

Selection into Credit Markets: Evidence from Agriculture in Mali

March 2020

Lori Beaman, Dean Karlan, Bram Thuysbaert, and Christopher Udry¹

Abstract

We use a two-stage experiment on agricultural lending in Mali to test whether selection into lending is predictive of heterogeneity. Understanding such heterogeneity, and the selection process which reveals it, is critical for guiding modelling of credit markets in developing countries, as well as for policy. We find such heterogeneity: returns to capital are higher for farmers who borrow than for those who do not. In our first stage, we offer loans in some villages and not others. In the second stage, we provide cash grants to a random subset of all farmers in villages where no loans were offered, and to a random subset of the farmers who do not borrow in villages where loans were offered. We estimate seasonal returns to the grant of 130 percent for borrowers, whereas we find returns near zero for the sample representative of those who had recently not borrowed. Critical for both theory and policy, this heterogeneity persists even after conditioning on a wide range of observed characteristics.

JEL: D21, D92, O12, O16, Q12, Q14

Keywords: credit markets; agriculture; returns to capital

¹ Lori Beaman: l-beaman@northwestern.edu, Northwestern University; Dean Karlan: karlan@northwestern.edu, Northwestern University, IPA, J-PAL, and NBER; bram.thuysbaert@ugent.be, Ghent University; and Christopher Udry: christopher.udry@northwestern.edu, Northwestern University. The authors thank partners Save the Children and Soro Yiriwaso for their collaboration. Thanks to Yann Guy, Pierrick Judeaux, Henriette Hanicotte, Nicole Mauriello, Diego Santa Maria, and Aissatou Ouedraogo for excellent research assistance and to the field staff of Innovations for Poverty Action – Mali office. We thank Dale Adams, Alex W. Cohen and audiences at Cambridge University, Columbia University, Dartmouth College, MIT, BU, University of Michigan, the Federal Reserve Bank of Chicago, Stanford, the University of California-Berkeley, University of California-San Diego, and the University of Maryland for helpful comments. All errors and opinions are our own.

1 Introduction

The return to investment in productive activities depends on a myriad of influences, reflecting both the realization of risk and underlying heterogeneity in the characteristics, effort and constraints of producers. Some of this variation may be apparent to outside observers; much may not. Some of this variation may be apparent to producers themselves; much may not. A primary role of financial markets is to help capital flow to the highest return activities.

In a two-stage randomized controlled trial of loans and grants for low-income farmers in rural Mali, we study how self-selection and peer-selection affect this allocation process. We designed a two-stage protocol specifically in order to tease out how selection into loans depends on heterogeneity and the predictability of returns to capital. The sample consists of likely liquidity constrained farmers in rural Mali, a capital-poor economy not well integrated into global financial markets. In stage one of the experiment (the loan stage), a microcredit organization (Soro Yiriwaso, “Soro”) identified 198 villages that were within their expansion plans but which they had not previously entered. Soro then offered group-liability loans to all women farmers in 88 villages, randomly selected from the 198 villages. In these loan treatment villages, some farmers choose, or are chosen by their peers, to borrow via group liability loans under a community association. In stage two (the cash grant stage), after first waiting for households and the associations to make their loan decisions, we announced and immediately gave cash grants (40,000 FCFA, about US\$140) to a random subset of households that did not borrow in the loan villages and of all households in the no-loan villages.

The first stage effectively creates two samples over which we will compare the returns to the stage two cash grants: 88 “loan villages” (where we measure returns to the cash grant for individuals who did not borrow) and 110 “no-loan” villages (where we measure returns to the cash grant for all individuals, i.e. including those who would have borrowed had they been offered a loan, as well as those who would not have borrowed). Comparing the average returns in these two samples allows us to test an important selection question: do those who do not borrow have lower average returns than those who do borrow?

We find large average increases in investment and agricultural profits for the non-selected population (i.e., grant recipients in no-loan villages). Specifically, the cash grants in no-loan villages led to a significant increase in land being cultivated (8.7%, se=3.3%), fertilizer use (18%, se=5%), and overall input expenditures (16%, se=4%). These households also experienced an

increase in the value of their agricultural output and in gross profit² by 13% (se=4%) and 12% (se=5%), respectively. Thus, we observe a statistically significant and economically meaningful increase in investments in cultivation and an increase in gross profit from relaxing capital constraints. This impact on gross profit even persists after an additional agricultural season. Thus in this environment, capital constraints are limiting investments in cultivation.³

However, we find low, indeed zero, average returns to the cash grants for those who did *not* borrow (i.e., the difference between randomly receiving a grant and not among non-borrowers in loan villages). In loan villages, households given grants did not earn any higher gross profits from the farm than households not provided grants. This contrasts sharply with households given grants in the no-loan villages who had large increases in gross profits relative to those not provided grants. Therefore, we conclude that households which borrowed, and were thus selected out of the sample frame in loan villages, had higher marginal returns than those who did not borrow. The differences in the impact of the grants between households who borrow and those who do not are substantial. We estimate that among borrowing households, farm output increases by \$222 (se=120) and farm gross profits increase by \$183 (se=96). In contrast, we estimate that among households who do not borrow, receipt of the grant generates only \$25 of additional output and \$1.04 additional gross profit (neither being statistically significantly different from zero).

Thus, putting the findings from the two samples together, we infer that farmers with particularly high returns to capital are much more likely to select – or be selected – into borrowing. This implies that some of the variation in returns is predictable *ex ante*, and that farmers are aware of this heterogeneity in expected returns.

Although 93% of non-borrowing households report farming as their primary source of income, perhaps the non-borrowers did not invest in farming because they had higher return

² We do not have a complete measure of profits, and thus are using the term “gross profits” as this is the value of agricultural output net of most, but not all, expenses. Most importantly, the value of own, family and unpaid labor is not included. We explain in detail in section 2.2.

³ The increase in investment contingent upon receipt of the grant is sufficient to reject neoclassical separation, but not to demonstrate the existence of binding capital constraints. For example, in models akin to Banerjee and Duflo (2012) with an upward-sloping supply of credit for each farmer, a capital grant could completely displace borrowing from high-cost lenders, lower the opportunity cost of capital to the farmer and induce greater investment even though the farmer could have borrowed more from the high cost lender and thus was not capital constrained in a strict sense. However, there is no evidence that these grants lowered total borrowing. Therefore, we refer to the range of capital market imperfections that could cause investment responses to cash grants simply as credit constraints.

opportunities elsewhere. To examine this, we also look at other outcomes such as livestock ownership and small business operations. However, we do not find evidence of grant recipients in loan villages investing the cash in alternative activities more than their counterparts in no-loan villages.

What is driving this selection effect? The two-stage experiment allows us to conclude that indeed those who borrow are able to use capital more productively than those who do not borrow. But is this selection predictable from information plausibly observable to the lender *ex-ante*? If the heterogeneity is predictable by information observable to the lender *ex-ante*, then the lender could use this information both for social purposes (to focus their marketing efforts on those who stand the most to gain, from a poverty alleviation perspective) as well as expand access to credit (i.e., risk-based pricing, to alleviate adverse selection problems). We find that even after conditioning on the rich set of characteristics in our data, and exploiting a machine learning algorithm (Athey and Imbens 2016; Wager and Athey 2018; Athey and Wager 2019; Athey, Tibshirani, and Wager 2019) to detect heterogeneity in treatment effects, the positive selection induced by the lending process remains strong.

But which aspects of the lending process create the positive selection? Is this driven by borrower self-selection, lender (in this case, peer) selection, or both? The experimental design itself does not allow us to separate these mechanisms, nor does the institutional setting of this credit market provide benefit or cost shifters that would permit estimates of the selection process using local instrumental variables methods as in Heckman (2010) or Eisenhauer et al. (2015). Nor do we elicit directly from peers their predictions on returns for others, as done by Hussam et al. (2020).

We instead provide a simple economic model of the selection process that makes an important prediction: peers will select individuals with high profit *levels* (due to concerns about repayment), whereas individuals will self-select to borrow if they have high *marginal* profits (due to interest in maximizing profits). To examine this, we compare the distribution of returns in no-loan villages (thus a representative sample of everyone) to loan villages (thus only to those selected *out* from borrowing, either by themselves or their peers). In no-loan villages we find no correlation between baseline gross profits and marginal returns to the grant. In the loan villages, however, baseline gross profits are *negatively* correlated with marginal returns to the grant; more specifically, those with higher baseline gross profits have close to zero marginal returns to the grant, whereas those with low baseline gross profits have positive marginal returns to the grant. We find both that high marginal, low average gross profit farmers are under-represented among borrowers, suggesting that they are screened out of borrowing by the lender; and that low marginal, high average gross profit farmers are under-represented among borrowers, suggesting self-selection. We see the same pattern when we use a causal forest algorithm to estimate

conditional average treatment effects (CATEs). CATEs based on a causal forest trained in the no-loan villages show a high density of farmers with high baseline profits and high CATEs. When the causal forest is trained in the loan villages, however, these farmers are notably less represented. Farmers who have both high marginal returns and high baseline profits are much more likely to be borrowers. Thus both mechanisms are at work.

Our experiment also speaks to three additional questions important to academic and policy questions regarding lending to low income farmers: we are able to examine whether loans versus grants generate different investment behavior (we do not find a difference); we are able to examine the efficacy of a microlending program that targets farmers (as compared to the more standard microenterprise focus of microlenders); and we are able to use a seven-year follow-up survey to measure the long-run effects of grants.

First, on comparing grants to loans, note that about 21% of households in our sample received loans (in loan villages), which is typical of other microcredit contexts but of course far below the 100% take-up rate of the grants. The average loan size was 32,000 FCFA (US\$113). Like the grants, we find that offering loans led to an increase in investments in cultivation, particularly fertilizer, insecticides and herbicides, and an increase in agricultural output. We do not detect, however, a statistically significant increase in gross profits. Therefore we observe farmers investing in cultivation when capital constraints are relaxed through credit. Our treatment on the treated estimates of the impact of borrowing on the cultivation activities and harvests of those who borrowed are large and consistent with our entirely separate estimates of the impact of grants on borrowers. Therefore, it does not appear that the lending process leads to dramatically different behavior on the part of farmers than cash grants.

Second, underlying our experiment is an estimate of the impact of an agriculture microcredit program: we find high returns, particularly when compared to experiments estimating the impact of microcredit designed for entrepreneurship.⁴ High average returns to agricultural investment could emerge when farmers lack capital and face credit and savings constraints. Microcredit organizations have attempted to relieve credit constraints, but most microcredit lenders focus

⁴ The evidence from traditional microcredit, targeting micro enterprises, is mixed: some randomized evaluations find an increase in investment in self-employment activity (Crépon et al. 2015; Angelucci, Karlan, and Zinman 2015) while others do not (Karlan and Zinman 2011; Attanasio et al. 2015; Augsburg et al. 2015; Banerjee, Duflo, et al. 2015; Tarozzi, Desai, and Johnson 2015). See Banerjee, Karlan and Zinman (2015) and Meager (2019) for an overview of the above seven studies. Rarely have randomized evaluations of microcredit found an increase in the profitability of small businesses as a result of access to microcredit, at least at the mean or median. These limited results from microcredit come despite evidence that the marginal returns to capital can be quite high for micro-enterprises (de Mel, McKenzie, and Woodruff 2008).

on small or micro business entrepreneurial financing. Furthermore, the typical microcredit loan requires frequent, small repayments and therefore does not facilitate investments in agriculture, where income comes as lump sums once or twice a year (see Karlan and Mullainathan 2007 for a discussion of this; see Fink, Jack, and Masiye 2018 for an experiment demonstrating the importance of this timing issue for farmers). By contrast, the loan product studied here is designed for farmers by providing capital at the beginning of the planting season and requiring repayment as a lump sum at harvest. However, lending may not be sufficient to induce investments in the presence of other constraints. Farmers may be constrained by a lack of insurance (Karlan et al. 2013), have time inconsistent preferences (Duflo, Kremer, and Robinson 2011), or face high costs of acquiring inputs (Suri 2011).

These loan impact results are in stark contrast to a long history of failed agricultural credit programs (Adams 1971), which often were implemented as government programs and thus plagued by politics (Adams, Graham, and Von Pischke 1984). In the expansion of microcredit in the 1980s and onward, we have seen several changes occur at once: a shift from individual to group lending processes (although now this trend is reversing (Giné and Karlan 2014; de Quidt, Fetzer, and Ghatak 2012)), a shift from balloon payments to high frequency repayment (Field et al. 2013 study a lending product that partially reverses this trend, with a delayed start to repayments), a shift from government to nongovernment (and now to for-profit) institutions, and a shift from agricultural focus to entrepreneurial focus (Karlan and Morduch 2009; Armendariz de Aghion and Morduch 2010). The loan impact component of this study effectively returns to this older question, but tests an agricultural lending model that is different than had been employed in the past, one with group liability, little to no subsidy, and no government involvement.

Third, we conducted a follow-up survey in 2017, almost seven years after the grants were made, to measure the long-term effects of the grants. We find no evidence of persistent effects of the grants over this extended period, which was marked by political upheaval and systematic changes in cropping patterns in the region, as well as the highly variable seasonal rainfall realizations typical of the West African semiarid tropics.

2 The Experimental Design and Data

2.1 The Experimental Design

Agriculture in most of Mali, and in all of our study area, is exclusively rainfed. Evidence from nearby Burkina Faso suggests that income shocks translate into consumption volatility (Kazianga and Udry 2006), so improving agricultural output can have important welfare consequences not

only on the level of consumption but also the household's ability to smooth consumption within a year. The main crops grown in the area include millet/sorghum, maize, cotton (mostly grown by men), rice and groundnuts (mostly grown by women). At baseline, about 40% of households were using fertilizer⁵, and 51% were using other chemical inputs (herbicides, insecticide).

The sample consists of 198 villages identified by Soro as villages they had not previously entered but were within their expansion plans. The villages are located in two *cercles* (an administrative unit larger than the village but smaller than a region) in the Sikasso region of Mali.⁶

Figure 1 demonstrates the design.

Stage One: Loans

Soro, a Malian microcredit organization and affiliate of Save the Children (an international nongovernmental organization based in the United States), marketed, financed, implemented and serviced the loans. After a baseline survey was completed (see below), we randomly assigned the 198 villages to either loan (88 villages) or no-loan (110 villages) status using a re-randomization technique ensuring balance on key variables.⁷ This stage one randomization was done at the village level (because that is how Soro marketed and implemented loans).

⁵ The government of Mali introduced heavy fertilizer subsidies in 2008. The price of fertilizer was fixed to 12,500 FCFA (US\$44) per 50 kg of fertilizer. This constituted a 20% to 40% subsidy, depending on the type of fertilizer and year. Initial usage of the subsidy was low in rural areas initially but has grown over time, helping to explain the increase in input expenses we observe in our data from baseline to endline (Druilhe and Barreiro-Huré 2012).

⁶ Bougouni and Yanfolila are the two cercles. Both are in the northwest portion of the region and were chosen because they were in the expansion zone of the Soro. The sample was determined by randomly selecting 198 villages from the 1998 Malian census that met three criteria: (1) were within the planned expansion zone of Soro Yiriwaso, (2) were not currently being serviced by Soro Yiriwaso, and (3) had at least 350 individuals (i.e., sufficient population to generate a lending group).

⁷ First, we ran a loop with a set of number of iterations which randomized villages to either loan or no-loan in each iteration, and then we selected the random draw that minimized the t-values for all pairwise orthogonality tests. This is done because of the difficulties stratifying using a block randomization technique with this many baseline and continuous variables. For village-level randomization of stage one loans, we used: village size, an indicator for whether the village was all Bambara (the dominant ethnic group in the area), distance to a paved road, distance to the nearest market, the percent of households having a plough, the percentage of women having a plough, fertilizer use among women in the village, average literacy rate, and the distance to the nearest health center. For household-level randomization of stage two grants, after first stratifying on stage one village loan status, we used: whether the household was part of an extended family; was polygamous; the primary female respondent's: land size, fertilizer use, and whether she had access to a plough; an index of the household's agricultural assets and other assets, and

Soro offered their standard agricultural loan product, called *Prêt de Campagne*, in the 88 loan villages. This product is given exclusively to women, but naturally money may be fungible within the household. Unlike most microloan products, the loan is designed specifically for farmers: loans are dispersed at the beginning of the agricultural cycle in May-July and repayment is required after harvest. Administratively the loan is administered to groups of women organized into village associations, and each individual woman then receives an informal contract with their village association. Qualitative interviews with members outside the study villages, prior to the intervention, revealed that the application process is informal with few administrative records at the village level. For example, there are records of neither loan applications nor denials. Nor is a record kept of more subtle, informal processes of “application” or “denial”, such as women who discuss the possibility of joining the group to get a loan but who are discouraged from joining (such data would have been helpful for ascertaining the extent of peer versus self-selection, for instance). The size of the group is not constrained by the lender: a group could add a member without decreasing the size of loan each woman received. The size of the loan to each woman is also determined through an informal, iterative process. Repayment is tracked only at the group level, and there is nominally joint liability. On average there are about 30 women per group and typically one, though up to three, associations per village. This is a limited liability environment since these households have few assets and the legal environment of Mali would make any formal recourse on the part of the bank nearly impossible. However, given that loans are administered through community associations, the social costs of default could be quite high. In practice we observe no defaults over the two agricultural cycles where we were collaborating with Soro.⁸

Soro offered loans in the loan villages for two years, the 2010 and 2011 agricultural seasons. The average loan size in 2010 was 32,000 FCFA (US\$113).⁹ The annual interest rate is 25% plus 3% in fees and a mandatory savings of 10%.

Women who borrowed are represented by the far left box in Figure 1.

Stage Two: Grants

per capita food consumption. See Bruhn and McKenzie (2009) for a more detailed description of the randomization procedure.

⁸ This is not atypical for Soro. In an assessment conducted by Save the Children in 2009, 0% of Soro’s overall portfolio for this loan product was at risk (> 30 days overdue) in years 2004-2006, rising to only 0.7% in 2007.

⁹ We use the 2011 PPP exchange rate with the Malian FCFA at 284 FCFA per USD throughout the paper.

Grants worth 40,000 FCFA (US\$140) were distributed by Innovations for Poverty Action (“IPA”), and with no stated relationship to the loans or to Soro, to about 1,600 female survey respondents in May and June of 2010 (i.e., planting time).

In the 110 no-loan villages, households were randomly selected to receive grants and a female household member – to parallel the loans – was always the direct recipient. This corresponds to the boxes on the right side of Figure 1. US\$140 is a large grant: average input expenses, in the absence of the grant, were US\$196 and the value of agricultural output was US\$522. The size of the grant was chosen to closely mimic the size of the average loan provided by Soro, though *ex post* the grant ended up being slightly larger on average than loans. In no-loan villages, we also provided some grants to a randomly selected set of men, but we exclude those households from the analysis in this paper.¹⁰

Figure 1 also depicts that in the 88 loan villages, grant recipients were randomly selected among survey respondents who did not take out a loan.¹¹ We attempted to deliver grants at the same time in all villages, but administrative delays on the loan side meant that most grants were delivered first in no-loan villages, and there is an average 20-day difference between when no-loan households received their grants from their counterparts in loan villages. We discuss the implications of this delay in section 3.2.3.

In order to minimize the possibility of dynamic incentives to not borrow, we informed recipients that the grants were a one-time grant, not an ongoing program, and also distributed some grants in loan villages to some borrowers who were not in the survey. Thus it was not obvious that borrowing precluded someone from being a grant recipient.

2.2 Randomization Balance Check and Attrition

We conduct different tests to verify that there are no important observable differences between the different groups in the sample, using variables not included in the randomization procedure. Appendix Table 1 looks at baseline characteristics across three comparisons: (i) loan to no-loan

¹⁰ The grants to men are intended for a separate paper analyzing household dynamics and bargaining, and we do not consider them useful for the analysis here since the loans were only given to women.

¹¹ We determined who took out a loan by matching names and basic demographic characteristics from the loan contracts between the client and Soro, which Soro shared with us on an ongoing basis. There were a few cases (67) where Soro allowed late applications for loans and households received both a grant and a loan. The majority (41 out of 67) of these cases occurred because there were multiple adult women in the household, and one took out a loan and another received a grant. We include controls for these households. The results are similar if these observations are excluded.

villages; (ii) grant to no-grant households in no-loan villages; and (iii) grant to no-grant households in loan villages. Few covariates are individually significantly different across the three comparisons, and an aggregate test in which we regress assignment to treatment on the set of 11 covariates fails to reject orthogonality for each of the 3 comparisons (p-value of 0.26, 0.91 and 0.67, respectively, reported at the bottom of the table).

Our attrition rate is low: approximately one percent each round. Regardless, Appendix Table 2 reports tests for differential attrition comparing the same groups as in Table A1, from baseline to the first follow-up and to the endline. For each of the three comparisons, we fail to reject that attrition rates are on average the same in the compared groups for both follow up years. In a regression of attrition on the nine covariates, treatment status, and the interaction of nine covariates and treatment status, a test that the coefficients on treatment status and the interaction terms are jointly zero fails to reject for all six regressions (results on bottom row of Appendix Table 2).

2.3 Data

A baseline survey was conducted in January-May 2010. A first follow-up survey was conducted after the first year of treatment and the conclusion of the 2010 agricultural season¹² in January-May 2011; a second follow-up survey was conducted after the second year of treatment and the conclusion of the 2011 agricultural season in January-May 2012; and a third follow-up survey was conducted seven years after the initial grant distribution in January-May 2017. In the four rounds, similar survey instruments covered a large set of household characteristics and socioeconomic variables, with a strong focus on agricultural data including cultivated area, input use and production output at individual and household levels. Throughout the paper we will refer to “gross profits” as a key outcome variable. We do not have a complete measure of profits. Gross profit is the value of agricultural output net of most, but not all, expenses. Specifically, gross profit is the value of harvest (whether sold, stored or consumed) minus the cost of fertilizer, manure, herbicide, insecticide, hired labor, cart and traction animal expenses (rental or maintenance), and seed expenses (although valuing last year’s seeds at zero). We do not subtract the value of own, family or other unpaid labor or the implicit rental value of land used, because

¹² We also conducted an “input survey” on a subsample of the sample frame right after planting in the first year (September-October 2010), in order to collect more accurate data on inputs such as seeds, fertilizer and other chemicals, labor and equipment use. This input survey covered a randomly selected two thirds of our study villages (133 villages) and randomly selected half of the households (stratifying by treatment status) to obtain a subsample of 2,400 households. We use the input survey if conducted, and if not we use the end of season survey. We also control for timing of the collection of the data in all relevant specifications.

both the labor and land markets are too thin to provide reliable guidance on these values. We will, however, examine the use of these inputs directly.

We also collected data on food and non-food expenses of the household as well as on financial activities (formal and informal loans and savings) and livestock holdings.¹³ The food expenditure module asked about consumption of over 50 food items over the previous seven days. We calculate prices using village-level reports in all sample villages. We use these sample-wide prices to convert consumption of all items into expenditures. It is important to note that there is a lot of consumption seasonality in Mali (Beaman, Karlan, and Thuysbaert 2014). Our measure of food expenditure reflects consumption in the post-harvest season only.

3 Selection into loans and the return to cash grants

We focus on agricultural outcomes, so consider agricultural output Q . $\{Q(0,0), Q(0,1), Q(1,0)\}$ represents the set of possible outputs in year 1 of households in our sample. $Q(0,0)$ is a random variable representing potential output if the household neither borrows nor receives a grant; $Q(1,0)$ and $Q(0,1)$ are similarly defined for households who receive a loan but not a grant, and for those who receive a grant but not a loan.¹⁴ ¹⁵ The joint distribution of potential outcomes is $F(Q(0,0), Q(0,1), Q(1,0))$, and the three marginal distributions are denoted $F_{NG}(Q(0,0))$, $F_G(Q(0,1))$ and $F_B(Q(1,0))$.

Define B as a binary variable representing the loan take-up outcome for each household. Thus $B = 1$ if the household would take up a loan if located in a loan village, and $B = 0$ if the household would not take up a loan if in a loan village. Our first goal is to identify the expected value of the effect on output of receiving a grant for households for which $B = 1$ versus those for which $B = 0$. The two-stage randomization provides straightforward identification of the expected value of the return to an exogenous cash grant for these two groups.

¹³ The survey instruments are all available upon request.

¹⁴ There is a fourth logically possible potential outcome, $Q(1,1)$ for households who both borrow and receive a grant, but this is irrelevant in our context because no one who receives a loan is ever assigned to the grant treatment.

¹⁵ This is a minor adaptation of the standard potential outcomes notation building on Rubin (1974); Heckman (1992; 1997); Imbens and Angrist (1994); Angrist et al. (1996); Heckman et al. (1997).

Define $G \in \{0,1\}$ and $L \in \{0,1\}$ as random variables designating a household's status in the grant treatment arm and in a loan treatment village, respectively. The first stage randomization of villages ensures

$$\{Q(0,0), Q(0,1), Q(1,0), B\} \perp L. \quad (1)$$

The second stage randomization of grants across the random sample when $L = 0$ and across non-borrowers when $L = 1$ ensures

$$\{Q(0,0), Q(0,1), Q(1,0), B\} \perp G|L = 0 \quad (2)$$

$$\{Q(0,0), Q(0,1)\} \perp G|(B = 0, L = 1). \quad (3)$$

There is 100 percent take-up of the offer of a grant, so in our sample of non-borrowers in loan villages, we observe

$$Q|(L = 1) = L(1 - B)[Q(0,1)G + Q(0,0)(1 - G)]$$

and in our sample of the full population of no-loan villages we observe

$$Q|(L = 0) = (1 - L)[Q(0,1)G + Q(0,0)(1 - G)]$$

Therefore, (1) and (2) imply that data from the full population of no-loan villages can be used to identify the conditional marginal distributions

$$F_G(Q(0,1)|L = 0, G = 1) = F_G(Q(0,1)|L = 0, G = 0) = F_G(Q(0,1)) \quad (4)$$

and

$$F_{NG}(Q(0,0)|L = 0, G = 0) = F_{NG}(Q(0,0)|L = 0, G = 1) = F_{NG}(Q(0,0)). \quad (5)$$

Similarly, (1) and (3) imply that data from the population of non-borrowers in loan villages can be used to identify

$$\begin{aligned} F_G(Q(0,1)|B = 0, L = 1, G = 1) &= F_G(Q(0,1)|B = 0, L = 1, G = 0) \\ &= F_G(Q(0,1)|B = 0) \end{aligned} \quad (6)$$

$$\begin{aligned} F_{NG}(Q(0,0)|B = 0, L = 1, G = 0) &= F_{NG}(Q(0,0)|B = 0, L = 1, G = 1) \\ &= F_{NG}(Q(0,0)|B = 0). \end{aligned} \quad (7)$$

The loan village population provides an estimate of $\mathbb{P}(B = 1|L = 1) = \mathbb{P}(B = 1)$, which together with (4) and (6); and together with (5) and (7) provides

$$F_G(Q(0,1)|B = 1) = \frac{F_G(Q(0,1)) - F_G(Q(0,1)|B = 0)(1 - \mathbb{P}(B = 1))}{\mathbb{P}(B = 1)} \quad (8)$$

$$F_{NG}(Q(0,0)|B = 1) = \frac{F_{NG}(Q(0,0)) - F_{NG}(Q(0,0)|B = 0)(1 - \mathbb{P}(B = 1))}{\mathbb{P}(B = 1)}.$$

If we define the effect of receiving a grant without a loan as $\Delta_G Q \equiv Q(0,1) - Q(0,0)$, then from (2), we have

$$\begin{aligned} & \mathbb{E}(Q(0,1)|B = 0, L = 1, G = 1) - \mathbb{E}(Q(0,0)|B = 0, L = 1, G = 0) \\ &= \mathbb{E}(Q(0,1) - Q(0,0)|B = 0, L = 1) \equiv \mathbb{E}(\Delta_G Q|B = 0, L = 1) \end{aligned} \quad (9)$$

and

$$\begin{aligned} & \mathbb{E}(Q(0,1)|L = 0, G = 1) - \mathbb{E}(Q(0,0)|L = 0, G = 0) \\ &= \mathbb{E}(Q(0,1) - Q(0,0)|L = 0) \equiv \mathbb{E}(\Delta_G Q|L = 0). \end{aligned} \quad (10)$$

(9), along with $\mathbb{P}(B = 1|L = 1) = \mathbb{P}(B = 1)$, can be estimated from the loan villages and (10) can be estimated with data from the no-loan villages. Equation (1) then implies that we can identify three average treatment effects of immediate interest:

$$\begin{aligned} & \mathbb{E}(\Delta_G Q) = \mathbb{E}(\Delta_G Q|L = 0) \\ & \mathbb{E}(\Delta_G Q|B = 0) = \mathbb{E}(\Delta_G Q|B = 0, L = 1) \\ & \mathbb{E}(\Delta_G Q|B = 1) = \frac{\mathbb{E}(\Delta_G Q) - \mathbb{E}(\Delta_G Q|B = 0)}{\mathbb{P}(B = 1)} + \mathbb{E}(\Delta_G Q|B = 0) \end{aligned} \quad (11)$$

which are the average effects of receiving a grant (without a loan) amongst the general population, amongst those who would not borrow if they were in a loan village, and amongst those who would borrow if they were in a loan village.

Finally, we are also interested in the expected treatment effect of borrowing on those who would borrow if loans were available: $E(\Delta_B Q|B = 1) \equiv E(Q(1,0) - Q(0,0)|B = 1)$. We have already noted that (1) implies that $E(Q(0,0)|B = 0)$ is identified from data on the output of non-borrowers in loan villages, and that $E(Q(0,0))$ is identified from average output in nonloan villages. From (1), data from the loan villages identifies $\mathbb{P}(B = 1|L = 1) = \mathbb{P}(B = 1)$. Then, in parallel with (11)

$$\mathbb{E}(Q(0,0)|B = 1) = \frac{\mathbb{E}(Q(0,0)) - \mathbb{E}(Q(0,0)|B = 0)}{\mathbb{P}(B = 1)} + \mathbb{E}(Q(0,0)|B = 0) \quad (12)$$

(12) and (1) imply that we can identify the average treatment effect on the treated of borrowing:

$$\begin{aligned} \mathbb{E}(Q(1,0)|L = 1, B = 1, G = 0) - \mathbb{E}(Q(0,0)|B = 1) \\ = \mathbb{E}(Q(1,0)|B = 1) - \mathbb{E}(Q(0,0)|B = 1) = \mathbb{E}(\Delta_B Q|B = 1). \end{aligned} \quad (13)$$

3.1 Observable characteristics of borrowers versus non-borrowers

Take-up of the loans, determined by matching names from administrative records of Soro with our sample, was 21% in the first agricultural season (2010-11) and 22% in the second (2011-2012). Despite the similarity in overall take-up numbers, there is a lot of turnover in clients. Only about 65% of clients who borrowed in year 1 took out another loan in year 2. This overall take-up figure is similar to other evaluations of group-based microcredit focusing on small enterprise (for analysis of randomized evaluations of group-based microcredit, see Angelucci, Karlan, and Zinman 2015; Attanasio et al. 2015; Banerjee, Duflo, et al. 2015; Crépon et al. 2015; Tarozi, Desai, and Johnson 2015; and for a summary discussion of these studies, see Banerjee, Karlan, and Zinman 2015).

Table 1 provides descriptive statistics from the baseline on households who choose to take out loans in loan villages, compared to non-clients in those villages. We provide information on the household as a whole, as well as the primary female respondent and primary male respondent. There is a striking pattern of selection into loan take-up: households that invest more in agriculture, have higher agricultural output and gross profits are more likely to take out a loan. Borrowers also have more agricultural assets and livestock. Figure 2 demonstrates that this holds across the whole distribution. Women in households who borrow are also more likely to own a business and are more “empowered” by three metrics: they have higher intra-household decision-making power, are more socially integrated, and are more engaged in community decisions.¹⁶ Households that borrow also have higher consumption at baseline than non-clients.

¹⁶ All three of these variables are indices, normalized by the no-grant households in no-loan villages. The household decision-making index includes questions on how much influence she has on decisions in the following domains: food for the household, children’s schooling expenses, their own health, her own travel within the village, and economic activities such as fertilizer purchases and raw materials for small business activities. The community action index includes questions on frequency she speaks with different village leaders, and different types of participation in village meetings and activities. The social capital index includes questions about seven other randomly selected community members from our sample and whether the respondent knows the person, are in the same organization, would engage in informal risk sharing and transfers with the person, and topics of their discussions (if any).

3.2 Returns to the grant in loan and no-loan villages

To isolate the role of selection into loans, we focus on the first year of the experiment.¹⁷ Table 2 shows the estimates from the following regression using the first follow up data on farm investments and output.

$$Y_{ijt} = \alpha + \beta_1 grant_i + \beta_2 grant_i \cdot loan_j + X_{ijt}\pi + \lambda_j + \epsilon_{ijt} \quad (14)$$

where $grant_i$ indicates individual i received a grant in May-June 2010, and $loan_j$ indicates that the MFI offered loans in village j . We include additional baseline controls (X) which include the baseline value of the dependent variable y_0 ¹⁸ plus its interaction with village type (loan village / no loan village) and the baseline variables used in the re-randomization routine (listed in the notes of table 2). λ_j are village fixed effects. β_1 and β_2 are the primary coefficients of interest. β_1 is the effect of the cash grant on the outcome Y_{ijt} in the no-loan villages, i.e., the average effect of the cash grant in a representative sample of the population. β_2 shows the differential impact of receiving a grant on the outcome Y_{ijt} for the households that did not borrow (in loan villages).

Table 2 shows the estimates from this regression for a variety of cultivation outcomes (inputs along with harvest output and gross profits), and Table 3 shows the analogous estimates for other, non-cultivation outcomes such as livestock, small business ownership, consumption, and female empowerment.

3.2.1 Agriculture

Columns (1)-(8) of Table 2 look at agricultural inputs and crop choice. We first focus on the first row of coefficients, β_1 , which capture the impact of the grant in no-loan villages – which are intended to be representative of the population in Southern Mali. We find that households who received a grant in no-loan villages, compared to those who did not, cultivated more land (0.18 ha, $se=0.07$). This is approximately a 8.7% increase (control mean = 2.07) compared to households who did not receive a grant in no-loan villages. Households also allocate their land to a different crop mix: column (2) shows that 0.07 more hectares ($se=0.02$) are dedicated to

¹⁷ The second year data is more difficult to interpret. In loan villages, a different set of households borrowed in year 2 than in year 1. In particular, we observe a positive, though modest, treatment effect of receiving a grant on taking out a loan in year 2. The impact of the grant in year 2 in loan villages is therefore a combination of mechanisms and does not isolate selection. The results in year 2 are shown in section 6.1.

¹⁸ In cases where the observation is missing a baseline value, we instead give the lagged variable a value of -9 and also include an indicator for a missing value.

growing rice and groundnuts, which are cash crops in the area. The grant also induced an increase in the use of the plough (6 percentage points, $se=1pp$), the quantity of seeds used (5 kg, $se=2.1$), and in hired labor days (2.7 days, $se=0.8$). While 2.7 days over the entire agricultural season is a small number, these households use little hired labor: the mean in the control in 2011 is only 17 days. We observe no change in family labor. Fertilizer and other chemical inputs increased by 18 percent (\$21, $se=6$). The grants therefore led to an overall increase in agricultural investment: Column (8) shows that on net, input expenses that we can put a dollar value on increased by \$30 ($se=8$).¹⁹

Columns (9)-(10) report statistically significant and economically meaningful increases in output and gross profits: output increased by \$66 ($se=19$) and gross profits increased by \$39 ($se=16$), equal to 13 and 12 percent increases, respectively. Overall, we see significant increases in investments and ultimately gross profits from relaxing capital constraints.²⁰

Critically, the coefficient on Grant * Loan village (β_2) demonstrates striking heterogeneity in the returns to the cash grant between no-loan and loan villages. The Grant * Loan village coefficient shows that the selected sample of households who did not take out a loan do not experience the same positive returns when capital constraints are relaxed.

Column (1) shows that households in loan villages who did not take out a loan did not increase the amount of land they cultivated when randomly selected to receive a grant ($\beta_2 = -0.16$ ha, $se=0.10$ and the p-value of the test that the sum of β_1 and β_2 is zero is 0.80). The interaction terms for family labor and fertilizer/other chemical expenses are also negative (-6.9 days, $se=6.5$ and -\$15, $se=9$, respectively). Households who received grants in loan villages did seem to increase some inputs, such as quantity of seeds and hired labor, although neither is statistically significant as shown in columns (2)-(6). On net, column (8) shows that total input expenses in loan villages increase in response to the grant by about \$20 (p value is 0.02), which is not statistically different from the estimate in no-loan villages of \$29.

¹⁹ Value of land and the shadow wage of family labor cannot be estimated given the extremely thin land and labor market in this area. In addition, only seeds that were purchased in the market and rental costs of a plough are included in total input expenses. The value of seeds used from the previous year's harvest, and the cost of using their own plough are also not included. See the notes in table 2 for more details.

²⁰ We are not estimating the marginal product of capital as in de Mel et al. (2008) but instead the "total return to capital" – i.e., cash. Beaman et al. (2013) also shows in this same area that labor inputs also adjust along with agricultural inputs, making it impossible to separate the returns to capital from the returns to labor without an additional instrument for labor inputs. We are therefore capturing the total change in profits and investment behavior when capital constraints are relaxed.

However, even though we observe increased inputs for the grant recipients in loan villages, we see no corresponding increase in agricultural output nor in gross profits. The β_2 interaction coefficient for output is similar in magnitude as β_1 but *negative* (-\$41, se=28), offsetting the increase in output in no-loan villages (\$66, se=19). The test that the sum of the two coefficients is different from zero is not rejected ($p=0.23$), indicating that the (intentionally) selected sample did not experience a statistically significant increase in output when given a grant. Similarly for gross profits, the total effect in loan villages is essentially zero (\$1.04), which is not significantly different from zero ($p=0.95$) and fairly precisely measured. Thus while there is some evidence that households who did not take out loans used some of the grant to increase agricultural inputs, there is no evidence of increases in agricultural output nor gross profits – in stark contrast to the random sample of households in no-loan villages.

These estimates imply that there is a great deal of heterogeneity in marginal returns to relaxing capital constraints across farmers, and that those who borrow are disproportionately those with high marginal returns. The return to the grant implied for would-be borrowers in no-loan villages is \$131 (se=68) in additional gross profits per \$100 of grant.²¹ In contrast, the return for non-borrowers is close to zero.

The analysis indicates that households who do not borrow are those without high returns in agriculture to cash transfers. In contrast to the literature on health products, where much of the evidence points towards limited screening benefits from cost sharing (Cohen and Dupas 2010; Ashraf, Berry, and Shapiro 2010; Tarozzi et al. 2014), we find that the repayment liability leads to lower return households being screened out. The design does not allow us to experimentally determine whether households are self-selecting (demand side) or being screened by the lender (supply side). We return to this question in section 5.

3.2.2 Other outcomes

Table 3 shows the estimates of equation (14) looking at outcomes other than agriculture. The most striking results are in columns (1) and (2): grant-recipient households in no-loan villages are more likely to own livestock (11 percentage points, se=1pp), and there is a large (\$166, se=71) increase in the value of total livestock compared to no-grant households. This represents a 14% increase in the value of household livestock, and is slightly larger than the value of the grant itself. Recall we saw in Table 2 that households also spent an extra \$30 on cultivation investments. The livestock value is measured several months after harvest; these results could indicate either that

²¹ Calculated as $(\beta_1 + 0.79(\beta_1 + \beta_2))/(0.21) * (100/140)$ where 0.21 is the loan take-up rate in loan villages, and the grant size is \$140.

post-harvest, households moved some of their additional farming profits into livestock, or may reflect measurement challenges.²² We also see that the grant increased the likelihood in no-loan villages that a recipient household had a small enterprise (Column 3: 4 percentage points higher, $se=2pp$).²³ Grant recipient households also consumed more, including 5.7% more food (Column 4: \$0.34 per day in adult equivalency, $se=0.14$) and 5.8% in non-food expenditures (Column 5: \$2.53 per month, $se=1.39$). Columns (6)-(9) show no main effect of the grant on whether the household has any financial savings, membership in rotating, savings and loans associations (ROSCAs), education expenses or medical expenses.²⁴

The investment and spending patterns among grant recipient households in loan villages for the most part echo those described above in no-loan villages. Column (1) shows that while grant recipients in loan villages were overall more likely to own livestock than their control counterparts, the magnitude of the effect is smaller than in the no-loan villages (interaction term is -4 percentage points, $se=2pp$). The remainder of the outcomes however show few differences.²⁵

Taken together, Table 3 shows that the grants benefited households in a variety of ways. However, we have no strong evidence that households in loan villages, who did not experience higher agricultural output and gross profits as in no-loan villages, used their grants to invest in alternative higher-return activities other than cultivation.

²² We may also over-value recently purchased livestock. At the household level, we collected data on the quantity of animals. We use village-level reports of livestock prices to value livestock quantities for all households. Therefore, if recently purchased livestock are younger or smaller in treatment household, leading to a large estimated treatment effect.

²³ Appendix Table 5 shows in Column (1) that despite increasing the extensive margin of small business, we do not measure an increase in business profits after year 1.

²⁴ Columns (2) through (4) of Appendix Table 5 also show no impact in year 1 on women's empowerment, involvement in community decisions nor social capital, respectively.

²⁵ The only outcome which suggests potential heterogeneity in behavior between loan and no-loan villages is medical expenses, in Column (9). Medical expenses (in the last 30 days) are marginally-significantly higher in loan grant households (\$5.01, $se=2.55$), since medical expenses may have declined (-\$2.58, $se=1.87$) among grant recipients in no-loan villages. The total effect in loan villages is not statistically different from zero ($p=0.16$). This is a difficult outcome to interpret because (i) having more resources could mean a household is more likely to treat illnesses they experience but by contrast (ii) they are also more able to invest in preventative care, which should lower medical expenses.

3.2.3 Robustness

Timing of delivery of grants

One concern about our interpretation of the results is a timing issue: households received grants in loan villages on average 20 days later than in no-loan villages because of delays in the administration of the loans. If farmers in loan villages received grants too late in the agricultural cycle to make productive investments, we would erroneously conclude that there is positive selection into agricultural loans, since we would observe more investments and returns in no-loan villages than in loan villages. This is particularly a concern since we observe farmers increase the amount of land they farm, which is a decision which occurs very early in the agricultural cycle. In Appendix Table 3, we look at land cultivated (i.e., an investment decision made early in the process) and an index of all the agricultural outcomes and find no relationship with the timing of the grant, among the grant-recipient households in no-loan villages. We look at two main specifications: one in which we include the date the grant was received linearly and with its square, and a second in which the sample is split into the first half of the grant period and the second half (since most of the grants in the loan-available villages were distributed in the second half). In both cases we control for whether this was the team's first visit to the village (revisit to village).²⁶

Spillovers

It is possible that households that received neither grants nor loans were *indirectly* affected by the study interventions. Spillovers could be either positive (if grants or loans were shared) or negative (through general equilibrium effects on locally determined prices or competition over land). We do not have a perfect way to estimate such spillovers. We do, however, have data from an additional 69 villages in the same administrative units (cercles) as our study villages.²⁷ Appendix Table 4 shows that no-grant households in no-loan villages had similar agricultural practices to households in villages where we did no intervention. There are no significant differences in hectares of land cultivated, suggesting that the increase in land cultivated among grant recipients was not zero-sum with households who did not get a grant. We also observe no

²⁶ Households who are revisited are those who were not available during the first visit to the village. They may be systematically different than households who are reached during a first-visit.

²⁷ Our partner organization would only commit to not enter 110 villages, which serve as our no-loan villages. The villages we use as no-intervention villages were leftover replacement villages and not entirely randomly selected. For example, the no-intervention villages have larger average population size but fewer children per household than study villages. Soro Yiriswaso may have offered loans in up to 15 of the 69 villages in year 1. Removing those 15 villages leaves Appendix Table 4 qualitatively unchanged.

change in land cultivated with rice or groundnuts (column 2). This is important since land used to grow rice, which needs to be in a flood plain, is more constrained than other types of land and is thus most likely to be crowded out by treated households. There are also no significant differences in total input expenses, value of the harvest, and gross profits (columns (6)-(8)). Column (9) suggests no changes in equilibrium prices, given we observe no difference in our price index. The one significant difference is the number of hired labor days (column 4). Non-grant recipients in no-loan villages hired more labor by 3.5 laborer days ($se=1.4$). While this is precisely estimated and a point estimate comparable to main treatment effect in Table 2, recall that this is four man-days over the entire course of the agricultural season and therefore unlikely to have affected total output and gross profits.

4 Unobservable versus observable predictors of marginal returns

4.1 Predicting returns based on observable characteristics

Table 1 demonstrated that loan-takers are systematically different at baseline than those who do not take out loans on a number of characteristics, including some which are surely important in cultivation: they have more land, spend more in inputs, and enjoy higher output and gross profits. Are these baseline characteristics enough to predict who could most productively use capital on their farm? Theoretically, the prediction is ambiguous: in many models those who have the highest returns are households who are the most credit constrained. But we observe that individuals who take out loans have on average *more* wealth in the form of livestock. It could be that those with lower returns to investments in cultivation instead invest in livestock. Several variables show that those who take-up are wealthier in general (more land, more livestock, higher consumption), and wealthier households may also have access to better technologies, like a plough, which could increase their returns to capital.

Here we examine whether the marginal returns from grants and the selection effect discussed above are predicted fully by characteristics observed in the baseline, or if there is additional selection that occurs based on unobservables.

We start by examining heterogeneity in returns by observable characteristics in no-loan villages only, in the unselected random sample of farmers. Columns (1) and (2) of Table 4 show that there is limited evidence of heterogeneity using the variables which we saw to be important in Table 1, including baseline gross profits, baseline land size, and baseline value of livestock.

Instead of relying on our intuition for choosing baseline characteristics, we can exploit a machine learning algorithm to estimate heterogeneity in treatment effects (Athey and Imbens 2016;

Wager and Athey 2018; Athey, Tibshirani, and Wager 2019). Researcher-chosen characteristics may (i) be subject to concerns about inference in light of multiple testing and simultaneously (ii) miss important heterogeneity which results from non-linear combinations of baseline characteristics. The causal tree algorithm of Athey and Imbens (2016) extends the basic intuition of decision trees like those used in random forests by selecting splits in order to maximize heterogeneity in treatment effects across leaves (less a penalty for the variance of treatment and control outcomes in each leaf). This ensemble of decision trees-based approach provides estimates of conditional average treatment effects (CATE) for each household.

We implement the generalized random forest method (Wager and Athey 2018; Athey and Wager 2019; Athey, Tibshirani, and Wager 2019) using the R package *grf* version 0.10.4 (Tibshirani et al. 2018; Athey, Tibshirani, and Wager 2019). We grow a rather large number of trees in each forest (250,000) and account for clustering at the village level. See appendix A1 for details on implementation of the causal forests methodology.

In column (3) of Table 4, we assess heterogeneity using the predicted treatment effects from the algorithm trained on the no-loan village data only. As in Chernozhukov et al. (2018), Davis and Heller (2017) and Davis and Heller (2019), we examine how well the estimated treatment effects - CATEs – predict how gross profits vary with treatment. The point estimate is positive, but noisy (0.33, se=0.58). This is suggestive but far from conclusive evidence of heterogeneity in no-loan villages.

Columns (1)-(3) demonstrate that if we had only implemented a cash grant experiment in randomly selected villages, without the experimental design which allows us to compare returns to non-borrowers, we would not have concluded on the basis of the characteristics we observe that there is substantial heterogeneity in the returns to investments in cultivation.

We also estimate CATEs from the causal forests algorithm trained on the selected sample of non-borrowers in loan villages. Table 4, Column 4 looks at this loan villages subsample. When we train a causal forest algorithm on this sub-sample, we find strong evidence of heterogeneous treatment effects. Grant * predicted causal effects is positive and significant at the 5% level (1.28, se=0.49). We return to this finding in section 5.2.

4.2 Does heterogeneity based on observables explain the heterogeneous treatment effects for borrowers and non-borrowers in the experiment?

Table 4 shows no strong evidence of heterogeneity based on observables in the agricultural returns to grants in the random sample of farmers in no-loan villages. We now explore the possibility that observable characteristics (which we have seen in Table 1 are correlated with loan

take-up) can account for the lower return to grants of non-borrowers in loan villages. To explore whether the experiment induces selection not picked up by observable characteristics, we use a specification that interacts baseline characteristics (Z) with an indicator for receiving a grant:

$$Y_{ijt} = \alpha + \beta_1 grant_i + \beta_2 grant_i \cdot loan_j + \gamma_1 grant_i \cdot Z_{ijt} + \gamma_2 Z_{ijt} + X_{ijt}\pi + \lambda_j + \epsilon_{ijt} \quad (15)$$

We structure our analysis by sequentially increasing the controls we include in the regression, by first focusing on Z variables which would be fairly observable to microcredit institutions (MFIs), then including variables which would be fairly observable to the community and therefore may be included in peer screening mechanisms (as in group-lending). Finally, we include the predicted treatment effects from the causal forest model trained on the no-loan villages. This should be a robust synthesis of many covariates, and their interactions.

Table 5 shows our empirical specification (15) with gross profits as the outcome, with different baseline household-level controls. Column (1) is identical to Column (10) in Table 2 and is included for ease of comparison. Column (2) includes Z variables measured at baseline, and their interactions with grant receipt, that an MFI may be able to easily observe: the household's landholdings (in hectares), the value of their own livestock, gross profits, an indicator for whether the household has six or more adults (the 90th percentile), an indicator for the presence of an extended family, and the number of children in the household. Column (2) shows that the estimates of the differential effect of the grant in loan versus no-loan villages is reduced in magnitude slightly (-\$33, $se=22$ compared to -\$38 without controls) but is qualitatively unchanged. We show the coefficients from the interactions between some of these Z variables and grant receipt. Strikingly, higher baseline gross profits do not predict higher returns to the grant, on average. We also do not observe a statistically significant relationship between baseline livestock value and returns to the grant. However, larger households do benefit more from the grants than smaller households, and households with larger baseline landholdings have lower returns.

Column (3) adds in additional information that would likely be known within the community and thus usable in a peer lending screening process: the primary female respondent's intra-household decision-making power, her engagement in community decision-making and her social capital. In all specifications, the estimates on the differential impacts of the grants in loan versus no-loan villages are slightly smaller in magnitude but qualitatively similar.

Column (4) includes the predicted treatment effects from the causal forest algorithm trained on no-loan villages and then used to predict CATEs for the entire sample. This table uses data from both no-loan and loan villages, but we continue to see no meaningful heterogeneity in returns based on a model trained on the no-loan village data.

We therefore conclude that our estimates of selection effects are not driven by the rich set of observables we measure at the baseline, but by characteristics more difficult for outsiders to observe, such as land productivity, access to complementary inputs, or farmer skill.

In the next section we examine whether the selection is a demand-side effect (people choosing whether to borrow or not) or a supply-side effect (lenders or peers choosing whether to let a farmer into their lending circle).

5 Is screening driven by supply-side or demand-side forces?

In section 3.2 we showed that providing cash grants to households who did not take out loans led to lower agricultural returns – and in fact zero returns – compared to households who were randomly selected in no-loan villages. The experimental design provided us with a transparent method for showing that the impact of the grants on gross profits in the random sample of households is greater than their impact in the selected sample of non-borrowers. In contrast, the experimental design itself does not allow us to differentiate how the screening itself occurs: it may be the result of self-selection on the part of farmers (demand-side) or due to lender screening on the part of the MFI or community associations (supply-side).²⁸

5.1 Theoretical framework

We begin with a simple model to illustrate what we mean by self-selection and by lender screening. In order to distinguish these concepts, the model requires three elements. First, there is liquidity constraint that generates a potential demand for credit. Second, some potential conflict of interest between the borrower and the lender is required if self-selection and lender screening are to be distinguished. In this model, the conflict emerges from limited liability. Third, multiple dimensions of heterogeneity across borrowers will generate patterns of self-selection that differ from those caused by lender screening. Self-selection will be largely driven by

²⁸ The MFI itself has little to no information about individual loan applicants. However, women must go through a community association – which in principle has joint liability for the loan – in order to get a contract with the MFI. It is therefore possible that the associations are screening out some farmers who want to borrow.

heterogeneity in marginal productivity; screening of borrowers by lenders will depend more on heterogeneity that affects the total value of output.

The heterogeneity was introduced in section 3 as a vector of characteristics of the household Z . We focus on two dimensions of this heterogeneity of endowments: $z_i = (\eta_i, \theta_i)$. θ_i is an average productivity shifter that affects output but not the marginal product of the input, and η_i affects the marginal product. We consider a situation of symmetric information – both the farmer and the lender know the farmer’s endowment before the loan is transacted, but neither knows the realization of ϵ_i , which we think of as a random shock to output realized after borrowing is completed. The lender provides loans normalized to size 1, at an interest rate of r (these parameters are set exogenously at a national level by Soro). At the start of the farming season, i chooses whether to borrow $B_i \in \{0,1\}$, and the lender chooses whether to lend $L_i \in \{0,1\}$. The loan is made if and only if $B_i L_i = 1$. We are assuming that the farmer has no alternative use for capital outside of agriculture.²⁹

Gross profit for borrowers is $q_i = Q_B(\theta_i, \eta_i, \epsilon_i)$ and for non-borrowers is $q_i = Q_N(\theta_i, \eta_i, \epsilon_i)$. A convenient specification for gross profits that satisfies the assumptions on η and θ is

$$q_i = \eta_i(f_N + (f_B - f_N)B_i L_i) + \theta_i + \epsilon_i \quad (16)$$

where constants $f_B \geq f_N$ reflect the liquidity constraint that generates the demand for credit. For shorthand, we refer to differences across farmers in θ as differences in average productivity and to differences in η as differences in marginal productivity.

Output is produced, and because of limited liability, the loan is repaid in full if and only if gross profit is sufficiently high.³⁰ The lender receives $\min\{(1+r)B_i L_i, q_i\} - (1+\rho)B_i L_i$ and the farmer keeps $\max\{q_i - (1+r)B_i L_i, 0\}$, where $\rho (< r)$ is the cost of funds to the lender. We assume that both farmer and the lender maximize expected profits.

As a consequence of limited liability, a borrowing farmer earns zero if $\epsilon \leq \bar{\epsilon}$, where

²⁹ This assumption implies that $Q_G \equiv Q_B$, that is, that the farmer uses the grant in the same way he/she would use the loan. This possibly unrealistic assumption can be generalized, at the cost of additional notation, while preserving the lessons we draw for patterns of selection in the following paragraphs. In section 7, we provide evidence that the uses of and returns to the grants are similar to those of the loans.

³⁰ We make assumptions to ensure that output is nonnegative. So $(\eta_i, \theta_i) \in P = [\eta^L, \eta^H] \times [\theta^L, \theta^H]$ with $\eta^L f_N + \theta^L \equiv q^L > 0$, and ϵ_i is drawn from a continuous density $g(\epsilon)$ with positive support on $[\epsilon^L, \epsilon^H]$ with $\epsilon^L \geq -q^L$. The expected value of ϵ is 0.

$$\bar{\epsilon} \equiv (1 + r) - \eta_i f_B - \theta_i \quad (17)$$

The farmer will want to take a loan if and only if

$$(\eta_i f_B + \theta_i - (1 + r))(1 - G(\bar{\epsilon})) + \int_{\bar{\epsilon}}^{\epsilon^H} \epsilon g(\epsilon) d\epsilon \geq \eta_i f_N + \theta_i \quad (18)$$

The set

$$B^{rej} = \left\{ (\eta, \theta) \mid (\eta_i f_B + \theta_i - (1 + r))(1 - G(\bar{\epsilon})) + \int_{\bar{\epsilon}}^{\epsilon^H} \epsilon g(\epsilon) d\epsilon < \eta_i f_N + \theta_i \right\}$$

defines the characteristics of farmers who would choose not to borrow. Define η_i^* as the level of marginal productivity such that a loan that will not be defaulted on has an expected return just equal to the interest cost:

$$\eta_i^*(f_B - f_N) = 1 + r \quad (19)$$

η_i^* is independent of θ_i . We can say something about the magnitude of η^* , because the marginal cost of borrowing is on the order of 30%. No household with expected returns under this magnitude will borrow to invest in agriculture unless the probability of default is positive.

Second, define $\theta_i^* \equiv (1 + r) - \eta_i^* f_B - \epsilon^L$. A borrowing farmer with an endowment $\{\theta_i^*, \eta_i^*\}$ never defaults, and (18) is satisfied with equality. For all $\eta_i \geq \eta_i^*$, a farmer with endowment $\{\theta_i^*, \eta_i\}$ chooses to borrow. Similarly, for all $\theta_i < \theta_i^*$, a farmer with endowment $\{\theta_i, \eta_i^*\}$ defaults with positive probability, and chooses to borrow because of the limited liability constraint.³¹ In Figure 3, we show the set of (η, θ) such that a farmer would choose $B_i = 1$. The solid curve labeled B partitions the space such that farmers with endowments to the southeast of B choose to borrow.³²

The lender will choose to make the loan if and only if expected profits are positive. So the lender is willing to lend to i if and only if

³¹ From (16) $(\eta_i^* f_B - (1 + r)) = \eta_i^* f_N$ and $\epsilon^L = (1 + r) - \eta_i^* f_B - \theta_i^*$. For $\theta < \theta_i^*$, $G((1 + r) - \eta_i^* f_B - \theta) > 0$. Therefore, $E\max(\eta_i^* f_B + \theta + \epsilon_i - (1 + r), 0) > \eta_i^* f_N + \theta$.

³² B is upward-sloping below $\{\theta_i^*, \eta_i^*\}$ by reasoning analogous to that in the preceding note. If $\{\tilde{\theta}, \tilde{\eta}\}$ is a point on B with a positive probability of default: then for all $\theta < \tilde{\theta}$, $G((1 + r) - \tilde{\eta} f_B - \tilde{\theta}) < G((1 + r) - \tilde{\eta} f_B - \theta)$ and the farmer endowed with $\{\theta, \tilde{\eta}\}$ strictly prefers to borrow.

$$(1+r)(1-G(\bar{\epsilon})) + (\eta_i f_B + \theta_i)G(\bar{\epsilon}) + \int_{\epsilon}^{\bar{\epsilon}} \epsilon g(\epsilon) d\epsilon \geq (1+\rho). \quad (20)$$

$L^{rej} = \{(\eta, \theta) | (1+r)(1-G(\bar{\epsilon})) + (\eta_i f_B + \theta_i)G(\bar{\epsilon}) + \int_{\epsilon}^{\bar{\epsilon}} \epsilon g(\epsilon) d\epsilon < (1+\rho)\}$ defines the set of borrower characteristics such that the lender would not be willing to lend to the borrower with those characteristics. Equation (20) is satisfied for farmer i with endowment $\{\theta_i^*, \eta_i^*\}$, because $r > \rho$. In Figure 3, the dashed curve labeled L partitions the space such that the lender is willing to make a loan to farmers with endowments to the northeast of L ; the set L^{rej} is the area to the southwest of L .

We can now consider the consequences of self-selection versus lender screening for the observed distribution of gross profits. In the no-loan villages, where grants were given to a random sample of the population, we have

$$\begin{aligned} q_i^{NG} &= \eta_i f_N + \theta_i + \epsilon_i \\ q_j^{Grant} &= \eta_j f_B + \theta_j + \epsilon_j \end{aligned} \quad (21)$$

for farmers i and j randomly selected into the no grant and grant treatment groups, respectively. Let $h(\eta, \theta)$ denote the joint density of η and θ in the population of our study area, then given our randomization, the distributions of q_j^{Grant} and q_i^{NG} simply reflect draws from the full density $h(\cdot)$. The left panel of Figure 4 depicts these distributions; as can be anticipated from our preceding results, the distribution of q_j^{Grant} lies to the right of that of q_i^{NG} over virtually the whole range.

In the loan villages, grants were given to a random sample of non-borrowers. Suppose that selection into borrowing is being driven by the simultaneous operation of both borrower side self-selection and by lender-side screening; that is, that the selection is driven jointly by equations (18) and (20). In this case, the joint density of η and θ in the population of non-borrowers is the truncated probability distribution

$$h^{rej}(\eta, \theta) = \frac{h(\eta, \theta)}{\text{prob}((\eta, \theta) \in \{L^{rej} \cup B^{rej}\})} \quad (22)$$

with support $(\eta, \theta) \in \{L^{rej} \cup B^{rej}\}$. As can be seen in Figure 3, the endowments of the approximately 80 percent of the population who do not borrow differ from the overall population in two ways. First, because of lender screening, the distribution of endowments in the selected population of non-borrowers has greater weight on low values of average productivity θ . Second, because of borrower self-selection, the selected population contains a higher proportion of farmers with low marginal productivity. Put differently, self-selection implies that among non-

borrowers with high average productivity, a disproportionately large share will have low marginal productivity. And lender screening implies that among non-borrowers with high marginal productivity, a disproportionate share will have low average productivity.

5.2 Empirical evidence of self and community selection

We start by first looking at the cumulative distribution functions (CDFs) of gross profits in no-loan and loan villages.³³ The right panel of Figure 4 depicts the distributions of q_j^{Grant} and q_i^{NG} for the randomly chosen grant recipients and non-grant recipients among the population of non-borrowers in the loan villages. There are two distinctive feature of this graph. First is the presence of a significant fraction of non-borrowers with relatively high gross profits (>\$500), but approximately zero marginal return from the grant. This pattern in the data is consistent with the mechanism of self-selection: farmers who have low returns to capital would opt out of loans and therefore show up in this sample. Second is the presence of a significant fraction of non-borrowers with high marginal productivity but low average productivity (measured by gross profits). This feature corresponds to the mechanism of lender screening. We infer that the realizations of q_j^{Grant} and q_i^{NG} are determined jointly by equations (18) and (7), so that non-borrower endowments are drawn from $h^{rej}(\eta, \theta)$. Both self-selection and lender screening are occurring in this credit market.

Correlations between observable characteristics of borrowers and non-borrowers and the return to grants are also informative of the nature of the selection process. We saw in Table 1 that there are a number of observable characteristics that are strongly (positively) correlated with loan take-up. Consider any such attribute, Z , that we *a priori* expect to be correlated with average productivity, θ . For example, baseline gross profits would be one such attribute. In Table 6, we report the results of estimating

$$Y_{ijt} = \alpha + \beta_1 grant_i + \beta_2 grant_i \cdot loan_j + \gamma_1 grant_i \cdot Z_{ijt} + \gamma_2 Z_{ijt} + \delta_1 grant_i \cdot Z_{ijt} \cdot loan_j + X_{ijt}\pi + \lambda_j + \epsilon_{ijt} \quad (23)$$

where we have augmented specification (15) with an additional interaction of Grant * Z * Loan village. This additional interaction permits us to examine whether the correlation between Z and the marginal return to the grant is different for the general population (γ_1) than for a selected population of non-borrowers ($\gamma_1 + \delta_1$). This helps illuminate whether the underlying mechanism is self-selection driven, lender driven, or both. The higher average productivity, θ , associated with

³³ Note that this is the same sample as we use in table 2, and therefore continues to exclude households who borrowed in loan villages.

the higher value of Z reduces the likelihood that the farmer has been screened out of borrowing by the lender, so non-borrowers with higher values of Z are more likely to have self-selected out of borrowing because they have low marginal productivity. Hence, among the population of non-borrowers in loan villages, higher values of Z are associated with lower values of η , relative to the association in the population in general.

Column (1) of Table 6 examines the association between baseline gross profits and the marginal return to the grant in the overall population and in the selected sample of non-borrowers. In accord with our model, households in loan villages have a significantly more negative correlation between baseline gross profits and the return to a grant than households in the overall population (Grant * Baseline gross profits * Loan village: $-\$0.18$, $se=0.07$). In the context of our model, both lender screening and self-selection are required to generate this pattern and the relationship demonstrated in Table 1. Screening by the lender generates the positive correlation between baseline gross profits and loan take-up, and households with low returns to additional liquidity self-select out of borrowing.

In columns (2)-(4), we report the estimates of equation (23) for three additional characteristics of households that are positively associated with loan take-up and plausibly with average productivity θ : baseline value of livestock holdings, baseline food consumption per capita (in USD), and baseline non-food expenditure per capita (in USD). In column (2), we find no significant difference in the correlation between baseline livestock and the return to the grant in the overall population and in the loan villages ($-\$0.015$, $se=0.013$), so this measure provides no evidence in support of the hypothesis that both dimensions of selection are operating. In column (3), we examine baseline harvest period expenditure on food (harvest period expenditure will minimize the likelihood of strong nutrition-productivity effects). The association between baseline food expenditure and the return to the grant in loan villages is much lower than the same correlation in the overall population villages ($-\$23$, $se=6$). Similarly, in column (4), we use baseline non-food consumption per capita as Z , hypothesizing that this quantity may be strongly positively correlated with a household's permanent income (and hence with θ) and less strongly correlated with the marginal product of additional agricultural investment. Again, we find a much lower association between non-food consumption and the effect of the grant on gross profits in loan villages than in the overall population ($-\$1.61$, $se=0.61$).

We return again to the causal forest analysis shown in Table 4. We can estimate CATEs using either an algorithm trained on no-loan villages only or on loan villages only. Table 7 explores whether the baseline characteristics which are associated with high predicted treatment effects (CATEs) are the same in both models. Table 7, Column 1 shows that in the general population of no-loan villages, households with high CATEs have *higher* baseline gross profits, more food and

non-food consumption, more livestock and more landholdings. In contrast, column 2 shows that in the sample of non-borrowers in loan villages, households with high CATEs have lower baseline gross profits, lower baseline food consumption and non-food expenditure, lower livestock values and smaller land holdings.

The comparison between columns (1) and (2) is striking as 6 out of the 8 characteristics have the opposite sign in their correlation with predicted treatment effects in the two models. This is further evidence consistent with our model of peer and lender selection. Among the selected sample in the loan villages who did not borrow, we see that those who are less poor—as proxied by food and non-food consumption – have lower returns. These are households that would be attractive to potential group members since they are less likely to default given a high θ . Therefore, their low returns are likely driven by self-selection: they did not borrow because they knew they had low returns (low η). In the full sample in column (1), we see a positive correlation between baseline food and non-food consumption and predicted returns. This is consistent with Table 1, where borrowers tended to be less poor than non-borrowers. To square this with column (2), the model would suggest that the less poor households with expected high returns choose to borrow, and left the sample that we used to train the model in the loan villages. Those that remain are the less poor households with low anticipated returns, generating the negative correlation in column (2).

This provides a potential explanation for why we do not observe strong evidence of heterogeneous returns in the random sample in Table 4, and why the limited heterogeneity found in the full sample in Table 5 did not explain the differential returns to grants of borrowers and non-borrowers in loan villages. There is unobserved heterogeneity within households with similar observable characteristics, i.e. there are characteristics not observed in our data that drive the selection that we uncover through the experiment.

Figure 5 demonstrates visually the effects of self- and lender-selection on the joint distribution of baseline gross profits and the return to grants. The x-axis is the quantiles of baseline gross profits, while the y-axis is the quantiles of the predicted treatment effects (CATEs). Figure 5a reports the results of the causal forest trained on and estimated in the no-loan village sample. The highest density of observations is in the upper right and there is an apparent positive correlation between baseline gross profits and the estimate CATE of a grant. Figure 5b reports the results of the causal forest trained on and estimated in the sample of non-borrowers in the loan villages. The high-baseline profit and high CATE quadrant of the population is much less represented: these are households that demand loans and are able to borrow.

6 Persistent effects of grants

We focus first on the impact of the grants in year 2 and then discuss the results from a longer-term follow up.

Agriculture

We observe a persistent increase in output and gross profits in the 2011-2012 agricultural season (year 2) from the grant given in 2010. In Table 8, column (8) shows that output is higher in grant recipient households by \$52 (se=23) and Column (9) demonstrates that gross profits was higher by \$49 (se=17). This is striking since we do not observe grant-recipient households spending more on inputs *that we can easily measure* in Column 8 (\$1.10, se=10.45). Recall that there are a number of inputs, such as land, seeds used from the previous year's harvest, and family labor, that we cannot value. Columns 2-4 provide evidence that grant recipients continued to make different investments than the control group. Grant recipients in no-loan villages planted 6.5% more land with rice and peanut crops in year 2. Rice and peanuts are high value crops. Grant recipients in no-loan villages were also 4.9% more likely to use a plough during land preparation (4 pp, se=1pp), and used 6.8% more seeds (6.1 kg, se=2.6).

We show the estimates of the interaction term of Grant * Loan village in year 2 in Table 8, but the interpretation of the results is challenging. In the second year of the experiment, the MFI offered loans again. Only about half of households who took out a loan in year 1 took out another loan. There were also households who did not borrow in year 1 who chose to borrow in year 2. Moreover, households who randomly received a loan in year 1 are more likely to receive a loan in year 2. With the caveats in mind, we see a similar negative interaction term on gross profits in Column 10 as in year 1 (-\$40, se=24). The lower gross profits may be a result of higher input use: Column 8 shows that, in loan villages, grant-recipient households spent more on agricultural inputs (\$30, se=17) than control households in 2012.

Other outcomes

Table 9 shows the persistent impacts of the grant in year 2 on non-agricultural outcomes. Columns 1 and 2 demonstrate that grant-recipient households are more likely to own livestock (9 percentage points, se=2pp) and continue to hold more livestock assets (\$184, se=102) than control households in no-loan villages. They are also more likely to own a business (3 percentage points, se=1pp).³⁴ There is not a significance increase in food consumption in year 2 (\$0.24, se=0.19) but an increase in monthly non-food expenditure (\$3.89, se=2.13). Households are also

³⁴ In results available from the authors, business profits increase by 18% (\$41, se=19) in year 2.

more likely to have financial savings (3.3 percentage points, se=1.9pp) and be members of rotating savings and loans associations (ROSCAs) (3.9 percentage points, se=1.9pp). Columns 8-9 show that there continues to be no measurable impact on educational expenses (\$0.39, se=3.76), or medical expenses (-\$0.72, se=1.82).

Table 9 shows that, similar to year 1, there is no evidence of households in no-loan villages using grants differently than those in loan villages across this set of non-agricultural outcomes (livestock ownership, owning a small business, and consumption) in year 2.

Longer-term follow up

In 2017, almost seven years after the grants were distributed, we conducted another round of data collection, interviewing 5,560 households of the original sample. Table 10 shows no evidence of a persistent effect of the grant on the key agricultural outcomes analyzed in the paper. The time period between 2012 and 2017 was a tumultuous time in Mali. There was a military coup in March 2012, followed by a French military intervention in the north of the country until 2014 (all of which were factors in why there was a large gap in our field work between the second and seven year follow-ups). Second, unrelated to the political instability, there was an expansion in cotton cultivation in the Segou region of Mali. From 2007 to 2010, it is estimated that there was between 200 and 244 million tonnes of cotton produced per year. However, in 2017 that figure had risen to 703 million tonnes (USAID, 2018). The increase largely came from increases in land dedicated to cotton cultivation. The state-owned Malian Textile Development Company (CMDT), which was re-structured starting in late 2010, provides fertilizer and credit to cotton farmers. This change in cultivation patterns could easily wash out any long-term benefits from a single cash transfer many years prior.

Note that we did not analyze if there is a difference in agricultural outcomes between loan and no-loan villages since our partner organization Soro was unable to provide any information on whether loans were disbursed in the treatment and/or control villages between 2012 and 2017.

7 Impact of the loans

We also show our estimates of the intent-to-treat (ITT) effects of being offered an agricultural loan on the same set of outcomes already discussed in section 3. In this analysis, we exclude all grant recipients, from both loan and ineligible villages. Tables 11 and 12 show the results of the loan intent-to-treat analysis. We use the following specification:

$$Y_{ijt} = \alpha + \beta_1 loan_j \cdot I\{t = 2011\} + \beta_2 loan_j \cdot I\{t = 2012\} + X_{ijt}\pi + \epsilon_{ijt} \quad (24)$$

where (X) includes the baseline value of the dependent variable y_0 , *cercle* (an administrative unit above a village and below a region) fixed effects, and the village stratification controls listed in the notes of the Table 2. The specification uses probability weights to account for the sampling strategy, which depends on take-up in the loan villages. See notes in table 11 for details.

Tables 11 and 12 show the ITT estimates for agricultural outcomes and broader outcomes, respectively. In Table 11, we observe an increase in input expenditures on family labor days (8.6, se=4.8) and in fertilizer and other chemicals expenses (\$14, se=7); total input expenses rose by \$20 (se=9) in villages offered loans. Land cultivated also increases but is not statistically significant at conventional levels (0.08 ha, se=0.06). The value of the harvest rose by \$34 (se=20), but we do not measure a statistically significant increase in gross profits (\$19, se=16).

Is the impact of the loan different from the impact of the grant? We calculate the Treatment on the Treated estimates for year 1 for the sub-population who take up loans.³⁵ Compared to the estimate of the impact of the grant from table 2, we do not reject the hypothesis that the per \$100 dollar effects of grants and loans are the same for any of the agricultural outcomes.³⁶ Taken as a whole, the grants and loans are having similar effects on agricultural outcomes.

Table 12 demonstrates that overall, the microcredit agricultural loans did not have broad impacts beyond agriculture. We do not detect an impact on outcomes such as food and non-food consumption, whether the household has a small business, nor educational expenses.³⁷ We observe a large but imprecisely estimated impact on livestock (columns 1-2). There is one significant impact: column (9) shows a *reduction* in medical expenses (-\$5.03, se=1.64). We are not, however, able to document any corresponding increase in preventative health care expenditures.

These results on the impact of loans stand in stark contrast to the recent experimental literature on the impact of entrepreneurially-focused credit (see Angelucci, Karlan, and Zinman 2015;

³⁵ See table notes of Table 11.

³⁶ The standard errors are calculated using a bootstrap routine: the difference in the impact of the grant and loan is estimated for 1,000 draws of households (with replacement), with probability weights for households calculated in each bootstrap sample for the loan impact estimation.

³⁷ Appendix Table 5 further shows no detectable effect on business profits, women's decision-making power within the household, women's involvement in community decisions, nor on women's social capital. This is similar to the existing evaluations of microcredit (finding no impact on these measures: Attanasio et al. 2015; Augsburg et al. 2015; Banerjee, Duflo, et al. 2015; Crépon et al. 2015; one exception is Angelucci, Karlan, and Zinman 2015 which does find some impacts on these outcomes). Soro Yiriwaso did not have any explicit component of the program emphasizing women's empowerment.

Attanasio et al. 2015; Augsburg et al. 2015; Banerjee, Duflo, et al. 2015; Crépon et al. 2015; Karlan and Zinman 2011; Tarozzi, Desai, and Johnson 2015, and an overview in Banerjee, Karlan, and Zinman 2015; in contrast, Breza and Kinnan 2018 finds noticeable general equilibrium effects as a consequence of a state-wide shutdown of the microcredit market). Analysis pooling these studies using a Bayesian hierarchical model approach, however, unravels evidence of positive treatment effect at higher quantiles, even though the average treatment effect is a fairly precise null (Meager 2018; 2019). An earlier agricultural lending literature also documented institutional failures, typically with high default rates (Adams, Graham, and Von Pischke 1984; Adams 1971), although a newer study in Zambia finds positive impacts from loans tailored to agriculture, much akin to here (Fink, Jack, and Masiye 2018).

The impact estimates are also promising from the perspective of the microfinance institution: repayment was 100%, and the retention to the following year (65%) is on par with typical client retention rates for sustainable, entrepreneurially-focused microcredit operations.

8 Conclusion

Capital constraints are binding for at least some farmers in Southern Mali, and we find that agricultural lending with balloon payments (i.e., with cash flows matched to those of the intended productive activity) can increase investments in agriculture. This is an important policy lesson since the majority of microcredit has focused on small enterprise lending, and the typical microcredit loan contract – where clients must start repayment after a few weeks – is simply ill-suited for agriculture. Field et. al. (2013) finds similar results merely from delaying the onset of high frequency repayment, within the context of microenterprises. In Mali, for example, Soro Yiriwaso is among very few microcredit organizations with a product specially designed for agriculture, despite the fact that the vast majority of households in rural Mali depend on agriculture for a sizeable part of their livelihood.

These results are also important for policy analysis. The random choice of communities into which to enter by the lender is sufficient for us to estimate ITT effects of the lending program, avoiding strong assumptions on the selection process. Our results provide evidence of quantitatively important selection on unobserved variables, which has methodological implications for impact evaluation. Had we matched borrowers to non-borrowers on observable characteristics to assess the impact of lending to farmers, we would have overestimated the impact of credit, since conditional on an unusually wide range of observed characteristics those who borrow have substantially higher returns to capital than those who do not borrow.

There are also lessons relevant for the targeting of social programs. Cash transfer programs are often means-tested and recent work suggests that both community targeting, where community members rank-order households to identify the poor, and ordeal mechanisms can be an effective way of generating screening on wealth/income in developing countries (Alatas et al. 2012; 2013). Price is the screening mechanism we look at here with agricultural loans. In a different agricultural setting, Jack (2013) finds that a willingness to accept mechanism can induce self-selection among landholders in Malawi, leading to improved project success for tree planting. Similarly, Maitra et al. (2020) examines alternative mechanisms for hiring agents to manage loans to farmers, and finds more impact on farmers when the agents had prior experience lending and transacting with farmers. We find that the lending process is a mechanism that generates positive selection so farmers who benefit the most from relaxing capital constraints are more likely to choose to borrow.

We find that the returns to capital in cultivation are heterogeneous and that higher marginal-return farmers self-select into borrowing more so than low marginal-return farmers. But there is also a set of high marginal return, extremely poor households that are unable to borrow. This has important implications for models of credit markets. In particular, our results provide rigorous empirical evidence for systematic selection into contracts, which is embedded in several models (e.g., Evans and Jovanovic 1989; Buera 2009; Moll 2014) but which has lacked clear empirical evidence. As recognized by Banerjee et al. (2015) and Kaboski and Townsend (2011), our results highlight the need to incorporate heterogeneity of returns in credit market models.

References

- Adams, Dale W. 1971. "Agricultural Credit in Latin America: A Critical Review of External Funding Policy." *American Journal of Agricultural Economics* 53 (2): 163–72. <https://doi.org/10.2307/1237428>.
- Adams, Dale W., Douglas H. Graham, and J. D. Von Pischke, eds. 1984. *Undermining Rural Development with Cheap Credit*. Westview Special Studies in Social, Political, and Economic Development. Boulder: Westview Press.
- Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A Olken, and Julia Tobias. 2012. "Targeting the Poor: Evidence from a Field Experiment in Indonesia." *The American Economic Review* 102 (4): 1206–1240.
- Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Olken, Benjamin, Ririn Purnamasari, and Matthew Wai_Poi. 2013. "Self-Targeting: Evidence from a Field Experiment in Indonesia."
- Angelucci, Manuela, Dean Karlan, and Jonathan Zinman. 2015. "Microcredit Impacts: Evidence from a Randomized Microcredit Program Placement Experiment by Compartamos Banco." *American Economic Journal: Applied Economics* 7 (1): 151–82.

- Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin. 1996. "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association* 91 (434): 444–55. <https://doi.org/10.2307/2291629>.
- Armendariz de Aghion, Beatriz, and Jonathan Morduch. 2010. *The Economics of Microfinance*. 2nd ed. Cambridge, MA: MIT Press.
- Ashraf, Nava, James Berry, and Jesse M Shapiro. 2010. "Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia." *American Economic Review* 100 (5): 2383–2413. <https://doi.org/10.1257/aer.100.5.2383>.
- Athey, Susan, and Guido Imbens. 2016. "Recursive Partitioning for Heterogeneous Causal Effects." *Proceedings of the National Academy of Sciences* 113 (27): 7353–60. <https://doi.org/10.1073/pnas.1510489113>.
- Athey, Susan, Julie Tibshirani, and Stefan Wager. 2019. "Generalized Random Forests." *The Annals of Statistics* 47 (2): 1148–78. <https://doi.org/10.1214/18-AOS1709>.
- Athey, Susan, and Stefan Wager. 2019. "Estimating Treatment Effects with Causal Forests: An Application." *Observational Studies* 5: 36–51.
- Attanasio, Orazio, Britta Augsburg, Ralph De Haas, Emla Fitzsimons, and Heike Harmgart. 2015. "The Impacts of Microfinance: Evidence from Joint-Liability Lending in Mongolia." *American Economic Journal: Applied Economics* 7 (1): 90–122. <https://doi.org/10.1257/app.20130489>.
- Augsburg, Britta, Ralph De Haas, Heike Harmgart, and Costas Meghir. 2015. "The Impacts of Microcredit: Evidence from Bosnia and Herzegovina." *American Economic Journal: Applied Economics* 7 (1): 183–203. <https://doi.org/10.1257/app.20130272>.
- Banerjee, Abhijit, Emily Breza, Esther Duflo, and Cynthia Kinnan. 2015. "Do Credit Constraints Limit Entrepreneurship? Heterogeneity in the Returns to Microfinance." *Working Paper*.
- Banerjee, Abhijit, and Esther Duflo. 2012. "Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program." *M.I.T. Working Paper*.
- Banerjee, Abhijit, Esther Duflo, Rachel Glennerster, and Cynthia Kinnan. 2015. "The Miracle of Microfinance? Evidence from a Randomized Evaluation." *American Economic Journal: Applied Economics* 7 (1): 22–53. <https://doi.org/10.1257/app.20130533>.
- Banerjee, Abhijit, Dean Karlan, and Jonathan Zinman. 2015. "Six Randomized Evaluations of Microcredit: Introduction and Further Steps." *American Economic Journal: Applied Economics* 7 (1): 1–21.
- Beaman, Lori, Dean Karlan, and Bram Thuysbaert. 2014. "Saving for a (Not so) Rainy Day: A Randomized Evaluation of Savings Groups in Mali." <http://www.nber.org/papers/w20600>.
- Beaman, Lori, Dean Karlan, Bram Thuysbaert, and Christopher Udry. 2013. "Profitability of Fertilizer: Experimental Evidence from Female Rice Farmers in Mali." *American Economic Review Papers & Proceedings*, May.
- Breza, Emily, and Cynthia Kinnan. 2018. "Measuring the Equilibrium Impacts of Credit: Evidence from the Indian Microfinance Crisis." *Working Paper*.
- Bruhn, Miriam, and David McKenzie. 2009. "In Pursuit of Balance: Randomization in Practice in Development Field Experiments." *American Economic Journal: Applied Economics* 1 (4): 200–232.

- Buera, Francisco J. 2009. "A Dynamic Model of Entrepreneurship with Borrowing Constraints: Theory and Evidence." *Annals of Finance* 5 (3): 443–64. <https://doi.org/10.1007/s10436-009-0121-2>.
- Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Ivan Fernandez-Val. 2018. "Generic Machine Learning Inference on Heterogenous Treatment Effects in Randomized Experiments." National Bureau of Economic Research.
- Cohen, Jessica, and Pascaline Dupas. 2010. "Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment *." *Quarterly Journal of Economics* 125 (1): 1–45. <https://doi.org/10.1162/qjec.2010.125.1.1>.
- Crépon, Bruno, Florencia Devoto, Esther Duflo, and William Pariente. 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 (1): 123–50. <https://doi.org/10.1257/app.20130535>.
- Davis, Jonathan M.V., and Sara B. Heller. 2017. "Using Causal Forests to Predict Treatment Heterogeneity: An Application to Summer Jobs." *American Economic Review* 107 (5): 546–50. <https://doi.org/10.1257/aer.p20171000>.
- . 2019. "Rethinking the Benefits of Youth Employment Programs: The Heterogeneous Effects of Summer Jobs." *The Review of Economics and Statistics*, June, 1–47. https://doi.org/10.1162/rest_a_00850.
- Druilhe, Z., and J. Barreiro-Huré. 2012. "Fertilizer Subsidies in Sub-Saharan Africa." *FAO ESA Working Paper* No 12-04.
- Duflo, Esther, Michael Kremer, and Jonathan Robinson. 2011. "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya." *American Economic Review* 101 (6): 2350–90. <https://doi.org/10.1257/aer.101.6.2350>.
- Eisenhauer, Philipp, Heckman James J., and Edward Vytlacil. 2015. "The Generalized Roy Model and the Cost-Benefit Analysis of Social Programs." *Journal of Political Economy* 123 (2): 413–43. <https://doi.org/10.1086/679498>.
- Evans, David S., and Boyan Jovanovic. 1989. "An Estimated Model of Entrepreneurial Choice under Liquidity Constraints." *Journal of Political Economy* 97 (4): 808–27. <https://doi.org/10.1086/261629>.
- Field, Erica, Rohini Pande, John Papp, and Natalia Rigol. 2013. "Does the Classic Microfinance Model Discourage Entrepreneurship Among the Poor? Experimental Evidence from India." *American Economic Review* 103 (6): 2196–2226. <https://doi.org/10.1257/aer.103.6.2196>.
- Fink, Günther, B. Kelsey Jack, and Felix Masiye. 2018. "Seasonal Liquidity, Rural Labor Markets and Agricultural Production." Working Paper 24564. National Bureau of Economic Research. <https://doi.org/10.3386/w24564>.
- Giné, Xavier, and Dean S. Karlan. 2014. "Group versus Individual Liability: Short and Long Term Evidence from Philippine Microcredit Lending Groups." *Journal of Development Economics* 107 (March): 65–83.
- Heckman, James. 1992. "Randomization and Social Policy Evaluation." In *Evaluating Welfare and Training Programs*, edited by Charles F. Manski and Irwin Garfinkel, 201–30. Cambridge, Mass: Harvard University Press.

- . 1997. “Instrumental Variables: A Study of Implicit Behavioral Assumptions Used in Making Program Evaluations.” *The Journal of Human Resources* 32 (3): 441. <https://doi.org/10.2307/146178>.
- Heckman, James J. 2010. “Building Bridges between Structural and Program Evaluation Approaches to Evaluating Policy.” *Journal of Economic Literature* 48 (2): 356–98. <https://doi.org/10.1257/jel.48.2.356>.
- Heckman, James J., Hidehiko Ichimura, and Petra E. Todd. 1997. “Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme.” *The Review of Economic Studies* 64 (4): 605–54. <https://doi.org/10.2307/2971733>.
- Hussam, Reshmaan, Natalia Rigol, and Benjamin N. Roth. 2020. “Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design In The Field.” *Working Paper*.
- Imbens, Guido W., and Joshua D. Angrist. 1994. “Identification and Estimation of Local Average Treatment Effects.” *Econometrica* 62 (2): 467. <https://doi.org/10.2307/2951620>.
- Jack, B Kelsey. 2013. “Private Information and the Allocation of Land Use Subsidies in Malawi.” *American Economic Journal: Applied Economics* 5 (3): 113–35.
- Kaboski, Joseph P., and Robert M. Townsend. 2011. “A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative.” *Econometrica* 79 (5): 1357–1406. <https://doi.org/10.3982/ECTA7079>.
- Karlan, Dean, and Jonathan Morduch. 2009. “Access to Finance.” In *Handbook of Development Economics*, edited by Dani Rodrick and M. R. Rosenzweig. Vol. 5. Elsevier.
- Karlan, Dean, and Sendhil Mullainathan. 2007. “Rigidity in Microfinancing: Can One Size Fit All?” *QFinance*, December. <http://www.qfinance.com/financing-best-practice/rigidity-in-microfinancing-can-one-size-fit-all?page=1>.
- Karlan, Dean, Isaac Osei-Akoto, Robert Darko Osei, and Christopher R. Udry. 2013. “Agricultural Decisions after Relaxing Credit and Risk Constraints.” *Quarterly Journal of Economics*, *Forthcoming*. <https://doi.org/10.2139/ssrn.2169548>.
- Karlan, Dean, and Jonathan Zinman. 2011. “Microcredit in Theory and Practice: Using Randomized Credit Scoring for Impact Evaluation.” *Science* 332 (6035): 1278–84. <https://doi.org/10.1126/science.1200138>.
- Kazianga, Harounan, and Christopher Udry. 2006. “Consumption Smoothing? Livestock, Insurance and Drought in Rural Burkina Faso.” *Journal of Development Economics* 79 (2): 413–46. <https://doi.org/10.1016/j.jdeveco.2006.01.011>.
- Maitra, Pushkar, Sandip Mitra, Dilip Mookherjee, and Sujata Visaria. 2020. “Decentralized Targeting of Agricultural Credit Programs: Private versus Political Intermediaries.” *National Bureau of Economic Research Working Paper* 26730 (February). <https://doi.org/10.3386/w26730>.
- Meager, Rachael. 2018. “Aggregating Distributional Treatment Effects: A Bayesian Hierarchical Analysis of the Microcredit Literature.” *LSE Working Paper*.
- . 2019. “Understanding the Average Impact of Microcredit Expansions: A Bayesian Hierarchical Analysis of Seven Randomized Experiments.” *American Economic Journal: Applied Economics* 11 (1): 57–91. <https://doi.org/10.1257/app.20170299>.

- Mel, Suresh de, David McKenzie, and Christopher Woodruff. 2008. "Returns to Capital in Microenterprises: Evidence from a Field Experiment." *Quarterly Journal of Economics* 123 (4): 1329–72.
- Moll, Benjamin. 2014. "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?" *American Economic Review* 104 (10): 3186–3221. <https://doi.org/10.1257/aer.104.10.3186>.
- Quidt, Jonathan de, Thiemo Fetzer, and Maitreesh Ghatak. 2012. "Group Lending Without Joint Liability." *London School of Economics Working Paper*.
- Rubin, Donald B. 1974. "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies." *Journal of Educational Psychology* 66 (5): 688–701. <https://doi.org/10.1037/h0037350>.
- Suri, Tavneet. 2011. "Selection and Comparative Advantage in Technology Adoption." *Econometrica* 79 (1): 159–209. <https://doi.org/10.3982/ECTA7749>.
- Tarozzi, Alessandro, Jaikishan Desai, and Kristin Johnson. 2015. "The Impacts of Microcredit: Evidence from Ethiopia." *American Economic Journal: Applied Economics* 7 (1): 54–89. <https://doi.org/10.1257/app.20130475>.
- Tarozzi, Alessandro, Aprajit Mahajan, Brian Blackburn, Dan Kopf, Lakshmi Krishnan, and Joanne Yoong. 2014. "Micro-Loans, Insecticide-Treated Bednets, and Malaria: Evidence from a Randomized Controlled Trial in Orissa, India." *American Economic Review* 104 (7): 1909–41. <https://doi.org/10.1257/aer.104.7.1909>.
- Tibshirani, Julie, Susan Athey, R. Friedberg, V. Hadad, L. Miner, Stefan Wager, and M. Wright. 2018. "Grf: Generalized Random Forests (Beta)." *ArXiv:1610.01271 [Econ, Stat]*. <https://github.com/grf-labs/grf> R package version 0.10.2.
- Wager, Stefan, and Susan Athey. 2018. "Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests." *Journal of the American Statistical Association* 113 (523): 1228–42. <https://doi.org/10.1080/01621459.2017.1319839>.

Appendix A1: Causal forest estimates

We implement a generalized causal forest to estimate conditional average treatment effects (CATE) at the observation level. This method has two clear advantages over standard linear regression methods. First, it allows the researcher to consider a relatively high-dimensional set of observable characteristics that may influence the effectivity of the treatment. Second, it accounts for the potentially non-linear relationship between the treatment effect and the predictors.

The causal tree algorithm of Athey and Imbens (2016) selects splits in order to maximize heterogeneity in treatment effects across leaves, less a penalty for the variance of treatment and control outcomes in each leaf. They propose an “honest” approach for estimation, using only one half of the sample (the training sample) to determine and cross-validate the splits. Then, each observation in the second half of the sample (the estimation sample) is assigned to a terminal leaf according to its observable characteristics, and the predicted CATE are calculated as the difference between the mean outcomes of treatment and control observations within each terminal leaf.

Wager and Athey (2018) builds on this method and propose a causal forest algorithm that assigns each individual observation the average of its predicted CATE across a large number of trees. Under this approach, each tree is estimated through the honest method described above, but using only a random sub-sample drawn without replacement. Only a random fraction of the available covariates is made available when determining each split.

We employ the generalized causal forest method proposed by Athey et al. (2019), which adapts the Generalized Random Forests method to the estimation of CATE. The algorithm has two basic steps. First, a causal forest is grown (with each tree based on a random sub-sample of the data, which is then split in half into a training sample to define leaves and an estimation sample to calculate CATE). Second, individual CATE is estimated using a set of kernel-based weights for all other observations in the sample. These weights are derived from the fraction of trees where each observation in the sample falls in the same terminal leaf as the target observation i .

Implementation

i. Preparing the dataset

Our sample for the estimation of the causal forests consists of all observations present at both the baseline and the first follow-up rounds of surveys. We estimate a different causal forest for the no-loan villages and the loan villages. The covariates are baseline net revenue, an indicator for the presence of an extended household, per capita food and non-food consumption, the value

of livestock owned, area of land cultivated, the value of agricultural assets owned by the household, the total days of labor used, and the index of social capital.

ii. The algorithm

We implement the algorithm using the R package *grf* version 0.10.4 (Tibshirani et al. 2018). Following Athey and Wager (2019), we allow the algorithm to tune the parameters through cross-validation using the “R-learner” objective function for heterogeneous treatment effects. The parameters that are determined through this method are the number of variables considered during each split, minimum node size, the fraction of the sample drawn for the construction of each tree, the percentage of observations assigned to the training and the estimations samples, the split balance parameters, and whether empty leaves are pruned from the estimated trees.

Regarding the number of trees in the forest, the literature suggests that it should be as large as possible given computational constraints. Hence, we tested different values for this parameter and noticed that the correlation between the predictions across different pairs of random seeds increases slightly with the number of trees in the forest until reaching 100,000 trees, after which it stabilizes. We verified that increasing the number of trees to 250,000, 500,000 or even 1,000,000 does not lead to meaningful changes in the distribution of the predictions or their stability. Therefore, we use 250,000 trees. The correlation between the predictions generated by different random seeds was consistently above 0.9 in the no-loan sample, and above 0.99 in the loan sample.

Finally, we account for the fact that the observations in our sample are grouped in unevenly sized clusters (i.e., villages in our setting). In practice, this modifies the causal forest algorithm in two ways. First, the training and estimation samples for each tree are determined by selecting a random subset of clusters, and then drawing an equal number of observations from each cluster. Second, the out-of-bag predictions for each observation i are generated using only the trees where no observation in the training or estimation samples belongs to the same cluster as the target observation i .

Since some clusters in our study have a very small number of observations, we follow Athey and Wager (2019) and increase the number of observations to be drawn from each cluster for the training and estimation samples (the default is the size of the smallest cluster). This improves the stability of the tree-growing algorithm substantially, at the cost of using fewer observations from the clusters that are below this threshold. Considering that our sample size is relatively large compared to other field experiments, we decided to fix this parameter at the 25th percentile of the distribution of cluster sizes.

iii. Assessing treatment heterogeneity

In this subsection, we evaluate whether our generalized casual forest algorithm succeeded in identifying treatment heterogeneity. We conduct this analysis separately for the no-loan and loan villages.

Although the out-of-bag predictions from our model exhibit considerable variation, Athey and Wager (2019) warns that this does not necessarily rule out the possibility that the obtained estimates might just be noisy due to overfitting. Therefore, we follow Davis and Heller (2019) and compare how the predicted CATE relate with the actual treatment effects. First, we group the observations in each sample in 20 bins according to their predicted CATE. Then we calculate the treatment effect for each bin, following the same specification as in our main results (i.e., we control for net revenue at baseline, village fixed effects, and stratification controls that are listed in the notes to Table 2). Finally, we plot the resulting treatment effect versus the mean predicted CATE per bin in Appendix Figure 1.

To provide a more robust test for heterogeneity, we employ a calibration test motivated by the best linear predictor of CATE method of Chernozhukov et al. (2018). Consider the no-loan villages. Let $B(Z)$ be the random forest predictor of $b_0(Z) \equiv E(Q(0,0)|Z)$, so $B(Z_{ijt})$ is the prediction from the random forest of the net output of a random household with characteristics Z_{ijt} that does not receive a grant or a loan. Similarly, let $S(Z)$ be the causal forest predictor of $s_0(Z) \equiv E(Q(1,0) - Q(0,0)|Z)$, so $S(Z_{ijt})$ is the predicted CATE for a household with characteristics Z_{ijt} . The probability of randomization into the grant treatment is $p(Z)$. We estimate

$$Y_{ijt} = \alpha_0 + \alpha_1 B(Z_{ijt}) + \beta_1 (grant_i - p(Z_{ijt})) + \beta_2 (grant_i - p(Z_{ijt})) (S(Z_{ijt}) - \bar{S}) + S(Z_{ijt}) + \epsilon_{ijt}$$

by weighted least squares using weights $\left(p(Z_{ijt})(1 - p(Z_{ijt}))\right)^{-1}$. Chernozhukov et al. (2018) shows that rejecting the hypothesis that $\beta_2 = 0$ also rejects the hypothesis that there are no heterogeneous treatment effects, and implies that $S(Z)$ is a relevant predictor of that treatment effect heterogeneity.

In the loan villages, $B(Z)$ is the predictor of $b_l(Z) \equiv E(Q(0,0)|Z, B = 0)$, so $B(Z_{ijt})$ is the prediction from the random forest of the net output of a random household with characteristics Z_{ijt} that does not borrow when in a loan village, and who receives neither a grant nor a loan. Similarly, in the loan villages, $S(Z)$ is the causal forest predictor of $s_l(Z) \equiv E(Q(1,0) - Q(0,0)|Z, B = 0)$.

We find a coefficient for β_2 of -0.03 for the households in the no-loan sample, and a coefficient of 1.05 (p-value = 0.009) for the loan sample. We note that these findings are in line with the patterns observed in Appendix Figure 1. Overall, the results suggest that the algorithm succeeded

in finding meaningful heterogeneity for the loan sample. For the no-loan sample, on the other hand, the evidence is weak and inconclusive.