

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

We thank Bruno Barsanetti, Ozgur Bozcaga, Janjala Chirakijja, Ofer Cohen, 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, 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 (SES 1156182). 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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

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? Could the small microcredit treatment effects often found for the average household mask important heterogeneity? In Hyderabad, India, 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 those in control neighborhoods. 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
50 Memorial Drive
Cambridge, MA 02142
and NBER
banerjee@mit.edu

Esther Duflo
Department of Economics, E52-544
MIT
50 Memorial Drive
Cambridge, MA 02142
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
8 Upper Campus Road
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 (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). However, rigorous empirical support for this idea has been elusive.¹

In this paper we document evidence that non-convexities² can and do give rise to poverty-trap dynamics. We also uncover an important reason why the search for such evidence has been challenging: households are heterogeneous, with some households 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. For example, some households are able to supervise several employees working with sewing machines—if they can buy or rent the machines and find space for them—while others are only able to sew by hand by themselves and face steeply diminishing returns.

In such a world, the impact of improved access to credit would be heterogeneous in a specific way. Some people are borrowing for consumption and are not interested in starting a business. Even among those who are willing to start a business, some may only have access to an inefficient technology with diminishing returns. Cheaper credit may prompt them to start a new business, but their target business size is small and therefore any revenue and profit effects will be small. In contrast, those who have access to the high return-high fixed cost technology, who we call gung-ho entrepreneurs (henceforth GEs) have a larger target business size and therefore will take full advantage of the additional credit. As a result, the revenue and profit effects will be large. Moreover, access to credit will allow some intermediate-wealth GE households to escape the poverty trap, and consequently their treatment effects will persist and increase over time, even if credit is withdrawn.

¹Kaboski and Townsend (2011) estimate a structural model assuming production non-convexities. Specifically, each period, households have access to a single, indivisible and illiquid project of stochastic, fixed scale, and households must decide whether to operate that specific business or work in the labor market. However, this assumption of non-convexities in production is not directly testable in their data, and they are unable to precisely detect any average impacts on investment. Instead, their structural model of consumption smoothing, precautionary savings, and investment has more power to explain the observed consumption patterns in the data. McKenzie and Woodruff (2006) argue that fixed costs of *starting* a business do not appear to constrain Mexican microentrepreneurs. (Note this is distinct from whether there are fixed costs of moving from one mode of production to another.)

²The 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, will be able to grow much larger. We will use these notions interchangeably.

We look for evidence of non-convexities and poverty traps using a new wave of data collection on the sample of households that were included in a microcredit RCT in Hyderabad, India (Banerjee et al., 2015). In 2006, a microfinance lender, Spandana, entered into 52 neighborhoods, randomly selected out of 104 in a matched-pair design. The main finding of our original study was that, on average, the effects of microfinance were very small, two years and four years after introduction. These results are broadly consistent with those of five other RCTs (Attanasio et al. (2015); Augsburg et al. (2015); Crépon et al. (2015); Karlan and Zinman (2009); Tarozzi et al. (2015); Angelucci et al. (2015)). However, for businesses that existed before Spandana entered, we found an expansion in businesses scale (sales, inputs and investment) after two years, and an average increase in profits of Rs. 2,206 in treatment areas, representing a 100% increase relative to the control mean of Rs. 2,000. The overall effect on an index of business outcomes was significant and positive (0.09 standard deviation, with a p-value of 0.057 after correcting for multiple inference). Immediately after the first endline, the control neighborhoods became eligible for microcredit. Two years later, the control group where as likely to have taken a microcredit loan, but the treatment group had borrowed for longer and had larger loans. The differences between business performance were not significant any more, even for the group that had a business before, although we still find that they had higher profits on average (p-value 0.125).

Meager (2019) shows that the difference between those with a pre-existing business and those without one is a common pattern in microfinance studies. In a meta-analysis she finds that that, on average and in 5 of the 6 sites she studies, existing business owners experience more-positive treatment effects than other households. Moreover, in a study using data from a Moroccan microcredit RCT leveraging machine learning methods to detect heterogeneous treatment effects, Chernozhukov et al. (2018) use double machine learning and find that the existence of a prior business is the only household level variable that consistently predicts a higher treatment effect. It makes economic sense: In the absence of microcredit, the cost of capital is high. Therefore, only those with high net returns – either due to high productivity, preferences for self-employment, or lower outside options – should select into entrepreneurship. We therefore identify the pre-existing business owners as the gung-ho entrepreneurs (GE). The rest of the sample is a mix of consumption borrowers and “reluctant entrepreneurs” (REs), who may start a business when credit is cheap but do not intend to grow it.

Shortly after our second endline survey, there was a severe crisis of the microfinance movement in Andhra Pradesh and the entire microfinance sector discontinued its activities in Hyderabad. (A timeline of events appears in Figure 2.) By that time, GE households in control neighborhoods had borrowed Rs. 2428 in total from Spandana on average, and GE households in treatment households had borrowed Rs. 4625, or 76% more on average. Two

years later, six years after the treatment neighborhoods were first exposed to microcredit, four years after the control neighborhoods got access, and two years after everyone lost access to microfinance, we went back to all the households included in the original study. While we still do not detect significant impacts on consumption on average, we now do 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), turnover (expenses and revenues) and employment (employees and wage bill).³ The 6-year business impacts we document here are more precise and larger than those seen 2 or 4 years after the initial intervention.

Crucially, while detectable on average, these positive impacts are driven by the GE subsample. Panel A of Figure 1 presents the GE treatment effects pictorially across the three survey waves for a subset of business outcomes. These firms' asset stocks, investment, self-employment hours, business expenses, revenues and profits are all higher in treatment neighborhoods by statistically and economically significant margins. Self-employment hours increase almost 20%, the stock and flow of business assets increase by almost 25% and 40%, respectively, business expenses increase by 80%, and revenues more than double, relative to GEs in control areas. We also find positive and significant effects on the average profits of GEs, with effects concentrated in the top tercile of the distribution. For the GE subsample, we also find large significant impacts on both business and non-business durables spending.⁴ They are also borrowing more from informal lenders. This is consistent with the presence of a non-convexity: in response to additional microcredit access, these households seek out even more credit from other sources. Finally, while we find no impact on average non durable consumption in this group, we do see positive impacts for most of the distribution, for the 30th to 80th quantiles. (See figure 3, panel E.) The results for this group demonstrate that the effects of a temporary influx of cheaper credit persist, and even increase over time, even though the control group got access as well, and even though both groups eventually lost access to microcredit.

In contrast, for the non-GEs (comprising REs and consumption borrowers), there is no evidence of impacts on most business outcomes or on consumption. Panel B of Figure 1 presents a subset of business treatment effects for the non-GE sample across the three survey waves for a subset of business outcomes. The effects are much smaller than those for the GEs, are generally not significantly different from zero, and show no tendency to increase over time. We also find that their informal borrowing went down. The lack of effect on non GEs is driven by two facts. First, only 20% of these households start

³The findings on wages and employment are consistent with the general equilibrium predictions laid out theoretically by Buera et al. (2017) and tested empirically in the context of India by Breza and Kinnan (2018).

⁴If household durables, which include both items like gold and those like television, are a combination of savings and consumption, this suggests that the income gains experienced by the GEs are partly saved and partly consumed.

businesses after 2006, either in treatment areas or in control.⁵ Many of these households use cheaper credit to replace existing, more expensive sources of credit (Banerjee and Duflo, 2014). Second, even the businesses they do start remain small.

Our analysis of the differences between the GE and non-GE samples poses several challenges. First, the GE businesses are, by definition, older on average than the non-GE businesses. Thus, it could be the case that the non-GE businesses simply need more time and/or experience before their trajectory starts to resemble that of the GEs---this would be a story not of heterogeneity, but of a head start. Related, it could be that the GEs are not inherently better, but instead are benefiting from a first-mover advantage (eg, setting up the first grocery store in the community instead of the fourth). A different set of challenges relates to the fact that, as mentioned above, only a small share of non-GEs (approximately 20%), start businesses after Spandana's entry. This raises the concern that the businesses which are started are equally productive as the GEs, but the effects are attenuated by a large number of zeros. This in turn makes it challenging to test the idea that the presence of microfinance induces some reluctant entrepreneurs to start new small businesses, which look negatively selected compared to the existing (ie, GE) businesses.

To overcome these challenges, we take advantage of the fact that Spandana entered different neighborhoods in a staggered fashion, over 13 months between April 2006 and April 2007. As a result, we observe businesses in different treatment areas that opened up at the exact same time (say, August of 2006): some opened before Spandana opened in its area (e.g., Spandana's branch may have opened in October); others opened after Spandana opened in its area (e.g., Spandana may have opened in May). Moreover, randomization was done at the matched pair level. Therefore, 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 for the youngest pre-Spandana firms. Comparing firms that opened in this common period but after Spandana arrived in a treatment area to those in the matched pair gives the combined selection treatment effects, again on the youngest firms. If the GEs are in fact advantageously selected on productivity (or alternatively, the non-GEs are negatively selected), then the treatment-control differences for the pre-Spandana firms should be larger than those for the post-Spandana firms created at the same point in calendar time.⁶ If,

⁵Among households that did not have a business in 2006 (ie, the non-GEs), 18% have opened a business by our third endline survey in 2012.

⁶The idea of comparing a set of business selected without microcredit access and that are then "shocked" with microcredit to identify a pure treatment effect on an advantageously selected sample is akin to a non-experimental version of the experiment in Beaman et al. (2015). In contrast to them, we lack a source of identification that allows us to estimate a pure selection effect (microfinance cannot be taken away from the businesses that selected in post microfinance entry). We are, however, able to follow the businesses over time to understand if and how the effects persist.

on the other hand, the GEs are simply benefiting from a head start, then within this “overlapping” sample we should see no differential treatment-control differences.

In fact, we find that, within this sample, the GEs demonstrate large positive treatment effects, while the non-GEs demonstrate “effects” (comprising treatment and selection effects) that are imprecise but negative. A back-of-the-envelope decomposition shows that, if the true treatment effect on the non-GE entrepreneurs (the “reluctant entrepreneurs”) 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.

Motivated by these large and persistent effect of early access to microcredit on the GE, we explore quantitatively how they could have arisen as a result of an intervention which provided only temporary differential exposure to microcredit for the treatment group. We use the data to structurally estimate a simple model of firm growth in the presence of credit constraints. The model allows for two different technologies, one with diminishing returns and one with constant returns, with a fixed cost required to adopt the better (ie, CRS) technology. There is also heterogeneity among the entrepreneurs—one group has access both technologies (we think of these as the GEs) while the other (the REs) can only access the diminishing returns technology.

This model generates a process in which the impact of temporary access to some additional credit cause a divergence among the GE firms—but not among the RE 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. Households with enough wealth, on the other hand, can afford the fixed cost of operating the constant returns technology and thereafter do not face diminishing returns. 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. (See Figure 9.) Thus the impacts of even a short-term headstart in microfinance access can persist because the extra wealth that households earned when they have access to microcredit can make them rich enough to continue to run the high technology even when microcredit is withdrawn.

The model does a good job in generating impacts very similar to what we observe in the data. We also show that 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 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, Angelucci et al. (2015); Banerjee et al. (2015); Meager (2019); ?), 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.⁷ 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 the heterogeneity is driven by selection, not experience or entrepreneur age. This also allows us to show that businesses started because of microfinance are indeed worse, holding age and experience fixed.

As we mention 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 two contemporary studies presenting evidence of non-convexities and poverty trap dynamics in different settings. Balboni et al. (2018) find evidence consistent with poverty traps among the ultra-poor in Bangladesh, following up on the evaluation of a large productive asset transfer analyzed in Bandiera et al. (2017). Kaboski et al. (2019) also investigate the possibility of non-convexities in 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.

These findings have a number of important implications for credit market policy. 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—then a more selective approach may be better. While we have no reason to question the fact that even the REs benefit from the loan (see the discussion of distributional effects below, 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. Second, it raises the issue of whether, from the point of view of growth, much bigger (and more selective) loans are desirable.⁸

⁷Additionally, Karaivanov and Yindok (2015) estimate a model which makes a distinction between “voluntary” and “involuntary” entrepreneurship using data from urban Thailand and examines heterogeneous responses to credit.

⁸La Porta and Shleifer (2008) make the case that most of the firms in the informal economy are marginal to the main story of growth. Related, Diao et al. (2016) 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.

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. (Hussam et al. (2018) also provide evidence that microentrepreneurs in India are aware of the marginal returns to additional capital of their own businesses, and those of their peers.) This has important implications in terms of designing a menu of credit contracts that can induce GEs and REs to select different options.⁹

The remainder of the paper is organized as follows. In section 2, we describe the experimental setup of Banerjee et al. (2015), the AP microfinance ordinance, and our 2012 survey. Section 3 presents results for the full sample. Section 4 presents results for a unique subsample which allows us to distinguish selection from other reasons GEs and REs could differ. In Section 5, we present our model and estimation results, and show how the parameters shed light on the poverty trap. Section 6 concludes.

2. DATA AND EXPERIMENTAL DESIGN

2.1. Setting and Previous Work. We build upon two existing rounds of panel data collected by Banerjee et al. (2015). As discussed in that paper, 104 neighborhoods in Hyderabad were randomized so that 52 received access to credit from Spandana, a large lender that was then moving into Hyderabad, starting in early 2006. The remaining neighborhoods only received access in mid-2008 after a round of data collection conducted in late 2007 - early 2008. A second round of data collection was conducted in mid-2010 to examine longer-term impacts of access to microfinance. Coincidentally, this second endline was completed just a few months before the microfinance landscape abruptly changed, as we discuss below. Figure 2 shows the timeline of the data collection as it relates to changes in microfinance access.

Banerjee et al. (2015) examined the effects of the intervention on outcomes measured in 2007-8 and in mid-2010.¹⁰ Key outcomes examined in that work include borrowing from various sources, consumption, business creation, and business income, as well as measures of human development outcomes such as education, health, and women’s empowerment.¹¹

⁹Related, Maitra et al. (2017) show that incentivized agents can identify productive and lower-risk borrowers in West Bengal.

¹⁰As described below, the survey instrument for this paper is based on that used in Banerjee et al. (2015) to facilitate comparisons across time, although new modules were added.

¹¹We 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. Borrowers, who were organized into joint liability 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. In fact, conversations with former borrowers in 2011 indicate that residents of Hyderabad viewed the lenders as exchangeable.

At the first (2007-8) endline, treated households did borrow more from microcredit institutions (though fewer than a third of treated households borrowed).¹² No significant difference was found on consumption, but there were significant positive impacts on investment in durables. Treated households started more businesses, and those whose businesses were already in existence before microcredit (ie, the GEs) invested more in those businesses. The average profits of these existing businesses increased, with particularly large gains at higher quantiles, while the median marginal new (ie, RE) business was both less profitable and less likely to have even one employee in treatment than in control areas.

At the second (mid-2010) endline, when microcredit was available both in treatment and control groups but treatment group households had the opportunity to borrow for a longer time, businesses in the treatment group had significantly more assets. But the average business was still small and not very profitable, though, once again, a tail of businesses appeared to experience gains from longer microfinance access. There was still no difference in average consumption. No effect was found on women's empowerment or human development outcomes either 2 or 4 years after the initial treatment.

These results hint at important heterogeneity. However, many unresolved issues remained. Since during the 2006-2010 period, treatment households always had access to microfinance, one question is whether the impacts seen, particularly those on business outcomes, are sustainable in the absence of continued access to new loans. Another question is whether newly created businesses will, given more time, catch up to the existing businesses, or whether they are on permanently different trajectories. These are among the questions we address in this paper.

2.2. Andhra Pradesh Microfinance Ordinance. The second round of endline data analyzed in Banerjee et al. (2015) was collected in mid-2010, only a few months before the Andhra Pradesh (AP) state government put forth a sweeping new regulation of the microfinance sector. 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. The government was worried about widespread over-borrowing by its citizens and alleged abuses by microfinance collection agents. The provisions of the Ordinance (promulgated as a law in December 2010) brought the activities of the MFIs in the state to a complete halt. Under the law, MFIs were not permitted to approach clients to seek repayment and were further barred from disbursing any new loans.¹³ In the months following the ordinance, almost 100% of microfinance borrowers in AP defaulted. Furthermore, Indian banks pulled back tremendously on their willingness to lend to any MFI across the country, and MFIs even outside of

¹²As discussed in Banerjee et al. (2015), other MFIs entered Hyderabad between 2006 and 2008, when the control group was treated. That the control group had access to microfinance before Spandana entered may make the initial treatment less powerful, but it does not invalidate the original experimental design.

¹³However, it was not illegal for borrowers to seek out their lenders to make payments.

Andhra Pradesh were forced to contract their lending activities, at least temporarily.¹⁴ In mid-2011, the Reserve Bank of India (RBI) issued new guidelines for the microfinance sector and established itself as the national regulator for the industry. In 2012, when we returned to collect an additional survey round, the environment for MFIs in the rest of India had improved in large part due to the RBI’s actions, but MFIs in AP were still not permitted to operate under state law and were unable to collect on their loans or issue new credit.

The respondents surveyed for the [Banerjee et al. \(2015\)](#) study experienced the direct consequences of the AP ordinance. Approximately one third of respondents reported having a loan outstanding at the time of the second endline survey in mid-2010, and close to 50% had taken at least one microloan from any lender between 2004 and 2010.¹⁵ 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 3.6, below, we provide evidence that will help disentangle those two effects (and strongly suggests that the effects are not driven by the windfall effects).

2.3. Follow-Up Data Collection. In mid-2012, we returned to the respondents of the 2010 survey round of [Banerjee et al. \(2015\)](#) and conducted a follow-up survey with 5,744 households located in 103 of the original 104 combined treatment and control neighborhoods.¹⁶ In addition to the outcomes analyzed in [Banerjee et al. \(2015\)](#), we added a module to capture the household’s worries, happiness, and time preferences, and retrospective questions about the household’s exposure to the AP crisis and desire to borrow from MFIs in the future. We also added survey questions about the respondent’s social network.¹⁷

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. 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.

¹⁴See [Breza and Kinnan \(2018\)](#) for an analysis of the impacts of the ordinance on lending outside of Andhra Pradesh.

¹⁵See Table 2, columns 3 and 4, respectively.

¹⁶One (treatment) area was dropped because it was used for piloting. It was crucial to pilot in an area where past waves of surveying had taken place since familiarity with surveyors significantly increases households’ willingness to respond accurately. All our results below control for strata dummies from the original strata assignment and therefore also omit the control area assigned to the same stratum.

¹⁷The network outcomes are analyzed more fully in a separate paper ([Banerjee et al., 2018b](#)), which focuses specifically on the impact of microfinance on social networks.

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 second endline (2010), 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.

3. REDUCED FORM RESULTS

We aim to use the empirical setting to explore the long-run, persistent impacts of microfinance. In the first part of our empirical analysis, we follow Banerjee et al. (2015) and investigate the intent to treat (ITT) comparisons between the initial treatment neighborhoods 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

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

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.¹⁸ For all specifications, standard errors are clustered at the area level.

While we are interested in tracking the average impacts of microfinance over the entire population, we are especially keen to understand the differential impacts for gung-ho entrepreneurs vs. other households. For these specifications, the regressions take the form

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

Here, we indicate that household i in area a 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 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 $\beta_1 + \beta_2$.

The following sections discuss results for intent-to-treat estimates of treatment effects on multiple sets of outcomes. For most, we present each set of results in a regression table with two panels: Panel A shows average treatment effects for each outcome variable (the first specification described above), while Panel B shows heterogeneous effects by entrepreneurial status (the second specification described above). We report the p-values

¹⁸Altogether, there were 52 strata, or pairs. Pairs were formed to minimize the sum across pairs A, B $(\text{area A avg loan balance} - \text{area B avg loan balance})^2 + (\text{area A per capita consumption} - \text{area B per capita consumption})^2$. Within each pair one neighborhood was randomly allocated to treatment.

of the total treatment effect $\beta_1 + \beta_2$ for the gung-ho entrepreneurs at the bottom of each table.

3.1. Effects on Microcredit Borrowing. We aim to identify the persistent, long-run impacts of microfinance two years after the withdrawal of microfinance from the entire state of Andhra Pradesh. Before doing so, it is important to first understand how the exposure to microfinance was affected by the initial treatment status. Over the course of the three survey rounds, we have collected a number of measures that capture the exposure to microfinance. Table 2 presents the treatment effects for a set of these measures.

A natural measure of exposure is the likelihood of ever borrowing from any MFI. Panels A and B contain regressions of indicators for past borrowing at different points in time on treatment status. In column 1 of panel A, the outcome is an indicator for having borrowed in the three years prior to endline 1 (in 2007/2008); since there was essentially no microfinance lending in Hyderabad prior to 2005 this is equivalent to an indicator for having ever borrowed. Treatment households were approximately 11 percentage points more likely to have ever borrowed than control households.¹⁹ Columns 2 and 3 measure the incidence of borrowing around the time of endline 2 (2010). Column 2 captures any borrowing from microfinance in the three years prior to endline 2, while Column 3 reports the effects of the initial treatment on having a loan outstanding in October 2010, immediately before the AP Ordinance and 3-11 months after endline 2 was administered.²⁰ There are no detectable differential impacts on the probability to have a current loan before or after the time of endline 2. This evidence suggests that by 2010, the control group had caught up to the treatment group in terms of current access to credit. However, the treatment group did get a head start, and they had a long time to try microfinance. In column 4, we consider an indicator for whether the household ever reported borrowing at any time in any survey round. This is the union of the outcomes from Columns 1-3 and a retrospective question asked at the time of endline 3 about loans outstanding in October, 2010 at the time of the crisis. While approximately 50% of the control group had ever borrowed before the AP ordinance, households in the treatment group were 4.4 percentage points more likely to have ever borrowed (a 9% increase).

The original treatment could have also affected households via the intensive margin, namely the number of loans taken over time, the number of MFIs from which the household borrowed, and the total amount of credit taken. Panels C and D focus on this intensive margin. All outcomes in these panels are snapshots at the time of endline 2.²¹ Here,

¹⁹This 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 endline 1, while the 11pp value includes loans already fully repaid.

²⁰This variable is measured using retrospective questions in endline 3.

²¹We would ideally also like to measure each household's total stock of microfinance taken between 2006 and October 2010 from all MFIs. However, this is infeasible because the value of loans taken and fully repaid between survey waves was not measured. However, the existence of such loans (though not the

Column 1 captures whether a household had an active loan at the time of endline 2. Columns 2 and 3 explore the number and the total value of the MFI loans outstanding at the time of the second endline survey. While the number of MFI loans²² is no different in treatment and control neighborhoods, the overall amount of credit is larger in treatment areas. The average treatment household reports Rs. 946 more borrowing than the average control household. This amounts to a 14% increase in credit over the control group. Because treatment group borrowers had earlier access to microfinance through Spandana, this effect captures the fact that most microlenders increase the loan size offered to clients over time.²³ In column 4, we report that treatment households are 50% more likely to have a Spandana loan than households in the control group, and column 5 shows that they had Rs. 1,132 more credit from Spandana in 2010.

We next ask whether the exposure effects vary between GEs and non-GEs. Panels B and D capture the heterogeneous treatment effects. On the extensive margin, we do not detect any significant differences in ever borrowing from a microfinance institution between gung-ho entrepreneurs and others (Panel B and cols 1 to 2 of Panel D). However, the point estimate for the differential impact on the total amount of MFI credit taken in 2 at endline 2, is large, although insignificant (Panel D, column 3). We do find a treatment effect on the amount borrowed from Spandana in 2010 of Rs. 800 for non-GEs. This treatment effect is twice as large for the GEs (Panel D, column 5).

Overall, households in the original treatment neighborhoods started borrowing earlier and were more likely to ever borrow from an MFI. They also had more credit outstanding before the AP crisis. Though we cannot measure the total value of loans ever taken from microfinance, the evidence also suggests that treatment households borrowed for longer (more loan cycles) and had a larger overall stock of microfinance credit. We also find some suggestive evidence that the gung-ho entrepreneurs were taking larger loans (but were equally likely to be borrowers) from microlenders a few months before the crisis.

Finally, we note that there is no single sufficient statistic that captures all of these effects. In the results that follow, we focus on the reduced form ITT treatment effects.

3.2. Total borrowing. We can also look at treatment effects on borrowing from non-microcredit sources, and borrowing overall. This is instructive for several reasons. First, households who are not constrained should simply refinance older, more expensive debt with newer, cheaper debt, with no net impact on total borrowing as long as they borrowed more from informal sources than was available from microfinance. Credit-constrained households, on the other hand, will increase total borrowing (Banerjee and Duflo, 2014).

amount) was measured, so we can construct a proxy for ever borrowing at any time, presented in Panels A and B.

²²This can be interpreted as the number of lenders since it is extremely rare for a borrower to have multiple loans from one lender at a point in time.

²³Increases of between Rs. 2,000 and Rs. 5,000 between loan cycles are common.

From a policy perspective, understanding effects on total borrowing are relevant: if new credit fully crowds out existing sources, the effect is not to increase total liquidity, but only to reduce the interest rate at which the marginal unit of capital is borrowed. Any crowd-in, especially among GEs, is also relevant for interpreting the magnitudes of the treatment effects on business inputs, revenues and profits. And, as we argue below, patterns of crowd-in or -out are informative about whether households are aware of their own “type.”

Table 3 reports treatment effects on households borrowing from sources other than microfinance. Panel A shows the average treatment effects. Averaging across GEs and others, total credit (col 1) and informal credit (col 2) increase, but not significantly. Informal loans taken for business does increase significantly.²⁴ As mentioned above, our survey also collected information on individuals’ social networks. The final two columns of table 3 examine these measures, since social networks are an important source of informal credit. Column 4 shows treatment effects on individuals’ overall network degree—the number of other households with whom they interact, or could interact if the need arose, across both financial interactions (eg, borrow or lend small amounts of money, cooking oil or kerosene) and non-financial interactions (eg, ask for advice, socialize together). Column 5 shows the measure examining only financial networks. For the full sample, overall network degree and financial network degree both fall significantly. This crowd-out of network relationships for the average borrower is further examined in Banerjee et al. (2018b).

Next, in Panel B, we examine whether these effects differ for GEs vs. others. The main treatment effect identifies the effects for the non-GEs. For these households, total borrowing, total informal borrowing, and total loans taken for business purposes show no significant changes. The differential effect for GEs, on the other hand, is large and significant for all of these outcomes. Summing the main and interaction effects, GEs in the treatment group had, on average, Rs. 12,400 more in outstanding informal debt relative to the control group (significant at the 5% level), and Rs. 12,000 more in outstanding informal debt used for business purposes (significant at the .1% level). The fact that increased access to microfinance increases demand for other sources of credit among the GEs is already at odds with a simple globally concave production function, and instead suggests that the GEs face a non-convexity in their choice set.²⁵

The final two columns of Panel B shed light on the differential effects on informal credit seen for GEs vs. non-GEs. Column 4 shows treatment effects on individuals’ overall network degree. The effect for non-GEs is negative, significant, and substantial in magnitude, representing roughly a 9% fall in total links. The differential effect for

²⁴To classify whether a loan was used for a business purpose, we code survey responses coding the primary purpose of the loan.

²⁵This argument assumes the increase in borrowing from other sources is purely demand driven. However, it is possible that the intervention may have also relaxed credit *supply* by either signaling borrower quality or because of the (real or perceived) junior status of microfinance debt. We discuss below a number of other pieces of evidence that point to a non-convexity as an explanation.

GEs, however, shows that for these households, there is no significant net change in social networks. Column 5 shows that, examining only financial networks, the pattern is similar. In short, non-GEs appear to experience shrinking social networks, while GEs' networks are maintained.

3.3. Business Outcomes. Table 4 reports treatment effects on outcomes related to household businesses. We find that the effects of microfinance on business creation described in Banerjee et al. (2015) persist even in the absence of ongoing microcredit: treatment households were 3.8% more likely to have a business, and own 0.056 more businesses on average, than control households (Panel A, columns 1 and 2). On average, businesses in the treatment group are larger as well. Households in the treatment group have over Rs. 1,500 more in business assets than households in the control group (column 4), and report 27.3 log points higher expenses and 31 log points higher revenues from their businesses than the control group (columns 5 and 6). Profits, reported in column 7 (in levels due to zero and negative values) are significantly higher on average, by just under 600 Rupees per month.²⁶

This finding—on average, microfinance improves business outcomes in a way that persists even once loans are no longer available—is novel (and contrasts with our own short term effects and that of the other studies), and is in itself important for policy and for future research examining the effect of microcredit. Even if policy makers or microfinance organizations could not distinguish between GE and non-GE applicants, these results are reassuring in that microfinance lending can durably improve business outcomes on average, at least in a context like Hyderabad.

While positive business effects are apparent for the treatment group as a whole, as Panel B shows, these results are driven almost entirely by effects on gung-ho entrepreneurs alone. GEs in the treatment group are 6.4% more likely to own a business and own, on average, 0.10 more businesses than those in control (columns 1 and 2). Their businesses are larger and more profitable, as well: GEs in treatment own over Rs. 3100 more in business assets (column 4) and report spending 83% more in business inputs (column 5), receiving 60.5 log points, or 104%, more in business revenue (column 6) and earning almost 1300 Rupees, or 28%, higher profits. As Panel A of Figure 1 shows, these treatment effects are also larger in magnitude than those measured at either endline 1 or endline 2.²⁷ In contrast, these

²⁶Recall that some of this increase in the business outcomes reported in columns 4 to 7 is driven by the extensive margin—treated households have more businesses. Below, when we focus on a sample entirely of business owners, we show that there are significant effects on the intensive margin as well—again, driven by the GEs.

²⁷The figure also shows that for several of the outcomes, the treatment effects appear to become smaller at endline 2 before expanding again at endline 3. We have one working hypothesis for why this pattern might emerge. If treated GE households spent the first few years post-microfinance access investing in the household business, they may have decided to take a small break and increase their consumption and/or spending on household durables at endline 2.

same outcomes for the non-GE households in the treatment group are no different than the outcomes for those in the control group.

One question that OLS results alone cannot answer is how the effects we document are distributed throughout the population. For instance, 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 3, Panel A plots the results of bootstrapped quantile regressions for business profits on treatment status for GEs. As this figure shows, 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. No portion of the distribution for non-GE households, on the other hand, experienced such results, as Panel B shows. Next we examine business assets: Figure 3 Panel C plots the quantile results for business assets for GEs. Here 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. Panel D shows that, as with profits, there is no effect on business scale anywhere in the distribution for the non-GEs, with the possible exception of at the very top (90 to 95th percentile): these may be true GEs who happened to start their businesses later. Similarly, there are surely some non-GEs who, due to wealth, easy access to credit, etc. happened to start their businesses before Spandana's entry. Misclassifying such firms will have the effect of making the two groups appear more similar and thus making it harder to reject the null that GEs and non-GEs have similar treatment effects.

3.4. Labor Demand. The potential for microfinance to affect local labor markets through its effect on firms' labor demand has been noted theoretically (Buera et al., 2017) and demonstrated empirically (Breza and Kinnan, 2018; Fink et al., 2018). Table 5 reports effects on employment in household firms, wage bills, and household labor supply. As Panel A shows, treatment households on average are more likely to have more than one and more than two workers (cols 1 and 2); have 0.21 more employees in their largest business (col 3); pay out Rs. 370 more in wages to employees each month, more than 100% of the control group mean; and work 2.75 more hours per week in their businesses (col 6) than do control group households. These results when taken together, show that treatment households on average increase their total labor demand, consistent with the effects observed for microfinance on labor markets in other contexts.

However, as Panel B reveals, there is significant heterogeneity in these treatment effects. Non-GEs in the treatment group have .174 more employees in their largest business (col 3) and pay out Rs. 275 more in wages than in the control group (col 4), but show no significant differences in their total or self-employment labor supply relative to the control

group. Gung-ho entrepreneurs, on the other hand, show multiple significant treatment effects, which are uniformly larger than the effects for non-GEs: GEs in treatment areas are 5.7pp more likely to own a business with multiple employees (column 1), 3.2pp more likely to own a business with over two employees (column 2), have .277 more employees in their largest business (col 3), and pay out Rs. 587 more in monthly wages to employees (col 4); work an additional 6.65 total hours per week (col 5), of which 5.827 hours are in self-employment (column 6). Thus, not only do GEs in treatment neighborhoods have larger businesses several years after the introduction of microcredit; they are also contributing more labor time to their businesses on a weekly basis.²⁸ Moreover, the fact that GEs' businesses increase their hiring suggests that the benefits to their businesses may spill over to non-GE households in the form of greater opportunities for employment.²⁹

3.5. Consumption. Table 6 shows intent-to-treat estimates for treatment effects on household expenditure. As Panel A, column 1 shows, we find no significant average effect of increased exposure to microfinance on monthly consumption per adult equivalent. Once again, this lack of a significant average treatment effect masks considerable heterogeneity, both between GEs and non-GEs and within each group of households. We find no significant average treatment effects on consumption for either GEs or non-GEs, as Panel B, column 1 shows. But as demonstrated in Panel E of Figure 3 (displaying the results of bootstrapped quantile regressions for per capita consumption for gung-ho entrepreneurs), among the GEs, more than half of the distribution of per capita consumption (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 (Figure 3, Panel F) do we find any significant positive (or negative) treatment effects: the effect is a fairly precise zero throughout the distribution.

Columns 3, 4, and 5 report results for annual household spending on durable goods, both in total and broken into spending on durables for business use and non-business use.³⁰ We find a marginally significant average treatment effect of Rs. 560 in increased total spending on durable goods (Panel A, column 3) and a highly significant, though small, average treatment effect of Rs. 24 in increased spending on durable goods for

²⁸As with the effects on business outcomes discussed above, some of the effects on labor demand are driven by the extensive margin—treated households are more likely to operate any business.

²⁹To the extent that this generates positive spillovers to non-GE households in treatment areas in the form of higher earnings/consumption, it makes the (already small and insignificant; see section 3.5) treatment effects on non-GEs an upper bound on the true effect.

³⁰Because of outliers in these distributions, we Winsorize data of reported spending on durables in each category at the 95th percentile of each distribution.

household businesses (Panel A, column 5). These results, as Panel B reveals, are driven entirely by gung-ho entrepreneurs. In the treatment group, GEs spent Rs. 1,937 more on durables and Rs. 61 more on business durables in the previous year than GEs in the control group, while non-GEs in treatment and control show no differences in either of these outcomes (columns 3 and 5). Moreover, GEs show a large and highly significant increase in spending on non-business durables: Rs. 1,540, or 18.9% of the mean for GEs in the control group (Panel B, column 4).

Consistent with [Banerjee et al. \(2015\)](#)'s results for the second endline survey (in 2010), we find no difference between treatment and control households - whether among GEs or non-GEs - in spending on festivals (column 6). But in contrast to that study, as column 2 shows, we also find no difference in spending on "temptation goods," goods that households in the baseline survey said that they would like to spend less on (alcohol, tobacco, betel leaves, gambling, and food consumed outside the home).³¹ Additionally, there is no difference between treatment and control households in monthly spending on education (column 7) and health (column 8).

One of the most disappointing features of the first wave of microfinance impact evaluations is the the lack of a positive effect on household consumption. [Banerjee et al. \(2015\)](#) and others do find an initial increase in durable consumption. However, they do not observe positive effects in overall consumption or in longer-run household durable consumption. Our results point to some optimism, at least when microfinance is directed toward gung-ho entrepreneurs. While in 2012, the GEs continued to invest their labor hours and capital in their businesses, we also observe that a sizable subset of the distribution does in fact enjoy a consumption increase, and that the average household is able to purchase more household durables. If the marginal returns to business capital are still high, then we might expect even larger consumption increases in the future. For these seasoned entrepreneurs, a high marginal value to an additional rupee of business investment may explain the absence of a short-run consumption effect.³²

On a less positive note, in Appendix Table 10, we consider whether access to microfinance in the past caused any differences in happiness and worries, as measured by responses to survey questions. We find that that treatment households are both more worried and less happy than control households. We are not able to statistically distinguish the impacts by GE status. These results are consistent with the results in [Banerjee,](#)

³¹The fact that we do not find an effect on temptation good spending may not be surprising in our setting where MFIs are no longer operating. One possible source of the initial [Banerjee et al. \(2015\)](#) temptation goods effect may have been individuals scaling back unnecessary consumption in order to make the weekly MFI loan repayment. When microfinance is no longer present, there is no need to come up with the weekly payment amount.

³²The fact that, at the very top of the distribution for the GEs, treatment effects on consumption appear (insignificantly) negative is suggestive evidence of the presence of high returns leading to earnings being re-invested. (See Panel E of Figure 3.)

Duflo, and Hornbeck (2018a), which show that business-owning households are willing to give up microfinance when the real price goes up slightly, despite large negative effects on their business: households clearly care about other things than profits, and business success is not welfare.

3.6. Threats to Validity.

Attrition. Given the long time horizon since the sampled households were first contacted, it is perhaps unsurprising that some individuals have attrited from the sample, i.e., could not be found. Appendix Table 11 shows that, relative to the first endline (2007), approximately 84% of households were located at endline 3 in 2012. 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 12 shows the results. With a few exceptions, the Lee bounds are informative, i.e. when the non-attrition-adjusted estimate is significantly different from zero, the bounds do not include zero.

Windfall Effects of the AP Ordinance. It is also important to understand the differential repercussions of the AP ordinance on the treatment and control groups. Note that the effects were twofold. First, all households—uniformly in treatment and control areas—lost access to future credit. Second, households with outstanding loans received an implicit write-off of the remaining principal and interest, i.e. a windfall equal to the amount they would otherwise have had to repay. Borrowers who had received a new loan just before the ordinance received a large windfall, while those who were close to fully repaying the loan and obtaining a new loan received a small windfall.³³ Again, this occurred in both treatment and control areas. We would like to interpret differences in the treatment versus the control group we find in this paper as coming through increased past exposure to microfinance and to nothing else. However, if the treatment-control differences are also driven by differential windfalls when microfinance was withdrawn, then the interpretation would be muddled.

As we describe above, in Section 3.1, there are no differences in borrowing on the extensive margin between treatment and control on the eve of the AP Ordinance. However, 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

³³In a companion paper (Banerjee et al., 2014) we consider the effects of the windfall on household consumption and investment.

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 2010.³⁴ 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 13 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 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 14 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

³⁴The 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.

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.³⁵ 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.

4. A TEST OF THE SELECTION MECHANISM: OVERLAPPING SAMPLE RESULTS

We have so far argued that the differences observed between the GEs and non-GEs reflect that fact that households differ in their underlying potential productivity as entrepreneurs, and that when microlenders lend to households who have not demonstrated entrepreneurial potential, they screen in those who are less well-positioned to benefit (at least in terms of marginal product of capital). However, other explanations could be driving our results—namely, the GE businesses are older, on average, and have more experience. It could be that, with time, the non-GEs would accumulate enough age/experience and would then look like the GEs. Moreover, the non-GEs comprise both consumption borrowers and reluctant entrepreneurs (REs), making it difficult to draw conclusions about the REs alone. Recall that only 20% of non-GE households (those who did not have a business in 2006) start new businesses by 2012.

To overcome these challenges, we use the fact that Spandana did not enter all treated neighborhoods at exactly the same time: branches opened in treatment areas between April 2006 and April 2007.³⁶ 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—Spandana may have opened first in the largest areas, those closest to its headquarters, etc. However,

³⁵Recall 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.

³⁶The timing of the first endline was such that no area was surveyed fewer than 12 months after Spandana entered.

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—in expectation, Spandana would have opened a branch there at the same time. We refer to the sample of businesses that opened during the time that Spandana was opening branches as the “overlapping sample.”

Figure 4 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 first endline. 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 , endline 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, benefit from a first-mover advantage, etc., then among this overlapping sample, those that opened pre-Spandana (because Spandana opened relatively late in their area) should have indistinguishable treatment effects from those that opened post-Spandana (because Spandana opened relatively early in their area), since both the pre- and post-Spandana businesses among this sample are the same in terms of age, experience, etc. If, on the other hand, microfinance induces businesses to enter that have lower returns than those who enter in the absence of microfinance, then the firms that opened pre-Spandana should have different (larger) long-term treatment effects than those that opened post-Spandana but at the same point in calendar time. We note that this sample is small—approximately 300 households who opened a business during the overlapping sample window—which will reduce our statistical power but, as we show below, the effects for the GEs (i.e., the pre-Spandana businesses within this sample) are strong enough to nonetheless be detectable.

Effects on business outcomes. Table 8 shows the results. Panel A shows the EL1 treatment effects on business outcomes for businesses opened pre-2006, before Spandana opened any branches in Hyderabad. These businesses are not in the overlapping sample but fit the definition of GEs. The treatment effects for this sample are large and positive, though in some cases imprecisely estimated; however, the index of business incomes is 0.071 standard deviations higher in treatment than control, significant at 10%. The standard errors partly reflect systematic heterogeneity in treatment effect. Next, Panel B shows EL1 treatment effects for businesses opened in the overlapping sample window, *before* Spandana had opened in their area (equivalent to B^T in Figure 4). The counterfactual is given by businesses opened in the same time frame in the control areas in the same

matched pairs (equivalent to B^C in Figure 4). The treatment effects for these businesses are similar to those for pre-2006 businesses (the older GEs) in Panel A. If anything, the effects are stronger: the effect on the index of business incomes is 0.148 standard deviations, significant at 5%.

Finally, Panel C shows the EL1 treatment effects 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 4). Again, the comparison is between businesses in treatment and control areas within a matched pair (equivalent to A^T vs. A^C in Figure 4). The treatment versus control differences for these businesses, while imprecisely estimated, are uniformly negative. The effect on the index of business incomes is -0.183; while this is not significantly different from zero, it is significantly different from the effect of plus 0.148 seen for the pre-Spandana (but same-aged) businesses.

Note that this comparison combines two separate effects – the treatment effect for the post-Spandana businesses and any selection effect. Given that the businesses are the same age as the businesses studied in Panel B, these results are strongly indicative of negative selection effects. In fact, we can use the results in Panels B and C to get a back-of-the-envelope estimate of how much worse the average RE is, due to selection. We use the index of business outcomes (column 6 of table 8) for this calculation. We make two simplifying assumptions. First, we assume that, for this sample, the true, causal treatment effect of microfinance for REs is zero, and second, that the treatment - control comparison for the post-Spandana sample is zero. (The point estimate is -0.183, but is not significantly different from zero; to be conservative, we assume it is zero.) Given that entry by microfinance induced 18% more businesses to start in treatment vs. control, this implies that those new businesses must be 0.67 standard deviations worse than the pre-Spandana businesses.

Figure 10 further investigates whether age or experience effects could be at play in generating the observed differences between the GE and RE samples. Businesses opened before 2006 are separated into quintiles of age (with quintile 1 being the oldest and 5 the newest), and treatment effects on the index of business outcomes are estimated separately for each quintile, using the corresponding quintile in control areas as the counterfactual. The dashed gray horizontal line shows the overall treatment effect for the pre-2006 businesses.³⁷ Panel A shows the EL1 results, which were summarized above. Panel B shows the results at EL2, which are quite similar: all age quintiles of pre-2006 businesses show treatment effects which are indistinguishable from each other. In Panel C, the EL3 effects are plotted. There is now more loss of precision, in part because some entrepreneurial households have now closed their businesses, but the qualitative pattern remains the same.

³⁷Note that, for some businesses, we know that they opened prior to 2006 (i.e., they meet the definition of GEs), but they did not report a specific year in which they opened. Such businesses are dropped from this figure since we cannot assign them an age quintile.

The fact that—among businesses of the same age—those opened pre-Spandana show significant, positive treatment effects while treatment-control comparisons for those opened post-Spandana show insignificant, negative coefficients; and the fact that, as Figure 10 shows, there is no systematic tendency of older businesses to have different treatment effects than newer businesses among the pre-2006 sample, both buttress our interpretation that the differential effects we observe for GEs vs. others are due to selection rather than age or experience.

Effects on total borrowing. Table 9 shows how treatment assignment affects households' demand for overall borrowing at the first endline, within the overlapping sample. We again allow the treatment effect to differ for those who entered entrepreneurship before MFI entry (ie, the GEs) vs. those who entered later (who are a mix of “reluctant entrepreneurs”, or REs, and some GE who happened to have entered later). We report effects on total borrowing (from any source, formal or informal); total borrowing which the respondent states was used for a business purpose³⁸; total informal borrowing (from family, friends, neighbors, and business associates such as suppliers and customers); and informal borrowing for business purposes. Column 1 shows that the REs' demand for overall borrowing actually appears to fall (though this effect is not significant due to the small sample size). For GEs, the point estimate suggests a large positive effect (again, not significant). In column 2 we focus on borrowing for business purposes. Again, the treatment effect is (insignificantly) negative for the non GEs. The differential effect for the GEs is large—approximately RS. 43,000—and, despite the small sample size, significant at the 10% level. Column 3 looks at total borrowing from informal sources. As with total borrowing, REs' demand for overall informal borrowing appears to fall (though this effect is not significant), while for GEs, the point estimate suggests a large (but insignificant) positive effect. In column 4, we examine informal loans taken for business purposes. For non GEs, there is a negative (insignificant) effect, while the differential effect for GEs is large (approximately Rs. 12,500) and significant at the 10% level.

These results have two important implications. First, for the gung-ho entrepreneurs, access to microcredit has a crowd-in effect, facilitating additional borrowing from informal sources as well. The opposite appears to be true for the REs: their total and informal borrowing does not increase and even appears to fall. These results also show 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. (We will argue below that access to microfinance increases their demand for non-microfinance credit due to a nonconvexity in the production function that GEs—but not REs—can access.) The REs, on the other hand, do not respond to

³⁸Households were asked, for all outstanding loans “For what purposes did you actually use the loan?” Purposes relating business investment and business expenses were classified as business purposes.

greater microcredit access by increasing their borrowing from other sources—if anything, other sources of borrowing appear to fall. This is consistent with the evidence from other settings that households are aware of their potential returns to capital (Beaman et al., 2015; Hussam et al., 2018).

5. MODEL AND ESTIMATION

Our reduced form analyses of the full and overlapping samples suggest that relaxing credit constraints through microfinance has large, persistent, and even divergent effects for the gung-ho entrepreneurs (GEs), and that the GEs appear to be advantageously selected relative to the reluctant entrepreneurs (REs) who entered in response to microcredit. We next present a simple framework to explore whether microfinance might unlock a poverty trap for a subset of GEs. Importantly, we are able to investigate whether the underlying production frontier available to the GEs exhibits non-convexities, which in turn, could lead to poverty trap dynamics.

We first write down a simple dynamic household optimization problem with borrowing limits. Our empirical treatment of the model then proceeds in four steps: A) we use the first endline 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.

5.1. Dynamic Household Optimization Problem. Households maximize the discounted sum of the utility from consumption

$$(5.1) \quad U(c_t)_{t=0}^{\infty} = \sum_{t=0}^{\infty} u(c_t)$$

subject to wealth and borrowing constraints, introduced below.

Production Function. We next consider the production frontier for the *Gung Ho* (GEs) entrepreneurs. We assume that type is persistent and that, as suggested by the evidence in section 4, individuals know their type. GEs have access to two distinct production technologies that turn inputs K_t into revenues Y_t . K_t measures the rupee value of total inputs used in production. The *Low* technology has revenues equal to:

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

We posit that this technology exhibits decreasing returns to scale with $\alpha < 1$.³⁹ Denote the optimal scale of this technology as K_L^* . The GEs also have access to a *High* technology.

³⁹In the estimation, we do not make this restriction and allow for any positive value of α .

This technology requires a minimum investment \underline{K} , but comes with a higher marginal product past the minimum scale. We further assume that this technology has constant, rather than decreasing, returns to scale.⁴⁰ The *High* technology's revenues are given by:

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

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.

Borrowing Constraint. We consider two regimes, $\tau \in \{1, 2\}$. (As explained below, the regime that is operative for a given household depends on both time and treatment status.) 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 finding of crowd-in for the GEs, in this regime they can borrow from microfinance and from informal lenders. This “social network” 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 project proceeds (including savings), so we abstract away from both distressed and strategic 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. The (gross) interest rate on this borrowing is R . All loans must be repaid at the end of each period.

Wealth Dynamics. We assume that individuals are heterogeneous in starting wealth W_0 . We also assume that all productive decisions are separable across time. Individuals enter each period t with wealth W_{t-1} and decide how much to invest. Any wealth not invested in the productive business can be saved within the period for a gross return ρ . In our context, many households save “under the mattress” or in a bank account where $\rho \ll R$. The entrepreneur also chooses whether to operate the high technology inside the business, denoted by the indicator D_t^H .

Given the production frontier and the borrowing costs, the GE has total end of period financial profits from running the business equal to:

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

⁴⁰The *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.

where δ is the depreciation rate for total capital. Given that the GE also saves an amount S , the total cash on hand at the end of the period is:

$$\pi(K_t) + \rho S_t$$

where $S_t = W_{t-1} + B - K_t$.

Consumption. Finally, at the end of each period households choose how much to consume, c_t , out of their cash on hand. Any cash on hand that is not consumed is passed on to the subsequent period: $W_t = \pi(K_t) + \rho S_t - c_t$.

At the end of the period, profits are realized, capital net of depreciation (δ) is liquidated, any loans are repaid, and the household chooses how to divide the profits between consumption c_t and future wealth, W_t .

Full Utility Maximization Problem. The utility maximization problem in its recursive form is:

$$V(W|\theta, \tau, g) = \max_{c, W', K, B, D^H} u(c) + \beta E(V(W')|\theta, \tau, g)$$

$$W' + c = D^H Y_H(K) + (1 - D^H) Y_L(K) + (1 - \delta)K + \rho S - RB$$

$$0 \leq B \leq \bar{b}_\tau$$

$$W' \geq 0$$

$$K \leq W + B$$

$$S = W + B - K$$

5.2. Production Function Estimation. We estimate four production parameters of the model (A_L, α, A_H, K). 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 first wave of survey data (EL1). 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 wealth, particularly non-financial wealth such as land. To make progress without a direct measure of baseline wealth, we make the assumption that the treatment effects in capital satisfy a monotonicity assumption – that

is, individuals at any given percentile of the treatment capital distribution are directly comparable to those in the control capital distribution.

Figure 5 illustrates where identification of the production function parameters comes 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 first stage treatment-versus-control comparison by quantile – how total capital changes across the distribution with access to microcredit. Using total capital as the first stage 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 CDFs hint at the idea that the treatment effects might be heterogeneous in starting wealth. Panel A suggests that the log impact on total capital is largest for values of log capital between roughly 5 and 9, which the effects on revenues are apparent for values of log capital between roughly 5 and 10. The fact that the effects decay less quickly in revenues than in capital is direct, albeit suggestive, evidence, of increasing returns. Moreover, for the lower quantiles, it appears that access to microfinance does not affect capital decisions (or revenues) at all. This is consistent with the amount of microfinance not being large enough to move some households out of the low production technology.

In the estimation we assume that the capital choice of each business satisfies a monotonicity assumption; namely 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. Under this assumption, comparing capital levels in treatment versus control at any quantile gives us the treatment effect on capital *for that quantile*. Given the exogenous shock in capital, *at each quantile* we can then construct predicted revenues for each individual in the sample according to

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

The predicted revenues at each value of $\{A_L, \alpha, A_H, \underline{K}\}$ can then be compared to the true revenues 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 within treatment group by quantile of capital.⁴¹ Note that we only assume monotonicity in capital, not in revenues. Moreover, given that some firms produce 0 revenues, we allow stochastic business closure within period to match the empirical rates in the data. For each bin in the capital distribution, assuming the average capital level for that bin, we can calculate the predicted revenues under each value of the parameter vector and compare it to the average empirical revenues in the same bin. In practice, we select the parameters $\{A_L, \alpha, A_H, \underline{K}\}$ that minimize

⁴¹Specifically, we create 15 quantile bins.

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 bootstrapped⁴² 95% confidence intervals in brackets.⁴³

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

The estimated fixed cost, Rs. 7,900 corresponds to roughly the median of the estimated baseline wealth distribution. In Figure 6, 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 7 shows the “gross profits” for the three possible technologies available to a household: the low technology, the high technology, and the savings technology. For the low and high technologies, the gross profits are equal to revenues plus proceeds of selling back undepreciated capital. The gross profits of the savings technology are equal to the 45-degree line, because it is simply a storage technology with a gross return of 1. 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. (Note that these 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.)

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

⁴²For the bootstrapping procedure, we draw with replacement from the full dataset of Pre-Spandana, young businesses.

⁴³When, in addition, the curvature parameter on the high technology is estimated, we obtain an estimate of unity, ie constant returns to scale. However, this is computationally intensive, so in the main estimation we constrain the high technology to exhibit CRTS).

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.).⁴⁴

5.3. Calibrating and Solving for the Policy Functions. Next, we solve the dynamic program given the production parameters from step 1. 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/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. 12,000.⁴⁵ This reflects the first stage on total credit (Table 3, column 3). 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 \(2012\)](#).⁴⁶ Below, we provide a sensitivity analysis to these preference parameter choices.

5.4. 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 intervention spans seven years (2006-2012, inclusive). Given a baseline year of 2006, we simulate the model 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

⁴⁴The 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.

⁴⁵Recall that $\bar{b}_1 = 0$.

⁴⁶Our 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 individuals take loans at a 25% interest rate to fund consumption.

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.⁴⁷ 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 .⁴⁸ 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 $\tau = 2$. 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 EL3 data. The model replicates the qualitative patterns observed in the EL3 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

⁴⁷Specifically, we discretize wealth in bins of Rs. 300.

⁴⁸In our core simulate-forward exercise we assume that $\nu=500$. We explore the sensitivity of the results to changes in ν in Appendix table 15.

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 15, 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 EL3 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 at EL3 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 iid noise distribution, the treatment effects remain large (see Panel B, Appendix table 15).

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 8 shows the implied treatment effects on expected⁴⁹ 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 3, 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 9 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

⁴⁹The expectation is taken with respect to the draws from the noise distribution; see Section 5.4.

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.

6. 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 third endline), 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 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 when they can only access 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. This explains why the treatment effects persist—and even grow—after microcredit is removed from both treatment and control areas.

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.⁵⁰ We find evidence for the “crowdout” hypothesis among the reluctant entrepreneurs: access to formal finance reduces these households’ takeup 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.

⁵⁰Karlan and Zinman (2018) find evidence of crowd-in across different (formal) microlenders in Mexico. However, to our knowledge we are among the first to examine the relationship between microcredit and informal credit access.

REFERENCES

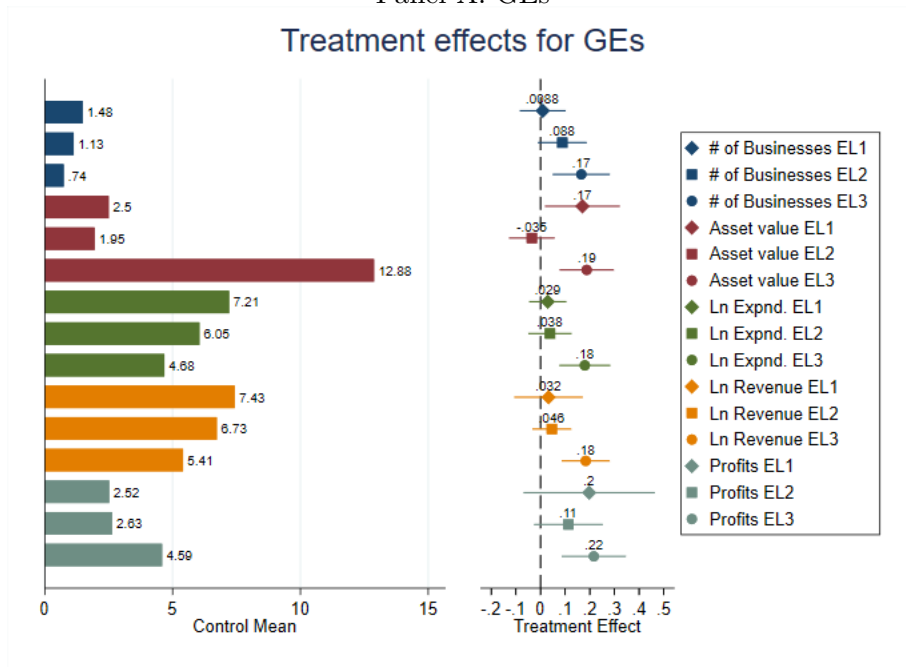
- AHGION, P. AND P. BOLTON (1997): "A Theory of Trickle-Down Growth and Development," *Review of Economic Studies*, 64, 151–172.
- ANDREONI, J. AND C. SPRENGER (2012): "Estimating time preferences from convex budgets," *American Economic Review*, 102, 3333–56.
- ANGELUCCI, M., D. KARLAN, AND J. ZINMAN (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 (2018): "Why do People Stay Poor?" Tech. rep.
- BANDIERA, O., R. BURGESS, N. DAS, S. GULESCI, I. RASUL, AND M. SULAIMAN (2017): "Labor markets and poverty in village economies," *The Quarterly Journal of Economics*, 132, 811–870.
- BANERJEE, A., E. BREZA, E. DUFLO, C. KINNAN, AND K. PRATHAP (2014): "Microfinance as commitment savings: Evidence from the AP crisis aftermath," Tech. rep., Working Paper.
- 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., E. DUFLO, AND R. HORNBECK (2018a): "How much do existing borrowers value Microfinance? Evidence from an experiment on bundling Microcredit and insurance," *Economica*, 85, 671–700.
- BANERJEE, A. V., E. BREZA, A. G. CHANDRASEKHAR, E. DUFLO, M. O. JACKSON, AND C. KINNAN (2018b): "Changes in social network structure in response to exposure to formal credit markets," *w*.
- BANERJEE, A. V. AND E. DUFLO (2005): "Growth theory through the lens of development economics," *Handbook of economic growth*, 1, 473–552.
- (2014): "Do firms want to borrow more? Testing credit constraints using a directed lending program," *Review of Economic Studies*, 81, 572–607.
- BANERJEE, A. V. AND A. F. NEWMAN (1993): "Occupational Choice and the Process of Development," *Journal of Political Economy*, 101, 274–298.

- BEAMAN, L., D. KARLAN, B. THUYSBAERT, AND C. UDRY (2015): "Selection into Credit Markets: Evidence from Agriculture in Mali," Tech. rep.
- BREZA, E. AND C. KINNAN (2018): "Measuring the equilibrium impacts of credit: Evidence from the Indian microfinance crisis," Tech. rep., National Bureau of Economic Research.
- BUERA, F. J., J. P. KABOSKI, AND Y. SHIN (2011): "Finance and development: A tale of two sectors," *American economic review*, 101, 1964–2002.
- (2017): "The Macroeconomics of Microfinance," .
- 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.
- DIAO, X., J. KWEKA, AND M. McMILLAN (2016): "Economic transformation in Africa from the bottom up: Evidence from Tanzania," Tech. rep., National Bureau of Economic Research.
- FINK, G., B. K. JACK, AND F. MASIYE (2018): "Seasonal Liquidity, Rural Labor Markets and Agricultural Production," Tech. rep., National Bureau of Economic Research.
- 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. ROTH (2018): "Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design In The Field," Tech. rep.
- KABOSKI, J., M. LIPSCOMB, V. MIDRIGAN, AND C. PELNIK (2019): "How Important are Nonconvexities for Savings and Development? Experimental Evidence from Uganda," Tech. rep.
- KABOSKI, J. AND R. TOWNSEND (2011): "A structural evaluation of a large-scale quasi-experimental microfinance initiative," *Econometrica*, 79, 1357–1406.
- KARAIVANOV, A. AND T. YINDOK (2015): "Involuntary Entrepreneurship: Evidence from Thai Urban Data," SFU Working Paper.
- 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.
- 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.
- TAROZZI, A., J. DESAI, AND K. JOHNSON (2015): "The Impacts of Microcredit: Evidence from Ethiopia," *American Economic Journal: Applied Economics*, 7, 54–89.

FIGURES

Panel A: GEs



Panel B: non-GEs




FIGURE 1. Treatment effects over time

Note: The left panel shows control means and the right panel shows treatment effects with 95% confidence intervals. The control means of assets and profits are in 1000s of rupees, while the treatment effects are in standard deviations; the # business variable is in units of businesses (both control mean and treatment effect), and the business expenditure and revenue outcomes are in log rupees (both control mean and treatment effect).




FIGURE 2. Timeline of Survey Activities and Microfinance Crisis

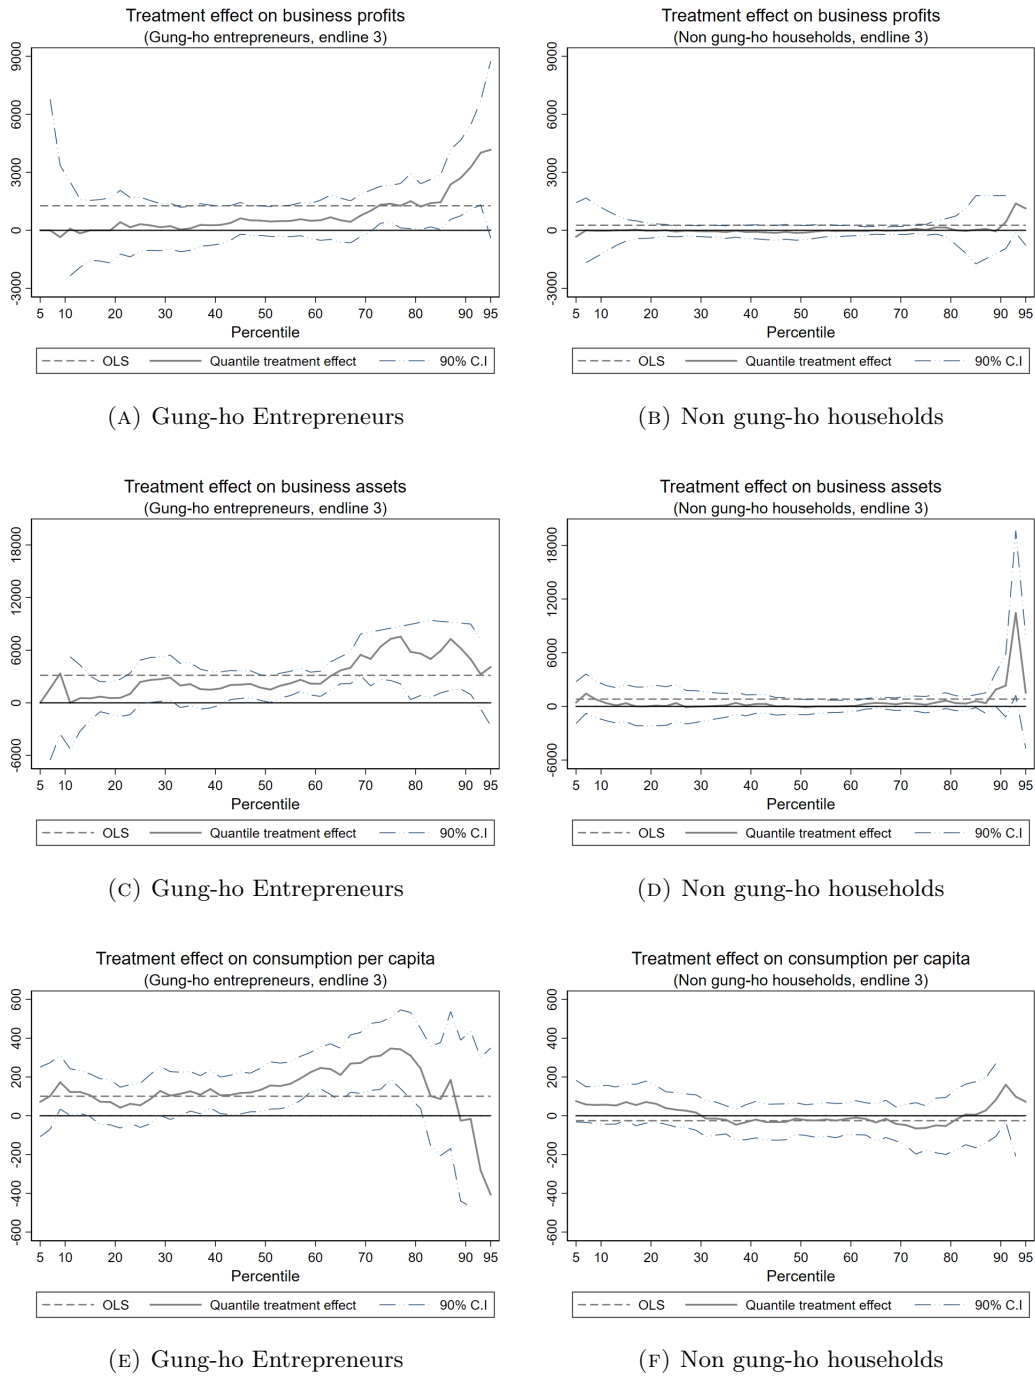


FIGURE 3. Quantile treatment effects

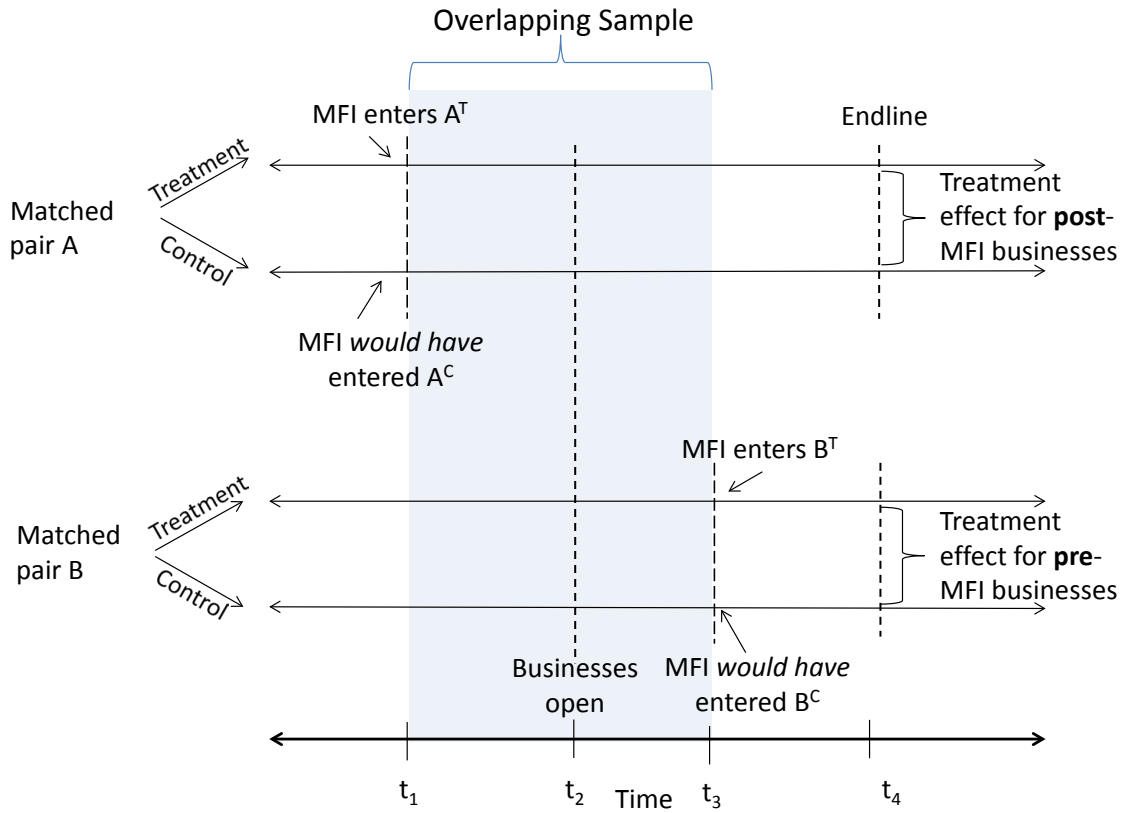


FIGURE 4. Overlapping sample identification




FIGURE 5. Distributions (CDFs) of Total Capital and Revenues by Treatment Status (EL1, overlapping sample only)




FIGURE 6. Production functions




FIGURE 7. Gross profit functions

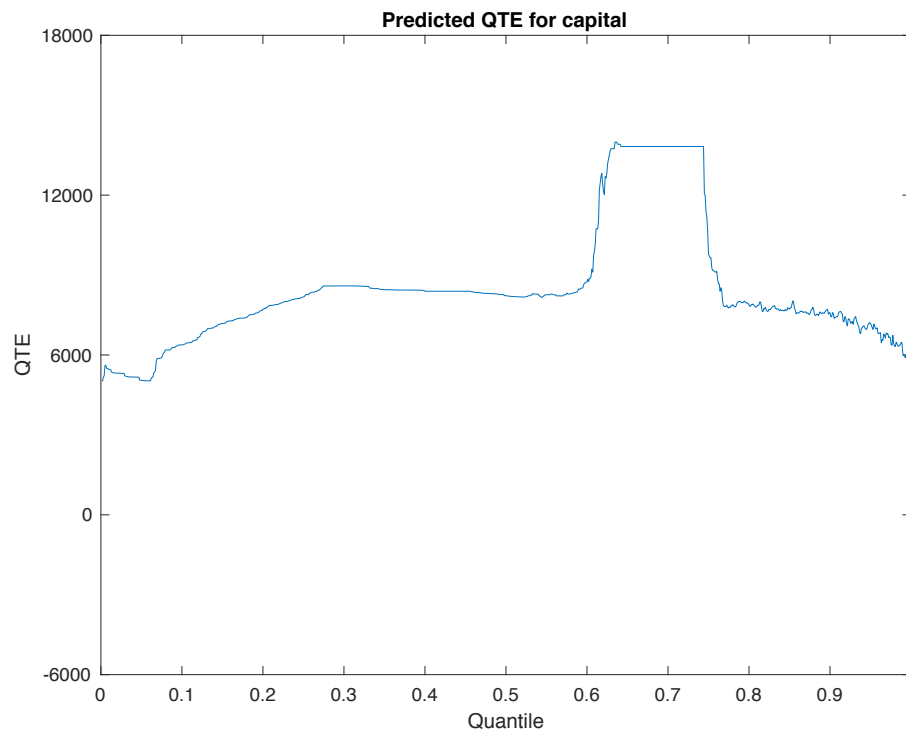


FIGURE 8. Model-implied quantile treatment effects for EL3 capital

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 9. Wealth transitions with and without credit access

TABLES

TABLE 1. Endline 3 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 (endline 2)</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 endline 2 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. Exposure to microfinance

	(1)	(2)	(3)	(4)	
	Borrowed from MFI in last 3 years (EL1 1)	Borrowed from MFI in last 3 years (EL2 2)	Outstanding MFI loan in 10/10 (EL3)	Borrowed from MFI between 2004 and 2010	
Panel A: Cumulative exposure to microcredit					
Treatment	0.109*** (0.022)	0.032 (0.022)	-0.009 (0.018)	0.044* (0.024)	
Control Mean	0.256	0.420	0.202	0.498	
Control Std. Dev.	0.436	0.494	0.402	0.500	
Observations	6804	6128	5745	5467	
Panel B: Cumulative exposure to microcredit by entrepreneurial status					
Gung-ho entrepreneur (GE)	0.163*** (0.023)	0.112*** (0.022)	0.035** (0.016)	0.110*** (0.022)	
Treatment	0.109*** (0.021)	0.029 (0.023)	-0.012 (0.019)	0.036 (0.026)	
Treatment × GE	-0.002 (0.030)	0.005 (0.030)	0.009 (0.025)	0.020 (0.032)	
Treatment + Treat × GE	0.107	0.034	-0.003	0.057	
P(Treat + Treat × GE ≠ 0)	0.001	0.312	0.899	0.091	
Control Mean (Non-GEs)	0.206	0.385	0.190	0.463	
Control Std. Dev. (Non-GEs)	0.404	0.487	0.392	0.499	
Observations	6804	6128	5745	5467	
	(1)	(2)	(3)	(4)	(5)
	Any MFI loan	Number of MFI loans	Total MFI loan amount	Any Spandana loan	Total Spandana amount
Panel C: Microcredit exposure as of endline 2					
Treatment	0.008 (0.020)	0.026 (0.038)	946.417** (474.365)	0.061*** (0.014)	1132.643*** (257.510)
Control Mean	0.332	0.530	6670.434	0.112	1897.522
Control Std. Dev.	0.471	0.937	13627.432	0.316	6768.526
Observations	6143	6143	6143	6143	6143
Panel D: Microcredit exposure as of endline 2 by entrepreneurial status					
Gung-ho entrepreneur (GE)	0.093*** (0.020)	0.173*** (0.049)	2557.957*** (671.712)	0.052*** (0.018)	798.113** (388.901)
Treatment	0.003 (0.021)	0.000 (0.038)	677.234 (508.180)	0.050*** (0.014)	800.099*** (267.354)
Treatment × GE	0.013 (0.031)	0.075 (0.073)	754.962 (929.289)	0.034 (0.024)	1036.985** (504.799)
Treat + Treat × GE	0.016	0.075	1432.197	0.083	1837.084
P(Treat + Treat × GE ≠ 0)	0.617	0.299	0.102	0.001	0.000
Control Mean (Non-GEs)	0.302	0.472	5812.723	0.096	1629.648
Control Std. Dev. (Non-GEs)	0.459	0.878	12661.459	0.294	6782.720
Observations	6143	6143	6143	6143	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. In Panels A and B, the column 1 outcome is an indicator for ever borrowing in the 3 years before endline 1 (in 2007/2008). The column 2 outcome is any borrowing from microfinance in the 3 years before endline 2. Column 3 reports the effects of the initial treatment on having a loan outstanding in October 2010, reported in endline 3. In column 4, the outcome is an indicator for whether the household ever reported borrowing at any time in any survey round. In Panels C and D, all outcomes are measured as of endline 2. The col 1 outcome is whether a household had an active loan EL 2. The outcome in cols 2 and 3 are the number and value of MFI loans outstanding; outcomes in cols 4 and 5 are the number and value of Spandana loans outstanding.

TABLE 3. Total borrowing (endline 3)

	(1)	(2)	(3)	(4)	(5)
	All loans	All loans (informal)	Loans for bus. (informal)	Network Degree	Net degree (financial)
Panel A: Treatment effects					
Treatment	1600.639 (5104.127)	2668.157 (3545.218)	4731.785*** (1500.096)	-0.369*** (0.134)	-0.307*** (0.090)
Control Mean	79283.288	57151.686	15062.722	5.948	4.372
Control Std. Dev.	1.56e+05	1.13e+05	43892.196	3.722	2.603
Observations	5744	5744	5744	5492	5492
Panel B: Treatment effects by entrepreneurial status					
Gung-ho entrepreneur (GE)	11801.721* (6119.067)	3647.067 (5833.084)	3798.873** (1766.060)	0.195 (0.163)	0.145 (0.106)
Treatment	-2808.930 (5662.034)	-1683.957 (4226.917)	1453.121 (1853.462)	-0.492*** (0.136)	-0.382*** (0.094)
Treatment \times GE	14174.488 (9651.940)	14085.007* (7387.176)	10598.425** (4289.273)	0.394* (0.228)	0.238 (0.144)
Treatment + Treat \times GE	11365.558	12401.050	12051.547	-0.098	-0.144
P(Treat + Treat \times GE \neq 0)	0.194	0.046	0.001	0.671	0.316
Control Mean (Non-GEs)	74493.933	55097.667	13677.772	5.903	4.328
Control Std. Dev. (Non-GEs)	1.56e+05	1.15e+05	41486.339	3.655	2.522
Observations	5744	5744	5744	5492	5492

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. 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. Network degree is based on questions regarding who the respondent would approach/be approached by for different types of assistance or social situations. Financial network degree examines only questions regarding financial transactions. See text for details.

TABLE 4. Business outcomes (endline 3)

	(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: Treatment effects							
Treatment	0.038*	0.056*	0.006	1565.222***	0.273*	0.311**	576.774***
	(0.020)	(0.031)	(0.005)	(426.789)	(0.140)	(0.157)	(179.375)
Control Mean	0.307	0.371	0.032	6680.551	2.293	2.637	2066.436
Control Std. Dev.	0.461	0.613	0.176	20448.064	3.776	4.180	6039.441
Observations	5744	5744	5744	5744	5724	5589	5580
Panel B: Treatment effects by entrepreneurial status							
Gung-ho entrepreneur (GE)	0.422***	0.525***	0.025***	8906.264***	3.361***	3.892***	3493.457***
	(0.020)	(0.026)	(0.008)	(973.087)	(0.174)	(0.180)	(350.655)
Treatment	0.024	0.031	0.009	816.198	0.101	0.126	263.906
	(0.018)	(0.024)	(0.006)	(526.966)	(0.114)	(0.128)	(168.567)
Treatment × GE	0.040	0.076**	-0.011	2325.597	0.503**	0.593**	1004.523**
	(0.028)	(0.035)	(0.013)	(1483.448)	(0.221)	(0.231)	(501.565)
Treatment + Treat × GE	0.064	0.107	-0.002	3141.795	0.605	0.719	1268.429
P(Treat + Treat × GE ≠ 0)	0.008	0.008	0.849	0.011	0.003	0.001	0.004
Control Mean (Non-GEs)	0.177	0.208	0.025	3974.639	1.258	1.452	988.890
Control Std. Dev. (Non-GEs)	0.382	0.483	0.155	17568.209	2.982	3.318	4065.047
Observations	5744	5744	5744	5744	5724	5589	5580

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 5. Labor market outcomes (endline 3)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	More than 1 worker in any business	More than 2 worker in any business	Workers in largest business	Total wages paid	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)	373.747*** (133.018)	2.170 (1.661)	2.752** (1.159)	0.008** (0.004)
Control Mean	0.102	0.035	0.507	348.367	87.490	15.400	0.027
Control Std. Dev.	0.303	0.184	1.292	4700.427	56.528	30.304	0.161
Observations	5738	5738	5738	5736	5744	5744	5744
Panel B: Treatment effects by entrepreneurial status							
Gung-ho entrepreneur (GE)	0.174*** (0.018)	0.055*** (0.009)	0.765*** (0.071)	488.639* (266.816)	4.798** (2.107)	23.537*** (1.587)	0.019** (0.008)
Treatment	0.017 (0.012)	0.009 (0.006)	0.174** (0.076)	275.264** (118.604)	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)	311.864 (368.366)	6.501* (3.321)	4.569** (1.962)	0.006 (0.012)
Treatment + Treat × GE	0.057	0.032	0.277	587.127	6.651	5.827	0.012
P(Treat + Treat × GE ≠ 0)	0.025	0.010	0.060	0.093	0.017	0.004	0.228
Control Mean (Non-GEs)	0.049	0.019	0.279	197.888	86.111	8.175	0.021
Control Std. Dev. (Non-GEs)	0.215	0.135	0.865	2496.403	55.490	22.456	0.144
Observations	5738	5738	5738	5736	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 6. Reduced form: consumption (endline 3)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Per capita consumption (monthly)	Temptation goods (monthly)	All durables	Non-business durables	Business durables	Festivals (annual)	Education (annual)	Health (annual)
Panel A: Treatment effects								
Treatment	-11.583 (58.735)	-19.609 (32.860)	559.362* (283.691)	351.696 (239.737)	23.900*** (8.242)	388.891 (304.751)	86.911 (88.338)	-45.121 (99.606)
Control Mean	2791.712	619.036	9264.343	8482.853	88.575	8932.731	1724.557	1775.509
Control Std. Dev.	1964.732	971.999	15748.713	14264.700	351.011	10950.345	3105.393	4107.905
Observations	5738	5739	5744	5744	5744	5737	5738	5738
Panel B: Treatment effects by entrepreneurial status								
Gung-ho entrepreneur (GE)	106.282 (77.229)	11.260 (29.245)	20.574 (658.347)	-513.234 (563.800)	92.895*** (16.248)	512.767 (551.631)	307.695** (153.525)	91.821 (177.110)
Treatment	-48.970 (61.965)	-21.350 (37.511)	-54.105 (385.449)	-175.322 (323.643)	6.847 (9.196)	351.483 (372.316)	111.138 (108.004)	-77.592 (122.781)
Treatment × GE	120.166 (107.356)	5.530 (48.037)	1991.273** (850.072)	1716.980** (725.416)	54.254** (26.063)	114.824 (663.376)	-82.426 (211.525)	104.366 (239.184)
Treatment + Treat × GE	71.196	-15.821	1937.168	1541.658	61.101	466.308	28.712	26.775
P(Treat + Treat × GE ≠ 0)	0.491	0.721	0.004	0.007	0.007	0.394	0.869	0.891
Control Mean (Non-GEs)	2759.209	614.207	9266.641	8635.084	60.866	8781.591	1631.834	1736.854
Control Std. Dev. (Non-GEs)	1886.292	986.241	15716.467	14350.628	293.414	9151.870	2663.565	4262.700
Observations	5738	5739	5744	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 7. Lending balance in October 2010 (pre-ordinance)

	(1)	(2)	(3)
	MFI Loan	Installments	Large windfall
Treatment	-0.012 (0.019)	0.039 (0.024)	0.004 (0.005)
Gung-ho entrepreneur (GE)	0.035** (0.016)	0.002 (0.029)	0.004 (0.006)
Treatment × GE	0.009 (0.025)	-0.016 (0.037)	0.004 (0.010)
Control Mean	0.202	0.386	0.027
Control Std. Dev.	0.402	0.297	0.161
Observations	5745	1095	5745

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.

TABLE 8. Business results, overlapping sample (endline 1)

	(1)	(2)	(3)	(4)	(5)	(6)
	Workers in largest business	Assets (stock)	Log expenses	Log revenue	Profit	Index of business variables
Panel A: Entered entrepreneurship pre-2006						
Treatment	0.024 (0.097)	391.360 (406.039)	0.316 (0.194)	0.434 (0.280)	2220.869** (946.406)	0.071* (0.037)
Control Mean	0.452	2614.149	7.339	7.716	2996.170	0.012
Control Std. Dev.	1.828	4873.170	2.844	3.127	14984.800	0.567
Observations	1305	1184	1273	1232	1232	1305
Panel B: 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)	0.148** (0.066)
Control Mean	0.033	2100.623	6.758	7.093	1164.737	-0.106
Control Std. Dev.	0.181	3961.748	2.853	2.959	6351.331	0.321
Observations	133	119	130	128	128	133
Panel C: 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)	-0.183 (0.112)
Control Mean	0.242	2539.005	6.377	6.719	1785.719	-0.021
Control Std. Dev.	1.110	4850.283	3.402	3.591	6797.191	0.490
Observations	164	145	158	154	154	164

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 9. Total borrowing, overlapping sample (endline 1)

	(1)	(2)	(3)	(4)
	All loans	Loans for business	All loans (Informal)	Loans for bus. (Informal)
Treatment	-14,584 (21,208)	-18,760 (16,807)	-3,831 (13,361)	-2,641 (5,808)
Treatment \times GE	36,556 (36,356)	43,283* (25,498)	16,156 (21,251)	12,579* (6,846)
Observations	275	275	275	275
FE	Stratum \times Post	Stratum \times Post	Stratum \times Post	Stratum \times Post
Control Mean (Non-GE)	80617	34555	46710	11977

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. GE means a business which entered before the MFI entered (or would have entered) a treatment (control) area. 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

TABLE 10. Reduced form: worries and happiness (endline 3)

	(1)	(2)	(3)	(4)	(5)
	Overall worries index	Financial worries index	Happiness scale	Financial security scale	Beaten in last month
Panel A: Treatment effects					
Treatment	-0.052*	-0.037	-0.082***	0.020	-0.002
	(0.030)	(0.031)	(0.027)	(0.053)	(0.004)
Control Mean	-0.000	-0.000	0.000	3.763	0.043
Control Std. Dev.	0.588	0.658	1.000	1.270	0.204
Observations	5717	5717	5716	5721	5702
Panel B: Treatment effects by entrepreneurial status					
Gung-ho entrepreneur (GE)	0.039	0.045*	0.015	0.076	-0.016*
	(0.025)	(0.024)	(0.034)	(0.060)	(0.009)
Treatment	-0.061*	-0.049	-0.088***	0.009	-0.000
	(0.032)	(0.033)	(0.030)	(0.055)	(0.006)
Treatment \times GE	0.029	0.041	0.018	0.033	-0.005
	(0.036)	(0.035)	(0.055)	(0.077)	(0.011)
Treatment + Treat \times GE	-0.033	-0.008	-0.069	0.042	-0.006
P(Treat + Treat \times GE \neq 0)	0.416	0.830	0.160	0.603	0.478
Control Mean (Non-GEs)	-0.014	-0.016	-0.005	3.745	0.049
Control Std. Dev. (Non-GEs)	0.587	0.654	1.010	1.285	0.217
Observations	5717	5717	5716	5721	5702

Notes: Respondents were asked whether they were worried about potential stressors along 17 dimensions. Six of the dimensions cover financial worries, such as worrying about not having enough money for food, rent, healthcare, etc.; being worried about debt; and worried about finding work. The remaining 11 dimensions cover non-financial concerns relating to health risks (illness, accidents, deaths, etc.) and various dimensions of conflict within and across households. For each of these dimensions, the respondent rated how worried he or she was on a scale of 1-4, where 1 is not at all worried and 4 is very worried. Haushofer and Shapiro (2013) show that responses to a set of questions similar to ours (and on which our questionnaire is modeled) correlate with levels of salivary cortisol, a physiological marker of stress. In addition to the “worries” questions, households were asked to rate themselves on Likert scales for overall happiness and financial security, and were asked about the incidence of domestic violence. All of the worries, happiness and financial security indexes are scaled to have units of standard deviations. Larger outcomes for the index and scale variables indicate less worried and happier households. 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 11. Attrition

Attrition in treatment vs. control (relative to endline 1)	
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 12. Main Results with Lee Bounds

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	Any	Informal loans	Businesses	Workers	Wages	Biz assets	Log inputs	Log revenue	Profits	Biz index	
	Amt.	Any	Num.	Wages	Workers	Biz assets	Log inputs	Log revenue	Profits	Biz index	
Panel A: All Households											
Treatment	-0.018	2311.846	0.033	0.049	0.186*	328.996**	1375.721***	0.236	0.272	509.051**	0.074***
	(0.013)	(3055.914)	(0.018)	(0.027)	(0.078)	(111.253)	(379.209)	(0.124)	(0.139)	(158.134)	(0.018)
Treatment	-0.040	-1708.982	0.009	0.024	0.122	246.880	800.622	0.095	0.126	346.544	0.054
lower											
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	11532.067*	0.054*	0.092*	0.246	533.710	2796.175*	0.524*	0.628**	1152.770**	0.111***
	(0.023)	(5450.123)	(0.026)	(0.039)	(0.152)	(327.819)	(1191.440)	(0.215)	(0.232)	(417.835)	(0.027)
Treatment	0.019	8323.170	0.035	0.062	0.198	529.864	2203.546	0.357	0.391	948.023	0.090
lower											
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*	-1892.724	0.019	0.024	0.151*	230.592	640.382	0.065	0.087	198.866	0.067**
	(0.017)	(3838.662)	(0.018)	(0.025)	(0.074)	(118.377)	(516.518)	(0.116)	(0.129)	(171.452)	(0.021)
Treatment	-0.066	-5994.197	-0.000	0.003	0.082	138.565	160.201	-0.071	-0.016	72.547	0.037
lower											
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

TABLE 13. 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 14. 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.

TABLE 15. Sensitivity of Simulated Treatment Effects to Preference Parameters and Method for Eliciting Baseline Wealth

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

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

Notes: Table shows the EL3 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.

APPENDIX B. SUPPLEMENTAL FIGURES




FIGURE 10. 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=yongest).

APPENDIX C. PRODUCTION FUNCTION ESTIMATION DETAILS

We do 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.)