE} = \beta^{PR+LE}]$	0.01	0.09	0.82	0.69
$p[\beta^{DR+HE} = \beta^{PR+HE}]$	0.07	0.04	0.15	0.09
$p[\beta^{DR+HE} = \beta^{DR+LE}]$	0.15	0.76	0.31	0.27
$p[\beta^{PR+HE} = \beta^{PR+LE}]$	0.38	0.59	0.06	0.02
1(debt reduction) × Z_i		✓		✓
1(payment reduction) × Z_i		✓		✓
Panel B. Z_i **= Quartiles of baseline remaining maturity f.e.**				
1(debt reduction)	-0.033	-0.006	-0.060	-0.047
	(0.050)	(0.072)	(0.061)	(0.095)
1(debt reduction) × 1(high baseline equity)	0.102	0.095	0.090	0.202
	(0.071)	(0.108)	(0.089)	(0.138)
1(payment reduction)	0.099**	0.183**	-0.046	0.040
	(0.050)	(0.076)	(0.061)	(0.097)
1(payment reduction) × 1(high baseline equity)	0.065	0.086	0.165*	0.267**
	(0.073)	(0.110)	(0.088)	(0.133)
1(high baseline equity)	0.187***	0.061	-0.226***	-0.360***
	(0.055)	(0.081)	(0.068)	(0.100)
Observations	3,186	3,186	3,178	3,178
R-squared	0.093	0.102	0.026	0.033
1(debt reduction) × Z_i		✓		✓
1(payment reduction) × Z_i		✓		✓

Notes: Each column presents OLS estimates from a single regression. Dependent variable is described in the column heading and measured at 12-months for repayment and over 12-months for effort (please see Section 2.3 and Appendix Section C.1 for details). To facilitate comparison, we first report (in Panel A) or reproduce (in Panel B) the results from Table 6 in Columns (1) and (3). Then in Columns (2) and (4) of Panel A, we include interactions between each of the treatment group indicators and standardized remaining maturity (mean 0 and standard deviation 1). Then in Columns (2) and (4) of Panel B, we include interactions between each of the treatment group indicators and quartiles of remaining maturity. Variable definitions are same as Table 6 and all specifications control for the eight strata fixed effects. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.7: HTEs by Baseline Vehicle Equity: Robustness to Dropping Remaining Maturity Outliers

	Sample after ...			
	... removing remaining maturity in top and bottom deciles		... removing remaining maturity $\leq$ 3 year or $\geq$ 6 year	
Outcome variable:	Repayment Index (1)	Effort Index (2)	Repayment Index (3)	Effort Index (4)
β^{DR+LE}: 1(debt reduction) $\times$ 1(low baseline equity)	-0.042 (0.056)	-0.070 (0.068)	-0.039 (0.054)	-0.064 (0.067)
β^{DR+HE}: 1(debt reduction) $\times$ 1(high baseline equity)	0.018 (0.057)	0.038 (0.073)	-0.027 (0.068)	0.039 (0.091)
β^{PR+LE}: 1(payment reduction) $\times$ 1(low baseline equity)	0.094* (0.057)	-0.074 (0.068)	0.103* (0.056)	-0.067 (0.067)
β^{PR+HE}: 1(payment reduction) $\times$ 1(high baseline equity)	0.148** (0.058)	0.137** (0.070)	0.140** (0.070)	0.183** (0.087)
1(high baseline equity)	0.192*** (0.061)	-0.207*** (0.074)	0.187*** (0.066)	-0.217** (0.085)
Observations	2,534	2,534	2,185	2,185
R-squared	0.080	0.022	0.072	0.021
$p[\beta^{DR+LE} = \beta^{PR+LE}]$	0.01	0.96	0.01	0.97
$p[\beta^{DR+HE} = \beta^{PR+HE}]$	0.02	0.15	0.02	0.10
$p[\beta^{DR+HE} = \beta^{DR+LE}]$	0.46	0.28	0.89	0.36
$p[\beta^{PR+HE} = \beta^{PR+LE}]$	0.51	0.03	0.69	0.02
Strata f.e.	✓	✓	✓	✓

Notes: Each column presents OLS estimates from a single regression. Dependent variable is described in the column heading and measured at 12-months for repayment and over 12-months for effort (please see Section 2.3 and Appendix Section C.1 for details). Variable definitions are same as Tables 6 and A.6. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.8: HTEs by Predicted Repayment Probability on Index Components

	Repayment Index Components			Effort Index Components			
	1(current)	Arrears Amount	Arrears (Scaled)	Distance driven	Time driven	Time Spent on job	Total days driven
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1(debt reduction)	-0.049	0.082	0.085	-0.110	-0.133**	-0.102	-0.069
	(0.044)	(0.062)	(0.064)	(0.068)	(0.068)	(0.066)	(0.067)
1(debt reduction) × 1(high repayment prob.)	0.131	-0.174**	-0.208**	0.189**	0.188*	0.194**	0.155*
	(0.080)	(0.083)	(0.083)	(0.097)	(0.096)	(0.093)	(0.092)
1(payment reduction)	0.087*	-0.177***	-0.023	0.022	-0.020	0.024	0.054
	(0.050)	(0.063)	(0.063)	(0.069)	(0.067)	(0.065)	(0.064)
1(payment reduction) × 1(high repayment prob.)	0.131	0.001	-0.116	0.022	0.038	0.054	0.029
	(0.086)	(0.084)	(0.085)	(0.098)	(0.095)	(0.092)	(0.092)
1(high repayment prob.)	0.151***	-0.357***	-0.287***	-0.063	-0.076	-0.071	-0.022
	(0.058)	(0.064)	(0.064)	(0.074)	(0.072)	(0.069)	(0.069)
Observations	3,178	3,178	3,178	3,178	3,178	3,178	3,178
R-squared	0.060	0.114	0.088	0.010	0.011	0.037	0.036
Strata f.e.	✓	✓	✓	✓	✓	✓	✓

Notes: Same as Table 7, except that here each dependent variable is a summary index component. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.9: When repayment and effort responses diverge: Additional HTEs by predicted repayment probability

	Repayment Index (1)	Effort Index (2)
1(debt reduction)	-0.072	-0.104
	(0.049)	(0.063)
1(debt reduction) × 1(high repayment prob.)	0.411***	0.128**
	(0.049)	(0.063)
1(payment reduction) | (1(control) × 1(high repayment prob.))	0.247***	-0.010
	(0.043)	(0.053)
Observations	3,178	3,178
R-squared	0.100	0.022
Strata f.e.	✓	✓

Notes: Each column presents OLS estimates from a single regression. Dependent variable is described in the column heading and measured at 12-months for repayment and over 12-months for effort (please see Section 2.3 and Appendix Section C.1 for details). These variables are standardized and thus have means of approximately zero and TE estimates in standard deviation units. The effort index is rescaled for comparability to monthly treatment effect estimates. 1(high repayment prob.) takes the value of one if the predicted value $\hat{\mathcal{L}}_i$ (i.e., probability of being current on the loan six months into the experiment) is at or above the median value in our experimental sample, and zero otherwise (please see Appendix Section D for details on estimating $\hat{\mathcal{L}}_i$). $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table A.10: ATEs on key externalities

	Outside Debt Repayment Index (1)	Outside Debt Repayment Index (Winsorized at 5%) (2)	Risky Driving Index (full sample) (3)	Risky Driving Index (insurance sample) (4)
β^{DR} : 1(debt reduction)	-0.048	-0.059*	0.017	0.007
	(0.033)	(0.033)	(0.020)	(0.020)
β^{PR} : 1(payment reduction)	-0.004	-0.014	0.009	-0.001
	(0.032)	(0.033)	(0.019)	(0.020)
Observations	2,947	2,947	3,176	2,903
R-squared	0.025	0.025	0.012	0.007
$p[\beta^{DR} = \beta^{PR}]$	0.15	0.17	0.66	0.63
Strata f.e.	✓	✓	✓	✓

Notes: Each column presents intent-to-treat (ITT) estimates of average treatment effects (ATEs) from a single OLS regression per equation 3. Dependent variable is described in the column heading and measured as an average over the 12 months experimental period (please see Section 5.4 and Appendix Section C.1 for additional details on variable definitions). All LHS variables are standardized and thus have means of approximately zero and TE estimates in standard deviation units.$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

B Model Appendix: Alternative assumptions about default costs (Online Appendix)

Section 3 presents two stylized models to guide the interpretation of our empirical findings. The goal is not to fully characterize the optimal contract or the detailed dynamics of repayment, but rather to illustrate key mechanisms that link debt obligations, effort, and default decisions in a tractable environment. In this appendix, we consider more general versions of these models to explore whether the presence of different assumptions qualitatively alters our key results. We start by recapping the no-debt and the standard debt-overhang cases presented in Section 3, so that the appendix here is self contained.

B.1 Preliminaries: Effort choice without debt

Consider a simple two-period setting in which an entrepreneur chooses the level of effort e to exert in running a business. Effort generates profits in the first period according to a strictly concave function $f(e)$, with $f' > 0$ and $f'' < 0$. In the second period, the business is sold at a price that depends on first-period profits: $V(f(e))$, where $V' > 0$ and $V'' \leq 0$. We adopt this interpretation for simplicity. Alternatively, the model can be viewed as one in which an entrepreneur chooses effort today, considering its effect on both current profits and the business's future continuation value.

The entrepreneur discounts future payoffs with a factor $\beta < 1$ and incurs a linear cost of effort.[1] The entrepreneur chooses e to maximize total discounted value:

$$\max_{e} f(e) + \beta V(f(e)) - e \tag{B.1}$$

The first-best level of effort e^{FB} satisfies the first-order condition:

$$f'(e^{FB})(1 + \beta V'(f(e^{FB}))) = 1 \tag{B.2}$$

This condition captures the fact that effort is chosen to balance the marginal cost of effort with the marginal benefits, which include both current profits and the future valuation of the business.

B.2 Standard Debt Overhang Case

Now consider an entrepreneur facing a fixed debt burden with a short-term payment D_1 and a long-term payment D_2, giving a present value of $D = D_1 + \beta D_2$. If the entrepreneur fails to make D_1, repossession follows in the second period regardless of subsequent actions.

In this model, the borrower chooses three actions: the effort level e and the two repayments P_1, P_2. We assume the borrower first selects the effort level e (period zero), then makes the first payment decision P_1, and finally makes the second payment decision P_2.

[1] Our results would be qualitatively unchanged if we take a more general approach and assume a standard increasing and convex cost function $c(e)$.

In Section 3, we focused on the simple case where we assume that default leads to the full loss of the business. We now consider a more general assumption. In particular, we now consider two costs related to default:

- In case of default, the asset is repossessed and sold by the lender. We assume that this action leads to a loss of equity for the borrower: in particular, the borrower only receives a share $1-\gamma$ of the equity.
- On top of the future asset loss, default leads to an immediate non-monetary cost ϕ for the borrower. For instance, this cost can capture the utility cost related to having a default flag in the credit report. This cost affects the borrower's utility at the time of default. We assume that $\phi \geq 0$.

Recall that Sections 3.2 and 3.3 focus on the special case where $\gamma = 1$ and $\phi = 0$, while Section 3.5 focuses on the special case where $\gamma = 1$ and $\phi > 0$.

Solution. Before discussing the solution, note that the setup here removes any incentive for the borrower to (1) pay anything different than the payment due or zero; (2) repay anything in period 2 if they do not pay D_1. This implies that there are only three potentially optimal payment choice pairs from the borrower's perspective: (a) $P_1, P_2 = D_1, D_2$, (b) $P_1, P_2 = D_1, 0$, and (c) $P_1, P_2 = 0, 0$.

We solve the model by backward induction, beginning with the optimal choice of P_2 given e^* and P_1. If the borrower skips the first payment ($P_1 = 0$), then $P_2 = 0$ necessarily. If the first payment is made in full ($P_1 = D_1$), the borrower pays $P_2 = D_2$ only if:

$$V(f(e^\star)) - D_2 \geq -\phi + (1-\gamma)(V(f(e^\star))) \tag{B.3}$$

The logic is straightforward: by repaying D_2, the business retains full equity; by defaulting, it forfeits a share γ of equity and incurs the fixed default cost ϕ. This trade-off reduces to the following condition:

$$\gamma V(f(e^\star)) - D_2 \geq -\phi \tag{B.4}$$

In the first period, the business will choose whether or not to repay, knowing its future choice P_2. Therefore, we need consider two scenarios, based on equation B.4. If that condition holds, then a borrower that pays in the first period will also pay in the second period. This happens only if the benefit of full repayment outweighs the benefit of not making any payments. In other words:

$$f(e^\star) - e^\star - D_1 + \beta(V(f(e^\star)) - D_2) \geq f(e^\star) - e^\star - \phi + (1-\gamma)\,\beta V(f(e^\star)) \tag{B.5}$$

This expression can be re-written as:

$$\beta\gamma(V(f(e^\star))) - D \geq -\phi \tag{B.6}$$

Therefore, a borrower repays the debt in full $(P_1, P_2) = (D_1, D_2)$ if both equations B.4 and B.6

hold. On this point, we note that equation B.6 implies equation B.4, if the regularity condition $D \geq D_2$ holds. As we discuss in detail below, this explains why we only consider equation B.6 in Section 3.

If equation B.6 does not hold, then the borrower will never pay in full but may decide to pay D_1 only. This can happen because paying $P_1 = D_1$ allows the borrower to postpone the costly ϕ, which benefits the borrower on net if the payoff to paying in the first period only $(P_1, P_2) = (D_1, 0)$ is higher than the payoff to paying nothing $(P_1, P_2) = (0, 0)$:

$$f(e^\star) - e^\star - D_1 + (1-\gamma)\,\beta V(f(e^\star)) - \phi\beta \geq f(e^\star) - e^\star - \phi + (1-\gamma)\,\beta V(f(e^\star)) \tag{B.7}$$

which can be re-written simply as:

$$\phi \geq D_1(1-\beta)^{-1} \tag{B.8}$$

We note that this condition holds irrespective of whether equation B.4 holds or not. In other words, as soon as the total debt level D is large enough (i.e., equation B.6 does not hold), the borrower will not pay in full, but may pay in the short run if immediate default is sufficiently costly (i.e., if equation B.8 holds).

Following our solution procedure, the last case to consider is when total debt is small enough (i.e., equation B.6 holds), but long-term debt is too large (i.e., equation B.4 does not hold). In this case, it is easy to observe that the borrower will always find optimal to repay $(P_1, P_2) = (D_1, 0)$. The intuition for this result is that when equation B.8 holds, the immediate default cost is always high enough to incentivize repaying D_1, even when D is otherwise high enough to produce the standard debt overhang effect.[2] One can think of this case as a delayed debt overhang effect on repayment.

The bottom line is therefore that if either equation B.6 or equation B.4 does not hold, then we do not have full repayment. Furthermore, in this case, if the immediate cost of default is sufficiently high (i.e., equation B.8 holds), then the borrower will repay D_1 to delay the immediate default cost but then default in period 2: $(P_1, P_2) = (D_1, 0)$. Otherwise, the standard debt overhang result holds: the borrower will not make any payment $(P_1, P_2) = (0, 0)$.

Continuing our backwards induction, our last step is to solve for the borrower's effort choice in each of the three repayment cases. In the full repayment case, the borrower retains its pledged asset and business, and therefore its objective function is:

$$\max_e \; f(e) \; - D_1 + \; \beta\,(V(f(e)) - D_2) \; - \; e$$

This problem is equivalent on the margin to the no-debt model, and therefore optimal effort is first-best: $e^\star = e^{FB}$. In the other two cases, the borrower's asset will be repossessed, and therefore

[2] More formally: if equation B.6 holds, then $\phi \geq D_1 - \beta(\gamma V(f(e^\star) - D_2)$. If equation B.4 does not hold, then $(\gamma V(f(e^\star) - D_2)$ is negative, which has to imply that $\phi \geq D_1$ and therefore equation B.8 always holds.

the two objective functions are each equivalent to the following:[3]

$$\max_{e} \; f(e) \; - P_1 + \; \beta(1-\gamma)\,(V(f(e))) \; - \; e$$

Now optimal effort is e^{SB}, which implicitly defined by $f'(e^{SB})\left(1 + \beta\,(1-\gamma)\,V'(f(e^{SB}))\right) = 1$. It is easy to observe that $e^{SB} < e^{FB}$ as long as $\gamma > 0$.

Summary. The key results from this model are the following:

1 The borrower exerts first-best effort $e = e^{FB}$ and pays the debt in full (i.e., $(P_1, P_2) = (D_1, D_2)$) if the overall debt burden D and the long-term debt D_2 are sufficiently low. In other words, if equations B.6 and B.4 hold, then $e = e^{FB}$.

2 If either B.6 or B.4 does not hold, the borrower will choose $e = e^{SB}$, with $e^{SB} < e^{FB}$ if $\gamma > 0$.

3 Even if the borrower exerts lower effort $e = e^{SB}$, it may still repay in the short run with $(P_1, P_2) = (D_1, 0)$. This occurs only if the immediate cost of default is sufficiently large: $\phi \geq D_1(1-\beta)^{-1}$.

4 The lender can induce higher effort and repayment by reducing D and D_2 so that the conditions discussed above hold. In particular, a modification that lowers long-term debt D_2 while keeping short-term payments D_1 constant should achieve this.

Discussion. This more general version of the model in Section 3 yields several noteworthy insights. First, as expected, it reduces to simpler setup presented in Section 3 if equity is completely lost in the event of default (i.e., $\gamma = 1$), there is no immediate default cost (i.e., $\phi = 0$), and we assume the regularity condition $D \geq D_2$. In Section 3 we restrict attention to $D \geq D_2$ for parsimony and exposition (not maximal realism): in that simpler setup, the model reduces to a single debt overhang condition, equation B.6, which automatically implies equation B.4.

Second, even in the more general setup here, the basic mechanism from Section 3 remains: high debt burden reduces the net value of continuation, leading to default and underinvestment in effort. The main difference from Section 3 is that here two conditions must hold to get first-best effort.[4] The extra condition requires long-term debt D_2 to be sufficiently low, further highlighting that the *key incentive problem with debt overhang comes from long-term debt.* Consistent with that idea, one arm in our experiment tries to address the standard debt overhang problem by reducing long-term debt (i.e., D_2) while keeping short-term payment D_1 constant.

[3] We stress that we literally mean equivalent, not identical: clearly, across the two scenarios, certain parameters of the objective functions can differ (e.g., the role of ϕ). However, those parameters do not affect the marginal incentives to exert effort, and therefore it is equivalent to exclude them.

[4] In Section 3 , we technically do still need two conditions to hold to achieve the first best. But under the regularity condition, equation B.4 is implied by equation B.6, so only equation B.6 needs to be explicitly required. Here, by contrast, we must explicitly require both conditions to hold.

The only requirement necessary to keep the basic mechanism intact is that $\gamma > 0$, which implies that repossession prevents the entrepreneur from fully recovering their marginal investment in the business.[5] The precise cutoff for repayment does depend on γ and ϕ, as we review next, but the qualitative result on effort holds for any ϕ so long as $\gamma > 0$.[6] The only qualitative change is related to repayment, and it is discussed in the next paragraph.

Third, the framework shows that effort and repayment can diverge in some cases. As discussed above, introducing $\phi > 0$ leaves effort incentives qualitatively unchanged: the entrepreneur still reduces effort when the debt burden is high. Repayment, however, may differ. With $\phi > 0$, the entrepreneur may repay in the short run while anticipating eventual default, since repayment postpones the immediate penalty. This occurs only if the value of delaying default exceeds the first-period payment, i.e., when $\phi \geq D_1(1-\beta)^{-1}$. We view $\phi > 0$ as a realistic assumption– reflecting reputational, psychological, or legal costs associated with missed payments– that is strongly supported by our empirical results in Section 5.3. And we conjecture that this mechanism becomes even more important in more realistic multi-period settings with uncertainty, where delaying default resembles exercising the option to continue earning profits from the business—or, more generally, utility from the pledged asset.

As discussed in the main text, this case illustrates why it is essential to consider effort as well as repayment behavior. When $\phi > 0$, repayment in the short run does not rule out debt overhang in longer-run. We return to this point below, after adding liquidity constraints to the general version of our debt overhang model.

B.3 Adding liquidity constraints

Setting. We now extend our general version of the standard debt overhang model to incorporate liquidity constraints. Specifically, we assume that the entrepreneur is constrained to use only the cash generated by their business net of C before making any repayment in the first period, and that outside financing is unavailable.[7] That is, they must satisfy:

$$f(e) - C - P_1 \geq 0 \tag{B.9}$$

This constraint may prevent repayment even when the entrepreneur's net continuation value is positive (i.e., the debt overhang constraint is not binding). So now repayment and effort are jointly determined by both the continuation value of the business and short-run liquidity.

[5]This assumption is empirically realistic in our setting, given repossession costs, the prevalence of negative vehicle equity, and the sunk permitting cost, all of which erode liquidation value in the repossession state. Furthermore, $\gamma > 0$ is also consistent with our experiment's findings that contract terms affect effort; in contrast, with full recovery for the borrower, there is no debt overhang (or its functional equivalent created by liquidity constraints), and changes to contract terms will not change effort.

[6]From a quantitative standpoint, pinning down the most accurate specification of default costs is probably crucial to match the data. However, our model is designed to have a purely qualitative interpretation.

[7]We assume that the cost arises independently from the business, and must be incurred irrespective of the business status. For instance, C could reflect basic living expenses. Please see footnote 27 in Section 3.3 for discussion.

Solution. The inclusion of C creates an extra condition in the problem discussed above. In particular, we start by considering three potential cases for how C compares to the other model components. The first is that the liquidity constraint is not binding, even if the borrower makes the full payment D_1 in the first period. Formally, we have:

$$f(e^{SB}) - C - D_1 \geq 0 \tag{B.10}$$

When this condition holds, the model is functionally equivalent to the standard debt overhang model: functionally speaking, there is no liquidity constraint.

The second case is also an extreme one and extremely straightforward: the liquidity constraint is binding irrespective of the borrower's effort. Formally:

$$f(e^{FB}) - C - D_1 < 0 \tag{B.11}$$

Here the borrower is unable to make the payments even if they exert first-best effort. It follows that the borrower makes zero payments $(P_1, P_2) = (0, 0)$ and sets $e = e^{SB}$.

The third case to consider is one where liquidity is sufficient to make repayment possible only if first-best effort is exerted. Formally:

$$f(e^{SB}) - C - D_1 < 0 \ \wedge \ f(e^{FB}) - C - D_1 \geq 0 \tag{B.12}$$

Interestingly, this rules out the delayed default strategy implied by the standard model above. In other words, when the debt overhang constraint also binds, a borrower will not pay D_1 but not D_2 and exert e^{SB}, regardless of how large ϕ is. The intuition is that the borrower cannot afford to make a payment D_1 if effort is second-best. In this case, second best effort is always tied to no payments in both the short- and long-run.

Note that we have focused only on cases where a binding liquidity constraint can lead to lower effort. We are not asserting that these are the only *feasible* cases– under certain assumptions about V and f, the borrower may find it optimal to exert effort e above e^{FB}. But we do hold that under-investment in effort captures the *relevant* set of cases for our purposes, for two reasons. First and conceptually, our overall objective with the model is to show how liquidity constraints, in conjunction with repossession risk, create a second incentive compatibility constraint that lenders must satisfy to prevent lower effort and non-repayment. Second and empirically, the evidence from our experiment runs counter to the predictions produced by a model with over-investment in effort. That model predicts that relaxing liquidity constraints will induce lower effort, while we find evidence to the contrary and consistent with our under-investment model's various distinct testable predictions.

Discussion. As discussed in the paper, the presence of the liquidity constraint implies that lower effort e^{SB} can be generated by either the traditional debt overhang constraint or liquidity issues. Formally, the entrepreneur will fail to repay and exert low effort if either the debt burden condition (i.e. if equations B.6 and B.4) or the liquidity condition fail. This highlights how debt overhang and liquidity interact in determining behavior. Addressing only one constraint may be insufficient if the other remains binding.

In particular, the presence of liquidity constraints changes the contract terms necessary to increase effort. If the entrepreneur is only constrained by the overall debt burden, reducing D_2 can restore repayment and effort. But if liquidity constraints are binding, reducing D_2 alone may not suffice. In such cases, easing short-run cash flow pressures—e.g., by lowering D_1—is necessary to support both repayment and effort. But if both constraints are binding (i.e., a liquidity constrained borrower with low equity), inducing more effort and higher payments may require lowering D_1 as well as the overall debt burden D.

Overall, a more general version of our model with liquidity constraints nests the Section 3's simpler version, and leaves it qualitative results and testable predictions unchanged under realistic assumptions.

C Data Appendix (Online Appendix)

C.1 Variable Definitions and Construction

This appendix details the definitions of key variables used in our analysis. It follows our pre-analysis plan (`https://www.socialscienceregistry.org/trials/13052`) unless otherwise noted and motivated.

We note several conventions upfront, to minimize repetition below:

1. Summary index component variables are standardized to mean zero and unit variance.
2. Summary indices are the unweighted average of their component variables.
3. We measure outcome variables over both monthly and 12-month horizons.
4. For stock variables, like loan performance or balances, the monthly versions are measured as of each month-end. The 12-month horizon version is measured at either 12-month-end or as an average of the 12 month-ends, as detailed below.
5. For flow variables, like effort, we use the average or sum over the requisite time period, as detailed below.
6. Loan amounts are in South African Rand (R).
7. Minibus loan terms and repayment status variables use the lender's definitions.

C.1.1 First-Stage Outcomes

Our first-stage outcomes — the loan terms directly varied by our experiment– are interest rate, monthly payment, and remaining maturity. These are pulled each month from the lender's book, based on contractual terms in force at that time.

We measure these terms only for months when the loan remains on the books, with no imputation applied after a loan leaves the book (see footnote 19 in Section 2.3 for discussion). 7.8% of loans leave the books during our 12-month window for estimating treatment effects, and these account for 1.8% of the potential loan-months in that window.

(i) *Monthly Installment Owed.* Minimum monthly debt repayment due, in R.

(ii) *Interest Rate.* Interest rate, in percentage points.

(iii) *Remaining Maturity.* Outstanding maturity, in months.

For the 12-month horizon version of these variables, we average the monthly versions.

C.1.2 Minibus Loan Repayment Index

This index is the equally-weighted average of three standardized components:

$$\textit{Repayment Index}_t = \frac{\textit{1(Current)}_t - \textit{Arrears Amount}_t - \textit{Scaled Arrears}_t}{3}$$

We define the components as:

(i) *1(Current).* An indicator taking the value of one if the loan has $\leq$ 100 ZAR in balance that is $>$ 30 days late, and zero otherwise.

(ii) *Arrears Amount.* Total amount $>$ 30 days past due.

(iii) *Scaled arrears.* Total arrears amount divided by the minimum monthly total installment owed in that month.

In the 1.8% of loan-months following a loan that leaves the book we code a paid-in-full loan as current, and a loan that is not paid-in-full as not current. Similarly, for loan-months following a loan that leaves the book: we consider arrears amount (scaled arrears) to be zero for a paid-in-full loan, and last observed arrears amount (scaled arrears) for not paid-in-full. For monthly definitions we utilize the snapshot at each month-end. For the 12-month horizon, we utilize the snapshot at end of the 12th month.

C.1.3 Entrepreneurial Effort Index

This index is the equally-weighted average of the four standardized components defined below and derived from daily GPS data provided by the lender's telematics data vendor (more on this below):

(i) *Distance driven:* Sum of the total distance covered.

(ii) *Time Driven.* Sum of total active driving hours.

(iii) *Days worked:* Count of days with positive driving time.

(iv) *Time spent on the job.* Duration (in hours) between the start of the first trip and the end of the last trip each day. Coded as zero for days with zero trips.

At the pre-analysis stage, we assumed we could obtain higher-frequency GPS data, enabling us to construct the above component measures from scratch. We specified cleaning steps as such. Instead, thus far we have only been able to access vendor-provided data aggregated to the daily level. Following our pre-analysis plan, for the monthly versions of these variables we sum over each day's total and standardize the variable. For the 12-month horizon version, we sum over each month and divide by 12.

We only have reliable data for a vehicle while its loan is still on the lender's books. We do not impute values for effort measures for the 1.8% of loan-months after a borrower leaves the books, given our focus on the principal-agent problem — the lender cares only about driver effort while the borrower is obligated to repay (see Section 2.3 for discussion).

C.1.4 Outside Credit Access Index

This index is the equally-weighted average of the following standardized components, each measured from monthly credit bureau snapshots: (i) 1(has an open credit card); (ii) 1(has an open credit line that is not a credit card); (iii) total number of open credit lines; (iv) 1(has an open installment loan from another lender).[8] Absence of data is coded as a zero, given the nature of credit bureau data. For the monthly versions, we use the snapshot provided by the bureau for each month. For the 12-month horizon versions, we utilize the 12th-month snapshot.

C.1.5 Outside Borrowing Index

This index is the equally-weighted average of the following standardized components, each measured from monthly credit bureau snapshots: (i) 1(non-zero debt with other lenders) ; (ii) credit line utilization rate (i.e., share of total available credit limit outstanding); (iii) total owed on installment loans from other lenders; (iv) total owed on revolving credit lines from other lenders. Absence of data is coded as a zero, given the nature of credit bureau data. For the monthly versions, we use the snapshot provided by the bureau for each month. For the 12-month horizon versions, we average across months.

[8] We pre-registered to also include: (i) an indicator for having a credit score and (ii) credit score (conditional on having one). But we were unable to obtain the requisite credit score variable.

C.1.6 Outside Debt Repayment Index

This index is the equally-weighted average of the following standardized components, each measured from monthly credit bureau snapshots: (i) 1(any account 30+ days past due); (ii) number of accounts 30+ days past due; (iii) total balances 30+ days past due.[9] We define these only for loans outstanding at time of treatment assignment; borrowers with no such loans are coded as missing. For the monthly versions, we use the snapshot provided by the bureau for each month. For the 12-month horizon versions, we average across months.

C.1.7 Risky Driving Measures

This index is the equally-weighted average of the following standardized components:

(i) *Number of accidents.* Based on insurance claims filed with the lender.

(ii) *Instances of excessive speeding.* Count of GPS-recorded instances with speed of > 120 km/hr.

For the monthly versions, we consider the count for that month. For the 12-month horizon versions, we average across the monthly versions.

We lack insurance data for the 9% of borrowers insured by someone other than our lender, so we must decide how to deal with that missing data. We construct two versions of the summary index: one includes only borrowers with nonmissing accidents, and the other version of the index includes all borrowers by using only the speeding component for borrowers with no accident data.

C.2 Equity measurement and decomposition

Here we detail our measurement and decomposition of vehicle loan-to-value (LTV), which we use in Section 4.2 to document the prevalence of (deeply) underwater borrowers in our experiment sample, and in Section 5.2 to estimate some of our HTEs.

C.2.1 Measuring baseline equity

We define loan-to-value ($LTV_{i,t=-1}$) for borrower i at experiment baseline ($t = -1$) as:

$$LTV_{i,t=-1} = \frac{\text{Total Loan Amount Owed to the Lender}_{i,t=-1}}{\text{Market Value of the Vehicle Model}_{i,t=-1}}$$

The numerator is the total amount owed to the lender at experiment baseline: total loan principal outstanding after capitalizing arrears, as we do for every loan in our experiment (Section 4.1). The denominator is the estimated market value of the financed vehicle model at baseline, specifically from October 2023, since our data source provides quarterly estimates. We source the data through our lender, which obtains it from a vendor for use in underwriting and repossession operations.

[9] We pre-registered to also include 90+ days past due information for these components, but were unable to obtain the requisite variables.

Recall that Figure A.6 shows the distribution of baseline LTV for delinquent borrowers in our experiment. For estimating some of the HTEs in Section 5.2, we define baseline LTV discretely: borrowers with high baseline equity are those with baseline LTV less than or equal to the median LTV in our experiment sample. For the purposes of interpreting HTEs on this measure, it is useful to understand the sources of variation in our continuous estimate of LTV (as discussed in Section 5.2.1), and we detail that next.

C.2.2 Decomposition of baseline equity

To quantitatively estimate which vehicle, borrower, and/or loan characeristics explain variation in baseline LTV, we apply an R-squared decomposition model based on the Shapley value framework (Huettner and Sunder 2012; Biasi and Ma 2022). The empirical model here includes a broad set of vehicle attributes ((log) price paid by borrower at loan origination and dummies for the vehicle's make and model, vintage, and fuel type), borrower characteristics (dummies for borrowers' age and gender), features of the loan at origination (interest rate, LTV, and dummies for loan maturity), and indicator variables for the remaining loan maturity as of the month preceding the experiment. Taken together, these ten (sets of) variables explain an estimated 94% of the variance in baseline LTV.

The key results for our purposes are illustrated in Appendix Figure A.11, which shows the share of variance in baseline LTV explained by each variable or set of variables. The key factor is remaining maturity, which alone explains about half of the variance. This makes sense given the lack of heterogeneity in vehicle characteristics, borrower characteristics, and loan terms documented in Section 2. It also points to the importance of loan amortization schedules wherein principal is repaid gradually over time, especially earlier in the loan term. The next most important contributor to LTV variation, vehicle purchase price, explains only about one-fourth as much variation as remaining maturity.[10] Initial loan terms– interest rate, initial maturity, and down payment as reflected in initial LTV– together account for only about 18% of the variance in borrowers' baseline LTV.

D Machine Learning Approach to Predicting Liquidity (Online Appendix)

We now detail how we predict borrower liquidity during our experimental period, for use in HTE tests of the second half of Prediction 3 and Prediction 4. Following Kent (2020), Chernozhukov, Demirer, Duflo and Fernández-Val (2025), and other references below, our statistical approach uses the control group to predict our counterfactual of interest, balanced folds to reflect our experimental design, iteration over folds to cross-validate and minimize overfitting, and a combination of estimates from different machine learning (ML) methods to improve prediction accuracy and robustness.

[10] Since the empirical model includes vehicle characteristics that should determine its cost, vehicle purchase price here likely reflects the contribution of heterogeneity in markups to variation in LTV.

D.1 Econometric Problem and Approach

We seek to predict each borrower's liquidity $\mathcal{L}_i$ during our experiment and are interested in doing so net of any treatment effects. In other words, we wish to predict counterfactual liquidity: liquidity in the absence of any treatment. As such, we predict counterfactual liquidity for each borrower by training machine learning models exclusively on control group observations to predict $\mathcal{L}_i$.

Lacking a more direct measure of available liquidity with sufficient variation (as discussed in Section 5.2.2), we predict the likelihood of repayment to our lender in the absence of treatment. This likely proxies well for the borrower's available liquidity, given our various other results on the importance of liquidity constraints as a key driver of repayment behavior.

D.2 Data Details

We estimate $\mathcal{L}_i$ by predicting the likelihood that each borrower i was current on their minibus taxi loan after month 6 of the experiment rollout.[11] Our predictors are 28 borrower and loan characteristics (Z_i), measured at baseline or loan origination, specifically: minibus taxi loan characteristics (interest rate at origination, maturity at origination, remaining maturity, LTV at origination, high or low LTV at baseline, scaled arrears just prior to the experiment, share of loan principal outstanding at baseline, log of total loan amount outstanding with the lender at baseline, share of owed installment made in the three months prior to random assignment), vehicle characteristics (old or new vehicle, fuel type), borrower characteristics and demographics (age, gender, high or low debt-to-income at baseline, indicator variable for borrower having a transport permit, number of loans borrower has outstanding with the lender at baseline), and outside borrowing from credit bureau data (credit access index, borrowing index, credit score, high or low credit utilization, high or low estimated debt to income, indicator for open credit card, indicator for other open credit line, indicator for other installment loan besides the loan with our lender, number of credit lines, whether the borrower has non-zero debt outside of minibus debt, total balance on installment loans, total balance on credit lines).

D.3 Balanced Folds with Cross-Validation: Details

Following Wager and Athey (2018), we create balanced folds that maintain the treatment/control ratio within each validation fold. This ensures that each fold contains representative samples from both treatment and control groups. Following the recommendation in Hastie, Tibshirani and Friedman (2009) for our control group sample size, we use five folds to ensure that training sets maintain sufficient control observations for model fitting.[12] Specifically:

[11] We decided on month 6, a midpoint of our experimental period, for three reasons: 1) Allow enough time to elapse for borrowers to relapse into delinquency (recall that everyone in our experiment is brought current at baseline, per the lender's standard modification), 2) Capture the construct of interest, which is liquidity *during* the experiment period, 3) Later in the 12-month window would likely sacrifice predictive power, since we are using baseline information to predict.

[12] There is a bias-variance tradeoff to consider when choosing the number of folds – more folds decrease variance but increase bias.

Step 1. Balanced fold creation. We begin by separating our experiment sample into treatment (T= Payment Reduction and Debt Reduction arms) and control (C) groups. We then create $K = 5$ sub-folds for each of the treatment group: $T_1, T_2, \ldots, T_5$, and for the control group: $C_1, C_2, \ldots, C_5$. Finally, to create Fold_k, we combine the data from the *k-th* treatment and control folds ($\text{Fold}_k = T_k \cup C_k$). We summarize this step in Algorithm 1:

Algorithm 1 Balanced Fold Creation

Split by assignment: $\mathcal{C} \leftarrow \{i : T_i = 0\}$ (controls), $\mathcal{T} \leftarrow \{i : T_i = 1\}$ (treated).
Create K folds for treatment group: $T_1, T_2, \ldots, T_K$ ($K = 5$)
Create K folds for control group: $C_1, C_2, \ldots, C_K$ ($K = 5$)
for $k = 1$ to K **do**
    $\text{Fold}_k = T_k \cup C_k$
end for

We then iterate over each fold, using the focal fold as a test set and the remaining folds as training sets. Using *only control observations* from the training sets, we train a machine learning (ML) algorithm to generate predicted liquidity $\hat{\mathcal{L}}_i \equiv \hat{\mathcal{L}}(Z_i)$ for *all observations* in the test set. We use the software package `R` for our estimation. Specifically:

Step 2. Prediction by Cross-Validation. For each fold $k \in \{1, 2, 3, 4, 5\}$, Fold_k serves as the test set, while the remaining folds $\{1, \ldots, K\} \setminus \{k\}$, in combination, serve as the training set. We then train each of four ML methods — Elastic Net Regression (EN), Gradient Boost Machines (GBM), Principal Component Neural Networks (pcaNNet), and Random Forests (RF) — using *only control observations* from the training set, and generate predictions for *all observations* — both control and treatment — in the test set. This procedure prevents overfitting by ensuring that the ML model never observes test data during training. Training only on the control observations thus maintains the validity of causal inference. Importantly, the cross-validation methodology ensures that all observations, including the control group, receive out-of-sample predictions. We summarize this step in Algorithm 2:

Algorithm 2 Prediction

Require: Data $\mathcal{D} = \{(L_i, Z_i, T_i)\}_{i=1}^{N}$ where L_i is "current after 6 months", Z_i are baseline covariates, $T_i \in \{0, 1\}$ is assignment (treated/control); number of folds $K = 5$; $\{\text{Fold}_j\}_{j=1}^{K}$ from Step 1; learner set $\mathcal{M} = \{\text{Elastic Net (EN)}, \text{GBM}, \text{pcaNNet}, \text{Random Forest (RF)}\}$
Ensure: Out-of-sample predictions $\widehat{L}_i^{(m)}$ for each learner $m \in \mathcal{M}$; median ensemble prediction $\widehat{L}_i$
for $k = 1$ to K **do**
    test $\leftarrow \text{Fold}_k$;   train $\leftarrow \bigcup_{\ell \neq k} \text{Fold}_\ell$
    for each $m \in \mathcal{M}$ **do**
        Fit $\widehat{g}^{(m,k)} : Z \mapsto L$ on *controls only* in train: $\{(L_i, Z_i) : i \in \text{train}, T_i = 0\}$
        For all $i \in$ test, set $\widehat{L}_i^{(m)} \leftarrow \widehat{g}^{(m,k)}(Z_i)$
    end for
end for

D.4 Combining Estimates from Different Machine Learning Methods

Our final step takes the median prediction across the estimates produced for each borrower i by the four different ML methods described below.[13] Specifically:

Step 3. Summary Prediction: $\widehat{L}_i \leftarrow \text{median}(\widehat{L}_i^{(\text{EN})}, \widehat{L}_i^{(\text{GBM})}, \widehat{L}_i^{(\text{pcaNNet})}, \widehat{L}_i^{(\text{RF})})$

Timmermann (2005) shows that combining predictions this way frequently dominates the single best learner in out-of-sample tests.

The four different ML methods each capture different aspects of the relationship between predictors and the predicted outcome.[14] Elastic Net Regression combines LASSO and Ridge regression penalties, providing both variable selection and coefficient shrinkage. This method is particularly effective when dealing with correlated predictors and it performs automatic feature selection. Gradient Boosting Machines builds models sequentially, where each new model corrects errors from previous models. This ensemble method excels at capturing non-linear relationships and interactions between variables. Principal Component Neural Networks combines dimensionality reduction through principal component analysis with neural network prediction, making it effective for high-dimensional data while maintaining interpretability. Random Forests creates multiple decision trees using bootstrap samples and random feature selection. This method provides natural handling of non-linearities and variable interactions while being robust to outliers.

These four methods' predictions are strongly correlated with each other in our data, with pairwise correlations ranging from 0.58 to 0.95. The median difference between the highest and lowest prediction is 11pp, with a standard deviation of 9pp.

Appendix References (For Online Publication)

Biasi, Barbara, and Song Ma. 2022. "The Education-Innovation Gap." *NBER Working Paper 29853.*

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val. 2025. "Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, With an Application to Immunization in India." *Econometrica*, 93(4): 1121–1164.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* . 2nd ed., New York:Springer.

Huettner, Frank, and Marco Sunder. 2012. "Axiomatic Arguments for Decomposing Goodness of Fit According to Shapley and Owen Values." *Electronic Journal of Statistics*, 6: 1239–1250.

Kent, et al. 2020. "The Predictive Approaches to Treatment Effect Heterogeneity (PATH) Statement." *Annals of Internal Medicine*, 172(1): 35–45.

[13]Specifically, we take the midpoint of a linear interpolation between the two middle estimates.

[14]The key criterion for including a method as an input to Summary Prediction is simply that it has good predictive performance using baseline data (Chernozhukov, Demirer, Duflo and Fernández-Val 2025).

Timmermann, Allan. 2005. "Forecast Combinations." In *Handbook of Economic Forecasting.* Vol. 1, , ed. Robert F. Engle and Halbert White, Chapter 4, 135–196. Elsevier.

Wager, Stefan, and Susan Athey. 2018. "Estimation and Inference of Heterogeneous Treatment Effects using Random Forests." *Journal of the American Statistical Association*, 113(523): 1228–1242.