

NBER WORKING PAPER SERIES

CAN MICROFINANCE UNLOCK A POVERTY TRAP FOR SOME ENTREPRENEURS?

Abhijit Banerjee
Emily Breza
Esther Duflo
Cynthia Kinnan

Working Paper 26346
<http://www.nber.org/papers/w26346>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2019, Revised July 2024

We thank Bruno Barsanetti, Ishan Bhatt, Ozgur Bozcaga, Janjala Chirakijja, Ofer Cohen, Ritesh Das, Harsh Dev Goyal, Leticia Donoso Pena, Harris Eppsteiner, Zoe Hitzig, Taylor Lewis, Cecilia Peluffo, Sneha Stephen, Laura Stilwell and Yuta Toyama for their excellent research assistance. We thank the Centre for Microfinance at the Institute for Financial Research and Management, especially Parul Agarwal, for their help with the survey implementation. We thank Paco Buera, Edward Glaeser, Rema Hanna, Dan Keniston, Asim Khwaja, Maggie McMillan, Rohini Pande, Michael Peters, K.B. Prathap, Natalia Rigol, Ben Roth, Neng Wang and Bilal Zia for their comments as well as numerous seminar and conference participants. We are grateful to the NSF for generous financial support. Previous title: “Does Microfinance Foster Business Growth? The Importance of Entrepreneurial Heterogeneity.” MIT IRB # 1203004973. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Abhijit Banerjee, Emily Breza, Esther Duflo, and Cynthia Kinnan. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Can Microfinance Unlock a Poverty Trap for Some Entrepreneurs?

Abhijit Banerjee, Emily Breza, Esther Duflo, and Cynthia Kinnan

NBER Working Paper No. 26346

October 2019, Revised July 2024

JEL No. D25,O1,O14,O16

ABSTRACT

Can microcredit help unlock a poverty trap for some people by putting their businesses on a different trajectory? In Hyderabad, India, where microfinance was found to have very little short-run impact on the average business, we find that “gung ho entrepreneurs” (GEs), households who were already running a business before microfinance entered, show persistent benefits that increase over time. Six years later, the treated GEs own businesses that have 35% more assets and generate double the revenues as comparable households in control neighborhoods. They also borrow substantially more from other (non-microfinance) sources. We find almost no effects on non-GE households. A model of technology choice in which talented entrepreneurs can access either a diminishing-returns technology, or a more productive technology with a fixed cost, generates dynamics matching the data. These results show that heterogeneity in entrepreneurial ability is important and persistent. For talented but low-wealth entrepreneurs, short-term access to credit can indeed facilitate escape from a poverty trap.

Abhijit Banerjee
Department of Economics, E52-540
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
banerjee@mit.edu

Esther Duflo
Department of Economics, E52-544
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
eduflo@mit.edu

Emily Breza
Harvard University
Littauer Center, M28
1805 Cambridge Street
Cambridge, MA 02138
and NBER
ebreza@fas.harvard.edu

Cynthia Kinnan
Department of Economics
Tufts University
177 College Ave #620
Medford, MA 02155
and NBER
cynthia.kinnan@tufts.edu

1. INTRODUCTION

The idea that non-convexities may play a role in the persistence of poverty has a long history (e.g., Dasgupta and Ray (1986); Banerjee and Newman (1993); Ahgion and Bolton (1997); Lloyd-Ellis and Bernhardt (2000); Banerjee and Duflo (2005); Buera et al. (2011)). This idea is supported by anecdotal evidence that, within the same sector, poor (micro)entrepreneurs use technologies that are inefficient but cheap, while wealthier entrepreneurs use technologies that are more efficient but have high up-front costs (Lewis et al., 2001). At the household level, Balboni et al. (2022) find evidence consistent with poverty traps among the ultra-poor in Bangladesh.

However, rigorous empirical support for this idea among businesses has been elusive. In this paper we document that non-convexities can and do give rise to poverty-trap dynamics.¹ We also uncover an important reason why the search for such evidence has been challenging: households are heterogeneous, with some having the potential to move up into a more-efficient, higher-fixed-cost mode of operation while others are unable or unwilling to do so.

In such a world, the impact of improved access to credit would be heterogeneous in a specific way. Some people borrow purely for consumption, with no interest in starting a business. Others might use increased access to credit to start a business, but only have access to an inefficient technology with diminishing returns, and thus choose a small target business size. For these types of borrowers, the impacts of a credit expansion on revenues and profits are likely small at best. In contrast, borrowers with access to a high return-high fixed cost technology, who we call “gung-ho entrepreneurs” (henceforth GEs) have a larger target business size and are better positioned to take full advantage of additional credit. As a result, the revenue and profit effects for these households are potentially large. Moreover, access to credit allows some GE households to escape a poverty trap, if it enables them to switch into higher-return, higher fixed-cost technologies.

In this paper, we combine reduced-form and structural methods to look for systematic evidence of heterogeneity, the existence of non-convexities in business production, and poverty trap dynamics. To do this, we use a new wave of data collection on the sample of households that were included in a microcredit RCT in Hyderabad, India previously studied in Banerjee et al. (2015). In 2006, a microfinance lender, Spandana, entered into 52 neighborhoods, randomly selected out of 104 in a matched-pair design. We analyze borrowing, business and consumption outcomes at two points in time: i) 2 years after the lender moved into treatment areas, but prior to the control group receiving access (as captured in Banerjee et al. (2015)); and ii) 6 years after treatment areas received access — 4 years after control areas had also been treated,

¹The literature examining the persistence of poverty and the possibility of multiple equilibria uses various concepts including non-convexities, local increasing returns, and fixed costs. It is challenging if not impossible to differentiate these concepts, which have similar implications – some firms will be ‘trapped’ at a small scale while others, with similar fundamentals but better initial conditions (e.g., initial wealth), will be able to grow much larger. We will use these notions interchangeably.

and two years microcredit access was discontinued everywhere in Andhra Pradesh following a policy crisis. As of the 6-year survey, no household had ongoing access to microcredit.

We first sketch a simple framework incorporating non-convexities in business investment decisions to better understand the circumstances under which poverty traps may arise and to generate testable reduced form predictions. With globally decreasing returns, we should expect modest short-run effects of a microcredit expansion; these effects will diminish after the control group is treated. In contrast, persistence in business impacts is consistent with regions of the production frontier having non-decreasing returns, such as a non-convexity from fixed costs of technology upgrading. However, this need not correspond to a setting with a poverty trap if microfinance simply accelerates asset accumulation that would have already happened in its absence. We use the model's predictions on informal borrowing to further diagnose non-convexities and point to possible poverty trap dynamics. Under a poverty trap, entrepreneurs might opt to borrow less than their full available credit supply if the largest attainable investment level is not large enough to profitably switch technologies (or benefit from the non-convex region of returns). If however entrepreneurs face a high return on investment relative to borrowing costs, they should exhaust their total available credit supply, including from informal sources. Thus, by allowing some entrepreneurs, who were previously too far below the nonconvexity to reach it through a combination of formal and informal borrowing, microfinance can crowd in other borrowing.

In addition, we hypothesize that these potential poverty trap dynamics are only relevant for GEs, who have the complementary inputs (such ability, motivation, skills, physical mobility or others) needed to expand their businesses when the opportunity arises. We proxy for GE status by whether a household selected into entrepreneurship in the absence of microfinance access, when credit constraints were tighter. We hypothesize that in contrast, "reluctant entrepreneurs" (henceforth REs) are more likely to opt into entrepreneurship when the credit constraints are looser (i.e., after microcredit entry). This distinction is motivated by previous papers' findings of heterogeneity in microfinance impacts along the dimension of baseline business ownership (Banerjee et al., 2015; Meager, 2019; Chernozhukov et al., 2018) and the literature on entrepreneurial entry and financial constraints (e.g., Kerr and Nanda 2009).

We present reduced form analysis and interpret the results using our motivating framework, applied specifically to GEs. First, we show that the relationship between initial assets and subsequent wealth GEs exhibits the characteristic *S*-shape. Moreover, the relationship appears shifted between treatment and control, with treated households requiring less initial wealth to access the steepest part of the *S*-shaped curve.

Second, contrary to short run results, after 6 years, we find positive, statistically significant average impacts on a number of key business outcomes including total entrepreneurship rates, profits, business scale (purchases and stock of assets), revenues, expenses and and employment. However, these positive persistent impacts are entirely driven by the GE subsample. These

firms business outcomes are all higher in treatment neighborhoods by statistically and economically significant margins. In particular, asset stocks increase by 40% , revenues more than double, relative to GEs in control areas, We also find positive and significant effects on the average profits of GEs, and large significant impacts on both business and non-business durable spending. These persistent results on business outcomes, for GEs only, are consistent with both the existence of heterogeneity and the presence of investment opportunities with non-decreasing returns, as outlined in our simple framework.

Third, we show that microfinance crowds in informal borrowing (especially business borrowing) for GEs, but not for non-GEs. In response to additional microcredit access, these households seek out even more business credit from other sources. As our framework highlights, households who can eventually reach the high-return non-convex production region even in the absence of microcredit would prefer exhausting all of their potential credit supply before microcredit became available, leaving no scope for crowd-in. Crowd-in of finance among the GEs thus demonstrates that some GEs were trapped in a low equilibrium before microcredit — an that microcredit helped some of these GEs to escape. In response to additional microcredit access, these households seek out even more business credit from other sources. As our framework highlights, households who can eventually reach the high-return non-convex production region even in the absence of microcredit would prefer exhausting all of their potential credit supply before microcredit became available, leaving no scope for crowd-in. Crowd-in of finance among the GEs thus demonstrates that some GEs were trapped in a low equilibrium before microcredit — an equilibrium that microcredit helped some of these GEs to escape.

Fourth, we rule out alternative explanations of the differences between GEs and the REs. One concern is that the non-GE sample is composed of both REs (who are induced to start a business by MF entry) and consumption borrowers. Another concern is that, even within the group who start businesses, REs and GEs differ on other margins, such as wealth or age. First, in a subsample where we have baseline wealth, we show that initial wealth differences do not account for the results. Second, we exploit the staggered entry of Spandana to show that they are not explained by business age either: Spandana entered Hyderabad in a staggered fashioned over 13 months, allowing us to observe businesses in different treatment areas that opened up at the exact same time, some of whom opened before Spandana entered, while others opened after. Because randomization was done at the matched pair level, for each treated area, we have a pre-identified control area which serves as a counterfactual. Comparing firms that opened in this common period before Spandana arrived in a treatment area to those that opened at the same time in the matched pair gives us the pure GE treatment effect. A similar comparison for firms that opened in this period, but *after* Spandana arrived, gives the combined selection and treatment effects for RE firms. If the REs are negatively selected, then the treatment-control differences for the pre-Spandana firms (GEs) should be larger than those for the post-Spandana firms (REs) created at the same point in calendar time. If, on the other hand, the GEs have

larger effect on average just due to their head start, then within this “overlapping” sample we should see no differential treatment-control differences.

We find that, within this sample, the young GE businesses experience large positive treatment effects, while the non-GEs demonstrate a large 0.18 standard deviation *decrease* in a business index (p -value = 0.105). This difference reflects both treatment and selection. A back-of-the-envelope decomposition shows that, if the true treatment effect on the REs is zero, then the negative selection effect is large: the non-GEs would be worse by two-thirds of a standard deviation along an index of business outcomes.

Next, we use a simple structural approach to quantitatively test for the presence of poverty traps among GEs. We build from our motivating framework, which allows for two different technologies: one with diminishing returns and the other with constant returns, but which requires a fixed cost. We first use the short run responses of investment and revenue to credit to identify the parameters governing the production possibility frontier and show that a model with two technologies, one of which includes a fixed cost, fits the data better than a single globally concave production function with no fixed cost. We then embed the production technology into a dynamic problem to characterize the transition diagram mapping current to future wealth.

At the estimated parameters, the model generates a process in which the impact of temporary access to some additional credit causes a divergence among the GE firms, thereby helping to explain the large persistent impacts on the GEs. This divergence occurs in large part because the model generates a poverty trap: without cheap credit, talented but low-wealth households cannot afford the minimum efficient scale of the better technology, and so remain stuck at the maximum efficient scale for the diminishing returns technology. The key role of microcredit in this model is to reduce the minimum wealth level at which households can switch into the better technology, allowing intermediate-wealth households to escape the poverty trap. Households escaping from the poverty trap are a quantitatively important driver of the persistent effects we observe: these households explain two-thirds of growth in revenues and almost three quarters of growth in capital stocks. The remainder is explained by households who were already out of the poverty trap zone further scaling up their businesses. This exercise helps clarify how large, persistent effects of early access to microcredit could have arisen as a result of an intervention which provided only temporary differential exposure to microcredit for the treatment group.

This paper builds on a large body of evidence studying the returns to microfinance (Attanasio et al. (2015); Augsburg et al. (2015); Crépon et al. (2015); Karlan and Zinman (2009); Tarozzi et al. (2015); Angelucci et al. (2015)). While we follow others in our focus on heterogeneity in the returns to credit (specifically, Banerjee et al. (2015); Angelucci et al. (2015); Meager (2019, 2022)), to our knowledge we are the first to document a pattern of divergence over time between treatment and control groups. In our case, this is driven by a predictable group – individuals

who opted into self-employment when credit constraints were tight.² Importantly, we combine two previously unstudied sources of variation from the original Spandana experiment – the date of establishment of household businesses and the timing of the lender’s roll-out through Hyderabad - to show that some of the heterogeneity is driven by selection, not experience or entrepreneur age. This also allows us to show that businesses started because of microfinance indeed have lower potential, holding age and experience fixed. Our results on wages and employment also relate to a small theoretical and empirical literature on the equilibrium impacts of microcredit (Buera et al., 2021; Breza and Kinnan, 2021).

As mentioned above, while many theoretical papers are built on the idea that non-convexities can be a source of poverty traps, finding concrete evidence has been very challenging, and this is the first paper that illustrates the mechanism for small-scale entrepreneurs outside of extreme poverty. Our paper joins a few other contemporary studies presenting evidence of non-convexities and multiple equilibria in different settings. In addition to Balboni et al. (2022), Kaboski et al. (2022) also investigate the possibility of non-convexities. Using an experiment in which Ugandan households were given a choice over riskier vs. safer lotteries, they find evidence of increasing returns, but driven by a completely different mechanism—land purchases. Bari et al. (2024) conduct an experiment in Pakistan which suggests that nonconvex capital adjustment costs can give rise to persistent effects of larger loans to finance asset purchases.

Finally, our findings are also related to a literature exploring why many businesses appear to be our so-called “reluctant entrepreneurs” De Mel et al. (2010), Schoar (2010) and Adhvaryu et al. (2019) suggest that microentrepreneurship may largely be used as an outside option when wage jobs are unavailable or as a tool for mitigating shocks. Relatedly, Breza et al. (2021) show that when aggregate labor supply exogenously decreases in Indian villages, households reduce their self-employment work by almost one-quarter to work in market-rate wage jobs.

The remainder of the paper is organized as follows. We first present a simple dynamic framework of entrepreneurial investment under potential non-convexities in section 2. In section 3, we describe the experimental setup the AP microfinance ordinance, and the timing of data collection. In Section 4, we generate testable predictions, and we provide reduced form tests in Section 5. In Section 6, we present our model and estimation results, and show how the parameters shed light on the poverty trap. Section 7 concludes.

2. FRAMEWORK

We present a simple dynamic framework for an entrepreneurial household with potential non-convexities in the production technology. This kind of model has been shown to give rise to poverty trap dynamics under certain circumstances. Our goal is to characterize reduced form

²Additionally, Karaivanov and Yindok (2022) estimate a model which distinguishes “voluntary” from “involuntary” entrepreneurship using data from urban Thailand and examines heterogeneous responses to credit.

and structural predictions in the presence of a poverty trap, using the entry of microfinance as an exogenous shock.

2.1. Model Basics. Households are endowed with initial wealth W_0 and maximize the discounted sum of utility from consumption:

$$(2.1) \quad U(c_t)_{t=0}^{\infty} = \sum_{t=0}^{\infty} u(c_t)$$

Households can generate income from running a business. Entrepreneurs have a production technology $Y(K_t)$ which transforms capital, K_t into output, and which may exhibit non-convexities. We do not model labor inputs; this is consistent with our setting in which few businesses have employees, and can also be interpreted as a function in which the labor decision has been concentrated out. Figure 1 plots sample production frontiers, with a globally concave relationship in Panel A and a production frontier with a nonconvexity in Panel B.³

Households may want to borrow to fund business investment. This borrowing may come from microfinance (when available) or informal loans (from friends, family and moneylenders) at gross interest rate R . However, their credit supply is capped at \bar{b} , and B_t must therefore satisfy

$$(2.2) \quad 0 \leq B_t \leq \bar{b}.$$

At the end of the period, the entrepreneur’s financial profits $\pi(K_t)$ are equivalent to the firm’s output net of borrowing costs,

$$(2.3) \quad \pi(K_t) = Y(K_t) - RB_t.$$

Finally, households are endowed with an outside option investment technology (e.g., savings under-the-mattress or in a bank) with return ρ . Note that in our context, the borrowing cost (both formal and informal) is substantially higher than interest on bank deposits, so $R \gg \rho$. Any wealth not used in the business S_t is invested in this outside option technology.

At the end of period t , the entrepreneur receives her business profits net of borrowing costs, liquidates her business capital net of depreciation δ , and receives a return on the outside option investment. This resulting cash on hand can either be consumed in period t , as C_t , or passed to the next period as household wealth W_t ,

$$(2.4) \quad \underbrace{\pi(K_t) + \rho S_t + (1 - \delta)K_t}_{\text{cash on hand}} = C_t + W_t$$

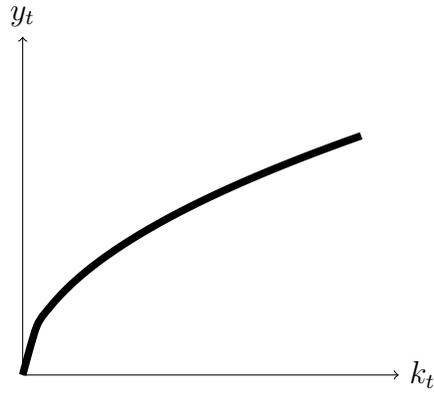
At the beginning of the next period, the entrepreneur decides how much to borrow, and how much to allocate the combination of her last period’s wealth W_t and next period’s borrowing,

³Consistent with our structural estimation in Section 6, we posit that the nonconvexity comes from a switch to a more productive technology requiring an upfront fixed cost investment; however the conclusion would be similar if the nonconvexity arose from a single production function with a region of increasing returns.

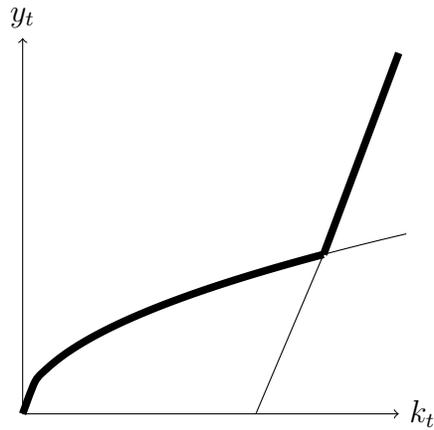
B_{t+1} between business capital K_{t+1} and the low-return outside option technology according to

$$(2.5) \quad S_{t+1} = W_t + B_{t+1} - K_{t+1}.$$

2.2. Wealth Dynamics and Poverty Traps. Solving the household's constrained optimization problem in each period yields policy functions for K_t , S_t and B_t , which depend on the state variable W_{t-1} and the model parameters. The solution to the household program also sheds light on the evolution of household wealth. Figure 2 shows three examples of possible wealth transition diagrams $F(W_{t-1})$.



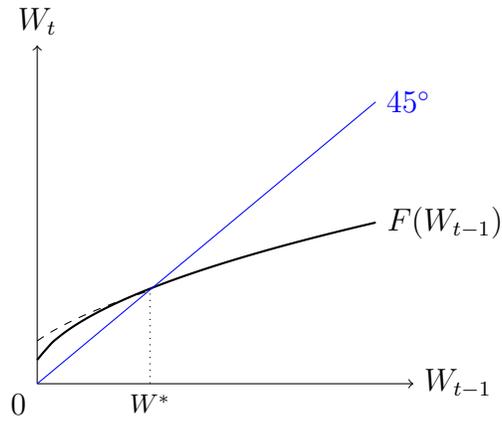
I: Concave production function



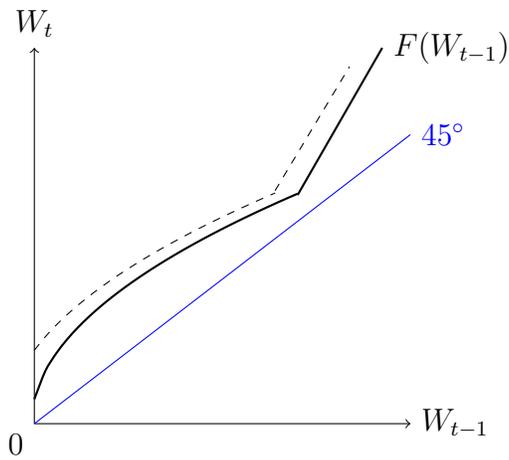
II: Non-convex production function

FIGURE 1. Possible production functions

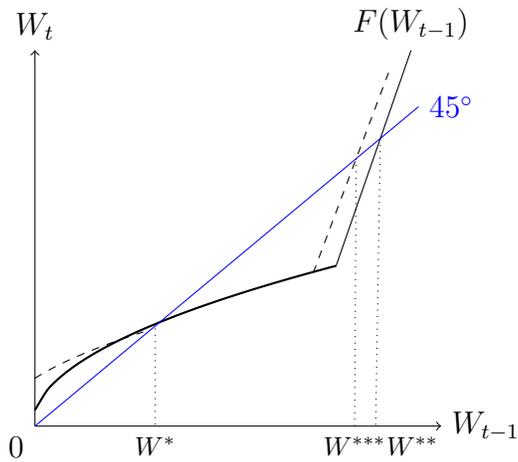
In Case A, which assumes a globally concave production frontier as in Figure 1 Panel I, the relationship between wealth today (W_t) and wealth yesterday (W_{t-1}) is globally concave, and there is a unique steady state, W^* where the curve crosses the 45 degree line. In this scenario,



A: Concave production function



B: Non-convexity, no poverty trap



C: Non-convexity and poverty trap

FIGURE 2. Possible transition dynamics

all individuals, regardless of initial wealth, will converge to wealth W^* . There is no poverty trap.⁴

In Cases B and C, we instead assume a production frontier with a nonconvexity, as in Figure 1 Panel II. This generates a non-convexity in the wealth transition diagram, giving it the well-known S-shape. However, this is not a sufficient condition for a poverty trap. In Case B, for example, the curve relating W_{t-1} to W_t has no steady state (i.e., the curve never crosses the 45 degree line). That means any household starting with non-zero levels of wealth will continue to get richer every period (We do not aim to differentiate settings where households continue to accumulate wealth forever from those where households eventually reach a higher steady state outside the support of the figure). There is no poverty trap.

In contrast, the wealth transition diagram in Case C does admit a poverty trap. The curve now crosses the 45 degree line multiple times. As in Subfigure A, there is a stable steady state at W^* , but also an unstable steady state at W^{**} . This means that anybody with initial wealth below W^{**} will end up converging to W^* . However, anybody with initial wealth greater than W^{**} will have increasing wealth. This scenario admits a poverty trap.

2.3. The role of formal credit. How is the transition diagram altered by access to formal credit? Formal credit allows individuals to borrow a larger amount than was previously possible, thereby allowing them to move to the better technology at a lower level of initial wealth (W_{t-1}). The dashed line segments in Figure 2 indicate this shift. In case C, with a poverty trap, the minimum level of wealth needed to escape the bad equilibrium falls from W^{**} to W^{***} . In case B, with a non-convexity but no poverty trap, the entire transition diagram shifts left, as even those with low starting wealth will take full advantage of additional credit in this case.⁵

Note that these different cases also have implications for the observed wealth and capital distributions in the data, assuming homogeneity in all parameters other than initial wealth. In Case A, which has a unique stable steady state, we would expect a unimodal distribution at low levels of wealth and capital. In contrast, in Case C, we should also expect a concentration of households at the low steady state, but we should also expect a larger mass of households at higher levels of assets and wealth. If there is also a higher steady state, then we should expect a bi-model distribution of wealth and assets.

2.4. Impacts of Microfinance. We next use this basic framework to analyze the impacts of microfinance, which raises the amount of credit supply \bar{b} , in each of the three cases. Specifically,

⁴The transition diagram does not pass through the origin because we assume that, even in the absence of formal credit, individuals have access to some informal borrowing.

⁵Another possibility is that access to microfinance creates a higher steady state where none existed before, by allowing households to borrow at lower interest rates and therefore shifting the transition diagram upward. Appendix Figure B1 depicts this case: without microcredit, the transition curve $F(W_{t-1})$ has a unique (low) steady state, but access to microcredit at a lower interest rate shifts the transition curve upward to $\bar{F}(W_{t-1})$, creating a second, higher steady state.

we aim to explore whether the pattern of impacts can identify the underlying regime. The dashed lines in Figure 2 plot the new transition diagrams with increased \bar{b} .

Business Outcomes. In Case A, relaxing the household's credit constraint will have limited impacts. Unconstrained households already at the steady state should experience no impact. Credit constrained households with initial wealth below the steady state, can use the extra credit to reach the steady state faster. These households should experience a short-run increase in capital investment and business profits. However, the businesses will stop growing once the new steady state is reached. Moreover, once the control group is treated, they should catch up quickly to the treatment group.

In Case B, all households in the range plotted in Figure 2 have strictly positive returns to borrowing more. Therefore, once microcredit relaxes households' borrowing constraints, we should see an increase in business investment and profits across the board – that is the transition diagram shifts and rotates to the left. If there is perpetual growth or if there is a new steady state at a substantially higher level of wealth, this increase should be more persistent. Due to the non-convexity in the production frontier and the wealth transition diagram, it is possible for the gap to widen between treatment and control households even after the control group is treated.

Finally, in Case C, relaxing the household borrowing constraint has three effects. First, for households that started with high enough wealth (e.g., those who were not stuck in the poverty trap), the impacts of microfinance access will be identical to Case B. Second, microfinance access will lower the threshold wealth W^{**} required to escape the poverty trap. These households now have increased incentives to invest in their business. We should also expect an increase to capital and business profits that persists over time, even once the control group is treated. Third, for households with wealth below W^* , the impacts will be modest and will look like Case A.

So, under a poverty trap, microfinance should have persistent positive impacts on business assets and profits, driven by households with higher initial wealth. However, note that Case B – with a non-convexity but no poverty trap – also shares the prediction of positive average treatment effects.

Consumption. While our simple model with non-convexities delivers clear predictions on business outcomes, the impacts on consumption are less straightforward. As in Kaboski and Townsend (2011), households may endogenously reduce consumption in the short run to facilitate larger investments in high-returning productive assets. While in the long-run we should expect an increase in consumption, the short-run consumption response is unsigned.

Non-Microfinance Borrowing. We next consider whether the response of the household's non-microfinance borrowing to microfinance can help identify between Case B and C.

In Case A, we should expect limited or no impact on non-microfinance borrowing. For unconstrained households, where there is margin to increase non-microfinance borrowing, access to microfinance has no impact.⁶ For constrained households, there is no scope to increase microfinance borrowing, as they are already maxing out their available credit supply.

In Case B, we also should expect a limited impact on non-microfinance borrowing. This is because all households have a strong incentive to use all of their credit supply before microfinance enters.

In Case C, we similarly should expect minimal effects on non-microfinance borrowing for those households who were already wealthy enough to avoid the poverty trap. However, the marginal household stuck in the poverty trap is *locally* unconstrained. Their credit demand is strictly less than supply. The entry of microfinance relaxes the household's *global* credit constraint, which is only valuable if it induces borrowing beyond the prior limit. This means that as long as microfinance credit supply isn't too large, there will be a set of households that both take microfinance and dramatically increase their non-microfinance credit demand. Thus, crowd-in of non-microfinance borrowing can distinguish cases B and C.

In Appendix Figure B1, we explore other cases that can generate these joint patterns of business expansion and crowd-in of borrowing. These cases tend to include some sort of non-convexity in the production frontier. Moreover, when microfinance enters, some individuals are moved into a better steady state or growth path. In Section 6, we use a structural approach to diagnose the patterns of steady states before and after the arrival of microfinance.

3. SETTING, EXPERIMENTAL DESIGN AND DATA

3.1. Experimental Design. We draw on variation generated by the randomized introduction of microfinance to neighborhoods in Hyderabad, India first studied by Banerjee et al. (2015). One hundred and four neighborhoods were randomized so that 52 received access to credit from Spandana, a large lender then moving into Hyderabad, starting in early 2006. The remaining 52 neighborhoods only received access in mid-2008, after a round of data collection conducted in late 2007 - early 2008. Spandana provided a relatively standard microfinance loan product. Borrowers, who were organized into groups, met on a weekly basis and made weekly installment payments. At the successful completion of a loan cycle, borrowers were offered larger loan sizes for subsequent cycles. Notably, within-group, these loans were largely homogenous in size and contract structure.

⁶If the microfinance borrowing cost is lower than the cost of non-microfinance borrowing, there could be a small crowd-out effect.

Banerjee et al. (2015) document, first, that demand for microcredit is not universal: at the first (2007-8) two-year follow-up, fewer than a third of treated households borrowed. Nonetheless, treated households did borrow more from microcredit institutions.⁷ Turning to downstream outcomes, both 2 and 4 years after the initial intervention, respectively, significant positive average impacts are seen on investment in durables. Treated households started more businesses, but there was no detectable effect on average business profits, consumption, women’s empowerment or human development outcomes.

Additionally, Banerjee et al. (2015) find that, in the short run, those whose businesses were already in existence before microcredit (the gung-ho entrepreneurs or GEs, in our taxonomy) invested more in those businesses. The average profits of these existing businesses increased, with particularly large gains at higher quantiles. In contrast, the average marginal new business performs worse, suggestively, in treatment versus control areas.⁸ Meager (2019) finds further support for the possibility of heterogeneity by baseline business ownership in a broader meta-analysis of Banerjee et al. (2015) and six additional microfinance RCTs. Our goal here is to further probe these results using additional, longer-run data and to interrogate the relevance of poverty traps in interpreting the results.

3.2. The AP Ordinance and Additional Data Collection. Coincidentally, the second wave of data collection was completed just a few months before the microfinance landscape abruptly changed, causing both treated and control areas to lose access to microcredit. On October 15, 2010, the AP government unexpectedly issued an emergency ordinance (The Andhra Pradesh Micro Finance Institutions Ordinance, 2010) to regulate the activities of MFIs operating in the state. Under the law, MFIs were not permitted to approach clients to seek repayment and were further barred from disbursing any new loans. In the months following the ordinance, almost 100% of microfinance borrowers in AP defaulted.

In 2012, roughly two years after the withdrawal of microcredit, we collected an additional, longer wave of endline data (six-year follow-up), allowing us to understand whether and how *past* access to credit affected household and business outcomes. At that time, MFIs in AP were still not permitted to operate under state law. Note that the AP Ordinance had two effects on borrowers in our sample – the default (windfall) effect and the effect of a reduction in future credit. In Section 5.5, we provide evidence that will help disentangle those two effects (and strongly suggests that the effects are not driven by the windfall effects).

3.3. Data. We supplement the data used in Banerjee et al. (2015) in two ways. First and most importantly, in mid-2012, we returned to the respondents of the 2010 survey round and

⁷We interpret the comparisons between treatment and control as measuring the effects of increased exposure to microfinance in general. The loans offered by Spandana were very similar to those of the competitors operating in Hyderabad at the time.

⁸In Section 5.4 below, we present a refined method for separating selection effects on productivity (those who were more marginal with respect to entrepreneurship were able to start businesses in treatment areas) from experience effects (new businesses are mechanically younger than the average GE).

conducted a six-year follow-up survey with 5,744 households located in 103 of the original 104 combined treatment and control neighborhood.⁹ The 2012 round of data collection also measured social network outcomes which are analyzed in (Banerjee et al., 2024).

At the time of the survey, it had been 6 years since the original treatment group was first exposed to microfinance and 4 years since the control group had gained access to microfinance from Spandana, the implementing partner. All of the respondents experienced a simultaneous withdrawal of microfinance from Hyderabad in response to the AP ordinance shortly after the 2010 survey round, 2 years before the 2012 survey.¹⁰ Therefore, when we compare outcomes between the original treatment and control groups, we measure the impacts of the intensity of *past* exposure to microfinance against a backdrop where microfinance is no longer available. Figure B2 shows the timeline of the experimental intervention, data collection, and changes to microfinance access.

Table 1 provides a description of the households surveyed in the 2012 round. The table displays the means of demographic, consumption, and business outcomes for households in the control group. We also include information about the borrowing behavior of these households at the time of the four-year (2010) follow-up survey, which is a close proxy for the household’s borrowing right before the AP crisis. Note that approximately 30% of the control group had an outstanding microloan at that time.

Appendix Table A6 shows that, relative to the two-year follow-up survey (2007), approximately 84% of households were located in the six-year (2012) survey round. The attrition rate is not differential across treatment vs. control. Nonetheless, to address whether our results might be sensitive to attrition, we compute Lee (2009) bounds for key outcomes (Appendix Table A7).

We additionally leverage a round of pre-intervention baseline data collected with 2,800 households prior to the rollout of Spandana branches. As noted in Banerjee et al. (2015), due to the need for a larger sample for the endline, the main analysis sample was largely re-drawn. However, approximately 500 households from the baseline were included in the two-year follow-up sample.¹¹ Crucially, for this Baseline-Endline sample we can observe measures of household wealth, a key state variable in our theoretical framework, *prior* to the treatment group getting access to microfinance.

4. TESTABLE PREDICTIONS

The model laid out in Section 2, combined with our empirical setting, generates several key predictions.

⁹One (treatment) area was dropped because it was used for piloting.

¹⁰The 2010 survey round, collected 4 years after the initial intervention, is analyzed in Banerjee et al. (2015). We only use this survey round here to document treatment-control differences in borrowing on the eve of the AP ordinance.

¹¹Due to a data issue, these households cannot be perfectly linked but instead must be manually linked using fuzzy name matching; 487 baseline households are successfully matched to the endline data.

4.1. Heterogeneity by Entrepreneur Type. In the framework in Section 2, we considered a single representative household. However, as we discuss above, the prior literature suggests that incumbent businesses benefit from an increase in credit supply, while those without an incumbent business do not appear to benefit from microfinance. Basic patterns in the data support the idea that borrowers are highly heterogeneous in terms of entrepreneurship preferences or potential. First, only approximately half (49.1%) of microfinance borrowers report having any business at all at the 2-year follow-up. Second, among borrowers who do report having a business, less than two thirds (63%) report using the loan proceeds for any business expense.

These patterns could potentially be explained by heterogeneity in wealth. However, we argue that type matters beyond wealth: some entrepreneurs are of the “Gung-ho” (GE) variety - they have the ability and motivation to grow their business beyond a subsistence level. Other entrepreneurs are of the “Reluctant Entrepreneur” (RE) variety — they are not able/interested in growing their business and may engage in entrepreneurship as a response to labor market frictions (Karaivanov and Yindok, 2022; Breza et al., 2021). One way to capture this difference in our conceptual framework is to postulate that REs face the concave production frontier in Figure 1 Panel A, while GEs additionally have access to the high-fixed cost, high=productivity technology in Panel B.

As discussed above, we proxy for GE/RE status by whether an entrepreneur opened her business before microfinance became available (a GE) or only after (an RE). Suggestive evidence in favor of such heterogeneity is shown in Figure 3. Under a poverty trap, there are multiple equilibria, so we should see bi-modality in wealth. We do, but, consistent with our assumptions, only for the GEs, for whom two distinct modes of baseline wealth are observed. For REs, there is a single mode, quite similar in magnitude to the lower of the two GE modes.¹²

4.2. Predictions.

Prediction 1: S-shaped Evolution of Assets with Wealth for GEs. Under non-convexities (cases B and C of Figure 2), the wealth transition of GEs should exhibit the characteristic *S* shape. Figure 4 plots 2012 business assets (the sum of business durable assets, business working capital, and wage bill) against baseline (2006; pre-treatment) assets — measured as the sum of business and non-business durable assets plus business inputs — separately for treatment and control. Two features are apparent. One, the relationship indeed exhibits the *S*-curve shape: for low levels of initial assets, the mapping to endline assets is very flat; then, at a certain point the relationship becomes markedly steeper, suggesting a more efficient transformation of baseline to endline assets.¹³ Secondly, the relationship appears left-shifted in treatment relative to control:

¹²We will use the term “Reluctant Entrepreneur” (RE) for households who opened a business after the entry of microcredit. When we refer to all households who are not GEs, whether or not they opened a business, we will use the term “non-GEs.”

¹³We do not attempt to draw a 45-degree line on this graph. Due to differences in the baseline and endline measures, as well as inflation (for cash savings) and depreciation (for real assets), drawing such a line is subject

the minimum level of initial wealth associated with the steeper portion of the S -curve shifts from nearly INR. 60,000 down to about INR. 45,000.

Prediction 2: Persistence of treatment effects on business outcomes for GEs. Under non-convexities (Cases B and C in Figure 2), a temporary increase in access to credit can lead to persistent increases in business investment, scale, and profits, relative to the more credit constrained regime. Our empirical setting, where the control group was treated with a two-year lag and where where microfinance was unexpectedly withdrawn four years after microfinance entered treated neighborhoods, allows us to directly test for persistence.

In the scenario with non-convexities and no poverty trap (Case B of Figure 2, the entrepreneur benefits from borrowing as much as possible as soon as credit is available. Because the business has large, sustained returns, any head start in borrowing and investment leads to persistent gains in profits and wealth, even after the control group is treated. While losing access to microfinance is likely costly to these entrepreneurs, the head start from longer access to microfinance is still preserved following microfinance's withdrawal. In the scenario with a poverty trap (Case C of Figure 2), the same exact logic applies for those entrepreneurs who already had enough wealth to escape the poverty trap prior to the arrival of microfinance.

As described above, in the Case C scenario, microfinance might also help some households make business investments large enough to benefit from the high-productivity technology and escape from the poverty trap. This should lead to short run positive treatment effects. When microfinance arrives in the control neighborhoods, a new set of entrepreneurs should also be able to escape the poverty trap, but with a lag. In contrast to the no poverty trap case, when microfinance is unexpectedly withdrawn, some entrepreneurs who were transitioning toward larger business sizes and higher levels of wealth may be pulled back into the poverty trap. Because treated GEs had more periods of borrowing, they should be less likely to suffer this fate than control GEs. These dynamics could actually lead to growth in the gap between treatment and control in the longer-run post-microfinance regime.

Prediction 3: Limited business treatment effects for REs. Given that REs only have access to the low-productivity, decreasing returns to scale, any short-run impacts of microfinance should be limited. Moreover, there should be even less scope for persistent benefits relative to control once the control group has also received microfinance.

Prediction 4: Crowd in of non-Microfinance Loans. A key implication of the presence of poverty trap dynamics (as in Case C) is that increased access to microcredit will crowd-in other sources of credit, for those who, without microcredit access, could not have escaped the poverty trap even by leveraging the credit available to them, but who are able to do so with the combination of microcredit and other financing. As we discuss above, other production technologies (e.g.,

to significant guesswork. Our point is simply to show the existence of nonlinear dynamics potentially consistent with a poverty trap.

Case A or Case B) do not predict such crowd-in. This prediction also holds in the longer-run, after microfinance is withdrawn unexpectedly. This is because households that have successfully escaped a poverty trap despite losing access to microcredit should still demand high levels of credit to invest in their high-return businesses, while those who fall back into the poverty trap should have lower credit demand.

We will test this prediction, both during and after the period when microcredit was available, by comparing the non-microcredit borrowing of GEs who exogenously received more vs. less access to microcredit due to their treatment assignment.

Prediction 5: Negative selection of RE businesses. If microfinance encourages the entry of RE businesses, these businesses should be negatively selected.

5. REDUCED FORM RESULTS

To test these predictions, we investigate the intent to treat (ITT) comparisons between the initial treated and control neighborhoods. We interpret the results of such comparisons as the impacts of having greater access to microfinance for four years instead of two (in the past). The average treatment effects regression takes the form

$$(5.1) \quad y_{in} = \alpha + \beta \times Treat_n + \delta_{s \ni n} + \varepsilon_{in}$$

where i indexes individuals and n indexes neighborhoods, y_{ia} are outcome variables (generally measured in 2012), $Treat_n$ is an indicator for treatment neighborhoods in the original study (where microfinance entered in 2006), and β is the coefficient of interest. $\delta_{s \ni n}$ is a stratum fixed effect.¹⁴ For all specifications, standard errors are clustered at the area level.

We are especially keen to understand the differential impacts for gung-ho entrepreneurs vs. other households. For these specifications, the regressions take the form

$$(5.2) \quad y_{in} = \alpha + \delta GE_{in} + \beta_1 \times Treat_n + \beta_2 GE_{in} \times Treat_{in} + \delta_{s \ni n} + \varepsilon_{in}$$

Here, we indicate that household i in neighborhood n is a gung-ho entrepreneur by setting $GE_{in} = 1$. The coefficient β_1 can be interpreted as the treatment effect on the non-GEs (who include both consumption borrowers and reluctant entrepreneurs (REs), who started a business because of the greater credit access), while the coefficient β_2 is the differential treatment effect for the GEs above and beyond the impact on the non-GEs. Thus, the total treatment effect for the GEs is $\beta_1 + \beta_2$.

5.1. Access to microfinance. Table 2, document how the intervention changed microfinance borrowing at different time horizons. Surveys were collected two years after treatment, 4 years after treatment, and 6 years after treatment. Panel A reports average effects of microfinance

¹⁴Altogether, there were 52 strata, or pairs. Pairs were formed to minimize the sum across pairs A, B (area A avg loan balance – area B avg loan balance)² + (area A per capita consumption – area B per capita consumption)². Within each pair one neighborhood was randomly allocated to treatment.

(pooling GEs and non-GEs) while Panel B reports heterogeneous effects by type. Column 1 shows treatment effect on likelihood of borrowing from any MFI two years post-treatment, before Spandana moved into the control neighborhoods. At this horizon, treatment households were approximately 11 percentage points more likely to have borrowed in the prior three years than control households.¹⁵ Column 2 shows the treatment effect measured at the survey four years post-treatment, after the control neighborhoods received access, and while microfinance is still active. It's a dummy for having borrowed at any point between 2004 and 2010. The expansion of microcredit is evident in the controls means : On half of the control group had borrowed by the 4-year survey. Households in the treatment group were were still 4.4 percentage points more likely to have ever borrowed (a 9% increase). The ordinance that led to the withdrawal of microfinance took place in October 2010, only a few months after the four-year follow up survey was collected.

The original treatment assignment could have also affected households via the intensive margin, namely the total amount of credit taken. Column 3 explores the total value of the MFI loans outstanding at the two-year follow-up survey. The average treatment household reports Rs. 1,294 more borrowing than the average control household, a 14% increase in credit over the control group. At the 4-year follow up (column 4) treatment households are borrowing Rs. 788 more, despite the entry of Spandana in the control areas. This is a combination of the extensive margin (they are still more likely to have ever borrowed) and intensive margin (older clients have higher credit limits). We cannot reject equality between GEs and non-GEs in any of the specifications, consistent with the idea that microfinance organizations do not screen borrowers.

5.2. Testing Predictions 2 and 3: Persistent effects on business outcomes, only for GEs. Table 3 reports treatment effects on key business and consumption outcomes measured at the six-year follow-up. Additional results can be found in Appendix Tables A1, A2, and A3. Columns 1 to 4 examine outcomes related to household businesses. We find significant average effects on businesses (Panel A), which persist even in the absence of ongoing access to microcredit: treatment households pay out Rs. 370 more in wages to employees each month, more than 100% of the control group mean (col 1); have over Rs. 1,500 more in business assets than households in the control group (column 2), and report 31 log points higher revenues from their businesses than the control group (column 3). Profits, reported in column 4 (in levels due to zero and negative values) are significantly higher on average, by just under 600 Rupees per month. Some of this increase in the business outcomes reported in columns 1 to 4 is driven by the extensive margin—treated households have more businesses. (See Appendix Table A1.)

¹⁵This is slightly larger than the 8.4pp treatment effect on having an MFI loan reported by Banerjee et al. (2015) because that number includes only loans outstanding at the two-year follow-up, while the 11pp value includes loans already fully repaid.

While positive effects on business outcomes are apparent for the treatment group as a whole, as Panel B shows, these results are driven largely by effects on gung-ho entrepreneurs alone. GEs' businesses are larger and more profitable, as well: GEs in treatment pay out almost Rs. 600 more in wages (column 1) and own over Rs. 3100 more in business assets (column 2) than GEs in control.¹⁶ GEs in treatment areas also have higher revenues by 72 log points (column 3), and earn almost 1300 Rupees, or 28%, higher profits (column 4). In contrast, although non-GEs pay out a statistically significant Rs. 275 more in wages in treatment than in control (column 1), there are no significant effects on assets, revenues or profits for the non-GE households (columns 2-4).

Impact on the distribution. Are the results for GEs driven by a few businesses experiencing extremely large effects on business scale, or are the effects distributed more broadly? Could the null effects for non-GEs be masking offsetting positive and negative effects at different points in the distribution? To shed light on this, Figure 5, Panel A plots the results of quantile regressions for business profits on treatment status for GEs. A large section of the distribution of households by business profits (from around the 75th to 95th percentiles) experienced significant positive treatment effects on their business profits. A similar pattern is seen for business assets (Figure 5 Panel C); the effects are even more broadly distributed: from roughly the 30th to the 90th percentile, GEs in treatment have significantly larger businesses than their counterparts in control areas.

In contrast, no portion of the distribution of either profit or business has shifted for non-GE households (panels B and D).

Persistent consumption effects of access to microcredit. Next, we examine effects on consumption. The effect of credit access on consumption will depend both on changes to household income (i.e., business profits), which matter via an income effect; and the return to reinvesting cash on hand in the business, which matters via a substitution effect. In our dynamic model, with non-convexities households initially may prefer to save than consume to reach the high-productivity portions of the production frontier. Prior work has found no detectable short-run impacts on household non-durable consumption, and Banerjee et al. (2015) finds no short-run impact on non-business, i.e. household, durables.

Table 3 columns 5 to 7 reports our 6-year results. While neither the pooled sample, nor GEs or non-GEs separately experience significant average increases in total consumption (column 5), durables spending increases on average (Panel A, column 6). Panel B shows that the entirety of this pooled increase in durable goods spending is driven by GEs, who spend over Rs. 1900 more on durable goods in treatment than in control (significant at the 1% level). When we look specifically at non-business durables in column 7, we see that the GEs are consuming more at the six-year follow up. Since household durable goods, which include both items like gold

¹⁶For these two outcomes, the differential effect for GEs is not significant, but in both cases the total treatment effect for GEs is significant—at the 10% level for wages and at the 5% level for assets.

and those like televisions, are a combination of savings and consumption, this suggests that the income gains experienced by the GEs are partly saved and partly consumed.

Turning to distributional effects, Figure 5, Panel E of Figure 5 shows that, among the GEs, more than half of the distribution (from around the 30th to the 85th percentile) experienced positive treatment effects on consumption. At the 75th percentile of the distribution, we find a gain of just under Rs. 350 in monthly household consumption per adult equivalent, an increase of 10.4% over the 75th percentile of consumption among GEs in the control group (Rs. 3325). However, at no point in the distribution of per capita consumption for non-GEs (Panel F) do we find any significant positive (or negative) treatment effects: the effect is a fairly precise zero throughout the distribution.

5.3. Testing Prediction 4: Crowd-in of non-Microfinance Loans for GEs. Table 4 investigates effects on non-microfinance borrowing both at the two-year (columns 1-2) and six-year (columns 3-4) horizons. We focus on total informal borrowing, as formal loans are rare in our setting. In columns 1 and 3, we consider borrowing for any purpose, and in columns 2 and 4 we restrict to loans that were used, at least in part, for any business-related purpose. We find that at both horizons, earlier access to microfinance crowds in other types of business borrowing on average.

Notably, these effects are driven almost completely by the GEs. There are no significant impacts for the REs; in fact several of the point estimates are negative in sign. Moreover, the differential treatment effects for GEs are large and positive. The short-run treatment effect on business borrowing for GEs is Rs. 9,269 larger than for REs, ($p = 0.038$) and the six-year effect for GEs, after the entry of microfinance into the control group and the total withdrawal of microfinance, is RS. 10,598 larger than for REs ($p = 0.015$).

These findings are consistent with prediction 4 and strongly imply the existence of production non-convexities. They also are inconsistent with Case II from Figure 2. As we discuss above, they are consistent with the existence of a poverty trap, but we cannot rule out all other models with the reduced form evidence alone (see Appendix Figure B1). We use the structural evidence below to provide further evidence of the relevance of poverty traps in our setting.

Is “GE” a proxy for baseline wealth? One concern with the differential treatment effects we document is that the difference between GEs and non-GEs could be proxying for baseline wealth. To address this, we leverage our matched Baseline-Endline sample. Figure B3 plots the kernel density of baseline wealth, separately for GEs and non-GEs. Reassuringly, there is significant overlap between the two distributions. However, the non-GE distribution does have more mass at low wealth levels. Therefore, to check whether this could be driving the difference in treatment effects, we re-estimate our heterogenous-effects specification (equation 5.2), weighting the non-GEs to have the same distribution of baseline wealth as the GEs, and vice versa. The outcome variable we use is total capital (the sum of business durable assets,

business working capital, and wage bill) measured at the six-year endline. The results appear in Table 5. The first column presents the unweighted results. As expected, we see a small and insignificant treatment effect for the REs. Despite the small sample, we see a large and significant differential effect for the GEs (for whom the total treatment effect is also significant.)

Column 2 weights REs to look like GEs in terms of baseline wealth; here, if REs were showing smaller treatment than GEs effects solely due to lower wealth, we would expect to see the main treatment effect increase; however we do not observe this. This suggests that RE are different, and not only because they are poorer.

Our model, however, predicts that not just type, but also initial wealth does matter, since there are possibility of poverty traps for GE. Column 3 weights GEs to have the same wealth distribution of GE. The differential effect of treatment on GE is now smaller, and no longer significant (although the total effect for the GEs does remain significant).

These two sets of reweighted results show that both type (GE vs non-GE) and the role of initial wealth are crucial to understand our results. Type matters in that even rich non-GEs cannot access the higher steady state, and baseline wealth matters insofar as the poorer GEs are too poor to escape the poverty trap even with access to credit.

5.4. Prediction 5: Negative selection of RE businesses. Another difference between GE and RE, other than baseline wealth, is their age. To test whether the longer tenure of GEs explains the results, we use the fact that Spandana did not enter all treated neighborhoods at the same time: branches opened in treatment areas between April 2006 and April 2007. As a result, we observe treatment-area businesses that opened at the exact same point in time, but some opened before Spandana’s entry to the area (because Spandana’s branch in that area opened relatively late), while others opened after Spandana’s entry (because Spandana’s branch in that area opened relatively early). Of course, Spandana’s decision of where to open early vs late was not random. However, because randomization was done at the matched pair level, for each treated area, the control area in the same matched pair serves as a counterfactual, under the assumption that Spandana would have opened a branch there at the same time. We refer to the sample of businesses that opened during Spandana’s roll-out as the “overlapping sample.”

Figure 6 shows a schematic illustrating the idea behind this overlapping sample. In Matched Pair A, Spandana entered the treated area A^T , at t_1 ; in Matched Pair B, Spandana did not enter the treated area B^T , until t_3 . In both pairs, Spandana did not enter the control areas, A^C and B^C , until after the two-year follow-up survey. In each of the 4 areas, there is a set of businesses that opened at time t_2 , after Spandana entered A^T but before it entered B^T . Finally, at t_4 , outcomes, y , are measured. The comparison $\bar{y}_{A^T} - \bar{y}_{A^C}$ identifies the treatment effect on businesses opened *after* Spandana’s entry, while the comparison $\bar{y}_{B^T} - \bar{y}_{B^C}$ identifies the treatment effect on businesses of the same age, but opened *before* Spandana’s entry.

If the differential treatment effects found for GEs are simply due to the fact that GEs are older or more experienced, then among this overlapping sample, those that opened pre-Spandana

should have indistinguishable treatment effects from those that opened later. If, on the other hand, microfinance induces businesses to enter that have lower returns (negative selection), then the firms that opened pre-Spandana should have larger treatment effects than those that opened post-Spandana but at the same point in calendar time.

Table 6 shows the results. Panel A shows the two-year treatment effects for businesses opened in the overlapping sample window, *before* Spandana had opened in their area (equivalent to B^T in Figure 6). We find that the effect on the index of business incomes is 0.150 standard deviations, significant at 5%. We show similar, but slightly smaller results in Appendix Table A5 for the set of older businesses that opened pre-2006, and Appendix Figure B4 further decomposes the results by business age.

Panel B shows the two-year treatment-control difference for businesses opened in the overlapping sample window, *after* Spandana had opened in the treatment area of the matched pair (equivalent to A^T in Figure 6). Those differences, while imprecisely estimated, are uniformly negative. The effect on the index of business outcomes is -0.183; while this is not significantly different from zero at conventional levels ($p = 0.105$), it is significantly different from the effect of plus 0.150 seen for the pre-Spandana (but same-aged) businesses. This comparison combines two separate effects – the treatment effect for the post-Spandana businesses and any selection effect. Column 7 of Table 6 shows consistent results for non-microfinance business borrowing. Again, we find large crowd-in of non-microfinance business borrowing, but only for the GEs.

We next conduct a back-of-the-envelope exercise with the results in Table 6 to ask how negatively selected is the average RE that enters because of microfinance. To do this we need to isolate the new businesses that started because of microfinance from those that would have opened in any case in the post-Spandana sample. By comparing the number of new businesses in treatment vs. control, we estimate that microfinance induced 18% more businesses to enter. Next we make three simplifying assumptions. First, we assume that the treatment effect for the 82% of businesses that would have opened in absence of the intervention in the post-Spandana sample have the same returns to credit as the pre-Spandana sample. Second, we assume that the true, causal treatment effect of microfinance for REs is zero, and third, that the treatment vs. control comparison for the post-Spandana sample is also zero. We then solve for the selection effect x on the index of business outcomes: $0 = .82 \times (.150) + .18x$. Thus, $x = -.68$, suggesting that those new businesses must be 0.68 standard deviations worse than the pre-Spandana businesses.

These results show that, consistent with our predictions, the behavior of GEs and REs of the same age and experience level diverge substantially. These results also suggest that entrepreneurs are aware of their type: the GEs appear to know that, once they have access to microcredit, they can productively invest even more capital than they are able to access from microlenders directly. This is consistent with the evidence from other settings that households are aware of their potential returns to capital (Beaman et al., 2023; Hussam et al., 2022).

5.5. Threats to Validity: Windfall Effects of the AP Ordinance. The 2010 AP Ordinance caused microlenders to cease operations and withdraw from the study neighborhoods. While the ordinance reduced future access to microcredit, it also led to the implicit deferral or forgiveness of the outstanding loan principal and interest. In Appendix C, we explore whether this windfall can potentially explain some of the long-run treatment impacts found in Table 3.

While the treatment group did have more credit outstanding than the control group in 2010, the year of the crisis, we calculate that the implied difference in windfall size between treatment and control was only approximately Rs. 470. This is a small amount relative to the differences in business outcomes for the GEs in 2012. Moreover, we can also estimate the impacts of windfall size on business outcomes directly. To do this we can compare borrowers who had received a new loan just before the ordinance and who received a large windfall with those who were close to fully repaying their prior loan and who received a small windfall. While the differences between these groups is large – roughly Rs. 8,000 – We find no evidence that receiving a larger windfall in 2010 is associated with better business outcomes in 2012, either on average or separately for GE or non-GE households.

6. MODEL ESTIMATION

We now estimate a version of the framework presented in Section 2. Our goal is to generate estimated analogs of Figure 2 using the simplest possible structure to capture the possibility of a poverty trap. We ask whether at the estimated parameters, the transition dynamics match the patterns in Panel C. We first lay out the key equations and constraints needed for estimation. Our empirical treatment of the model then proceeds in four steps: A) we use the two-year follow-up data from the overlapping sample to estimate the production frontier; B) we solve the household’s dynamic program and obtain the implied policy functions; C) we ask whether the resulting policy functions give rise to a wealth-based poverty trap; D) we simulate the model forward for the GEs in our data in both treatment and control areas and compare the simulated long run treatment effects to the empirical findings presented above.

In all that follows, we only consider the GE households in our analysis.

6.1. Dynamic Household Optimization Problem. We build on the model outlined in Section 2, specifically equations 2.1, 2.2, 2.3, 2.4, and 2.5 and estimate it with the experimental variation arising from treatment assignment.

Production Function. It is well understood that the shape of the production function directly affects the evolution of wealth and the emergence of poverty trap dynamics. There are also many potential ways to generate such non-convexities. We model this by allowing GEs to choose between two distinct production technologies to convert inputs K_t into revenues Y_t : a *Low* technology with decreasing returns and a *High* technology requiring a fixed cost, but with

higher marginal returns above that minimal scale.¹⁷ While we think this type of technology switching is reasonable given the setting, we cannot rule out other microfoundations for non-convexities in the production frontier.

We define K_t to include the rupee value of total inputs used in production. The *Low* technology has revenues equal to:

$$(6.1) \quad Y_L(K_t) = A_L K_t^\alpha$$

We posit that this technology exhibits decreasing returns to scale, i.e. $\alpha < 1$.¹⁸ Denote the optimal scale of this technology as K_L^* . In contrast, we assume that the *High* technology has constant returns to scale, with revenues given by:

$$(6.2) \quad Y_H(K_t) = A_H(K_t - \underline{K}).$$

While the goal of this modeling exercise is to capture the dynamics for the *GE* entrepreneurs, implicit here is the assumption that the REs only have access to the *Low* technology, consistent with the finding of no short- or long-term effects on business capital or business outcomes for the REs.

Borrowing Constraint. We consider two regimes, $\tau \in \{1, 2\}$. For simplicity, in regime $\tau = 1$, households do not have access to credit and must finance business investment from their wealth. This can be thought of as a borrowing limit of zero: $\bar{b}_1 = 0$.

In the second regime, households have the ability to borrow. Motivated by the reduced form finding of crowd-in for the GEs, in this regime they can borrow from microfinance and from informal lenders. This informal borrowing comes from input suppliers, shop keepers, moneylenders, friends, and relatives. We assume that all project returns are deterministic and that lenders have a claim to the entrepreneur's resources, so we abstract away from default. This credit line, however, is not infinite, and all households are restricted to choose $0 \leq B_t \leq \bar{b}_2$, where \bar{b}_2 is the total amount that can be borrowed in regime 2. All loans must be (and are) repaid at the end of each period.

Within-Period Profits and Cash on Hand. Implicit in Equation 2.4 is the simplifying assumption that all productive decisions are separable across time. Individuals enter each period t with wealth W_{t-1} and decide how much to invest in the business (K_t) and in the outside option technology (S_t). Within the business, the entrepreneur chooses whether to select the *High* technology, denoted by the indicator D_t^H .

Given the technology choice, the profit function now becomes

¹⁷The *High* technology may have a greater span of control, allowing for the use of hired labor and avoiding decreasing returns due to fixed household labor, may correspond to a different production technology (e.g., mechanical sewing machine vs. sewing by hand), or may correspond to a business facing a less localized market.

¹⁸In the estimation, we do not impose this restriction; we allow for any positive value of α .

$$(6.3) \quad \pi(K_t) = D_t^H Y_H(K_t) + (1 - D_t^H) Y_L(K_t) - RB_t$$

Full Utility Maximization Problem. Given the above elements, the utility maximization problem in its recursive form is:

$$\begin{aligned}
 V(W|\tau) &= \max_{c, W', K, B, D^H} u(c) + \beta E(V(W')|\tau) \\
 W' + c &= D^H Y_H(K) + (1 - D^H) Y_L(K) + (1 - \delta)K + \rho S - RB \\
 0 &\leq B \leq \bar{b}_\tau \\
 W' &\geq 0 \\
 K &\leq W + B \\
 S &= W + B - K
 \end{aligned}$$

6.2. Production Function Estimation. Given the time separable nature of the production decision, investment choices are the solution to a static profit maximization problem conditional on starting wealth. In our first estimation step, we estimate the four production function parameters of the model $(A_L, \alpha, A_H, \underline{K})$.¹⁹

For this exercise we use data only on the overlapping sample households who opened businesses prior to Spandana’s entry (the GEs) from only the two-year wave of survey data. While this sample is small, it gives us a relatively homogeneous group of GEs for which to estimate the production function. Moreover, these businesses are young; this is useful because early on the business is likely to not yet be in a long-run steady state, allowing for the estimation of a larger support of the production function.

Estimating the production function requires heterogeneity in baseline wealth. Ideally, we would have an empirical estimate of W_0 for each household and feed that directly into the model. However, we do not have such a measure, both because baseline surveys were not collected for the majority of households in the original Banerjee et al. (2015) study and because it is challenging to elicit the value of total wealth, particularly non-financial wealth such as land. To make progress, we make the assumption that the treatment effects in capital satisfy a monotonicity assumption.

Figure 7 illustrates our identification strategy from by plotting the CDFs for the treatment and control groups separately for total capital (Panel A) and revenues (Panel B). Panel A captures the treatment-versus-control investment comparison by quantile – how total capital changes across the distribution with access to microcredit. Using total capital as the “first

¹⁹We also estimated a version of the model allowing for diminishing returns in the high technology, but we find that we cannot reject linearity.

stage” impact of credit allows for arbitrary crowd-in or crowd-out of outside credit. Panel B shows the reduced form – how business revenues change across the distribution with access to microcredit.

The monotonicity assumption requires that the household’s rank in the capital distribution within the treatment group is the same as the rank the household would have attained within the control group. Note that we only assume monotonicity in capital, not in revenues. Under this assumption, comparing capital levels in treatment versus control at any quantile gives us the treatment effect on capital *at that quantile*.

Given the exogenous shock in capital, for any candidate $(A_L, \alpha, A_H, \underline{K})$, at each quantile we can then construct predicted revenues for each individual in the sample in the treatment and the control group according to

$$Y_t(K_t) = D_t^H Y_H(K_t; A_H, \underline{K}) + (1 - D_t^H) Y_L(K_t; A_L, \alpha).$$

The predicted difference in revenues between treatment and control group for each quantile, and at each value of $\{A_L, \alpha, A_H, \underline{K}\}$ can then be compared to the true difference in revenues at the corresponding quantile in the data. This clearly also requires an exclusion restriction-like assumption – that microfinance only impacts revenues through changes in total capital and through the parametric form we specify above.

Given the relatively small sample, we bin the data into 15 groups by quantile of capital within treatment and control groups, and compute predicted revenues for each group in the 15 bins for a grid of discretized value of the parameters of interest. Given that some firms produce 0 revenues, we allow stochastic business closure within period to match the empirical rates in the data.

We select the parameters $\{A_L, \alpha, A_H, \underline{K}\}$ that minimize the GMM objective function, using revenues as our moment condition. For details, see Appendix C.

The estimated parameters of the production technologies are the following with bootstrapped 95% confidence intervals in brackets.

$$A_L = 45 [15, 90], \alpha = 0.4 [0.2, 0.6], A_H = 1 [0.06, 2], \underline{K} = 7900 [100, 14400]$$

The estimated fixed cost, Rs. 7,900 corresponds to roughly the median of the estimated baseline wealth distribution. In Figure 8, we plot the log gross revenues of the technologies as function of log capital. Note that while the revenue functions cross at $K = 9414$ (95% confidence interval [403,14765]), this is not the point at which a household would optimally switch into the high technology, because operating the low technology at this scale is dominated by the option of operating it at its optimal scale and investing any remaining wealth into the savings technology. Figure 9 shows the “gross profits” for the three possible technologies available to a household: the low technology, the high technology, and the savings technology, where we

set $\rho = 1$. For the low and high technologies, the gross profits are equal to revenues plus proceeds of selling back undepreciated capital. The point at which the gross profits of the low technology intersect those of the savings technology represents the optimal scale of the low technology (approximately Rs. 650); above this point it is optimal to save additional wealth until the point at which the savings technology function is intersected from below by that of the high technology. This intersection, which occurs at approximately Rs. 13,500, represents the minimum scale at which it is efficient to operate the high technology.²⁰

Note that the production frontier we estimate here is consistent with papers such as McKenzie and Woodruff (2006), which find no evidence of fixed costs of *entry* into small-scale entrepreneurship. In our context, it is very easy to get started selling vegetables or prepared foods from outside one's home. However, getting to a more substantial scale with higher marginal returns involves making lumpy investments (e.g., renting a formal storefront with a minimum scale, acquiring an asset, hiring an employee, etc.).²¹

6.3. Calibrating and Solving for the Policy Functions. Next, we solve the dynamic program given the estimated production parameters. We assume that the instantaneous utility function takes the standard CRRA form $u(c) = \frac{1}{1-\sigma}c^{1-\sigma}$. The other parameters from the model (return on outside option/depreciation/borrowing parameters $\rho, R, \bar{b}_1, \bar{b}_2, \delta$; and preference parameters σ, β) are either calibrated or estimated separately from our survey data.

As noted above, the gross return on savings (the “under-the-mattress” technology), ρ , is set to one. The borrowing rate R is set equal to the microfinance interest rate ($R = 1.25$) and the borrowing cap \bar{b}_2 is set to Rs. 12,000.²² This reflects the first stage on total credit (Table 2, column 4). The depreciation rate is calibrated to $\delta = 0.4$. This is based on the average split between working and fixed capital in the data, and the assumption that labor and variable working capital inputs depreciate fully and fixed capital (assets) does not depreciate at all. We calibrate the preference parameters as follows: $\beta = 0.85, \sigma = 0.8557$. We obtain the estimate of σ using the method proposed by Andreoni and Sprenger (2012).²³ Below, we provide a sensitivity analysis to these preference parameter choices.

6.4. Forward Simulation and Simulated Treatment Effects. Given the policy functions, we can simulate the model forward for all GE households in the data. The time frame of the

²⁰These gross profit functions do not account for borrowing costs; that is, they represent the decision of a household using only their own wealth. For a household who needs to borrow at the gross interest rate of 1.25, the minimum efficient scale of the high technology is Rs. 18,500.

²¹The nonconvexity could also arise from other factors that are not strictly fixed costs, such as adopting new management practices that require learning-by-doing before they yield high returns.

²²Recall that $\bar{b}_1 = 0$.

²³Our implementation of the Andreoni and Sprenger (2012) method involved having participants make decisions between payouts at several different points in time: 2 days, 32 days and 62 days. (There was no 0 day wait to avoid confounding trust with patience.) As is common in these types of exercises (Frederick et al., 2002), the implied annual discount rates from annualizing elicited monthly discount rates are implausibly high. Because of this issue, we use $\beta = 0.85$. Choosing a relatively low discount factor (ie, high discount rate) is consistent with the fact that some (mostly non-GE) individuals take loans at a 25% interest rate to fund consumption.

intervention spans 2006-2012. Given a baseline year of 2006, we simulate the model forward for 6 years and compare the implied average treatment effects at the end of the final year.

State Variable. The forward simulation requires initial values of the state variable – wealth (W_0) – for each household, an object that we do not observe in the data. However, given that we have the policy function mapping wealth to investment, and we observe investment in the data, we can infer the former from the latter by inverting the policy function to recover baseline wealth. If investment decisions for the GEs were one-to-one in wealth, this would be straightforward. However, our estimated production function (which is the upper envelope of the Low and High technologies) exhibits a region where increasing the scale of the business is dominated by savings (e.g., but selecting the optimal scale of the low technology and saving any remaining wealth). This means that individuals investing at the optimal scale of the low technology will represent a range of wealth levels. Moreover, in the data, we do not observe stark clustering at the optimal low technology. This is to be expected – in reality, individuals may experience idiosyncratic shocks to productivity, to borrowing costs, or to their outside value of capital.

To address these two issues, we do the following. We first start with a discretized grid of possible values of baseline wealth.²⁴ We next calculate the capital choice under the model for each potential level of wealth. We then perturb this value with a mean zero, iid noise shock with standard deviation ν .²⁵ This perturbation captures the fact that the empirical capital choice might not exactly match the predicted value due to elements of noise that we do not model directly (e.g., measurement error, exogenous productivity shocks, and optimization frictions). This will also generate dispersion around the optimal scale of the low technology and eliminate bunching. We draw 100 times from the noise distribution for each value of wealth in the grid. In order to assign a level of baseline wealth for each household we observe in the data, we match the empirical capital choice back to the closest value of capital predicted under the model and perturbed by the iid shock. This means that households in the data are potentially matched to different starting wealth values across the 100 draws from the shock distribution. We simulate the model forward for each of these 100 noise draws and average across draws. When simulating the model forward for treatment and control neighborhoods, we use counterfactuals based on the baseline wealth distribution in the control group alone.

Borrowing Regimes Across Time. As mentioned above, we consider two different borrowing regimes τ . It is important to note that, consistent with our empirical setting, each regime change is a surprise to all households. In years 1 and 2, the control areas are in the no borrowing regime ($\tau = 1$), while the treatment areas are in the borrowing regime ($\tau = 2$). In years 3 and 4, the control areas also have access to borrowing, so all household samples are in regime $\tau = 2$.

²⁴Specifically, we discretize wealth in bins of Rs. 300.

²⁵In our core simulate-forward exercise we assume that $\nu=500$. We explore the sensitivity of the results to changes in ν in Appendix table A10.

Finally, in years 5 and 6, microfinance is no longer available anywhere, so $\tau = 1$ for the full sample. During any regime τ , households believe that this regime will continue to be the status quo forever, which is reasonable in this context given the rapid entry of microfinance and the unanticipated nature of the AP Crisis.

Simulated Treatment Effects. We next compare the treatment effects from the simulations with those from that actual six-year endline data. The model replicates the qualitative patterns observed in the data quite well. The difference between the mean level of capital in the treatment group versus that in the control group, in the simulated period 6 data, gives us the simulated year 6 treatment effect on capital. In levels, the implied treatment effect is Rs. 8,053 at the estimated and calibrated parameters described above; in the data it is Rs. 7,342. In Panel A of Appendix table A10, we show sensitivity of this estimated treatment effect to different assumptions about the preference parameters (β, σ) . Specifically, we consider $\beta \in \{0.85, 0.90, 0.95\}$ and $\sigma \in \{0.86, 1.00, 1.25, 1.50\}$, to cover the range of values typically used in the macro literature. We find that for all of the parameter combinations, the six-year differences between treatment and control are large, ranging from 4,008 to 8,053. This demonstrates that the model can predict capital differences between treatment and control six years post-intervention of the same order of magnitude as we find in the data—quantities not targeted or used in the estimation. We also show that our estimated treatment effects are also robust to the assumptions used to recover the baseline wealth distribution. Specifically, when we change the variance of the noise distribution, the treatment effects remain large (see Panel B, Appendix table A10).

In keeping with our examination of effects throughout the distribution, we next compute the implied distribution of quantile treatment effects from the simulated data. Figure 10 shows the implied treatment effects on expected year 6 capital across the distribution, ordered by the level of expected year 6 capital in control. The empirical counterpart to this figure is Figure 5, Panel C. A comparison of the two reveals striking similarities: both show the treatment effects concentrated in roughly the upper third of the distribution, with modestly positive effects lower in the distribution; the effects at the very top are also modest. Note that, at the lower quantiles, we observe a positive predicted treatment effect because in some draws of the capital shock, even a low wealth individual could be pushed out of the poverty trap region by the extra capital access resulting from treatment.

To understand through the lens of the model why the largest effects are seen at the top (but not very top) of the distribution, Figure 11 plots the model-implied wealth transition diagrams for the credit and no-credit regimes. The upward-sloping lines show these transitions: taking a given level of starting wealth and reflecting the optimal technology, borrowing and consumption/savings decisions, what will next period's wealth be? In the no-credit regime, households with wealth roughly below the long-short dashed line will be unable to raise enough cash on hand (wealth plus borrowing) to invest in the high-returns technology. As a result,

facing low marginal returns and impatience ($\beta = 0.85$), these households never accumulate enough wealth to access the high technology and instead converge to the optimal scale of the Low technology. Only households whose wealth is above this value are able to operate the High technology.

In the credit regime, in contrast, all households whose wealth is roughly above the long-long dashed line will be able to access the high technology. The reason that intermediate-wealth households can now do so is, of course, that they can borrow. Once households are able to access the high technology, its constant marginal returns implies that their wealth will continue to grow over time. Finally, the richest households were not in the poverty trap zone, even in the absence of credit; they borrow to expand since the constant marginal return is above the interest rate.

The values of the long-short and long-long dashed lines correspond to the 75th and 50th percentiles of the estimated baseline wealth distribution, respectively. Households starting out in the third quartile of wealth are those who benefit most from access to microcredit.

Poverty trap vs. scale up? Our simple model emphasizes the fact that in the presence of fixed costs, talented but low-wealth households may be caught in a poverty trap, and that microcredit can allow (some of) them to escape this trap. However, both in the model and in practice, increased access to credit can have another effect as well: allowing households who were already out of the poverty trap zone to scale up their high-return businesses. With constant marginal returns, households operating the high technology will benefit from the extra liquidity and hence larger business scale. This raises the question of whether the bulk of the effects we observe are coming from intermediate-wealth households escaping the poverty trap and moving to the high technology, or from higher-wealth households scaling up their already high-return activities.

To shed light on this, we split households into four groups. First, households whose wealth is so low that, without borrowing, they cannot operate the low technology at its efficient scale (group 1). These households, even by borrowing the full Rs. 12,000 that is possible when microcredit is available, cannot reach the minimum efficient scale of the high technology under borrowing (Rs. 18,500). They will borrow only to reach the optimal scale of the low technology, which is approximately Rs. 650. Since the efficient scale of the low technology is so small, these households will have very small treatment effects. Second, households who can reach the optimal scale of the low technology without borrowing, but even by borrowing the full Rs. 12,000 cannot reach the minimum efficient scale of the high technology under borrowing (group 2). Group 2 households will not borrow at all (and so will have a zero treatment effect). Next are households who, without borrowing, could not reach Rs. 13,500 of capital but, with borrowing, can reach Rs. 18,500 (group 3). Group 3 comprises the households who are pushed out of the poverty trap: they borrow the full Rs. 12,000 and move into the high technology. Finally, there are households who were already able to reach at least Rs. 13,500 of capital.

They will optimally borrow the full Rs. 12,000 and scale up their business accordingly, taking advantage of the constant marginal returns (group 4). These households are not pushed out of the poverty trap by credit; they were already out of the poverty trap zone. Group 4 households simply benefit from running their business at a larger scale.

To decompose what share of the total treatment effects are coming from each of these groups, we sort households on the basis of their initial wealth into the four groups using the mapping between wealth and capital in the no-credit regime. We then calculate group-specific treatment effects by taking the difference in period 6 capital between treatment (who were exposed to microcredit for 4 periods before it was withdrawn) and control (who were exposed to microcredit for just 2 periods before it was withdrawn) and scaling these effects by the share of our sample who fall into each group. Recall that Group 1 has very small treatment effects and group 2 has a zero treatment effect, so the relevant question is the share of the overall effect coming from group 3 (who escape the poverty trap) vs. group 4 (who scale up their businesses).

We find that the households in our sample are roughly evenly distributed between groups 2, 3 and 4: 33% are in group 2, 37% in group 3 and 30% are in group 4. (Group 1 is a negligible share.) Of the effect on total capital, 73% is driven by escape from the poverty trap (group 3) and 27% is driven by accelerated growth in group 4. For revenues, 68% is driven by the poverty trap effect and 32% by the scale-up effect. (The increase in revenues is a bit more skewed to the already productive businesses because, unlike those escaping the poverty trap, they have already paid the fixed cost.) The effect of microcredit working through allowing households to escape from a poverty trap is a quantitatively significant phenomenon, accounting for roughly two thirds of the increase in revenues, and three quarters of the increase in business investment.

7. CONCLUSION AND DISCUSSION

We use the long-run, *persistent* effects of a randomized microfinance evaluation to shed light on the old, but elusive, question of whether fixed costs can give rise to poverty traps. The universal withdrawal of microfinance from the entire study area in 2010 (two years before we surveyed respondents for the longer-run follow-up), overlaid on a setting with randomized variation in microcredit access, create a setting uniquely well-suited to addressing this question. We show that the effects of access to formal credit through microfinance are highly heterogeneous. Essentially all of the benefits of credit access accrue by increasing entrepreneurship on the *intensive* margin: for those individuals with an existing business before the entry of microfinance (who we call gung-ho entrepreneurs or GEs), we find economically meaningful, positive effects on household businesses and consumption. Within this group, the bulk of effects come from households who escape from the fixed-cost-driven poverty trap and move into a more productive technology (with the remainder coming from already-productive businesses scaling up to exploit constant returns). The rest of the sample—those who start new businesses, or who never start a business at all—exhibit essentially zero impact of credit access. Notably, for this

group the effect is a fairly precise zero throughout the distribution: while these reluctant entrepreneurs and consumption borrowers do not experience detectable benefits from microcredit access, neither do they appear to experience harm.

Using the randomized variation to estimate production parameters for the gung-ho entrepreneurs, we find that the data are consistent with the co-existence of multiple production technologies, one of which only becomes optimal at sufficiently high levels of capital. We model this as a fixed cost: individuals with sufficient capital can pay the fixed cost to operate a technology with higher marginal returns; while those with less capital must make due with a technology that does not have an upfront cost, but has lower TFP and decreasing returns. Embedding these production parameters into a dynamic model of consumption and investment reveals that the estimated parameters give rise to a poverty trap: low-wealth GE households remain poor because saving up for the better technology is not possible using only the returns from the decreasing return technology. Moreover, the size of the credit infusion that we observe for the treatment group in our data (which comprises both microcredit and the crowd-in of informal credit) is sufficient to allow a significant share of intermediate-wealth GEs to escape this poverty trap before the unexpected withdrawal of microcredit due to the AP ordinance.

One important ingredient in understanding the disparate impacts of microfinance is the informal credit market. We are among the first to use experimental variation to study this interaction between access to formal and informal finance (Karlan and Zinman (2018) find evidence of crowd-in across different (formal) microlenders in Mexico), and we document important heterogeneity. We find evidence for the “crowdout” hypothesis among the reluctant entrepreneurs: access to formal finance reduces these households’ take-up of informal credit. This crowdout effect may explain why impacts of microfinance are minimal among certain sub-populations: while microfinance may reduce borrowing costs, overall demand for credit may change very little for some groups. In contrast, microfinance *crowds in* other sources of borrowing among the GEs and is key to replicating the magnitudes of the long run impacts in the model simulations. It is therefore essential for policymakers to understand these interactions when designing financial inclusion policies and when targeting financial products to specific groups.

Of course, some firms will be able to escape the poverty trap without the intervention, due to wealth, luck, talent or a combination of all three. However, our results show that short-term access to credit allows a significantly higher proportion of talented entrepreneurs to scale up their businesses. In sum, it appears that there are indeed sizable benefits from microfinance for some people, but it takes time for these benefits to accumulate. And it is important to look for the impacts in the right place. The disappointing effect of microfinance overall may be related to lenders’ lack or willingness (or ability) to identify the right beneficiaries.

A number of implications for credit market policy emerge from our results. First, microcredit organizations often emphasize the non-selective nature of their lending as an advantage. But

if most of the business growth comes from a minority of firms—who in turn provide employment opportunities which may help others avoid “reluctant” entrepreneurship (Adhvaryu et al., 2019; Karaivanov and Yindok, 2022; Breza et al., 2021)—then a more selective approach may be better. While we have no reason to question the fact that even the REs and consumption borrowers benefit from the loan (see the discussion of distributional effects here, and in Angelucci et al. (2015) who carefully explore the possibility that some groups end up doing worse from microcredit), there may be a case for focusing more energy on identifying the GEs and helping them grow. A small literature explores how to achieve this type of positive selection on productive investments. Beaman et al. (2023) show evidence for positive *self*-selection in the context of agricultural microloans in Mali. Hussam et al. (2022) provide evidence that microentrepreneurs in India are aware of the marginal returns to additional capital of their own businesses, and those of their peers. Moreover, our results on the crowd-in of other credit sources for GEs—and the opposite for REs—demonstrate that households are aware of their ability (or lack thereof) to scale up their business.

Second, our findings raise the issue of whether, from the point of view of growth, much bigger (and more selective) loans are desirable.²⁶ Recent work by Bari et al. (2024) and Bryan et al. (2023) show that when well-targeted, substantially larger loans than the microfinance status quo lead to sizable increases in profits. Finally, the fundamental heterogeneity in borrowing needs and productive capabilities among potential borrowers has important implications in terms of designing a menu of financial contracts that can induce GEs, REs and non-entrepreneurs to select different options.²⁷

²⁶La Porta and Shleifer (2008) make the case that most firms in the informal economy are marginal to the main story of growth. Related, Diao et al. (2018) show that, in Tanzania, a small subset of firms experience growth in employment and labor productivity, suggestive of positive returns to capital, while the remainder do not.

²⁷Related, Maitra et al. (2017) show that incentivized agents can identify productive and lower-risk borrowers in West Bengal.

REFERENCES

- ADHVARYU, A., N. KALA, AND A. NYSHADHAM (2019): “Booms, Busts, and Household Enterprise: Evidence from Coffee Farmers in Tanzania,” *The World Bank Economic Review*.
- AHGION, P. AND P. BOLTON (1997): “A Theory of Trickle-Down Growth and Development,” *Review of Economic Studies*, 64, 151–172.
- ANDREONI, J. AND C. SPRENGER (2012): “Estimating time preferences from convex budgets,” *American Economic Review*, 102, 3333–56.
- ANGELUCCI, M., D. KARLAN, AND J. ZINMAN (2015): “Microcredit Impacts: Evidence from a Randomized Microcredit Program Placement Experiment by Compartamos Banco,” *American Economic Journal: Applied Economics*, 7, 151–82.
- ATTANASIO, O., B. AUGSBURG, R. DE HAAS, E. FITZSIMONS, AND H. HARMGART (2015): “The Impacts of Microfinance: Evidence from Joint-Liability Lending in Mongolia,” *American Economic Journal: Applied Economics*, 7, 90–122.
- AUGSBURG, B., R. DE HAAS, H. HARMGART, AND C. MEGHIR (2015): “The Impacts of Microcredit: Evidence from Bosnia and Herzegovina,” *American Economic Journal: Applied Economics*, 7, 183–203.
- BALBONI, C., O. BANDIERA, R. BURGESS, M. GHATAK, AND A. HEIL (2022): “Why Do People Stay Poor?” *Quarterly Journal of Economics*, 137.
- BANERJEE, A., E. BREZA, A. G. CHANDRASEKHAR, E. DUFLO, M. O. JACKSON, AND C. KINNAN (2024): “Changes in social network structure in response to exposure to formal credit markets,” *Review of Economic Studies*, 91, 1331–1372.
- BANERJEE, A., E. DUFLO, R. GLENNERSTER, AND C. KINNAN (2015): “The Miracle of Microfinance? Evidence from a Randomized Evaluation,” *American Economic Journal: Applied Economics*, 7, 22–53.
- BANERJEE, A. V. AND E. DUFLO (2005): “Growth theory through the lens of development economics,” *Handbook of economic growth*, 1, 473–552.
- BANERJEE, A. V. AND A. F. NEWMAN (1993): “Occupational Choice and the Process of Development,” *Journal of Political Economy*, 101, 274–298.
- BARI, F., K. MALIK, M. MEKI, AND S. QUINN (2024): “Asset-based microfinance for microenterprises: Evidence from Pakistan,” *American Economic Review*, 114, 534–574.
- BEAMAN, L., D. KARLAN, B. THUYSBAERT, AND C. UDRY (2023): “Selection into credit markets: Evidence from agriculture in Mali,” *Econometrica*, 91, 1595–1627.
- BREZA, E., S. KAUR, AND Y. SHAMDASANI (2021): “Labor rationing,” *American Economic Review*, 111, 3184–3224.
- BREZA, E. AND C. KINNAN (2021): “Measuring the equilibrium impacts of credit: Evidence from the Indian microfinance crisis,” *The Quarterly Journal of Economics*, 136, 1447–1497.
- BRYAN, G., A. OSMAN, AND D. KARLAN (2023): “Big loans to small businesses: predicting winners and losers in an entrepreneurial lending experiment,” *American Economic Review*.

- BUERA, F. J., J. P. KABOSKI, AND Y. SHIN (2011): "Finance and development: A tale of two sectors," *American economic review*, 101, 1964–2002.
- (2021): "The macroeconomics of microfinance," *The Review of Economic Studies*, 88, 126–161.
- CHERNOZHUKOV, V., M. DEMIRER, E. DUFLO, AND I. FERNANDEZ-VAL (2018): "Generic machine learning inference on heterogenous treatment effects in randomized experiments," Tech. rep., National Bureau of Economic Research.
- CRÉPON, B., F. DEVOTO, E. DUFLO, AND W. PARIENTÉ (2015): "Estimating the Impact of Microcredit on Those Who Take It Up: Evidence from a Randomized Experiment in Morocco," *American Economic Journal: Applied Economics*, 7, 123–50.
- DASGUPTA, P. AND D. RAY (1986): "Inequality as determinant of malnutrition and unemployment: theory," *The Economic Journal*, 96, 1011–1034.
- DE MEL, S., D. MCKENZIE, AND C. WOODRUFF (2010): "Who are the microenterprise owners? Evidence from Sri Lanka on Tokman versus De Soto," in *International differences in entrepreneurship*, University of Chicago Press, 63–87.
- DIAO, X., J. KWEKA, AND M. MCMILLAN (2018): "Small firms, structural change and labor productivity growth in Africa: Evidence from Tanzania," *World Development*, 105, 400–415.
- FREDERICK, S., G. LOEWENSTEIN, AND T. O'DONOGHUE (2002): "Time discounting and time preference: A critical review," *Journal of economic literature*, 40, 351–401.
- HUSSAM, R., N. RIGOL, AND B. N. ROTH (2022): "Targeting high ability entrepreneurs using community information: Mechanism design in the field," *American Economic Review*, 112, 861–898.
- KABOSKI, J. AND R. TOWNSEND (2011): "A structural evaluation of a large-scale quasi-experimental microfinance initiative," *Econometrica*, 79, 1357–1406.
- KABOSKI, J. P., M. LIPSCOMB, V. MIDRIGAN, AND C. PELNIK (2022): "How Important are Investment Indivisibilities for Development? Experimental Evidence from Uganda," Tech. rep., National Bureau of Economic Research.
- KARAIVANOV, A. AND T. YINDOK (2022): "Involuntary entrepreneurship—Evidence from Thai urban data," *World Development*, 149, 105706.
- KARLAN, D. AND J. ZINMAN (2009): "Expanding credit access: Using randomized supply decisions to estimate the impacts," *The Review of Financial Studies*, 23, 433–464.
- (2018): "Long-run price elasticities of demand for credit: evidence from a countrywide field experiment in Mexico," *The Review of Economic Studies*, 86, 1704–1746.
- KERR, W. AND R. NANDA (2009): "Democratizing entry: Banking deregulations, financing constraints, and entrepreneurship," *Journal of Financial Economics*, 99(1), 124–149.
- LA PORTA, R. AND A. SHLEIFER (2008): "The Unofficial Economy and Economic Development," *Brookings Papers on Economic Activity*.

- LEE, D. S. (2009): "Training, wages, and sample selection: Estimating sharp bounds on treatment effects," *The Review of Economic Studies*, 76, 1071–1102.
- LEWIS, B., N. AGRAWAL, C. GADI, D. GOYAL, J. KULKARNI, A. TAWAKLEY, S. VISWANATHAN, A. WADHWANI, A. AUGEREAU, V. BHALLA, ET AL. (2001): "India: The growth imperative," *The McKinsey Global Institute*, 3.
- LLOYD-ELLIS, H. AND D. BERNHARDT (2000): "Enterprise, inequality and economic development," *The Review of Economic Studies*, 67, 147–168.
- MAITRA, P., S. MITRA, D. MOOKHERJEE, A. MOTTA, AND S. VISARIA (2017): "Financing smallholder agriculture: An experiment with agent-intermediated microloans in India," *Journal of Development Economics*, 127, 306–337.
- MCKENZIE, D. AND C. WOODRUFF (2006): "Do entry costs provide an empirical basis for poverty traps? Evidence from Mexican microenterprises," *Economic development and cultural change*, 55, 3–42.
- MEAGER, R. (2019): "Understanding the average impact of microcredit expansions: A Bayesian hierarchical analysis of seven randomized experiments," *American Economic Journal: Applied Economics*, 11, 57–91.
- (2022): "Aggregating distributional treatment effects: A Bayesian hierarchical analysis of the microcredit literature," *American Economic Review*, 112, 1818–1847.
- SCHOAR, A. (2010): "The divide between subsistence and transformational entrepreneurship," *Innovation policy and the economy*, 10, 57–81.
- TAROZZI, A., J. DESAI, AND K. JOHNSON (2015): "The Impacts of Microcredit: Evidence from Ethiopia," *American Economic Journal: Applied Economics*, 7, 54–89.

FIGURES

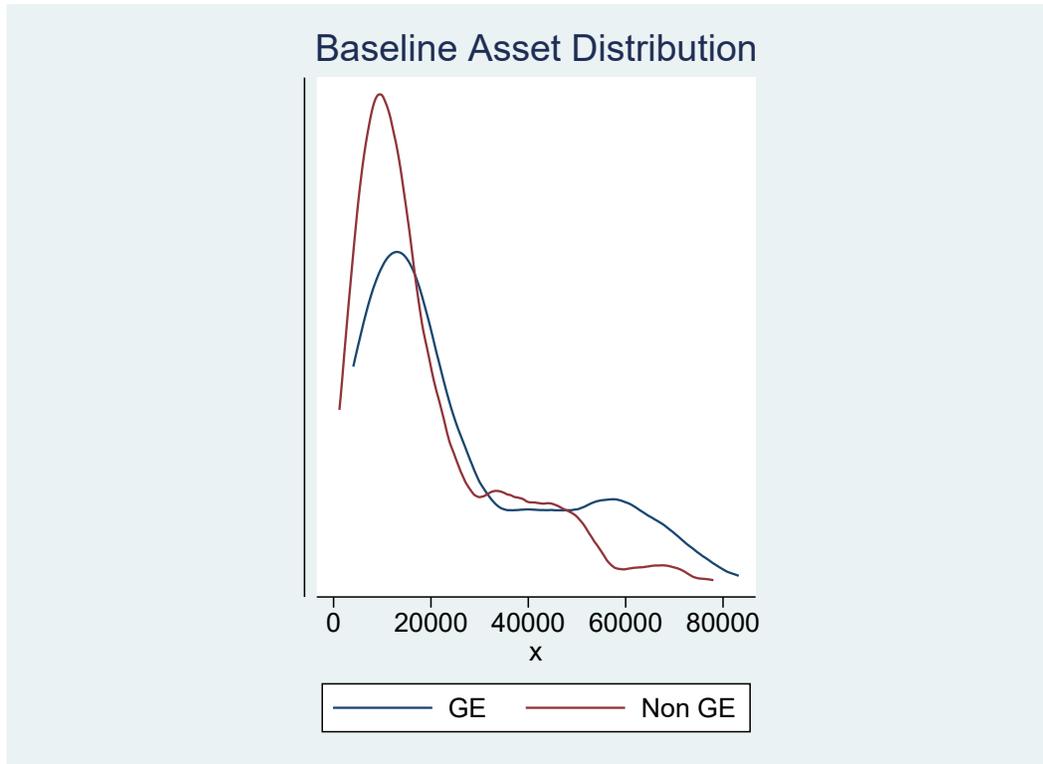


FIGURE 3. Baseline wealth, by type

Note: The figure plots the distribution of baseline (2006; pre-treatment) assets, measured as the sum of business and non-business durable assets plus business inputs, by type (Gung-ho Entrepreneur vs. Non-Gung-ho).

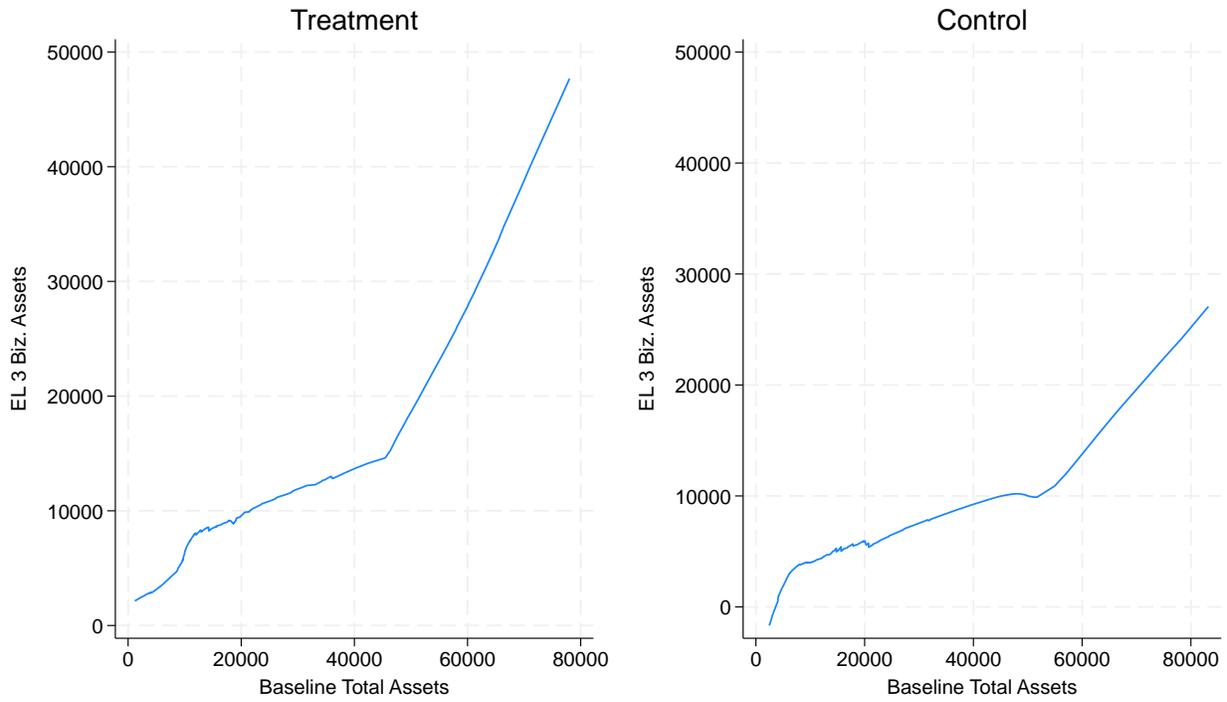


FIGURE 4. Six-year follo-up assets vs. baseline wealth, by treatment

Note: The measure on the x-axis is baseline (2006; pre-treatment) assets, measured as the sum of business and non-business durable assets plus business inputs. The y-axis measure is six-year (2012) business assets (the sum of business durable assets, business working capital, and wage bill). See Section 4.2 for details.

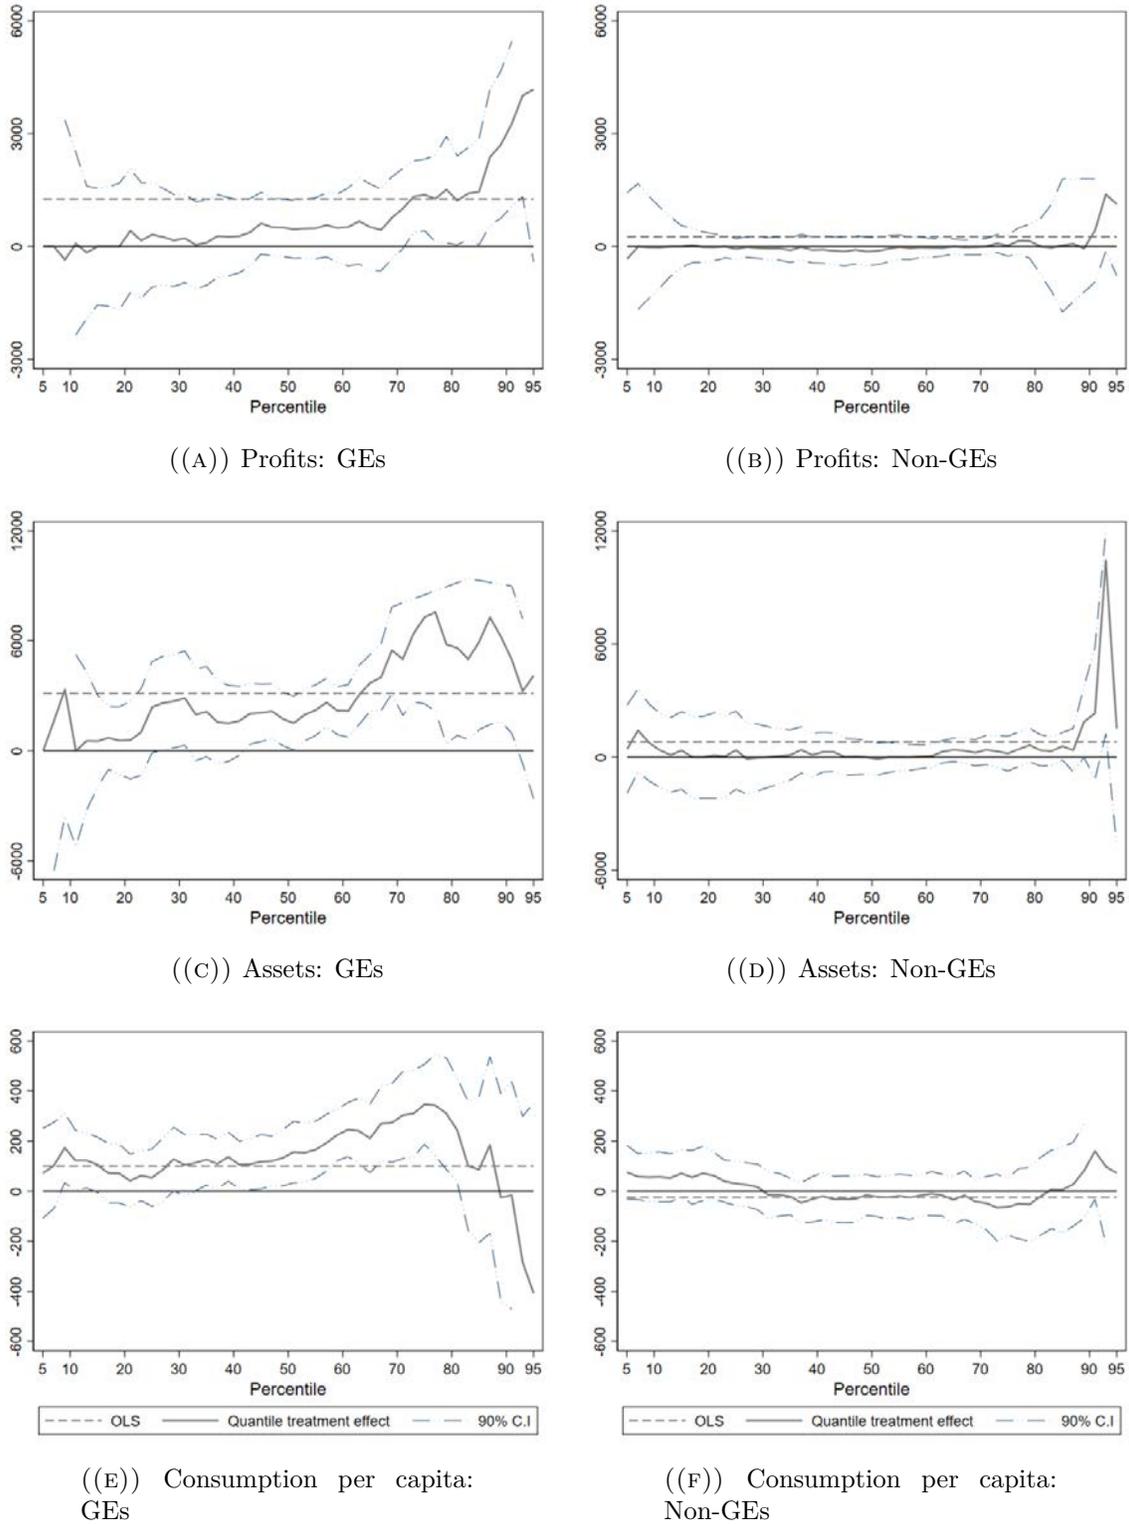


FIGURE 5. Quantile treatment effects on business outcomes (6 year endline)

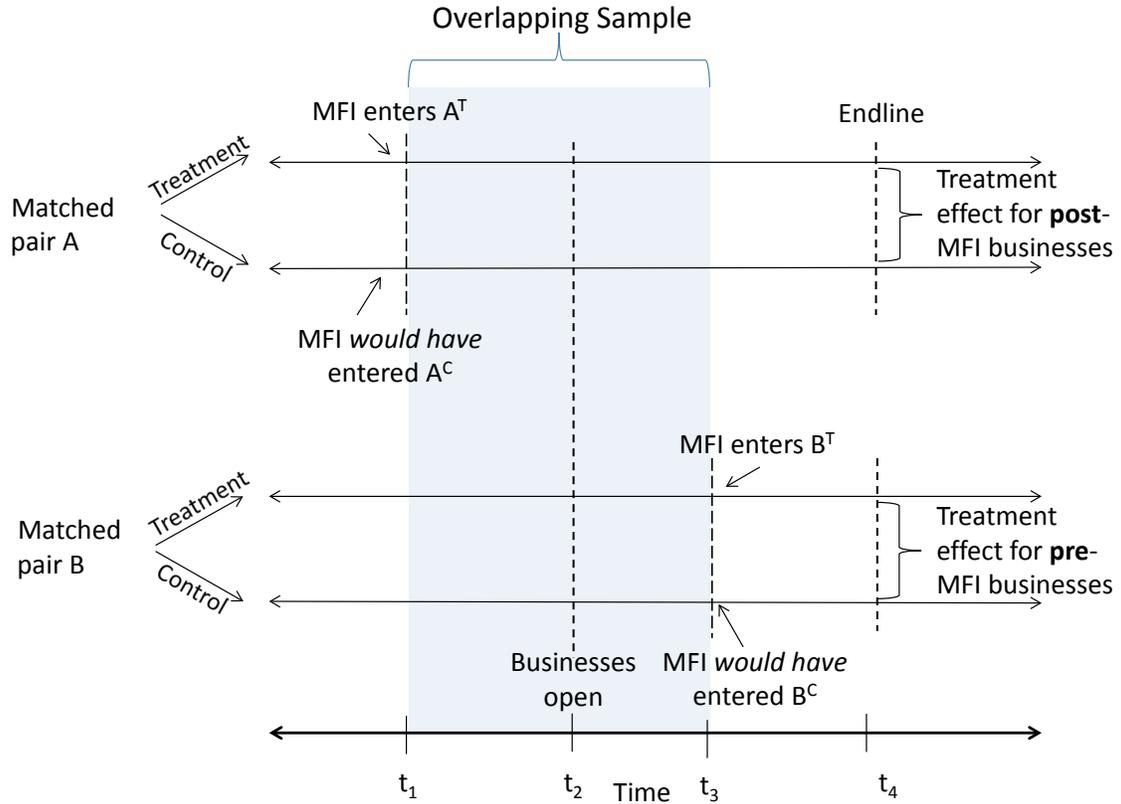
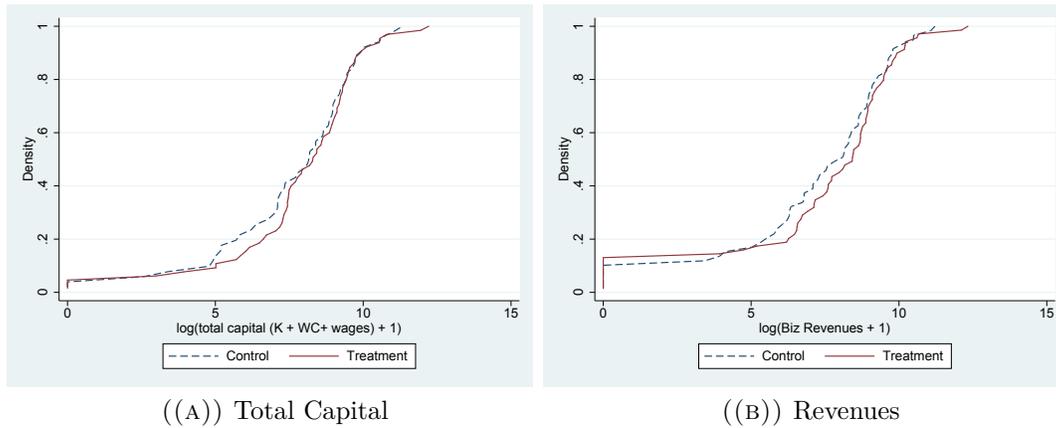


FIGURE 6. Overlapping sample identification



((A)) Total Capital

((B)) Revenues

FIGURE 7. Distributions (CDFs) of Total Capital and Revenues by Treatment Status (Two-year follow-up, overlapping sample only)

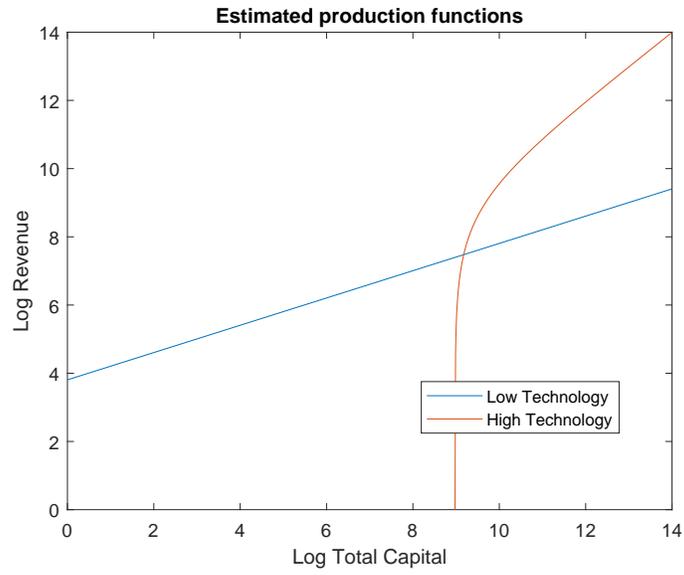


FIGURE 8. Production functions

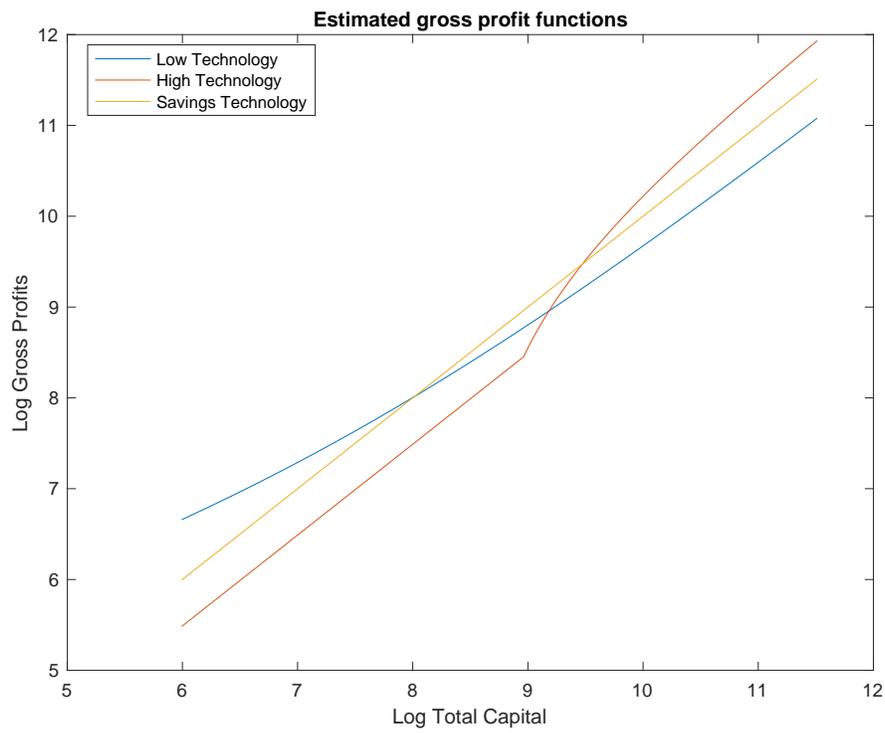


FIGURE 9. Gross profit functions

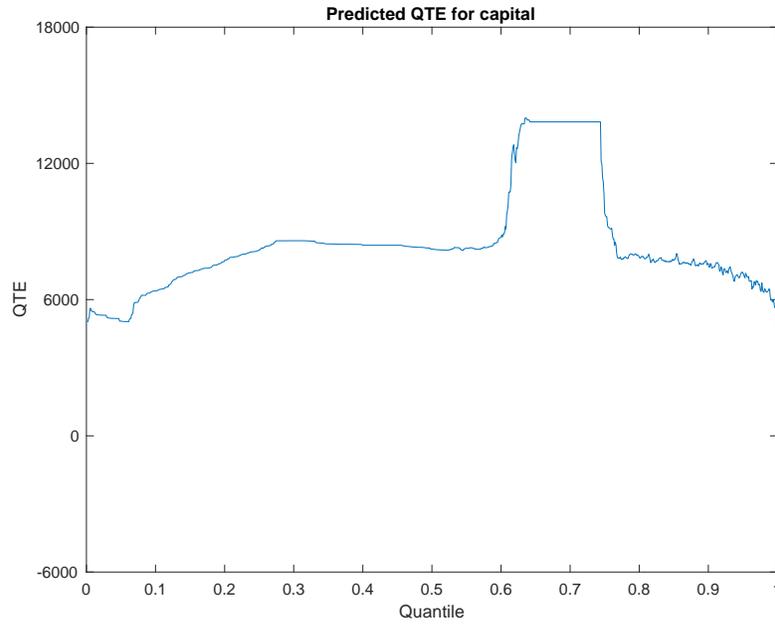


FIGURE 10. Model-implied quantile treatment effects for capital 6 years post intervention

Note: This figure plots the implied treatment effect on expected capital in year 6; see text for details. The lower quantile exhibit a positive predicted treatment effect because in some draws of the capital shock, even a low wealth individual could be pushed out of the poverty trap region by the extra capital access resulting from treatment.

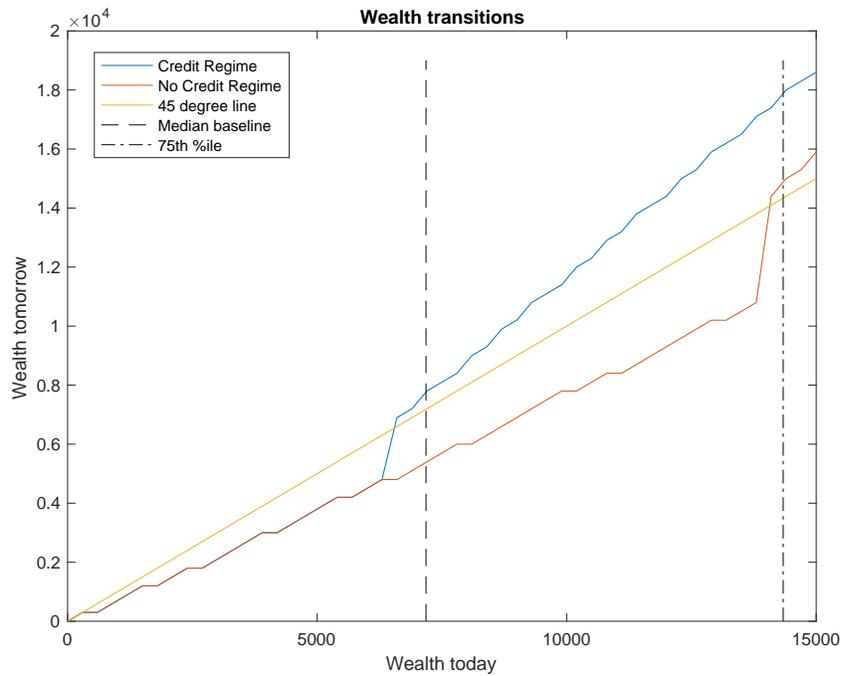


FIGURE 11. Wealth transitions with and without credit access

TABLES

TABLE 1. Six-year follow up summary household and business statistics (control group)

	Obs	Mean	Std. Dev
<u>Household composition</u>			
# members	2785	6.894	2.978
# adults (≥ 16 years old)	2785	4.221	1.975
# children (< 16 years old)	2785	1.638	1.368
Male head	2784	0.765	0.424
Head's age	2784	44.379	9.990
Head with no education	2784	0.334	0.472
<u>Access to credit (4-year follow-up)</u>			
Loan from Spandana	2946	0.112	0.316
Loan from other MFI	2946	0.268	0.443
Loan from a bank	2946	0.073	0.260
Informal loan	2946	0.603	0.489
Loan from Self-Help Group or other savings group	2946	0.092	0.290
Any type of loan	2946	0.905	0.293
<u>Amount borrowed at four-year follow-up from (Rs.):</u>			
Spandana	2946	1898	6769
Other MFI	2946	4773	10731
Bank	2946	5951	39247
Informal loan	2946	32252	76606
Self-Help Group or other savings group	2946	1003	5223
Total	2946	88244	144194
<u>Businesses</u>			
Has a business	2785	0.307	0.461
Gung-ho entrepreneur (GE)	2786	0.304	0.460
# of businesses	2785	0.371	0.613
# of businesses managed by women	2785	0.173	0.417
Share businesses managed by women	854	0.466	0.475
Sales (Rs.)	802	25240	80867
Expenses (Rs.)	849	16300	70729
Investment (Rs.)	854	3496	30499
More than 1 worker in any business	850	0.335	0.472
More than 2 workers in any business	850	0.115	0.320
# worker in largest business	850	1.660	1.884
Total work hours (hrs/week)	854	46.310	47.898
<u>Consumption (per household per month)</u>			
Consumption (Rs.)	2781	13077	9907
Non-durables cons (Rs.)	2781	11960	8455
Durables cons (Rs.)	2785	1115	3362
Asset index	2785	2.705	0.831

TABLE 2. Results on microcredit borrowing

	(1)	(2)	(3)	(4)
	Borrowed from MFI in last 3 Years (2y)	Ever borrowed from MFI before 2010 (4y)	Total MFI loan amount (2y)	Total MFI loan amount (4y)
Panel A: Cumulative exposure to microcredit				
Treatment	0.109*** (0.022)	0.044* (0.024)	1293.874*** (266.767)	788.346** (395.136)
Control Mean	0.256	0.498	2367.164	5556.331
Control Std. Dev.	0.436	0.500	6645.564	11351.364
Observations	6804	5467	6811	6143
Panel B: Cumulative exposure to microcredit by entrepreneurial status				
Treatment	0.109*** (0.021)	0.036 (0.026)	1121.479*** (255.969)	564.122 (423.303)
Treatment \times GE	-0.002 (0.030)	0.020 (0.032)	540.448 (559.421)	628.868 (774.078)
Gung-ho entrepreneur (GE)	0.163*** (0.023)	0.110*** (0.022)	1919.390*** (432.075)	2130.724*** (559.522)
Treat + Treat \times GE	0.107	0.057	1661.927	1192.990
P(Treat + Treat \times GE \neq 0)	0.001	0.091	0.003	0.102
Control Mean (Non-GEs)	0.206	0.463	1731.722	4841.876
Control Std. Dev. (Non-GEs)	0.404	0.499	4958.057	10546.729
Observations	6804	5467	6811	6143

Notes: Standard errors, clustered at the area level, reported in parentheses. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level. All regressions control for stratum FEs. The column 1 outcome is an indicator for ever borrowing in the 3 years before the 2-year endline (in 2007/2008). In column 2, the outcome is an indicator for whether the household ever reported borrowing at any time in any survey round (2, 4 or, retrospectively, 6 years). The column 3 outcome is total outstanding MFI borrowing at the 2-year endline. The column 4 outcome is total outstanding MFI borrowing at the 4-year endline (just prior to the withdrawal of microfinance).

TABLE 3. Reduced form results (Six-year follow-up)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Total wages paid	Total Assets (stock)	Log revenue	Profit (level)	Per capita consumption (monthly)	All durables	Non- business durables
Panel A: Treatment effects							
Treatment	373.747*** (133.018)	1565.222*** (426.789)	0.311** (0.157)	576.774*** (179.375)	-11.583 (58.735)	559.362* (283.691)	351.696 (239.737)
Control Mean	348.367	6680.551	2.637	2066.436	2791.712	9264.343	8482.853
Control Std. Dev.	4700.427	20448.064	4.180	6039.441	1964.732	15748.713	14264.700
Observations	5736	5744	5589	5580	5738	5744	5744
Panel B: Treatment effects by entrepreneurial status							
Treatment	275.264** (118.604)	816.198 (526.966)	0.126 (0.128)	263.906 (168.567)	-48.970 (61.965)	-54.105 (385.449)	-175.322 (323.643)
Treatment × GE	311.864 (368.366)	2325.597 (1483.448)	0.593** (0.231)	1004.523** (501.565)	120.166 (107.356)	1991.273** (850.072)	1716.980** (725.416)
Gung-ho entrepreneur (GE)	488.639* (266.816)	8906.264*** (973.087)	3.892*** (0.180)	3493.457*** (350.655)	106.282 (77.229)	20.574 (658.347)	-513.234 (563.800)
Treat + Treat × GE	587.127	3141.795	0.719	1268.429	71.196	1937.168	1541.658
P(Treat + Treat × GE ≠ 0)	0.093	0.011	0.001	0.004	0.491	0.004	0.007
Control Mean (Non-GEs)	197.888	3974.639	1.452	988.890	2759.209	9266.641	8635.084
Control Std. Dev. (Non-GEs)	2496.403	17568.209	3.318	4065.047	1886.292	15716.467	14350.628
Observations	5736	5744	5589	5580	5738	5744	5744

Notes: Standard errors, clustered at the area level, reported in parentheses. Total wages paid winsorized at the 95th percentile. Assets, expenses, revenues and profits are monthly and winsorized at the 0.5 and 99.5 percentiles. Log is $\log(x+1)$. s* significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level. All regressions control for stratum FEs.

TABLE 4. Effects on non-microcredit borrowing

	(1)	(2)	(3)	(4)
	Informal borrowing (2y)	Total business borrowing (2y)	Informal borrowing (6y)	Total business borrowing (6y)
Panel A: Non-microcredit borrowing				
Treatment	1327.308 (2086.617)	2757.482** (1389.416)	2668.157 (3545.218)	4731.785*** (1500.096)
Control Mean	40976.748	11169.293	57151.686	15062.722
Control Std. Dev.	77957.755	44445.289	1.13e+05	43892.196
Observations	6811	6067	5744	5744
Panel B: Non-microcredit borrowing by entrepreneurial status				
Treatment	-780.059 (2168.330)	-323.452 (1454.283)	-1683.957 (4226.917)	1453.121 (1853.462)
Treatment \times GE	6798.341 (5326.896)	9269.086** (4403.218)	14085.007* (7387.176)	10598.425** (4289.273)
Gung-ho entrepreneur (GE)	9186.418*** (3052.111)	13437.287*** (1953.008)	3647.067 (5833.084)	3798.873** (1766.060)
Treat + Treat \times GE	6018.282	8945.634	12401.050	12051.547
P(Treat + Treat \times GE \neq 0)	0.211	0.020	0.046	0.001
Control Mean (Non-GEs)	37947.476	6899.960	55097.667	13677.772
Control Std. Dev. (Non-GEs)	70481.529	37457.208	1.15e+05	41486.339
Observations	6811	6067	5744	5744

Notes: Standard errors, clustered at the area level, reported in parentheses. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level. All regressions control for stratum FEs. The column 1 outcome is total informal borrowing at the 2-year endline (in 2007/2008). In column 2, the outcome is business borrowing at the 2-year endline. The column 3 outcome is total informal borrowing at the 6-year endline (in 2012). The column 4 outcome is total business borrowing at the 2-year endline. Informal borrowing is loans from family, friends, neighbors, and business associates such as suppliers and customers. Loans for business are loans which the household reported they devoted to a business purpose.

TABLE 5. Effects on Capital: Reweighted regressions

	(1)	(2)	(3)
Effects on capital (6-year follow-up):			
	Unweighted	RE to GE	GE to RE
Treatment	6106.015 (4057.010)	6078.759 (3946.512)	6230.361 (4030.445)
Treatment \times GE	11629.804* (6077.825)	11127.548* (6007.240)	6148.310 (4080.228)
Treatment + Treat \times GE	17735.820	17206.307	12378.672
P(Treat + Treat \times GE \neq 0)	0.036	0.038	0.025
Control Mean (Non-GEs)	1757.200	1855.897	1757.200
Observations	481	481	481

Notes: Standard errors, clustered at the area level, reported in parentheses. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level. All regressions control for stratum-by-type FEs.

TABLE 6. Business results, overlapping sample (two-year follow-up)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Workers in largest business	Assets (stock)	Log expenses	Log revenue	Profit	Index of business variables	Informal borrowing for business
Panel A: Entered entrepreneurship post-2006, pre-Spandana							
Treatment	0.212* (0.123)	900.734 (829.002)	0.488 (0.456)	1.008* (0.567)	2801.011** (1293.561)	0.150** (0.067)	9937.288** (4543.011)
Control Mean	0.033	2100.623	6.758	7.093	1164.737	-0.105	1415.094
Control Std. Dev.	0.181	3961.748	2.853	2.959	6351.331	0.325	5914.240
Observations	133	119	130	128	128	133	118
Panel B: Entered entrepreneurship post-2006, post-Spandana							
Treatment	-0.288 (0.265)	-1500.608 (1159.788)	-0.566 (0.648)	-1.007 (0.823)	-1400.672 (1286.628)	-0.184 (0.113)	-2641.412 (5708.852)
Control Mean	0.242	2539.005	6.377	6.719	1785.719	-0.020	11976.563
Control Std. Dev.	1.110	4850.283	3.402	3.591	6797.191	0.494	36731.617
Observations	164	145	158	154	154	164	157

Notes: Standard errors, clustered at the area level, reported in parentheses. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level. All regressions control for stratum FEs. Col. 6 is a equally-weighted index of outcomes in cols 1-5. Loans for business are loans which the household reported they devoted to a business purpose; informal borrowing is loans from family, friends, neighbors, and business associates such as suppliers and customers.

APPENDIX A. SUPPLEMENTAL TABLES AND FIGURES

TABLE A1. Effects on Additional Business Outcomes

	(1)	(2)	(3)	(4)
	Has a business	Number of business	Started a business in last 12 months	Log expenses
Panel A: Treatment effects				
Treatment	0.038*	0.056*	0.006	0.273*
	(0.020)	(0.031)	(0.005)	(0.140)
Control Mean	0.307	0.371	0.032	2.293
Control Std. Dev.	0.461	0.613	0.176	3.776
Observations	5744	5744	5744	5724
Panel B: Treatment effects by entrepreneurial status				
Treatment	0.024	0.031	0.009	0.101
	(0.018)	(0.024)	(0.006)	(0.114)
Treatment \times GE	0.040	0.076**	-0.011	0.503**
	(0.028)	(0.035)	(0.013)	(0.221)
Gung-ho entrepreneur (GE)	0.422***	0.525***	0.025***	3.361***
	(0.020)	(0.026)	(0.008)	(0.174)
Treatment + Treat \times GE	0.064	0.107	-0.002	0.605
P(Treat + Treat \times GE \neq 0)	0.008	0.008	0.849	0.003
Control Mean (Non-GEs)	0.177	0.208	0.025	1.258
Control Std. Dev. (Non-GEs)	0.382	0.483	0.155	2.982
Observations	5744	5744	5744	5724

Notes: Standard errors, clustered at the area level, reported in parentheses. Assets, expenses, revenues and profits are monthly and winsorized at the 0.5 and 99.5 percentiles. Log is $\log(x+1)$. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1 level. All regressions control for stratum FEs

TABLE A2. Effects on Additional Business Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)
	More than 1 worker in any business	More than 2 worker in any business	Workers in largest months	Total weekly labor hrs	Total hrs in self employment	Closed a business in last yr
Panel A: Treatment effects						
Treatment	0.030** (0.015)	0.016** (0.006)	0.208** (0.087)	2.170 (1.661)	2.752** (1.159)	0.008** (0.004)
Control Mean	0.102	0.035	0.507	87.490	15.400	0.027
Control Std. Dev.	0.303	0.184	1.292	56.528	30.304	0.161
Observations	5738	5738	5738	5744	5744	5744
Panel B: Treatment effects by entrepreneurial status						
Treatment	0.017 (0.012)	0.009 (0.006)	0.174** (0.076)	0.150 (2.021)	1.259 (0.859)	0.006 (0.004)
Treatment × GE	0.040 (0.024)	0.023* (0.013)	0.102 (0.143)	6.501* (3.321)	4.569** (1.962)	0.006 (0.012)
Gung-ho entrepreneur (GE)	0.174*** (0.018)	0.055*** (0.009)	0.765*** (0.071)	4.798** (2.107)	23.537*** (1.587)	0.019** (0.008)
Treatment + Treat × GE	0.057	0.032	0.277	6.651	5.827	0.012
P(Treat + Treat × GE ≠ 0)	0.025	0.010	0.060	0.017	0.004	0.228
Control Mean (Non-GEs)	0.049	0.019	0.279	86.111	8.175	0.021
Control Std. Dev. (Non-GEs)	0.215	0.135	0.865	55.490	22.456	0.144
Observations	5738	5738	5738	5744	5744	5744

Notes: Standard errors, clustered at the area level, reported in parentheses. Durables variables (cols 3-5) winsorized at the 95th percentile. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level. All regressions control for stratum FEs.

TABLE A3. Appendix, consumption

	(1)	(2)	(3)	(4)	(5)	(6)
	Temptation goods (monthly)	Non-business durables	Business durables	Festivals (annual)	Education (annual)	Health (annual)
Panel A: Treatment effects						
Treatment	-19.609 (32.860)	351.696 (239.737)	23.900*** (8.242)	388.891 (304.751)	86.911 (88.338)	-45.121 (99.606)
Control Mean	619.036	8482.853	88.575	8932.731	1724.557	1775.509
Control Std. Dev.	971.999	14264.700	351.011	10950.345	3105.393	4107.905
Observations	5739	5744	5744	5737	5738	5738
Panel B: Treatment effects by entrepreneurial status						
Treatment	-21.350 (37.511)	-175.322 (323.643)	6.847 (9.196)	351.483 (372.316)	111.138 (108.004)	-77.592 (122.781)
Treatment × GE	5.530 (48.037)	1716.980** (725.416)	54.254** (26.063)	114.824 (663.376)	-82.426 (211.525)	104.366 (239.184)
Gung-ho entrepreneur (GE)	11.260 (29.245)	-513.234 (563.800)	92.895*** (16.248)	512.767 (551.631)	307.695** (153.525)	91.821 (177.110)
Treatment + Treat × GE	-15.821	1541.658	61.101	466.308	28.712	26.775
P(Treat + Treat × GE ≠ 0)	0.721	0.007	0.007	0.394	0.869	0.891
Control Mean (Non-GEs)	614.207	8635.084	60.866	8781.591	1631.834	1736.854
Control Std. Dev. (Non-GEs)	986.241	14350.628	293.414	9151.870	2663.565	4262.700
Observations	5739	5744	5744	5737	5738	5738

Notes: Standard errors, clustered at the area level, reported in parentheses. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level. All regressions control for stratum FEs.

TABLE A4. Effects on Borrowing: Reweighted regressions

	(1)	(2)	(3)
Effects on business borrowing (six-year follow-up)			
	Unweighted	RE to GE	GE to RE
Treatment	5969.886 (4010.710)	5135.230 (4325.828)	6239.750 (3987.709)
Treatment × GE	14472.436 (9640.823)	14558.119 (9684.004)	7836.070 (9365.067)
Treatment + Treat × GE	20442.321	19693.350	14075.819
P(Treat + Treat × GE ≠ 0)	0.039	0.043	0.123
Control Mean (Non-GEs)	10451.613	12539.550	10451.613
Observations	384	384	384

Notes: Standard errors, clustered at the area level, reported in parentheses. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level. Business borrowing is winsorized at the 95th percentile. The sample size in Panel B is lower than in Panel A due to non-response to questions about loan purpose. All regressions control for stratum-by-type FEs.

TABLE A5. Business results, pre-Spandana sample (two-year follow-up)

	(1)	(2)	(3)	(4)	(5)	(6)
	Workers in largest business	Assets (stock)	Log expenses	Log revenue	Profit	Index of business variables
Entered entrepreneurship pre-2006, pre-Spandana						
Treatment	0.024	389.934	0.316	0.434	2220.869**	0.071*
446.668	(0.097)	(393.948)	(0.194)	(0.280)	(946.406)	(0.036)
(2343.492)						
Control Mean	0.452	2591.982	7.339	7.716	2996.170	0.012
8585.175						
Control Std. Dev.	1.828	4760.308	2.844	3.127	14984.800	0.568
37485.339						
Observations	1305	1184	1273	1232	1232	1305
1190						

Notes: Standard errors, clustered at the area level, reported in parentheses. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level. All regressions control for stratum FEs.

TABLE A6. Attrition

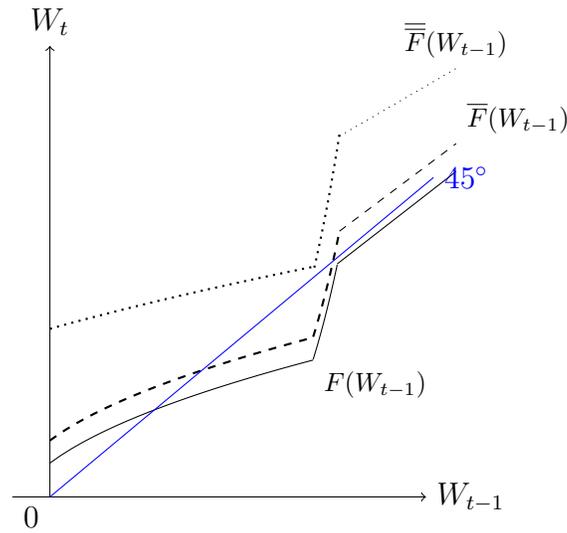
Attrition in treatment vs. control (relative to two-year follow-up)	
Found in endline 3, in treated	0.8315
Found in endline 3, in control	0.8602
<i>p-value of difference</i>	0.265

Notes: Standard errors, clustered at the area level, reported in parentheses. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level. All regressions control for stratum FEs.

TABLE A7. Main Results with Lee Bounds

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	Informal loans Any	Amt.	Businesses Any	Num.	Workers	Wages	Biz assets	Log inputs	Log revenue	Profits	Biz index
Panel A: All Households											
Treatment	-0.018 (0.013)	2311.846 (3055.914)	0.033 (0.018)	0.049 (0.027)	0.186* (0.078)	328.996** (111.253)	1375.721*** (379.209)	0.236 (0.124)	0.272 (0.139)	509.051** (158.134)	0.074*** (0.018)
Treatment											
lower	-0.040	-1708.982	0.009	0.024	0.122	246.880	800.622	0.095	0.126	346.544	0.054
upper	0.000	16562.409	0.057	0.108	0.322	660.302	4059.850	0.512	0.571	1377.160	0.166
Observations	6863	6863	6863	6863	6857	6855	6863	6843	6708	6699	6863
Panel B: GE Households Only											
Treatment	0.035 (0.023)	11532.067* (5450.123)	0.054* (0.026)	0.092* (0.039)	0.246 (0.152)	533.710 (327.819)	2796.175* (1191.440)	0.524* (0.215)	0.628** (0.232)	1152.770** (417.835)	0.111*** (0.027)
Treatment											
lower	0.019	8323.170	0.035	0.062	0.198	529.864	2203.546	0.357	0.391	948.023	0.090
upper	0.047	19155.102	0.063	0.122	0.411	457.400	4175.455	0.655	0.778	2079.608	0.149
Observations	2088	2088	2088	2088	2085	2083	2088	2078	2010	2006	2088
Panel C: Non-GE Households Only											
Treatment	-0.041* (0.017)	-1892.724 (3838.662)	0.019 (0.018)	0.024 (0.025)	0.151* (0.074)	230.592 (118.377)	640.382 (516.518)	0.065 (0.116)	0.087 (0.129)	198.866 (171.452)	0.067** (0.021)
Treatment											
lower	-0.066	-5994.197	-0.000	0.003	0.082	138.565	160.201	-0.071	-0.016	72.547	0.037
upper	-0.021	13434.927	0.051	0.083	0.257	420.839	3125.577	0.391	0.445	804.496	0.168
Observations	4775	4775	4775	4775	4772	4772	4775	4765	4698	4693	4775

APPENDIX B. SUPPLEMENTAL FIGURES



Non-convexity, multiple equilibria due to microcredit. If microcredit lowers borrowing costs slightly (from $F(W_{t-1})$ to $\bar{F}(W_{t-1})$; long dashed line) a new, better equilibrium is created. If microcredit lowers borrowing costs more significantly (to $\bar{\bar{F}}(W_{t-1})$; small dashed line) a new growth path is created.

FIGURE B1. Possible transition dynamics (cont.)

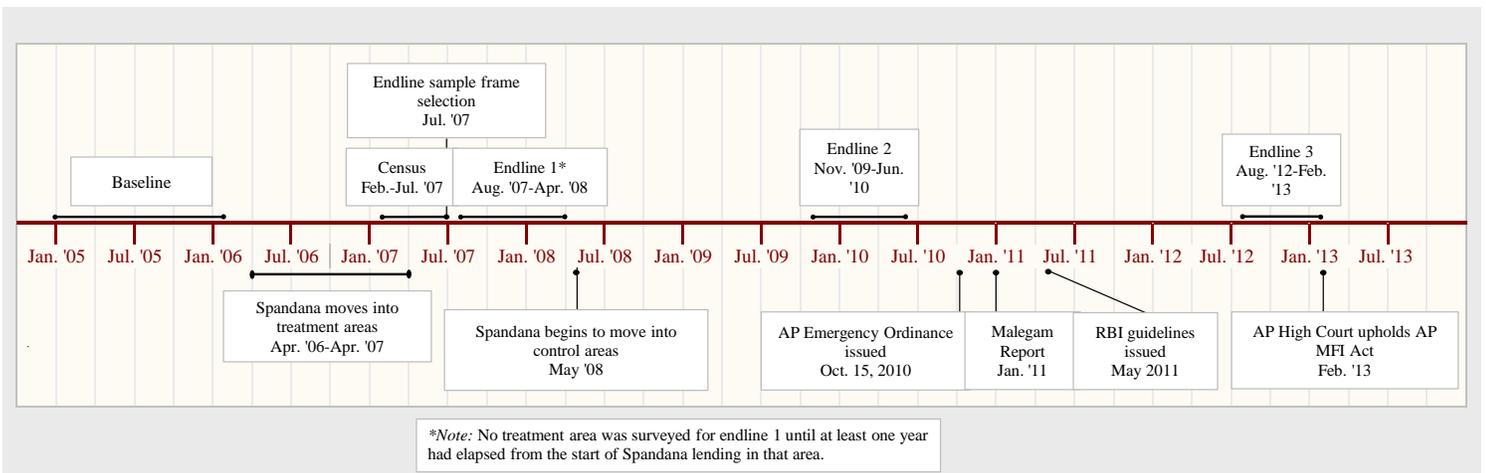


FIGURE B2. Timeline of Survey Activities and Microfinance Crisis

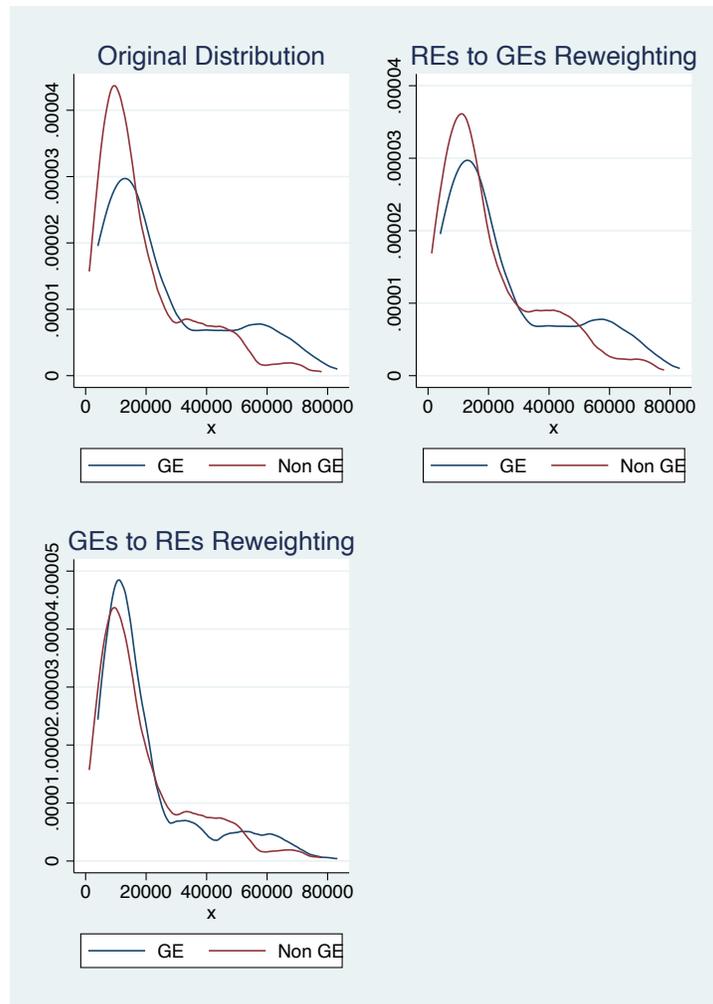


FIGURE B3. Overlap in the baseline wealth distributions of GEs and non-GEs

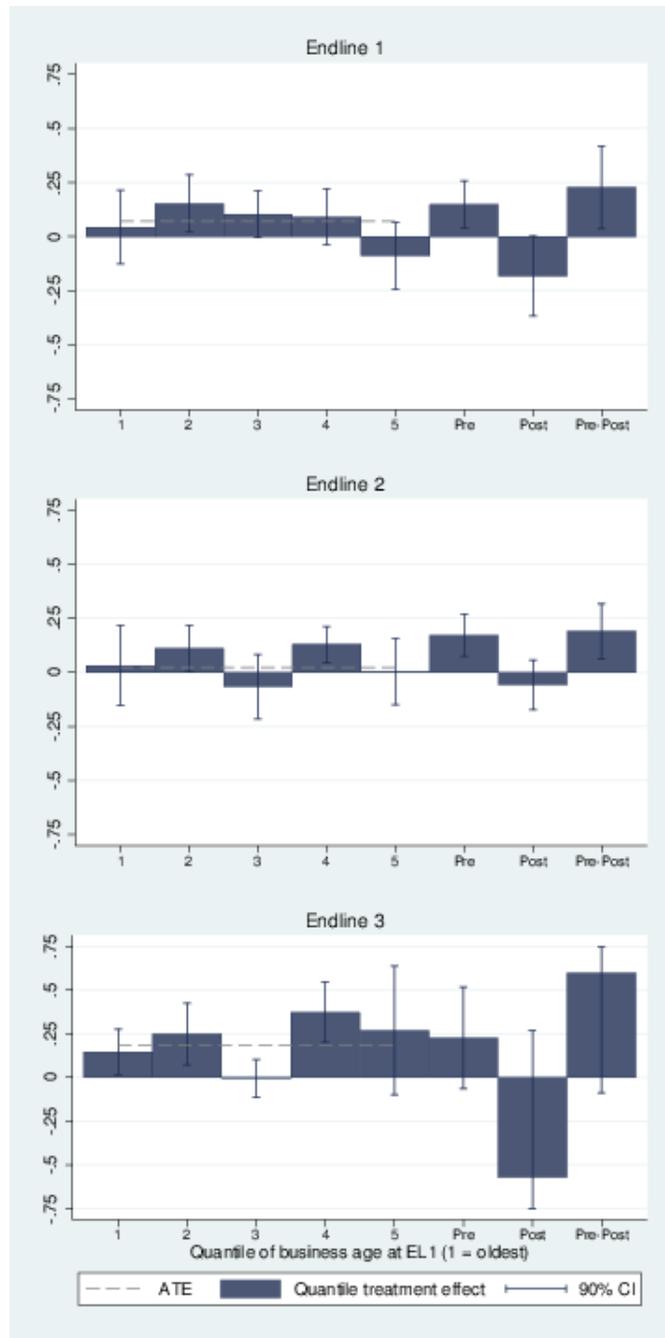


FIGURE B4. Experience vs. Selection: Treatment effects on index of business outcomes
 Note: The figure plots the treatment effects for gung-ho entrepreneurs, by quintile of business age (1=oldest, 5=youngest).

APPENDIX C. POTENTIAL WINDFALL EFFECTS OF THE AP ORDINANCE

Here, we explore the potential that larger windfalls in the treatment group can help to explain the persistent impacts on businesses and borrowing that we find.

As we describe in Section 5.3, there are differences in total microcredit balance coming through the extensive margin. Panel C of Table 2 shows that treated households had Rs. 946 more microcredit at endline 2. This implies that the treatment group did most likely have a slightly larger windfall (i.e., balance outstanding in October 2010). One should note that the size of the windfall was likely significantly smaller than the Rs. 946, which is the treatment effect on the initial loan amount, not the principal remaining at the time of the AP crisis. Given that microfinance installment payments are typically of equal size, occurring weekly over 50 weeks, the implied average difference in windfall would have been approximately half that amount, or Rs. 473.

A priori, a roughly Rs. 470 difference in windfall in 2010 is unlikely to drive the outcomes we observe in 2012. However, we are also able to directly examine the extent to which larger windfall differences had any impacts on business outcomes. To do this, we use data collected at endline 3 about each household's microfinance loans in October 2010. Specifically, households reported the size of their loans and how many installments of the loan had not yet been repaid as of October 2010.²⁸ This allows us to compare households who are otherwise very similar, but who received very different windfalls. Imagine two households, both of whom took loans in late 2009. One who borrowed in, say, September 2009 would have finished repaying that loan and gotten a fresh loan just before the crisis. They would have only made a few installment payments on this new loan prior to the ordinance and therefore received a large windfall. Another household, who borrowed in, say, November 2009, would have repaid almost that full loan, but not yet received a new loan, at the time of the crisis. Therefore, they received a small windfall. Thus, small differences in timing of the initial loan disbursement (e.g., September vs. November 2009) can lead to large differences in the size of the windfall, allowing a regression discontinuity design to identify the causal effect of receiving a larger windfall.

Appendix Table A8 shows that, among the sample with a loan maturity near the time of the crisis, having a slightly earlier maturity is associated with a large and significant increase in the share of the loan balance outstanding at the time of the AP crisis. Column 1 uses a broader measure of having a maturity near the crisis, i.e. households within +/- 10 weeks of the loan maturity date. Column 2 uses a narrower measure, households within +/- 8 weeks of the loan maturity date. In either case, those with the later maturity dates—who have almost but not quite finished repaying the previous loan—have only 4 to 4.5% of their loan balance outstanding; this is the amount they received as a windfall. In contrast, those whose

²⁸The AP Ordinance was very salient to all of the respondents at this time. This is especially true given that at the time of the survey, households were unable to take new loans from any microlender. Many households even showed us their microfinance loan cards when answering these questions.

maturity date was earlier—and had therefore repaid their loan and obtained a new loan—have approximately 80-85 additional percentage points of the loan balance outstanding; they received a large windfall. In other words, the loan timing instrument has a first stage that is quite large and highly significant.

Next, Appendix Table A9 presents the corresponding reduced form, comparing the main business outcomes across individuals with large versus small windfalls. In each specification, we control for an indicator of whether the respondent had a loan outstanding in October 2010, and therefore had any windfall (this would include both the Sept 2009 and Nov 2009 borrowers in the example above). Our coefficient of interest is an indicator for having received a large windfall (this would be 1 for the Sept 2009 borrower and 0 for the Nov 2009 borrower). This regressor measures the difference in outcome for large windfall recipients, relative to those who received a small windfall. In panels A and B, a large or small windfall is defined as someone with less than 10 or more than 40 weeks remaining in the loan cycle, i.e. households within +/- 10 weeks of their loan maturity date. In panels C and D, we consider a narrower window of plus or minus 8 weeks. In all cases, we find no evidence that receiving a larger windfall in 2010 is associated with better business outcomes in 2012, either on average or separately for GE or non-GE households. In fact, many of the point estimates are negative.

These results show that even large differences in windfall amount—of roughly 80% of loan balance, or approximately Rs. 8,000—led to no improvement in business outcomes two years later. Recall that the “experiment” that identifies this effect does not exploit treatment assignment, but instead relies on the discontinuity generated by the timing of the crisis. This makes it highly unlikely that the much smaller difference in windfall between treatment and control (approximately Rs. 470, or less than 5% of the balance of a Rs. 10,000 loan) can explain any portion of the long run results.

TABLE A8. Windfall first stage

	Share of loan outstanding in Oct. 2010	
	(1)	(2)
Large windfall (broad def.)	0.802*** (0.007)	
Large windfall (narrow def.)		0.848*** (0.007)
Low WF Mean	0.081	0.075
Low WF Std. Dev.	0.044	0.038
Observations	1095	1095

Notes: Broad windfall measure is loans with maturities from 10 weeks before to 10 weeks after the crisis; narrow windfall measure is loans with maturities from 8 weeks before to 8 weeks after the crisis. All regressions control for an indicator for receiving either a small (loan duration less than 8 or 10 weeks as of the crisis) or large (loan duration more than 42 or 40 weeks as of the crisis) windfall. Standard errors, clustered at the area level, reported in parentheses. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level. All regressions control for stratum FEs.

TABLE A9. Windfall effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Has a business	Number of business	Started a business in last 12 months	Total Assets (stock)	Log expenses	Log revenue	Profit (level)
Panel A: Windfall effects (broad windfall measure)							
Large windfall (broad def.)	-0.025 (0.048)	-0.024 (0.064)	-0.001 (0.020)	618.578 (2294.784)	-0.322 (0.428)	-0.283 (0.471)	334.441 (728.523)
No WF Mean	0.315	0.387	0.035	7044.668	2.307	2.659	2240.111
No WF Std. Dev.	0.465	0.632	0.183	20854.783	3.803	4.205	6623.068
Observations	5744	5744	5744	5744	5724	5589	5580
Panel B: Windfall effects by entrepreneurial status (broad windfall measure)							
Large windfall (broad def.)	0.029 (0.048)	0.034 (0.058)	-0.001 (0.022)	2684.679 (3308.989)	-0.134 (0.375)	-0.034 (0.445)	196.712 (699.033)
Large windfall (B) X GE	-0.143* (0.078)	-0.151 (0.101)	-0.001 (0.033)	-5503.608 (4797.283)	-0.510 (0.638)	-0.799 (0.696)	252.115 (1357.417)
P(LWF + LWF × GE ≠ 0)	0.090	0.212	0.948	0.368	0.308	0.192	0.718
Panel C: Windfall effects (narrow windfall measure)							
Large windfall (narrow def.)	-0.015 (0.055)	-0.025 (0.074)	-0.009 (0.024)	322.944 (2889.827)	-0.254 (0.524)	-0.204 (0.556)	448.723 (907.899)
No WF Mean	0.315	0.387	0.035	7044.668	2.307	2.659	2240.111
No WF Std. Dev.	0.465	0.632	0.183	20854.783	3.803	4.205	6623.068
Observations	5744	5744	5744	5744	5724	5589	5580
Panel D: Windfall effects by entrepreneurial status (narrow windfall measure)							
Large windfall (narrow def.)	0.030 (0.056)	0.022 (0.059)	-0.008 (0.027)	3706.244 (4453.758)	-0.160 (0.492)	-0.146 (0.533)	-307.100 (685.606)
Large windfall (N) X GE	-0.175* (0.105)	-0.196 (0.119)	0.001 (0.026)	-9291.611 (5762.282)	-0.768 (0.806)	-0.998 (0.913)	1016.813 (1567.896)
P(LWF + LWF × GE ≠ 0)	0.095	0.127	0.783	0.104	0.212	0.153	0.649
No WF Mean (Non-GEs)	0.182	0.219	0.028	4083.416	1.238	1.434	1108.522
No WF Std. Dev. (Non-GEs)	0.386	0.509	0.164	17327.776	2.979	3.319	4721.134
Observations	5744	5744	5744	5744	5724	5589	5580

Notes: Broad windfall measure is loans with maturities from 10 weeks before to 10 weeks after the crisis; narrow windfall measure is loans with maturities from 8 weeks before to 8 weeks after the crisis. LWF is large windfall; B is broad and N is narrow. GE is gung-ho entrepreneur. All regressions control for an indicator for receiving either a small (loan duration less than 8 or 10 weeks as of the crisis) or large (loan duration more than 42 or 40 weeks as of the crisis) windfall. Assets, expenses, revenues and profits winsorized at the 0.5 and 99.5 percentiles. Standard errors, clustered at the area level, reported in parentheses. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level. All regressions control for stratum FEs.

APPENDIX D. MODEL ESTIMATION DETAILS

D.1. Production function estimation details. We conduct the GMM estimation as follows:

- (1) Split the data into 15 quantiles (approx. 20 firms each): this trades off flexibility vs. sample size.
- (2) Define a grid for parameters $A_1, A_2, \alpha_1, \bar{K}$ and a shock matrix for business survival (ρ).

- (3) Calculate revenues under each technology, \hat{y}_1 and \hat{y}_2 , for every value of the parameters in the grid, for each observed value in K_T^* and K_C^* . With probability $1-\rho$ both y_1 and y_2 are 0 for a given value of K_T^* or K_C^* .
- (4) Take the max of \hat{y}_1 and \hat{y}_2 to get \hat{y} , the revenue given the optimal technology choice.
- (5) Calculate the difference between treatment and control for every observed capital value by taking $\hat{y}_T - \hat{y}_C$.
- (6) Aggregate the difference calculated above into 15 bins and take the average of each bin.
- (7) Calculate the difference between the quantiles of \hat{y} for treatment and control (“tdiff_pred”).
- (8) Find the values in the grid that minimize the squared difference between tdiff_pred and tdiff_data. (Where tdiff_data is the observed difference, from the data.)

D.2. Supplemental Model Figures.

TABLE A10. Sensitivity of Simulated Treatment Effects to Preference Parameters and Method for Eliciting Baseline Wealth

		σ			
		0.86	1.00	1.25	1.50
Panel A: Preference Parameters	β	0.85	0.90	0.95	
		8053.40	7643.26	7005.75	
		6955.73	6836.85	5808.94	
		5238.80	4897.19	4938.21	4007.89
		4007.89	4008.76	4020.33	

		ν			
		500	1500	3500	5000
Panel B: Variance in iid Noise	Predicted Treatment Effect: K	8053.40	7608.05	7228.27	7662.29

Notes: Table shows the 6 year treatment effects from the simulated data varying the preferences parameters in Panel A and the iid noise parameter in Panel B. β is the discount factor and σ is the CRRA curvature parameter. ν is the variance of the iid noise applied to the predicted capital decisions under the model.

REFERENCES

ADHVARYU, A., N. KALA, AND A. NYSHADHAM (2019): “Booms, Busts, and Household Enterprise: Evidence from Coffee Farmers in Tanzania,” *The World Bank Economic Review*.

AHGION, P. AND P. BOLTON (1997): “A Theory of Trickle-Down Growth and Development,” *Review of Economic Studies*, 64, 151–172.

ANDREONI, J. AND C. SPRENGER (2012): “Estimating time preferences from convex budgets,” *American Economic Review*, 102, 3333–56.

- ANGELUCCI, M., D. KARLAN, AND J. ZINMAN (2015): “Microcredit Impacts: Evidence from a Randomized Microcredit Program Placement Experiment by Compartamos Banco,” *American Economic Journal: Applied Economics*, 7, 151–82.
- ATTANASIO, O., B. AUGSBURG, R. DE HAAS, E. FITZSIMONS, AND H. HARMGART (2015): “The Impacts of Microfinance: Evidence from Joint-Liability Lending in Mongolia,” *American Economic Journal: Applied Economics*, 7, 90–122.
- AUGSBURG, B., R. DE HAAS, H. HARMGART, AND C. MEGHIR (2015): “The Impacts of Microcredit: Evidence from Bosnia and Herzegovina,” *American Economic Journal: Applied Economics*, 7, 183–203.
- BALBONI, C., O. BANDIERA, R. BURGESS, M. GHATAK, AND A. HEIL (2022): “Why Do People Stay Poor?” *Quarterly Journal of Economics*, 137.
- BANERJEE, A., E. BREZA, A. G. CHANDRASEKHAR, E. DUFLO, M. O. JACKSON, AND C. KINNAN (2024): “Changes in social network structure in response to exposure to formal credit markets,” *Review of Economic Studies*, 91, 1331–1372.
- BANERJEE, A., E. DUFLO, R. GLENNERSTER, AND C. KINNAN (2015): “The Miracle of Microfinance? Evidence from a Randomized Evaluation,” *American Economic Journal: Applied Economics*, 7, 22–53.
- BANERJEE, A. V. AND E. DUFLO (2005): “Growth theory through the lens of development economics,” *Handbook of economic growth*, 1, 473–552.
- BANERJEE, A. V. AND A. F. NEWMAN (1993): “Occupational Choice and the Process of Development,” *Journal of Political Economy*, 101, 274–298.
- BARI, F., K. MALIK, M. MEKI, AND S. QUINN (2024): “Asset-based microfinance for microenterprises: Evidence from Pakistan,” *American Economic Review*, 114, 534–574.
- BEAMAN, L., D. KARLAN, B. THUYSBAERT, AND C. UDRY (2023): “Selection into credit markets: Evidence from agriculture in Mali,” *Econometrica*, 91, 1595–1627.
- BREZA, E., S. KAUR, AND Y. SHAMDASANI (2021): “Labor rationing,” *American Economic Review*, 111, 3184–3224.
- BREZA, E. AND C. KINNAN (2021): “Measuring the equilibrium impacts of credit: Evidence from the Indian microfinance crisis,” *The Quarterly Journal of Economics*, 136, 1447–1497.
- BRYAN, G., A. OSMAN, AND D. KARLAN (2023): “Big loans to small businesses: predicting winners and losers in an entrepreneurial lending experiment,” *American Economic Review*.
- BUERA, F. J., J. P. KABOSKI, AND Y. SHIN (2011): “Finance and development: A tale of two sectors,” *American economic review*, 101, 1964–2002.
- (2021): “The macroeconomics of microfinance,” *The Review of Economic Studies*, 88, 126–161.
- CHERNOZHUKOV, V., M. DEMIRER, E. DUFLO, AND I. FERNANDEZ-VAL (2018): “Generic machine learning inference on heterogenous treatment effects in randomized experiments,” Tech. rep., National Bureau of Economic Research.

- CRÉPON, B., F. DEVOTO, E. DUFLO, AND W. PARIENTÉ (2015): “Estimating the Impact of Microcredit on Those Who Take It Up: Evidence from a Randomized Experiment in Morocco,” *American Economic Journal: Applied Economics*, 7, 123–50.
- DASGUPTA, P. AND D. RAY (1986): “Inequality as determinant of malnutrition and unemployment: theory,” *The Economic Journal*, 96, 1011–1034.
- DE MEL, S., D. MCKENZIE, AND C. WOODRUFF (2010): “Who are the microenterprise owners? Evidence from Sri Lanka on Tokman versus De Soto,” in *International differences in entrepreneurship*, University of Chicago Press, 63–87.
- DIAO, X., J. KWEKA, AND M. MCMILLAN (2018): “Small firms, structural change and labor productivity growth in Africa: Evidence from Tanzania,” *World Development*, 105, 400–415.
- FREDERICK, S., G. LOEWENSTEIN, AND T. O’DONOGHUE (2002): “Time discounting and time preference: A critical review,” *Journal of economic literature*, 40, 351–401.
- HUSSAM, R., N. RIGOL, AND B. N. ROTH (2022): “Targeting high ability entrepreneurs using community information: Mechanism design in the field,” *American Economic Review*, 112, 861–898.
- KABOSKI, J. AND R. TOWNSEND (2011): “A structural evaluation of a large-scale quasi-experimental microfinance initiative,” *Econometrica*, 79, 1357–1406.
- KABOSKI, J. P., M. LIPSCOMB, V. MIDRIGAN, AND C. PELNIK (2022): “How Important are Investment Indivisibilities for Development? Experimental Evidence from Uganda,” Tech. rep., National Bureau of Economic Research.
- KARAIVANOV, A. AND T. YINDOK (2022): “Involuntary entrepreneurship—Evidence from Thai urban data,” *World Development*, 149, 105706.
- KARLAN, D. AND J. ZINMAN (2009): “Expanding credit access: Using randomized supply decisions to estimate the impacts,” *The Review of Financial Studies*, 23, 433–464.
- (2018): “Long-run price elasticities of demand for credit: evidence from a countrywide field experiment in Mexico,” *The Review of Economic Studies*, 86, 1704–1746.
- KERR, W. AND R. NANDA (2009): “Democratizing entry: Banking deregulations, financing constraints, and entrepreneurship,” *Journal of Financial Economics*, 99(1), 124–149.
- LA PORTA, R. AND A. SHLEIFER (2008): “The Unofficial Economy and Economic Development,” *Brookings Papers on Economic Activity*.
- LEE, D. S. (2009): “Training, wages, and sample selection: Estimating sharp bounds on treatment effects,” *The Review of Economic Studies*, 76, 1071–1102.
- LEWIS, B., N. AGRAWAL, C. GADI, D. GOYAL, J. KULKARNI, A. TAWAKLEY, S. VISWANATHAN, A. WADHWANI, A. AUGEREAU, V. BHALLA, ET AL. (2001): “India: The growth imperative,” *The McKinsey Global Institute*, 3.
- LLOYD-ELLIS, H. AND D. BERNHARDT (2000): “Enterprise, inequality and economic development,” *The Review of Economic Studies*, 67, 147–168.

- MAITRA, P., S. MITRA, D. MOOKHERJEE, A. MOTTA, AND S. VISARIA (2017): “Financing smallholder agriculture: An experiment with agent-intermediated microloans in India,” *Journal of Development Economics*, 127, 306–337.
- MCKENZIE, D. AND C. WOODRUFF (2006): “Do entry costs provide an empirical basis for poverty traps? Evidence from Mexican microenterprises,” *Economic development and cultural change*, 55, 3–42.
- MEAGER, R. (2019): “Understanding the average impact of microcredit expansions: A Bayesian hierarchical analysis of seven randomized experiments,” *American Economic Journal: Applied Economics*, 11, 57–91.
- (2022): “Aggregating distributional treatment effects: A Bayesian hierarchical analysis of the microcredit literature,” *American Economic Review*, 112, 1818–1847.
- SCHOAR, A. (2010): “The divide between subsistence and transformational entrepreneurship,” *Innovation policy and the economy*, 10, 57–81.
- TAROZZI, A., J. DESAI, AND K. JOHNSON (2015): “The Impacts of Microcredit: Evidence from Ethiopia,” *American Economic Journal: Applied Economics*, 7, 54–89.