

Selection into Credit Markets: Evidence from Agriculture in Mali

August 2021

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 heterogeneous returns to capital. Understanding this 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% for would-be borrowers, whereas we find returns near zero for the sample representative of non-borrowers. We also provide evidence that there are some farmers – particularly those that are poor at baseline – that have high returns but do not receive a loan.

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. Paper previously circulated as “Self-selection into Credit Markets: Evidence from Agriculture in Mali”. 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 and Alex W. Cohen 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; some 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 demonstrate positive selection into borrowing with respect to marginal returns to capital. We designed a two-stage protocol specifically to test whether returns to capital are heterogeneous and sufficiently predictable that high return agents receive loans. 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 (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 of the experiment (the cash grant stage), after first waiting for households and the associations to make their loan decisions from stage one, 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 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. 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 to a grant than the implied returns to a grant among farmers who did borrow?

We find large average increases in investment and agricultural profits for the non-selected population (i.e., grant recipients vs. non-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 gross profit from relaxing capital constraints. This impact on gross profit even persists after an additional agricultural season. In this environment, therefore, 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, non-borrower households given grants did not earn any higher gross profit from the farm than non-borrowing households not provided grants. This contrasts sharply with households given grants in the no-loan villages who had large increases in gross profit 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 would borrow and those who do not are substantial. We calculate that among borrowing households, farm output would have increased by US\$222 (se=120) and farm gross profit by US\$183 (se=96) had those households received grants. In contrast, we estimate that among households who do not borrow, receipt of the grant generates only US\$25 of additional output and US\$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 non-borrowers did not invest in farming because they had higher return opportunities elsewhere. To examine this, we also look at other outcomes such as livestock ownership and

² We do not have a complete profit measure, and use instead the term “gross profit” for agricultural revenue net of most, but not all, expenses. Importantly, the value of family and unpaid labor is not subtracted. See section 2.3.

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. We therefore refer to capital market imperfections that could cause investment responses to cash grants simply as credit constraints.

small business operations. However, we do not find evidence that non-borrowers in loan villages invest the grant in alternative activities more than their counterparts in no-loan villages.

Farmers with high returns to grants are differentially selected into borrowing from Soro. But how efficient is this selection? In particular, are there identifiable women with high return investment opportunities who do not borrow? The two most salient concerns in our setting are that the womens' groups screen out potential borrowers based on their ability to repay; or that borrower risk aversion induces self-selection. We find evidence that within our selected sample of non-borrowers in loan villages, there are some households with high return projects but with low baseline levels of gross profits, food and non-food consumption, farm size and livestock holdings. In no-loan villages (thus a representative sample of everyone), returns to the grant are positively correlated, or uncorrelated with each of these baseline characteristics. However, in loan villages (thus only those selected *out* from borrowing, either by themselves or their peers), these same characteristics are negatively correlated with returns to the grant. We interpret these findings through the lens of a simple model. We conclude that the selection into borrowing of farmers with high return projects is more complete among farmers with high values of these baseline variables – i.e. nearly all of these high return farmers who receive a loan – than among farmers with low values of these baseline variables. Because these characteristics are plausibly associated with both a borrower's ability to repay and with her level of risk aversion, we cannot say if this excess selection is driven by borrower self-selection or by lender screening.

The heterogeneity in returns to loans that we discover is consistent with Meager (2020), which uses Bayesian hierarchical modeling of the quantiles of response to seven different microcredit interventions with RCTs to show evidence of strongly positive returns for a small set of borrowers, but near zero returns to borrowing for the large majority. Crépon et al. (2020) also finds a great deal of heterogeneity in the returns to loans (and grants) among microentrepreneurs in Egypt. Our finding that farmers are aware of these heterogenous returns is similar to that of Hussam et al. (2020), which finds that businesses (in their case, nonfarm enterprises in urban India) have widely varying marginal returns to grants, and that entrepreneurs themselves and community members are able to distinguish between those with relatively high and low returns. In a different setting (enterprise business plan competitions in Nigeria and in Ghana), McKenzie (2018; 2015), McKenzie and Sansone (2019), and Fafchamps and Woodruff (2017) provide evidence of the difficulty in predicting who will be the most successful, although average estimated returns are high.

Our experiment also speaks to three additional questions important to academia and policy: First, do loans generate different investment behavior than grants? Second, what is the impact of a

microlending program that targets farmers (as compared to the more standard microenterprise focus of microlenders)? Third, are the impacts of the cash grants persistent after seven years?

First, on comparing grants to loans, 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, 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 profit. 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 in magnitude 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. This is consistent with Crépon et al. (2020).

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 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 a lump sum once or twice a year (see Fink, Jack, & Masiye, 2018 for an experiment demonstrating the importance of this timing issue for farmers; see Karlan & Mullainathan, 2007 for a discussion of this). 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. Maitra et al. (2020) also finds positive impacts from an agricultural microcredit program on farm value-added in India for one version of the program, though not for a version which targeted the program differently. However, lending may not be sufficient to induce investments in the presence of other constraints. Farmers may be

⁴ The evidence from traditional microcredit, targeting micro enterprises, is more mixed; some randomized evaluations find an increase in investment in self-employment activity (Angelucci, Karlan, & Zinman, 2015; Crépon, Devoto, Duflo, & Pariente, 2015) while others do not (Attanasio, Augsburg, De Haas, Fitzsimons, & Harmgart, 2015; Augsburg, De Haas, Harmgart, & Meghir, 2015; Banerjee, Duflo, Glennerster, & Kinnan, 2015; Karlan & Zinman, 2011; Tarozzi, Desai, & 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, & Woodruff, 2008).

constrained by a lack of insurance (Karlan, Osei-Akoto, Osei, & Udry, 2013), have time inconsistent preferences (Duflo, Kremer, & 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 subsidized government programs and thus plagued by politics (Adams, Graham, & Von Pischke, 1984). In the expansion of microcredit in the 1980s and onward, we had seen several mostly simultaneous shifts: group instead of individual lending (de Quidt, Fetzer, & Ghatak, 2012; although now this trend is reversing, e.g. see Giné & Karlan, 2014); high frequency repayment instead of one-time balloon payments (see Field, Pande, Papp, & Rigol, 2013 for an important test, demonstrating the potential benefits to delayed-start repayment); nongovernment (and now for-profit) lending instead of government; and, enterprise targeted loans instead of agricultural (Armendariz de Aghion & Morduch, 2010; Karlan & Morduch, 2009). The loan impact component of this study tests a new model of agricultural credit with group lending, balloon payment, and nonprofit management (with little to no subsidy).

Third, we conducted a follow-up survey in 2017, almost seven years after the grants, to measure their long-term effects. We find no evidence that the grants had a persistent effects over this extended period, which was marked by political upheaval and systematic changes in cropping patterns, as well as the highly variable seasonal rainfall 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 rain fed. Evidence from nearby Burkina Faso suggests that income shocks translate into consumption volatility (Kazianga & Udry, 2006), so improved credit markets can have important welfare consequences from both increasing average production and insulating consumption from output volatility. The main crops grown in the area include millet/sorghum, maize, cotton (mostly grown by men), and 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).

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

The sample consists of 198 villages identified by Soro as villages that they had not previously entered but that were within their expansion plans. These are villages in which households have limited access to formal financial institutions: only 5% of households report receiving a formal loan at baseline.⁶ 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.⁷

Figure 1 presents 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.⁸ This stage one randomization was done at the village level (because that is how Soro marketed and implemented loans).

Soro offered their standard agricultural loan product, called *Prêt de Campagne*, in the 88 loan villages. There was no screening of the villages by Soro: loans were offered to women's

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 & Barreiro-Huré, 2012).

⁶ There is, instead, informal sources of financing and risk sharing through bilateral exchanges between households and through community savings and loans associations (ASCAs) and cooperatives. About 50% of households reported receiving a loan from family or friends at baseline.

⁷ Bougouni and Yanfolila are the two *cercles*, both in the northwest portion of the region and within the expansion zone of 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 being serviced by Soro Yiriwaso, and (3) contained at least 350 individuals (i.e., sufficient population to generate a lending group).

⁸ First, we ran a loop with a set number of iterations that 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 the following: village size, whether the village was all Bambara (the dominant ethnic group in the area), distance to a paved road, distance to the nearest market, percent of households with a plough, percent of women with a plough, frequency of fertilizer use among women in the village, literacy rate, and 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 the following: whether the household was part of an extended family; whether the household was polygamous; an index of the household's agricultural assets, other assets, and per capita food consumption; and, the primary female respondent's land size, fertilizer use, and plough access. See Bruhn and McKenzie (2009) for a more detailed description of the randomization procedure.

associations which were formed by the women in the village for the purpose of borrowing in each of the 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. 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. Soro itself was not directly involved in selecting who would receive a loan. 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 nominally there is 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. We observe no defaults over the two agricultural cycles during which we were collaborating with Soro.⁹

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).¹⁰ The annual interest rate is 25% plus 3% in fees and a mandatory savings rate¹¹ of 10%.

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

Stage Two: Grants

⁹ 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 (more than 30 days overdue) in years 2004–2006, rising to only 0.7% in 2007.

¹⁰ We use the 2011 PPP exchange rate with the Malian FCFA (284 FCFA per USD) throughout the paper.

¹¹ The mandatory savings are removed from the loan at the time of disbursement and held at the MFI. This deposit requirement may serve as part of a screening mechanism based on wealth or liquidity, as discussed in section 3.3

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—to parallel the loans—a female household member 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 approximate the average loan size provided by Soro, though *ex post* the grant is slightly larger on average than the 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.¹²

In the 88 loan villages, grant recipients were randomly selected among survey respondents who did not take out a loan (see Figure 1).¹³ 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 difference of 20-days between when grants were received by households in no-loan villages and 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 not to borrow, we informed recipients that the grants were a one-time grant, not an ongoing program, and also distributed an additional 104 grants (one or two per village) to loan village women not in our sample. It was therefore not obvious to those in the study that borrowing precluded someone from being a grant recipient.

2.2 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 season¹⁴ in January–

¹² 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.

¹³ 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.

¹⁴ We also conducted an “input survey” on a sub-sample 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

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. The four rounds used similar survey instruments, which 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 both the individual and the household level.

Throughout the paper we refer to “gross profit” 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 either the value of unpaid labor (own, family or other) or the implicit rental value of land used, because 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.¹⁵ 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 expenditure. It is important to note that there is a lot of consumption seasonality in Mali (Beaman, Karlan, & Thuysbaert, 2014). Our measure of food expenditure reflects consumption in the post-harvest season only.

2.3 Randomization Balance Check and Attrition

We conduct several tests to verify that we cannot reject the orthogonality of treatment assignment to baseline characteristics and attrition. Appendix Table 1 examines baseline characteristics across three comparisons: (i) loan to no-loan 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 statistically significantly different across the three comparisons, and an aggregate test in which we regress assignment to treatment on the set of 11 covariates fails

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 sub-sample of 2,400 households. We use the input survey if conducted, and we use the end of season survey if not. We also control for timing of the collection of the data in all relevant specifications.

¹⁵ The survey instruments are all available upon request.

to reject orthogonality for each of the three comparisons (p-value of 0.26, 0.91 and 0.67, respectively, reported at the bottom of the table).

Our attrition rate is low at approximately one percent each round.¹⁶

3. Identification

We focus on agricultural outcomes. Let $\{Q^{NG}, Q^G, Q^B\}$ represent the set of potential gross profits in year 1 of households in our sample, where Q^{NG} is a random variable representing potential profit if the household neither borrows nor receives a grant, and Q^G and Q^B are similarly defined for households who receive a grant but not a loan, and for those who receive a loan but not a grant, respectively.¹⁷ The joint distribution of potential outcomes is $F(Q^{NG}, Q^G, Q^B)$, and the three marginal distributions are denoted $F_{NG}(Q^{NG})$, $F_G(Q^G)$ and $F_B(Q^B)$.

Define $G \in \{0,1\}$ and $L \in \{0,1\}$ as random variables that designate a household's status in the grant treatment arm and in a loan treatment village, respectively. Not all women in loan communities borrow. Define $C = 1$ (for *complier*) if the household would borrow if its village is a loan village, and $C = 0$ if it would not borrow if located in a loan village.

Therefore, we can write a binary indicator of borrowing as

$$B = CL. \tag{1}$$

Furthermore, define the effect on profit of receiving a grant as $\Delta_G Q \equiv Q^G - Q^{NG}$. We seek to identify the expected value of the effect on profit of receiving a grant for households for which $C = 1$ versus those for which $C = 0$, that is $\mathbb{E}(\Delta_G Q | C = 1)$, and $\mathbb{E}(\Delta_G Q | C = 0)$. The two-stage randomization provides identification of these expected treatment effects.

3.1 Returns to grants for borrowers and non-borrowers

The first stage randomization of villages ensures

¹⁶ Despite the low attrition rate, we report differential attrition tests in Appendix Table 2. We compare 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, we fail to reject orthogonality for all six regressions (results on bottom row of Appendix Table 2).

¹⁷ 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).

$$\{Q^{NG}, Q^G, Q^B, C\} \perp L \quad (2)$$

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

$$\{Q^{NG}, Q^G, Q^B, C\} \perp G | L = 0 \quad (3)$$

$$\{Q^{NG}, Q^G\} \perp G | (C = 0, L = 1) \quad (4)$$

There is 100% take-up of the offer of a grant, so in our sample of the full population of no-loan villages and in our sample of non-borrowers in loan villages, we observe

$$Q = Q^G G + Q^{NG} (1 - G)$$

Therefore, (2) and (3) imply that data from grant recipients in no-loan villages can be used to identify the marginal distribution of Q^G in the population, in both loan and no-loan villages:

$$\begin{aligned} F_G(Q^G | L = 0, G = 1) &= F_G(Q^G | L = 0, G = 0) = F_G(Q^G | L = 1, G = 1) \\ &= F_G(Q^G | L = 1, G = 0) = F_G(Q^G) \end{aligned} \quad (5)$$

Similarly, (2) and (3) imply that data from non-grant recipients in no-loan villages identifies the marginal distribution of Q^{NG} (dropping the intermediate equalities for brevity in (6) and (8))

$$F_{NG}(Q^{NG} | L = 0, G = 0) = F_{NG}(Q^{NG}) \quad (6)$$

(2) and (4) imply that data from non-borrowers in loan villages can be used to identify the conditional (on $C = 0$) marginal distributions of the profits of those who receive and do not receive a grant, respectively:

$$\begin{aligned} F_G(Q^G | C = 0, L = 1, G = 1) &= F_G(Q^G | C = 0, L = 1, G = 0) \\ &= F_G(Q^G | C = 0, L = 0, G = 1) = F_G(Q^G | C = 0, L = 0, G = 0) \\ &= F_G(Q^G | C = 0) \end{aligned} \quad (7)$$

$$F_{NG}(Q^{NG} | C = 0, L = 1, G = 0) = F_{NG}(Q^{NG} | C = 0) \quad (8)$$

The loan village sample provides an estimate of $\mathbb{P}(C = 1 | L = 1)$, which with (2) identifies the share of compliers in the population $\mathbb{P}(C = 1 | L = 1) = \mathbb{P}(C = 1 | L = 0) = \mathbb{P}(C = 1)$. Therefore, we have identified the marginal distributions of profits for grant recipients and non-recipients among the selected population of those who would borrow:

$$F_G(Q^G | C = 1) = \frac{F_G(Q^G) - F_G(Q^G | C = 0)(1 - \mathbb{P}(C = 1))}{\mathbb{P}(C = 1)} \quad (9)$$

$$F_{NG}(Q^{NG}|C = 1) = \frac{F_{NG}(Q^{NG}) - F_{NG}(Q^{NG}|C = 0)(1 - \mathbb{P}(C = 1))}{\mathbb{P}(C = 1)}$$

With these marginal distributions identified, we can calculate the average effects of receiving a *grant* 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.

$$\begin{aligned} \mathbb{E}(Q^G) - \mathbb{E}(Q^{NG}) &\equiv \mathbb{E}(\Delta_G Q) \\ \mathbb{E}(Q^G|C = 0) - \mathbb{E}(Q^{NG}|C = 0) &\equiv \mathbb{E}(\Delta_G Q|C = 0) \\ \mathbb{E}(Q^G|C = 1) - \mathbb{E}(Q^{NG}|C = 1) &\equiv \mathbb{E}(\Delta_G Q|C = 1) \end{aligned} \quad (10)$$

We provide estimates of these three expectations in section 4.2 by estimating

$$Y_i = \alpha + \beta_1 grant_i + \beta_2 grant_i \cdot loan_{v(i)} + X_i \pi + \lambda_{v(i)} + \epsilon_i \quad (11)$$

where $\hat{\beta}_1$ is our estimate of $\mathbb{E}(\Delta_G Q)$ and $\hat{\beta}_1 + \hat{\beta}_2$ is our estimate of $\mathbb{E}(\Delta_G Q|C = 0)$ when the outcome Y_i is gross profit of farmer i in village $v(i)$, $\lambda_{v(i)}$ is a village fixed effect and X_i is a vector of baseline controls to be discussed below.

3.2 Average return to borrowing

Similarly, define the effect on profit of borrowing as $\Delta_B Q \equiv Q^B - Q^{NG}$. We also identify the expected treatment effect of borrowing on those who would borrow if loans were available: $\mathbb{E}(\Delta_B Q|C = 1)$. (2) implies that data from the population of borrowers in loan villages can be used to identify the conditional marginal distribution

$$F_B(Q^B|C = 1, L = 1) = F_B(Q^B|C = 1, L = 0) = F_B(Q^B|C = 1) \quad (12)$$

We have already noted that (2), (3) and (4) imply that $F_{NG}(Q^{NG}|C = 0)$ is identified from data on the profits of non-borrowers who do not receive a grant in loan villages, and shown in (9) that combining this with estimates of $\mathbb{P}(C = 1)$ and $F_{NG}(Q^{NG})$ identify $F_{NG}(Q^{NG}|C = 1)$. Thus we can identify the average treatment effect on the treated (TOT) of borrowing:

$$\mathbb{E}(Q^B|C = 1) - \mathbb{E}(Q^{NG}|C = 1) \equiv \mathbb{E}(\Delta_B Q|C = 1). \quad (13)$$

Note that we needed no assumption about whether farmers make the same investment decisions with a loan as they would with a grant in order to identify either (10) or (13).

3.3 Selection and the efficiency of the allocation of capital

Can we compare the selection into borrowing that we have identified with what would be optimal? Our experimental design does not allow us to do this directly. This section provides a theoretical framework that will guide our empirical tests of efficient allocation of capital.

We define optimal as the allocation of loans such that aggregate gross profits are maximized. Suppose that the opportunity cost of funds (of a fixed loan size) to the lender is ρ . Aggregate gross profits are maximized if all households with $\Delta_B Q \geq \rho$ borrow, while other households do not. However, in an environment of imperfect enforcement, incomplete information and uninsured risk, there may be potential borrowers with projects that could generate high returns, but who do not receive loans. Screening by the lender, self-selection, or both could drive this “excess selection”.

Among these frictions, the two most salient in our setting are (i) lender (more specifically, women’s group) screening on ability to repay and (ii) borrower risk aversion. In Appendix A1, we present two simple canonical models to provide guidance as to why certain high expected return borrowers do not take loans, and how excess selection can be detected in our setting. In the first, poorer or less collateralized potential borrowers with high marginal returns may be unable to make a credible repayment commitment. In the second, risk aversion may deter poorer or more risk averse farmers with high-expected return projects from borrowing for those projects. In both cases, the frictions imply that there will be non-borrowers for which their marginal return exceeds the opportunity cost of funds, and that the extent of this wedge decreases as a farmer’s baseline gross profits, collateral, or wealth increase. Our empirical tests that the allocation of loans maximizes profit are based on these common implications of the two models; we do not, therefore, distinguish between self-selection (based on risk aversion) and lender-selection (based on limited liability) as the source of the frictions that result in an inefficient allocation of loans.

We illustrate the basic predictions of the models in Figure 3. The efficient allocation is depicted in the left panel of Figure 3: the horizontal curve E defines the boundary in $(Q^{NG}, \Delta_B Q)$ between those who borrow and those who do not in a profit-maximizing allocation assigning credit exclusively to all farmers with a sufficiently profitable investment opportunity. A farmer i with values of $(Q_i^{NG}, \Delta_B Q_i)$ in the region $C = 1$ does not borrow because her returns are too low; her no-grant level of profits is irrelevant to the allocation. In panel B, the curve F defines the boundary in an allocation constrained by limited liability. The set of values of $(Q_i^{NG}, \Delta_B Q_i)$ such that a farmer does not borrow expands due to the friction. The dashed curve in Panel C depicts the boundary in the allocation in the presence of farmers’ decreasing absolute risk aversion (DARA). With either friction, the wedge between the lender’s cost of funds and the farmer’s

required expected return from the loan (weakly) decreases with the no-grant gross profit of the farmer. The wedge exists when a limited liability constraint binds, but this constraint is relaxed by increases in no-grant gross profits. Similarly, if a farmer has decreasing absolute risk aversion, then the expected return from borrowing she requires to accept the additional risk associated with borrowing declines with her no-grant gross profit.

Sections 3.1 and 3.2 demonstrated that our experimental design gives us clean identification of the returns to grants and the returns to loans. To evaluate the efficiency of the allocation of credit, however, we must consider the relationship between the two. In an efficient allocation, $\Delta_G Q = \Delta_B Q$, because both maximize profits. However, risk aversion generates selection across projects of a farmer as well as across farmers. The project chosen by a risk averse borrower who is given a grant will have an expected return (weakly) greater than the project that that farmer would have chosen to implement with the loan. Figure 3C also illustrates this: the (solid line) boundary in $(E(Q^{NG}), E(\Delta_G Q))$ between those who borrow and those who do not lies above that (dashed line) boundary in $(E(Q^{NG}), E(\Delta_B Q))$, and with DARA preferences the difference between the boundaries declines as $E(Q^{NG})$ rises. The key takeaway is that if we observe farmers in the non-borrower sample who demonstrate that they have high return projects (from their returns to the grants), we have evidence of excess selection. We also test empirically the hypothesis that the expected agricultural returns to grants for those who would borrow are equal to the expected agricultural returns to a loan for those who do borrow in section 6, but our interpretation of the evidence does not rely on farmers choosing the same projects in the loan versus grant treatment arms.

We take two complementary approaches to investigate empirically the extent to which there is excess selection out of borrowing by poor households with high return projects. First, if we assume that treatment effects of the grants are monotonic, the comparison of the gap between the distributions of profits of grant recipients and non-grant recipients in the no-loan villages with the analogous gap in the selected sample of non-borrowers in loan villages is informative that borrowing frictions exist. At high enough levels of non-grant gross profits in the loan villages, the only non-borrowers eligible to receive grants would be those farmers without high-return projects. Thus there will be small differences between gross profits of grant recipients and non-recipients, for sufficiently high Q^{NG} . If at low levels of non-grant profits Q^{NG} , even farmers with high return projects are unable to borrow, then the distribution of gross profits for grant recipients will be shifted to the right compared to that of non-recipients in loan villages, similar to the pattern in the no-loan villages. Therefore, we expect $F_{NG}(Q|C=0) - F_G(Q|C=0) - (F_{NG}(Q) - F_G(Q))$ to be declining in Q if there is excess selection. We examine these marginal distributions in section 5.

Second, we relax the monotone treatment effect assumption. Excess selection can be distinguished from efficient selection via their different implications for heterogeneous treatment effects. We start by noting a simple extension of (10) that implies that we can estimate the conditional (on any observable characteristic X^k) average treatment effects $\mathbb{E}(\Delta_G Q | X^k = x^k)$ and $\mathbb{E}(\Delta_G Q | X^k = x_i^k, C = 0)$. We estimate linear approximations to these conditional expectations in section 5 via the regression

$$Y_i = \alpha + \beta_1 grant_i + \beta_2 grant_i \cdot loan_{v(i)} + \gamma_1 grant_i \cdot X_i^k + \gamma_2 grant_i \cdot X_i^k \cdot loan_{v(i)} + X_i \pi + \lambda_{v(i)} + \epsilon_i \quad (14)$$

From this regression we construct $\hat{\beta}_1 + \hat{\gamma}_1 \cdot X_i^k$ as our estimate of $\mathbb{E}(\Delta_G Q | X^k = x_i^k)$, and $\hat{\beta}_1 + \hat{\beta}_2 + (\hat{\gamma}_1 + \hat{\gamma}_2) \cdot x_i^k$ as our estimate of $\mathbb{E}(\Delta_G Q | X^k = x_i^k, C = 0)$ when the dependent variable Y is gross profits Q .

If selection into borrowing is efficient, then conditional on any observed characteristic, the average return to grants should be higher in no-loan villages than in loan villages, because the distribution of $\Delta_G Q | X^k$ is truncated from above at ρ in the loan villages. Efficient selection also has implications for patterns of heterogeneity. Suppose that $\gamma_1 > 0$: that is, along observable dimension k , expected returns to grants are increasing in X^k in the no-loan villages (the argument is symmetric around zero). We show in appendix A1 that expected returns to grants in loan villages must also be increasing in X^k , but that the slope is attenuated towards zero. In loan villages, as X^k increases, the selection into non-borrowing becomes more severe due to the corresponding increase in expected returns, partially offsetting the increase in expected returns to the grant among non-borrowers in loan villages. So $\gamma_1 > \gamma_1 + \gamma_2 > 0$. Similarly, if $\gamma_1 < 0$, then $\gamma_1 < \gamma_1 + \gamma_2 < 0$.

With excess selection, it remains the case that conditional on any observed characteristic, the average return to grants should be higher in the no-loan villages. But the frictions generating excess selection imply that a borrowing farmer's marginal return to a loan exceeds ρ , and we define $\rho + h(X^k)$ to be the minimum marginal return required for a farmer with observed characteristic X^k to borrow. We choose X^k such that the hypothesized friction is declining in X^k : this wedge decreases as a farmer's no-grant gross profits, collateral, or wealth increase. There is excess selection if borrowing is determined by $\Delta_G Q_i > \rho + h(X^k)$ with $h(X^k) > 0$ and decreasing in X^k . Excess selection always reduces the slope of the relationship between average returns to grants and any X^k that is negatively correlated with borrowing frictions. If $\gamma_1 > 0$ (expected returns to the grant are increasing in X^k in the random sample) then $\gamma_1 > \gamma_1 + \gamma_2$. Recall that in the case of efficient selection, this effect could only attenuate the heterogeneity.

By contrast, if excess selection is sufficiently strong, the sign of the slope can change so that $\gamma_1 + \gamma_2 < 0 < \gamma_1$.¹⁸

We also examine the joint and potentially non-linear effects of a vector of baseline observables X that might be associated with efficient or excess selection. We implement a causal forest algorithm (Athey et al., 2019; Athey and Wager, 2019) to estimate conditional average treatment effects (CATEs) flexibly (see Appendix A2 for details on the methodology). We use the algorithm trained on no-loan villages to predict the CATE of a grant for farmer j , $E(\Delta_G Q | X = x_j)$. We use the algorithm trained on non-borrowers in loan villages to estimate the CATE of a grant for non-borrowing farmer i , $E(\Delta_G Q | X = x_i, c_i = 0)$. Efficient and excess selection into borrowing have the same observable implications for the relative slopes of these CATEs estimated in no-loan and loan villages.

We examine the hypothesis that the selection we observe is efficient in section 5 by focusing on a series of observable characteristics that are plausibly correlated with the salient borrowing frictions of ability to repay and borrower risk aversion (baseline gross profits, livestock ownership, food consumption or non-food expenditure at baseline¹⁹) and by using the causal forest algorithms.

4. Selection into loans and the return to cash grants

4.1 Observable characteristics of borrowers versus non-borrowers in loan villages

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 turnover in clients. 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, & Zinman, 2015; Attanasio, Augsburg, De Haas, Fitzsimons, & Harmgart, 2015; Banerjee, Duflo, Glennerster, & Kinnan, 2015; and for a summary discussion of these studies, see Banerjee, Karlan, & Zinman, 2015; Crépon, Devoto, Duflo, & Pariente, 2015; Tarozzi, Desai, & Johnson, 2015).

¹⁸ Similarly, if $\gamma_1 < 0$, $\gamma_1 + \gamma_2 < \gamma_1 < 0$ with sufficiently strong excess selection.

¹⁹ We attempted to measure risk aversion in the baseline survey. However, the data is very noisy and feedback from the field suggests the survey respondents did not understand well the questions. As seen in Table 1, loan takeup is not correlated with our measure of risk aversion, and poor quality data is one possible reason. We therefore focus on proxies of risk aversion: variables which would suggest households are close to subsistence.

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, and have higher agricultural output and gross profits are more likely to take out a loan. Borrowers also have more agricultural assets and livestock. The causal forest algorithm trained on data from no-loan villages provides estimates of the conditional average effect of the grant treatment (CATEs) given baseline characteristics of a household. We apply that model to all households in the loan villages to obtain predicted treatment effects for borrowers and non-borrowers. The final row of Panel A of Table 1 shows that households that borrow have higher predicted CATEs from the grant treatment than do non-borrowing households. 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.²⁰ Households that borrow also have higher consumption at baseline than non-clients.

4.2 Experimental results on returns to grants in loan and no-loan villages

Next, we present the estimated returns to receiving a grant 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 (equation 13). To isolate the role of selection into loans, we focus on the first year of the experiment.²¹ Table 2 shows the estimates from regression (11) using the first follow-up data on farm investments and output. Loan recipients are removed from the analysis sample. The baseline controls (X), which include the baseline value of the

²⁰ 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 the woman 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 the 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, is in the same organization, would engage in informal risk-sharing and transfers with the person, and topics of their discussions (if any).

²¹ 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 7.

dependent variable y_0 ²² and the baseline variables used in the re-randomization routine (listed in the notes of table 2). Standard errors are clustered at the village level. We also provide randomization inference p -values (Young, 2019) that account for both the re-randomization routine used to assign treatment status and multiple comparisons within families of outcomes (details discussed in table notes).

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 non-cultivation outcomes such as livestock, enterprise, consumption, and female empowerment.

4.2.1 Agriculture

Columns (1)–(8) of Table 2 examine agricultural inputs and crop choice. We first focus on the first row of coefficients, β_1 , which captures the impact of the grant in no-loan villages. We find that households who received a grant in no-loan villages cultivated more land than those who did not (0.18 ha, se=0.07). This is approximately an 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 growing rice and groundnuts, which are cash crops in the area. The grant also induced an increased use of the plough (6 percentage points, se=1), 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 was only 17 days. We observe no change in family labor. Fertilizer and other chemical inputs increased by 18% (US\$21, se=6). The agricultural inputs and crop choice variables in columns (1)–(7) are grouped together as a family of outcomes for the randomization-c p values (Young 2019). The adjusted p -values are qualitatively similar to our simple tests. Moreover, the omnibus test indicates a statistically significant ($p<.001$) experimental effect.

The grants led to an overall increase in agricultural investment: column (8) shows that measured input expenses increased by US\$30 (se=8).²³ Columns (9)–(10) report statistically significant and economically meaningful increases in output and gross profits: output increased by US\$66

²² 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.

²³ The 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.

(se=19) and gross profits increased by US\$39 (se=16), equivalent to 13% and 12% increases, respectively. Overall, we see statistically significant increases in investments and ultimately gross profits from relaxing capital constraints.²⁴

Critically, the coefficient on Grant * Loan village (β_2) demonstrates heterogeneity in the returns to the cash grant between households in no-loan villages and non-borrower households in loan villages. The β_2 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 non-borrower households in loan villages 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 -US\$15, se=9, respectively). Non-borrower 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). Column (8) shows that total input expenses among nonborrowers in loan villages increase in response to the grant by about US\$20 (p-value=0.02), which is not statistically different from the estimate in no-loan villages of US\$29. Note, however, that the inputs that are measured with the most precision—fertilizer and chemical expenses in column (7)—demonstrate a statistically significant difference in the impact of the grant on investment choices between loan and no-loan villages.

However, even though we observe increased inputs for the (non-borrower) grant recipients in loan villages, we see no corresponding increase in either agricultural output or in gross profits. The β_2 interaction coefficient for output is similar in magnitude to β_1 but *negative* (-US\$41, se=28), offsetting the increase in output in no-loan villages (US\$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, the total effect on gross profits among non-borrowers in loan villages is essentially zero (US\$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—in stark contrast to the

²⁴ 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) shows that labor inputs 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.

random sample of households in no-loan villages—no evidence of increases in either agricultural output or gross profits.

These point estimates imply that there is important 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 US\$131 (se=68) in additional gross profits per US\$100 of grant.²⁵ 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 (Ashraf, Berry, & Shapiro, 2010; Cohen & Dupas, 2010; Tarozzi et al., 2014), we find that the repayment liability leads to lower return households being screened out. Appendix A3 explores this in depth, and demonstrates that we are unable to predict either the returns to the grants or the heterogeneity in returns using baseline characteristics in no-loan villages (see Table A5).

4.2.2 Other outcomes

Table 3 shows the estimates of equation (17) on non-agricultural outcomes. 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=1), and there is a large (US\$166, se=71), statistically significant (but rather imprecisely measured) 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 US\$30 on cultivation investments. The livestock value is measured several months after harvest; these results could indicate that households moved some of their additional farming profits into livestock post-harvest, or they may reflect measurement challenges.²⁶ 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, se=2, control group mean =0.83). Grant recipient households also consumed more, including 5.7% more food (column (4)); US\$0.34 per

²⁵ Calculated as $(\beta_1 - 0.79(\beta_1 + \beta_2))/(0.21) * (100/140)$. The average return in the entire village is β_1 . The take-up rate of loans is 21%, so 79% of households in the village would be non-borrowers and would have earned a return of $(\beta_1 + \beta_2)$. The return is then scaled to be per US\$100, so we divide by the grant size of US\$140/100.

²⁶ We may 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.

day in adult equivalency, $se=0.14$, control group mean = 5.96) and 5.8% in non-food expenditures (column (5); US\$2.53 per month, $se=1.39$, control group mean = 43.81). Columns (6)–(8) show no statistically significant main effect of the grant on whether the household has any financial savings, education expenses or medical expenses.

The investment and spending patterns among grant recipient (non-borrower) households in loan villages for the most part echo those described above in no-loan villages. Column (1) shows that while non-borrower grant recipients in loan villages were overall more likely to own livestock than their non-borrower, no grant counterparts, the magnitude of the effect is smaller than in the no-loan villages (interaction term is -4 percentage points, $se=2$). The remainder of the outcomes however show few differences.²⁷

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

4.2.3 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 and investment decisions 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

²⁷ 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 (US\$5.01, $se=2.55$), since medical expenses may have declined (-US\$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 able to treat illnesses, but (ii) having more resources could lead to higher preventative care, which should lower total medical expenses.

and squared, 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 (rather than a revisit).²⁸

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.²⁹ 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 statistically 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 statistically significant 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 statistically significant differences in total input expenses, value of the harvest, and gross profits (columns (6)–(8)). The number of hired labor days (column (4)) is the one statistically significant difference: 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. Column (9) suggests no statistically significant changes in equilibrium prices. This makes sense since villages in Mali are small. Households engage in market activities in local weekly markets, which bring multiple villages together (Ellis & Hine, 1998).

²⁸ 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.

²⁹ 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 villages not used for the primary study, but the selection of villages into the experimental study sample was not explicitly randomized. For example, the no-intervention villages have larger average population size but fewer children per household than study villages. Also 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.

We note that this analysis cannot speak directly to the possibility of spillovers in loan villages. The dynamics of sharing a grant with others in a village in which loans are available may differ, and the direction is difficult to predict. There may be pressure to share or hide “free” money when others recently borrowed; on the other hand, those who needed capital would have received a grant and therefore grant recipients may share less. Moreover, recent evidence by Banerjee et al (2020) also highlight how the introduction of formal credit can alter existing informal risk sharing arrangements. The grants were a one-time transfer while the introduction of a formal microfinance organization may also mean that spillovers are different in loan and no-loan villages.

5. Evidence of inefficient selection

If there is excess selection of poorer farmers out of borrowing, we expect to observe two patterns in the data. First, the gap between the distributions of observed gross profits of grant recipients and non-recipients will differ in no-loan villages from that among non-borrowers in loan villages. Second, the gap between the average returns to grants in no-loan and among non-borrowers in loan villages is positive, but this gap is attenuated at sufficiently low levels of observed baseline gross profits (or any other observable correlated with the friction generating the excess selection). If the excess selection is sufficiently strong, an observable characteristic that is positively correlated with average grant returns in the full population can be negatively correlated with average grant returns in the selected sample of non-borrowers in loan villages, a sign change that does not occur with efficient selection.

First, the left panel of Figure 4 depicts the distributions of gross profits of grant recipients and non-recipients in no-loan villages. As anticipated from our preceding results, $F_G(Q^G)$ lies to the right of $F_{NG}(Q^{NG})$ over virtually the whole range. However, in the loan villages, grants were randomly allocated only within the selected sample of non-borrowers. In the right panel of Figure 4, above a certain relatively high level of gross profits ($> \$500$), grant recipients and non-recipients have identical profits.³⁰ Under the assumption of monotone treatment effects, these farmers have approximately zero marginal return from the grant. This pattern is broadly consistent with an efficient allocation: farmers who have low returns to capital do not borrow and therefore show up in this sample. However, at lower levels of gross profits, $F_G(Q^G|C = 0)$ lies to the right of $F_{NG}(Q^{NG}|C = 0)$. These are non-borrowers with high returns to the grant but low gross profits. This feature corresponds to the exclusion of poor farmers who experience borrowing

³⁰ Note that this is the same sample as we use in table 2, and therefore continues to exclude households who borrowed in loan villages.

frictions. This suggests there are some potential borrowers with high return projects who do not receive capital, highlighting excess selection.

Second, we analyze how observable characteristics of borrowers and non-borrowers are correlated with the return to grants. We saw in Table 1 that there are observable characteristics that are strongly correlated with loan take-up. Consider any such attribute, X^k , that we *a priori* expect might be negatively correlated with farmer-specific borrowing frictions. For example, baseline gross profits would be one such attribute. In Table 4, we report the results of estimating (15), which includes the interaction term $\text{Grant} * X^k * \text{Loan village}$. This additional interaction permits us to examine whether the correlation between X^k and the marginal return to the grant is different for the general population (γ_1) than for a selected population of non-borrowers ($\gamma_1 + \gamma_2$). The lower frictions associated with the higher value of X^k reduces the likelihood that the farmer has been screened out of borrowing by concerns of default or risk aversion. Non-borrowers with higher values of X^k are therefore more likely to have selected out of borrowing because they have low marginal productivity. Hence, among the population of non-borrowers in loan villages, higher values of X^k are associated with lower values of $\Delta_G Q$, relative to the association in the population in general.

Column (1) of Table 4 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 the overall population, there is no significant correlation between baseline gross profits and the return to grants. However, in accord with borrowing frictions that decline with baseline gross profits, households in loan villages have a statistically significantly negative correlation between baseline gross profits and the return to a grant than households in the overall population (= -US\$0.12, se=0.04). The negative correlation is evidence of excess selection.

In columns (2)–(5), we report the estimates of equation (15) for four additional characteristics of households that are positively associated with loan take-up and plausibly farmer-specific borrowing frictions: baseline value of livestock holdings, baseline food consumption per capita (in USD), baseline non-food expenditure per capital (in USD) and the baseline index of social integration. The point estimates for each show a positive correlation with returns to the grant in the overall population for the first three, although this is not statistically significant for livestock. Column (2) reports the results for the baseline value of livestock holdings. The correlation between livestock holdings and the returns to the grant for non-borrowers (those in loan villages) is not significantly negative (-US\$0.01, se=0.01). Thus, this provides no evidence in support of the hypothesis that farmers with low livestock holdings are subject to higher borrowing frictions. Next we examine the same but for food consumption (column (3)) and non-food expenditures (column (4)), hypothesizing that these may be strongly positively correlated with a household's

permanent income (and hence negatively with borrowing frictions). Here we do find statistically significant negative correlations: the returns to the grant are lower for those with both higher food consumption (-US\$11.72, se=3.92) and non-food consumption (-US\$1.04 se=.46) for the non-borrowers, despite the fact in the general population both correlations are positive. This sign change distinguishes excess selection from efficient selection. In contrast, in column (5) we find that the index of social integration is not significantly correlated with returns to the grant. Nor is there a significant difference in this correlation when we compare farmers in the no-loan villages with non-borrowing farmers in the loan villages. We do not find evidence, therefore, that our measure of social integration is correlated with borrowing frictions that generate excess selection.

We next estimate $E(\Delta_G Q | \mathbf{X} = \mathbf{x}_j)$, the predicted treatment effect (also known as the conditional average treatment effect or CATE) of a grant to a farmer with characteristics \mathbf{x}_j using a causal forest trained on data from the no-loan villages. We also estimate $E(\Delta_G Q | \mathbf{X} = \mathbf{x}_i, c_i = 0)$, the predicted treatment effect of a grant for non-borrowing farmer i using the algorithm trained on non-borrowers in loan villages. In order to perform inference with these estimates, we follow the method by Chernozhukov et al. (2018) which is compatible with any machine learning algorithm used to estimate heterogeneous treatment effects, including the causal forest algorithm we use. In Appendix A3, we show that there is robust evidence of heterogeneity in grant treatment effects among the selected sample in loan villages, and little evidence of observable heterogeneity in no-loan villages. However, our empirical setting provides a second way to see if the model detects meaningful heterogeneity: as discussed in section 4.1, we use the model estimated from the no-loan villages ($E(\Delta_G Q | \mathbf{X} = \mathbf{x}_j)$) to predict the CATES for borrowers and non-borrowers in loan villages. Table 1 shows that the predicted CATES are positively correlated with loan takeup.

Finally, we compare the CATEs estimated in the no-loan villages to those estimated among non-borrowers in loan villages net of village fixed effects. Table 5, column (1) shows that at baseline 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, more landholdings, and use more labor. The pattern we see is that less poor households have higher treatment effects from grants.

As in (17), efficient selection into borrowing implies that the response of the CATE to a change in any of the eight dimensions of X will be attenuated in the selected sample of non-borrowers relative to borrowers. However, only if there is excess selection, with the poorer households subject to higher borrowing frictions, can the correlation between these observables and the treatment effects of the grants turn negative in the selected sample in loan villages. Column (2)

shows that in the causal forest estimated in this selected sample, all of these correlations are significantly negative. Among the selected sample in the loan villages who did not borrow, we see that those who are less poor have lower returns. These are households that would be less likely to default, or to be less risk averse. This is consistent with Table 1, where borrowers tended to be less poor than non-borrowers. The less poor households with expected high returns 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, and poorer households, many with high returns who do not borrow due to borrowing frictions, generating the negative correlations in column (2).

The exception to this pattern is the social integration index. There is no statistically significant difference in this correlation between farmers in the no-loan villages and non-borrowing farmers in the loan villages, and the point estimate of the correlation among the selected sample of non-borrowers is between zero and the estimate in the no-loan villages, so this provides no evidence that borrowing frictions are associated with our measure of social integration.

Figure 5 demonstrates visually that the sign changes reported in Table 5 are not artifacts of linearity. The vertical axis of each figure is the local linear regression smoothed estimate of the CATE of a grant; the horizontal axis is the (5th through 95th percentiles) of each of the eight baseline characteristics of households, all net of village fixed effects. In each case, excepting the measure of baseline social integration, we see a positive (or near zero) relationship between the baseline measure of wealth and the estimated treatment effect of a grant in the no-loan villages. And in each case, we find a negative relationship between baseline wealth and the estimated treatment effect of a grant in the selected sample of non-borrowers in loan villages.

Average agricultural returns to the grants for non-borrowers in loan villages are zero, as shown in column (10) of Table 2, while they are on average high for the random sample in no-loan villages. However, Figure 5 demonstrates that average agricultural returns to grants for non-borrowers with low baseline values of profits, food consumption, non-food consumption, livestock, farm size or total labor are relatively large. Indeed, Appendix Table 7 shows that among non-borrowers in the first tercile of the distribution of baseline food and non-food consumption, average returns to the grant are at least as high as the average returns in no-loan villages. We refer to these households as “the most poor”. Thus, it appears that among the most poor, there are households with high returns to grants that are not borrowers, implying an inefficient allocation of loans.

Would these most poor farmers use the loans in a similar way to the grants? We cannot observe the returns to the grant for any individual farmer, of course. But we do observe the *ex-post* gross profits of grant recipients. Among the most poor households, it is not possible to reject the

hypothesis that the distribution of profits among those who receive grants in the loan villages (the non-borrowers) is the same as for those who receive grants in the no-loan villages. However, this may be due to low power. Among these most poor households, the median, second tercile and third quartile of the distributions of profits among those nonborrowers who receive grants is greater than or equal to those of the distribution of profits among grant recipients in no-loan villages (although none of the differences is statistically significant).³¹ The distribution of observed profits for grant recipients among the most poor, therefore, is consistent with the existence of high return households among the non-borrowers. There is no evidence of selection of high return farmers into borrowing amongst the most poor; all of the selection is occurring among the less poor. There are farmers with high returns who do not borrow.

6. 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. We use the following specification:

$$Y_{ijt} = \alpha + \beta_1 loan_j \cdot I\{t = 2011\} + \beta_2 loan_j \cdot I\{t = 2012\} + X_{ijt}\pi + \epsilon_{ijt} \quad (14)$$

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 6 for details.

Table 6 and Appendix Table 8 show the ITT estimates for agricultural outcomes and broader outcomes, respectively. In Table 6, we observe an increase in input expenditures on family labor days (8.6, se=4.8) and in fertilizer and other chemicals expenses (US\$14, se=7); total input expenses rose by US\$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 US\$34 (se=20), but we do not measure a statistically significant increase in gross profits (US\$19, se=16).

Loans have to be repaid, while grants do not. Concerns about the costs of default or risk could deter borrowers from investing in the highest return activities; loan recipients to use loans differently from the way in which they use grants, and to realize different returns for loans than grants. The selection effect we have identified, in which women with high agricultural returns to grants are strongly selected into borrowing, may not imply that these same women have high

³¹ These results are available upon request.

agricultural returns to loans. We calculate the Treatment on the Treated estimates for year 1 for the sub-population who take up loans.³² Compared to the estimate of the impact of the grant from table 2, we do not reject the hypothesis that the per US\$100 dollar effects of grants and loans are the same for any of the agricultural outcomes.³³ Taken as a whole, the grants and loans are having similar effects on agricultural inputs and outcomes.³⁴

Appendix Table 8 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, or educational expenses.³⁵ We observe a large but imprecisely estimated impact on livestock (columns (1)–(2)). We do find a statistically significant *reduction* in medical expenses (column (9), -US\$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 et al., 2015; Attanasio et al., 2015; Augsburg, De Haas, Harmgart, & Meghir, 2015; Banerjee, Duflo, et al., 2015, and an overview in Banerjee, Karlan, et al., 2015; in contrast, Breza & Kinnan, 2018 finds noticeable general equilibrium effects as a consequence of a state-wide shutdown of the microcredit market; Crépon et al., 2015; Karlan & Zinman, 2011; Tarozzi et al., 2015). Analysis pooling these

³² See table notes of Table 6. Interest charges and fees, plus the cost of the 10 percent deposit requirement imply that a \$100 loan would need to generate \$131 in additional revenue to be profitable to the farmer. We estimate that \$89 (se=37) of each \$100 loaned is used for higher farm expenses, generating additional farm output valued at \$145 (se=82). The remaining \$11 of the loan are likely invested in livestock (see Appendix table 8), which appears to generate an even higher return. These ToT estimates are very noisy, but consistent with the high estimated returns to grants estimated for borrowers.

³³ 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.

³⁴ Note that we do not remove the cost of the loan, i.e. interest payments, from gross profits. The true difference in take home profits between the grant and loan would be larger. We do not include the interest because the goal is to see if the behavior of farmers, in terms of investments and the associated agricultural output, differs between the grants and the loans. We see that there is no evidence that the fact that they must pay interest leads to different investment choices.

³⁵ Columns (9)–(11) of Appendix Table 8 further shows no detectable effect on women’s decision-making power within the household, women’s involvement in community decisions, or women’s social capital. This is similar to the existing evaluations of microcredit (one exception is Angelucci et al., 2015; finding no impact on these measures: Attanasio et al., 2015; Augsburg et al., 2015; Banerjee, Duflo, et al., 2015; Crépon et al., 2015). Soro Yiriwaso did not have any explicit component of the program emphasizing women’s empowerment.

studies using a Bayesian hierarchical model, however, uncovers evidence of positive treatment effect at higher quantiles, even though the average treatment effect is a fairly precise null (Meager, 2019, 2020). An earlier agricultural lending literature also documented institutional failures, typically with high default rates (Adams, 1971; Adams et al., 1984), although a newer study in Zambia finds positive impacts from agricultural loans, similar to those found here (Fink et al., 2018).

The impact estimates are also promising from the perspective of the microcredit 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.

7. Persistent effects of grants

We focus first on the impact of the grants in year 2 and then on the impact at the longer-term follow-up in year 7.

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 Panel A of Table 7, column (8) shows that output is higher in grant recipient households by US\$52 (se=23) and column (9) demonstrates that gross profit was higher by US\$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) (US\$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=1), 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 7, 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) of Panel A as in year 1 (-US\$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 (US\$30, se=17) than control households in 2012.

Other outcomes

Appendix 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, $se=2$) and continue to hold more livestock assets (US\$184, $se=102$) than control households in no-loan villages. They are also more likely to own a business (3 percentage points, $se=1$).³⁶ There is no significant increase in food consumption in year 2 (US\$0.24, $se=0.19$) but monthly non-food expenditure does increase (US\$3.89, $se=2.13$). Households are also more likely to have financial savings (3.3 percentage points, $se=1.9$). Columns (7)-(8) show that there continues to be no measurable impact on educational expenses (US\$0.39, $se=3.76$), or medical expenses (-US\$0.72, $se=1.82$).

Appendix Table 9 also shows that, similar to year 1, there is no evidence of households in no-loan villages using grants differently to 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 of the original sample households. Panel B of Table 7 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 between 200 and 244 million tonnes of cotton were produced per year. In 2017, that figure had risen to 703 million tonnes (USAID, 2018). The increase largely came from an increase in the 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.

³⁶ In results available from the authors, business profits increase by 18% (US\$41, $se=19$) in year 2.

8. Conclusion

Capital constraints are binding for at least some farmers in Southern Mali, and 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 ill-suited to agriculture. 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. Given the lackluster average estimated impact of entrepreneurial microcredit (Banerjee, Karlan, et al., 2015; Meager, 2019), our results suggest a path for microcredit lenders looking to shift their model towards a product that generates higher average returns for borrowers without increasing default. Naturally, further experimentation would be fruitful in order to test, for example, whether each of the three changes from the more “normal” microcredit model (group liability, agricultural focus, balloon repayment) was necessary.

These results are also important for policy analysis and program evaluation. The random choice of communities into which to enter by the lender enables 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, e.g. a quasi-experimental approach, 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, Banerjee, Hanna, Olken, & Tobias, 2012; Alatas et al., 2013). We look at a price-based screening mechanism, since agricultural loans charge a positive interest rate that induces selection. 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 take up agricultural microfinance loans more 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, as well as social policy that aims to relax liquidity constraints for the most vulnerable. In particular, our results provide rigorous empirical evidence for systematic selection into contracts, which is embedded in several models (Buera, 2009; e.g., Evans & Jovanovic, 1989; 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, D. W. (1971). Agricultural Credit in Latin America: A Critical Review of External Funding Policy. *American Journal of Agricultural Economics*, 53, 163–172.
- Adams, D. W., Graham, D. H., & Von Pischke, J. D. (Eds.). (1984). *Undermining rural development with cheap credit*. Boulder: Westview Press.
- Alatas, V., Banerjee, A., Hanna, R., Olken, B. A., & Tobias, J. (2012). Targeting the Poor: Evidence from a Field Experiment in Indonesia. *The American Economic Review*, 102, 1206–1240.
- Alatas, V., Banerjee, A., Hanna, R., Olken, Benjamin, Purnamasari, R., & Wai_Poi, M. (2013). *Self-Targeting: Evidence from a Field Experiment in Indonesia*.
- Angelucci, M., Karlan, D., & Zinman, J. (2015). Microcredit Impacts: Evidence from a Randomized Microcredit Program Placement Experiment by Compartamos Banco. *American Economic Journal: Applied Economics*, 7, 151–182.
- Angrist, J. D., Imbens, G. W., & Rubin, D. B. (1996). Identification of Causal Effects Using Instrumental Variables. *Journal of the American Statistical Association*, 91, 444–455. JSTOR.
- Armendariz de Aghion, B., & Morduch, J. (2010). *The Economics of Microfinance* (2nd ed.). Cambridge, MA: MIT Press.
- Ashraf, N., Berry, J., & Shapiro, J. M. (2010). Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia. *American Economic Review*, 100, 2383–2413.
- Athey, S., & Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113, 7353–7360.
- Athey, S., Tibshirani, J., & Wager, S. (2019). Generalized random forests. *The Annals of Statistics*, 47, 1148–1178.
- Athey, S., & Wager, S. (2019). Estimating Treatment Effects with Causal Forests: An Application. *Observational Studies*, 5, 36–51.
- Attanasio, O., Augsburg, B., De Haas, R., Fitzsimons, E., & Harmgart, H. (2015). The Impacts of Microfinance: Evidence from Joint-Liability Lending in Mongolia. *American Economic Journal: Applied Economics*, 7, 90–122.

- Augsburg, B., De Haas, R., Harmgart, H., & Meghir, C. (2015). The Impacts of Microcredit: Evidence from Bosnia and Herzegovina. *American Economic Journal: Applied Economics*, 7, 183–203.
- Banerjee, A., Breza, E., Duflo, E., & Kinnan, C. (2015). Do credit constraints limit entrepreneurship? Heterogeneity in the returns to microfinance. *Working Paper*.
- Banerjee, A., & Duflo, E. (2012). Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program. *M.I.T. Working Paper*.
- Banerjee, A., Duflo, E., Glennerster, R., & Kinnan, C. (2015). The Miracle of Microfinance? Evidence from a Randomized Evaluation. *American Economic Journal: Applied Economics*, 7, 22–53.
- Banerjee, A., Karlan, D., & Zinman, J. (2015). Six Randomized Evaluations of Microcredit: Introduction and Further Steps. *American Economic Journal: Applied Economics*, 7, 1–21.
- Banerjee, A., E. Breza, A. Chandrasekhar, E. Duflo, M. Jackson, and Cynthia Kinnan. Changes in social network structure in response to exposure to formal credit markets. No. w28365. National Bureau of Economic Research, 2021.
- Beaman, L., Karlan, D., & Thuysbaert, B. (2014). *Saving for a (not so) Rainy Day: A Randomized Evaluation of Savings Groups in Mali*. Retrieved from <http://www.nber.org/papers/w20600>
- Beaman, L., Karlan, D., Thuysbaert, B., & Udry, C. (2013). Profitability of fertilizer: Experimental evidence from female rice farmers in Mali. *American Economic Review Papers & Proceedings*.
- Breza, E., & Kinnan, C. (2018). Measuring the equilibrium impacts of credit: Evidence from the Indian microfinance crisis. *Working Paper*.
- Bruhn, M., & McKenzie, D. (2009). In Pursuit of Balance: Randomization in Practice in Development Field Experiments. *American Economic Journal: Applied Economics*, 1, 200–232.
- Buera, F. J. (2009). A dynamic model of entrepreneurship with borrowing constraints: Theory and evidence. *Annals of Finance*, 5, 443–464.
- Chernozhukov, V., Demirer, M., Duflo, E., & Fernández-Val, I. (2018). Generic Machine Learning Inference on Heterogenous Treatment Effects in Randomized Experiments. *NBER Working Paper*, 24678.
- Cohen, J., & Dupas, P. (2010). Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment *. *Quarterly Journal of Economics*, 125, 1–45.
- Crépon, B., Devoto, F., Duflo, E., & Pariente, W. (2015). Estimating the Impact of Microcredit on Those Who Take It Up: Evidence from a Randomized Experiment in Morocco. *American Economic Journal: Applied Economics*, 7, 123–150.
- Crépon, B., El Komi, M., & Osman, A. (2020). Is It Who You Are or What You Get? Comparing the Impacts of Loans and Grants for Microenterprise Development. *Working Paper*. Retrieved from <https://www.adam-osman.com/wp-content/uploads/2020/05/Loans-vs-Grants.pdf>
- Davis, J. M. V., & Heller, S. B. (2017). Using Causal Forests to Predict Treatment Heterogeneity: An Application to Summer Jobs. *American Economic Review*, 107, 546–550.

- Davis, J. M. V., & Heller, S. B. (2019). Rethinking the Benefits of Youth Employment Programs: The Heterogeneous Effects of Summer Jobs. *The Review of Economics and Statistics*, 1–47.
- de Mel, S., McKenzie, D., & Woodruff, C. (2008). Returns to Capital in Microenterprises: Evidence from a Field Experiment. *Quarterly Journal of Economics*, 123, 1329–1372.
- de Quidt, J., Fetzer, T., & Ghatak, M. (2012). Group Lending Without Joint Liability. *London School of Economics Working Paper*.
- Druilhe, Z., & Barreiro-Huré, J. (2012). Fertilizer subsidies in sub-Saharan Africa. *FAO ESA Working Paper, No 12-04*.
- Duflo, E., Kremer, M., & Robinson, J. (2011). Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya. *American Economic Review*, 101, 2350–2390.
- Ellis, S. D., & Hine, J. L. (1998). The Provision of Rural Transport services. *Sub-Saharan Africa Transport Policy Program Working Paper*, 37, 70.
- Evans, D. S., & Jovanovic, B. (1989). An Estimated Model of Entrepreneurial Choice under Liquidity Constraints. *Journal of Political Economy*, 97, 808–827.
- Fafchamps, M., & Woodruff, C. (2017). Identifying Gazelles: Expert Panels vs. Surveys as a Means to Identify Firms with Rapid Growth Potential. *The World Bank Economic Review*, lhw026.
- Field, E., Pande, R., Papp, J., & Rigol, N. (2013). Does the Classic Microfinance Model Discourage Entrepreneurship Among the Poor? Experimental Evidence from India. *American Economic Review*, 103, 2196–2226.
- Fink, G., Jack, B. K., & Masiye, F. (2018). *Seasonal Liquidity, Rural Labor Markets and Agricultural Production* (Working Paper No. 24564). National Bureau of Economic Research.
- Giné, X., & Karlan, D. S. (2014). Group versus individual liability: Short and long term evidence from Philippine microcredit lending groups. *Journal of Development Economics*, 107, 65–83.
- Heckman, J. (1992). Randomization and Social Policy Evaluation. In C. F. Manski & I. Garfinkel (Eds.), *Evaluating welfare and training programs* (pp. 201–230). Cambridge, Mass: Harvard University Press.
- Heckman, J. (1997). Instrumental Variables: A Study of Implicit Behavioral Assumptions Used in Making Program Evaluations. *The Journal of Human Resources*, 32, 441.
- Heckman, J. J., Ichimura, H., & Todd, P. E. (1997). Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme. *The Review of Economic Studies*, 64, 605–654. JSTOR.
- Hussam, R., Rigol, N., & Roth, B. N. (2020). Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design In The Field. *Working Paper*.
- Imbens, G. W., & Angrist, J. D. (1994). Identification and Estimation of Local Average Treatment Effects. *Econometrica*, 62, 467.
- Jack, B. K. (2013). Private information and the allocation of land use subsidies in Malawi. *American Economic Journal: Applied Economics*, 5, 113–35.
- Kaboski, J. P., & Townsend, R. M. (2011). A Structural Evaluation of a Large-Scale Quasi-Experimental Microfinance Initiative. *Econometrica*, 79, 1357–1406.
- Karlan, D., & Morduch, J. (2009). Access to Finance. In D. Rodrick & M. R. Rosenzweig (Eds.), *Handbook of Development Economics* (Vol. 5). Elsevier.

- Karlan, D., & Mullainathan, S. (2007). Rigidity in Microfinancing: Can One Size Fit All? *QFinance*. Retrieved from <http://www.qfinance.com/financing-best-practice/rigidity-in-microfinancing-can-one-size-fit-all?page=1>
- Karlan, D., Osei-Akoto, I., Osei, R. D., & Udry, C. R. (2013). Agricultural Decisions after Relaxing Credit and Risk Constraints. *Quarterly Journal of Economics*, *Forthcoming*. <https://doi.org/10.2139/ssrn.2169548>
- Karlan, D., & Zinman, J. (2011). Microcredit in Theory and Practice: Using Randomized Credit Scoring for Impact Evaluation. *Science*, *332*, 1278–1284.
- Kazianga, H., & Udry, C. (2006). Consumption smoothing? Livestock, insurance and drought in rural Burkina Faso. *Journal of Development Economics*, *79*, 413–446.
- Maitra, P., Mitra, S., Mookherjee, D., & Visaria, S. (2020). Decentralized Targeting of Agricultural Credit Programs: Private versus Political Intermediaries. *National Bureau of Economic Research Working Paper*, 26730. <https://doi.org/10.3386/w26730>
- McKenzie, D. (2018). Can Business Owners Form Accurate Counterfactuals? Eliciting Treatment and Control Beliefs About Their Outcomes in the Alternative Treatment Status. *Journal of Business & Economic Statistics*, *36*, 714–722.
- McKenzie, D. J. (2015). *Identifying and spurring high-growth entrepreneurship: Experimental evidence from a business plan competition* (Policy Research Working Paper Series No. 7391). The World Bank.
- McKenzie, D., & Sansone, D. (2019). Predicting entrepreneurial success is hard: Evidence from a business plan competition in Nigeria. *Journal of Development Economics*, *141*, 102369.
- Meager, R. (2019). Understanding the Average Impact of Microcredit Expansions: A Bayesian Hierarchical Analysis of Seven Randomized Experiments. *American Economic Journal: Applied Economics*, *11*, 57–91.
- Meager, R. (2020). Aggregating Distributional Treatment Effects: A Bayesian Hierarchical Analysis of the Microcredit Literature. *LSE Working Paper*.
- Moll, B. (2014). Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation? *American Economic Review*, *104*, 3186–3221.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, *66*, 688–701.
- Suri, T. (2011). Selection and Comparative Advantage in Technology Adoption. *Econometrica*, *79*, 159–209.
- Tarozzi, A., Desai, J., & Johnson, K. (2015). The Impacts of Microcredit: Evidence from Ethiopia. *American Economic Journal: Applied Economics*, *7*, 54–89.
- Tarozzi, A., Mahajan, A., Blackburn, B., Kopf, D., Krishnan, L., & Yoong, J. (2014). Micro-Loans, Insecticide-Treated Bednets, and Malaria: Evidence from a Randomized Controlled Trial in Orissa, India. *American Economic Review*, *104*, 1909–1941.
- Tibshirani, J., Athey, S., Friedberg, R., Hadad, V., Miner, L., Wager, S., & Wright, M. (2018). grf: Generalized Random Forests (Beta). *ArXiv:1610.01271 [Econ, Stat]*. Retrieved from <https://github.com/grf-labs/grf> R package version 0.10.2
- USAID. (2018). *On the functioning of agricultural markets in Mali*. Retrieved from https://cdn.ymaws.com/www.andeglobal.org/resource/resmgr/research_library/2018-11_MIFP_Study_on_Agricu.pdf

- Wager, S., & Athey, S. (2018). Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *Journal of the American Statistical Association*, *113*, 1228–1242.
- Young, A. (2019). Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results. *The Quarterly Journal of Economics*, *134*, 557–598.

Online appendix – not for publication

Appendix A1: Loan allocation with frictions

i. Limited liability

Consider a simple limited liability model of credit. An efficient allocation maximizes the gain in gross profits from loans, net of the cost of capital to the lender (ρ). The efficient allocation is defined by the function $B(\Delta_B Q, Q^{NG})$ chosen to maximize

$$\int_{\underline{\Delta_B Q}}^{\overline{\Delta_B Q}} \int_{\underline{Q^{NG}}}^{\overline{Q^{NG}}} B(\Delta_B Q - \rho) \tilde{f}(\Delta_B Q, Q^{NG}) dQ^{NG} d\Delta_B Q \quad (16)$$

where $\tilde{f}(\Delta_B Q, Q^{NG})$ is the joint density of marginal returns to borrowing ($Q^B - Q^{NG}$) and Q^{NG} implied by the joint distribution of potential outcomes $F(Q^{NG}, Q^G, Q^B)$ defined in section 3.

In this efficient allocation, $B(\Delta_B Q, Q^{NG}) = 1$ if $\Delta_B Q \geq \rho$, and $B = 0$ otherwise.

However, suppose there is limited liability. Because of limited liability, the maximum repayment that the lender can obtain from a borrower i is $\Delta_B Q$ if $\underline{c} \leq Q^{NG}$, $Q^{NG} + \Delta_B Q - \underline{c}$ if $Q^{NG} \leq \underline{c} \leq Q^{NG} + \Delta_B Q$, and 0 if $Q^{NG} + \Delta_B Q \leq \underline{c}$. The breakeven constraint of the lender, therefore, is

$$\begin{aligned} & \int_{\underline{\Delta_B Q}}^{\overline{\Delta_B Q}} \int_{\underline{c}}^{\overline{Q^{NG}}} B(\Delta_B Q) \tilde{f}(\Delta_B Q, Q^{NG}) dQ^{NG} d\Delta_B Q \\ & + \int_{\underline{\Delta_B Q}}^{\overline{\Delta_B Q}} \int_{\underline{c} - \Delta_B Q}^{\underline{c}} B(\Delta_B Q - \underline{c}) \tilde{f}(\Delta_B Q, Q^{NG}) dQ^{NG} d\Delta_B Q \\ & \geq \int_{\underline{\Delta_B Q}}^{\overline{\Delta_B Q}} \int_{\underline{Q^{NG}}}^{\overline{Q^{NG}}} \rho B(\Delta_B Q - \rho) \tilde{f}(\Delta_B Q, Q^{NG}) dQ^{NG} d\Delta_B Q. \end{aligned} \quad (17)$$

The left hand side of the breakeven constraint is the revenue generated by the lending, which is equal to the full gain in output for farmers not subject to the limited liability constraint plus the constrained payments from those farmers subject to the limited liability constraint (which are zero for all farmers with $Q_i^{NG} + \Delta_B Q \leq \underline{c}$). The RHS is the cost of all loans. The constrained

efficient allocation is the function $B(\Delta_B Q, Q^{NG})$ that maximizes (16) subject to the breakeven constraint (17).

If the breakeven constraint does not bind when $B_i = 1$ for all farmers i with $\Delta_B Q_i \geq \rho$, and $B_i = 0$ for all farmers with $\Delta_B Q_i < \rho$, then the unconstrained efficient allocation remains feasible. The breakeven constraint may not bind at the unconstrained efficient allocation if the distribution of farmers is such that the surplus generated by farmers for whom limited liability does not bind is sufficient to cover the losses from borrowers who are (at least partially) defaulting. In this case

$$\int_{\rho}^{\overline{\Delta_B Q}} \int_{\underline{c} + \rho - \Delta_B Q}^{\overline{Q^{NG}}} (\Delta_B Q - \rho) \tilde{f}(\Delta_B Q, Q^{NG}) dQ^{NG} d\Delta_B Q + \int_{\rho}^{\overline{\Delta_B Q}} \int_{\underline{Q^{NG}}}^{\underline{c} + \rho - \Delta_B Q} (\max(Q^{NG} + \Delta_B Q - \underline{c}, 0) - \rho) \tilde{f}(\Delta_B Q, Q^{NG}) dQ^{NG} d\Delta_B Q \geq 0.$$

The first term is the surplus generated from high-return farmers ($\Delta_B Q \geq \rho$) who pay the cost of their loans in full ($\Delta_B Q + Q^{NG} \geq \underline{c} + \rho$). The second term are the losses from high return farmers ($\Delta_B Q \geq \rho$) who are too poor to fully repay the cost of their loans ($\Delta_B Q + Q^{NG} < \underline{c} + \rho$). In this case, the allocation remains efficient.

However, if (17) is violated at the unconstrained efficient allocation, then it remains the case that $B_i = 1$ for all farmers with both $Q_i^{NG} + \Delta_B Q_i \geq \underline{c}$ and $\Delta_B Q_i \geq \rho$ (because such loans relax the breakeven constraint and increase net gain in output), and $B_i = 0$ for all farmers with $\Delta_B Q_i \leq \rho$ because such loans decrease the net gain in output and tighten the breakeven constraint. However, not all farmers with high marginal returns and low base output Q_i^{NG} can receive loans. The allocation of these remaining loans is determined by the function $B(\Delta_B Q, Q^{NG})$ to maximize

$$\int_{\rho}^{\overline{\Delta_B Q}} \int_{\underline{Q^{NG}}}^{\underline{c} + \rho - \Delta_B Q} B(\Delta_B Q - \rho) \tilde{f}(\Delta_B Q, Q^{NG}) dQ^{NG} d\Delta_B Q \quad (18)$$

subject to

$$\int_{\rho}^{\overline{\Delta_B Q}} \int_{\underline{Q^{NG}}}^{\underline{c} + \rho - \Delta_B Q} B(\max(Q^{NG} + \Delta_B Q - \underline{c}, 0)) \tilde{f}(\Delta_B Q, Q^{NG}) dQ^{NG} d\Delta_B Q \leq \int_{\rho}^{\overline{\Delta_B Q}} \int_{\underline{c} + \rho - \Delta_B Q}^{\overline{Q^{NG}}} (\Delta_B Q - \rho) \tilde{f}(\Delta_B Q, Q^{NG}) dQ^{NG} d\Delta_B Q \quad (19)$$

The RHS of (19) is a constant, the surplus generated by lending to high return farmers who repay the full cost of their loans. The problem is to allocate that fixed budget across the set of high-return farmers who cannot fully repay their loans to maximize (18).

The increase in (18) from lending to farmer j is $\Delta_B Q_j - \rho$, while the cost is $\rho - \max(Q_j^{NG} + \Delta_B Q_j - \underline{c}, 0)$. Therefore, farmers are allocated loans in order of decreasing ratios of benefit to cost: if $B_j = 1$ and $B_k = 0$, then $\frac{\Delta_B Q_j - \rho}{\rho - \max(Q_j^{NG} + \Delta_B Q_j - \underline{c}, 0)} \geq \frac{\Delta_B Q_k - \rho}{\rho - \max(Q_k^{NG} + \Delta_B Q_k - \underline{c}, 0)}$, and the boundary between $B(\Delta_B Q, Q^{NG}) = 1$ and $B(\Delta_B Q, Q^{NG}) = 0$ for farmers who partially repay their loans is characterized by $\frac{\Delta_B Q - \rho}{\rho - \max(Q^{NG} + \Delta_B Q - \underline{c}, 0)} = k$ for some constant $k > 0$. Therefore, the boundary between borrowers and nonborrowers in a constrained efficient allocation is downward sloping in $(Q^{NG}, \Delta_B Q)$. Thus, some farmers with high returns to capital may not receive loans, while similar farmers with the same marginal productivity but higher baseline output do borrow.

ii. Risk aversion

Alternatively, consider expected utility-maximizing farmers with decreasing absolute risk aversion. They are presented with an opportunity to borrow a fixed amount at cost ρ , with full enforcement. The loan would finance a risky project with random return $\Delta_B Q$ over baseline gross profit Q^{NB} . Suppose $E(\Delta_B Q) \geq \rho$ and that there is a farmer i indifferent between taking the loan to finance the project or not. Then any farmer with a higher no-grant gross profit with the same preferences and investment opportunity would strictly prefer to take the loan, and indeed would take a loan to finance a strictly inferior investment opportunity, with returns that are first order stochastically dominated by the project with return $\Delta_B Q_i$.³⁷ Farmers with lower no-grant gross profits require higher expected returns to be willing to accept the additional risks associated with borrowing. Risk aversion and self-selection also generates a boundary in $(E(Q^{NG}), E(\Delta_B Q))$ between those who borrow and those who do not that is downward sloping as in the dashed line Figure 3c.

Risk averse farmers will in general select different projects to finance with grants and loans. Suppose a farmer receiving a loan is indifferent between two risky projects with returns $\eta^1 \equiv \Delta_B Q^1$ and $\eta^2 \equiv \Delta_B Q^2$ with $E(\eta^1) > E(\eta^2)$. That farmer would strictly prefer the riskier, higher expected return project 1 if offered a grant rather than a loan. Therefore, the project chosen by the marginal borrower who is given a grant instead will have an expected return (weakly) greater

³⁷ For i , $EU(\Delta_B Q_i + Q_i^{NG} - \rho) = EU(Q_i^{NG})$. Then farmer j with $Q_j^{NG} > Q_i^{NG}$ with the the same project has $EU(\Delta_B Q_i + Q_j^{NG} - \rho) > EU(Q_j^{NG})$. So there is a constant $\epsilon_j > 0$ with $EU(\Delta_B Q_i - \epsilon_j + Q_j^{NG} - \rho) > EU(Q_j^{NG})$.

than the project that that farmer would have chosen to implement with the loan. Risk aversion generates selection across projects of a farmer as well as across farmers. Therefore, in Figure 3C, the solid line boundary in $(E(Q^{NG}), E(\Delta_G Q))$ between those who borrow and those who do not lies above that boundary in $(E(Q^{NG}), E(\Delta_B Q))$, and with DARA preferences the difference between the boundaries declines as $E(Q^{NG})$ rises.³⁸ Within-farmer selection of projects implies $E(\Delta_G Q|C = 1) \geq E(\Delta_B Q|C = 1)$. Since we have shown (in (10) and (13)) that each of these quantities are identified by our experimental design, in section 6 we examine the evidence that farmers may be selecting among projects.

iii. Distinguishing efficient and excess selection

The actual return to the grant for farmer i (which is unobserved to us, but perhaps is known to the farmer) is

$$\Delta_G Q_i = \mathbb{E}(\Delta_G Q|\mathbf{X} = \mathbf{x}_i) + \mu_i, \quad (20)$$

with $\mathbb{E}(\mu_i|\mathbf{X} = \mathbf{x}_i) = 0$ in the general population. If selection into borrowing is efficient, then borrowing is determined by $\Delta_G Q_i > \rho$.³⁹ This implies that non-borrowers have realizations of μ_i less than a threshold $\bar{\mu}_i \equiv \rho - \mathbb{E}(\Delta_G Q|\mathbf{X} = \mathbf{x}_i)$ and $\mathbb{E}(\mu|\mathbf{X} = \mathbf{x}_i, \mu < \bar{\mu}_i) \leq 0$. Therefore, taking expectations of (20) over the non-borrowers in loan villages, we have $\mathbb{E}(\Delta_G Q|\mathbf{X} = \mathbf{x}_i, C = 0) \leq \mathbb{E}(\Delta_G Q|\mathbf{X} = \mathbf{x}_i)$. Conditional on any observed characteristic, the average return to grants should be higher in the no-loan villages.

The treatment effect heterogeneity along dimension k in loan villages depends on (a) how that variable (say, baseline gross profit) is correlated with expected returns to the grant in the full population; and (b) how changes in those expected returns affect selection, $\mu < \bar{\mu}_i$. Assuming μ is independent of X^k , heterogeneity along dimension k among non-borrowers is related to that in the random sample by

³⁸ This discussion may raise the possibility that farmers borrowing with a limited liability constraint may also choose different projects than they would with a grant. In this case, the convexity of returns generated by the limited liability could induce borrowers to take *more* risk. However, this would imply some default in equilibrium, and we observe no instance of a defaulted loan.

³⁹ This abstracts from risk and thus implies that the farmer knows μ_i . To permit risk we let $\mu_i = v_i + \xi_i$ with v_i known to the farmer and $E(\xi_i | \mathbf{x}_i, v_i) = 0$. Now efficient borrowing is determined by $E(\Delta_G Q | v_i) > \rho$ and the following argument proceeds as stated, with the added notation.

$$\begin{aligned} \frac{d\mathbb{E}(\Delta_G Q | \mathbf{X} = \mathbf{x}_i, C = 0)}{dX^k} &= \frac{d\mathbb{E}(\mathbb{E}(\Delta_G Q | \mathbf{X} = \mathbf{x}_i) + \mu_i | \mathbf{X} = \mathbf{x}_i, \mu < \bar{\mu}_i)}{dX^k} \\ &= \left(1 - \frac{\partial \mathbb{E}(\mu | \mu < \bar{\mu}_i)}{\partial \bar{\mu}_i}\right) \frac{d\mathbb{E}(\Delta_G Q | \mathbf{X} = \mathbf{x}_i)}{dX^k}. \end{aligned} \quad (21)$$

Suppose that the grant treatment effect in no-loan villages is increasing in X^k (the argument is symmetric around zero). Then, farmers in the no-loan villages with higher values of X^k have higher expected returns $\left(\frac{d\mathbb{E}(\Delta_G Q | \mathbf{X} = \mathbf{x}_i)}{dX^k} = \gamma_1\right)$.⁴⁰ This implies that in loan villages, the selection into non-borrowing becomes more severe: the increase in expected returns reduces the critical value $\bar{\mu}_i^k$, partially offsetting the increase in expected returns to the grant among non-borrowers in loan villages $\left(0 \leq \frac{\partial \mathbb{E}(\mu | \mu < \bar{\mu}_i)}{\partial \bar{\mu}_i} \leq 1\right)$.⁴¹ So, $\gamma_1 > \gamma_1 + \gamma_2 > 0$.

With excess selection, the frictions generating excess selection imply that a farmer's marginal return to a loan exceeds ρ , and that this wedge decreases as a farmer's no-grant gross profits, collateral, or wealth increase. Consider an observed characteristic correlated with farmers' no-grant gross profits. There is excess selection if borrowing is determined by $\Delta_G Q_i > \rho + h(\mathbf{X})$ with $h(\mathbf{X}) > 0$ and decreasing in X^k . Non-borrowers, then, have $\mu_i < \rho + h(\mathbf{x}_i) - \mathbb{E}(\Delta_G Q | \mathbf{X} = \mathbf{x}_i) \equiv \bar{\mu}_i$. Treatment effect heterogeneity among the selected sample of non-borrowers is

$$\begin{aligned} \frac{d\mathbb{E}(\Delta_G Q | \mathbf{X} = \mathbf{x}_i, C = 0)}{dX^k} &= \left(1 - \frac{\partial \mathbb{E}(\mu | \mu < \bar{\mu}_i)}{\partial \bar{\mu}_i}\right) \frac{d\mathbb{E}(\Delta_G Q | X^k = x_i^k)}{dX^k} \\ &\quad + \frac{\partial \mathbb{E}(\mu | \mu < \bar{\mu}_i)}{\partial \bar{\mu}_i} \frac{dh(\mathbf{x}_i)}{dX^k}. \end{aligned} \quad (22)$$

The additional third term is always negative. Average returns to grants are larger in the general population than among the non-borrower subpopulation. But this gap is attenuated at sufficiently low levels of X^k , because even farmers with high return projects are not borrowing due to the high wedge generated by the friction. Excess selection always reduces the slope of the relationship between average returns to grants and any X^k that is negatively correlated with borrowing frictions. If $\gamma_1 > 0$ (expected returns to the grant are increasing in X^k in the random sample) then $\gamma_1 > \gamma_1 + \gamma_2$. Recall that in the case of efficient selection, this effect could only

⁴⁰ Recall γ_1 is the regression coefficient on $grant_i \cdot X_i^k$ in specification (14).

⁴¹ The first inequality is always true. If the distribution of μ_i^k is has a normal, power, double exponential or Pareto distribution, then the second follows.

attenuate the heterogeneity. By contrast, if excess selection is sufficiently strong, the sign can change $\gamma_1 + \gamma_2 < 0 < \gamma_1$.⁴²

Appendix A2: 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 CATEs 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 CATEs 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 CATEs. 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 CATEs). Second, each 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

⁴² Similarly, if $\gamma_1 < 0$, $\gamma_1 + \gamma_2 < \gamma_1 < 0$ with sufficiently strong excess selection.

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), and we allow the algorithm to tune the parameters through cross-validation using the “R-learner” objective function for heterogeneous treatment effects. This regularization method is not a standard cross-validation technique like “leave one out” or k-fold cross validation. It was developed by the authors specifically for generalized random forests. Intuitively, it picks random combinations of parameters to train multiple “mini forests”, then uses the out-of-bag predictions to estimate the objective function (the “R-objective”) for each forest, and picks the combination that minimizes it. This is explained in detail in section 1.3 of Athey & Wager (2019).

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. We used the “tune all” option in the algorithm (instead of manually selecting which parameters to tune) as done in the application in Athey and Wager (2019).

Regarding the number of trees in the forest, the documentation to the *grf* algorithm recommends “that users grow trees in proportion to the number of observations”. Davis and Heller (2019) use 100,000 trees. We tested different number of trees 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. The depth of the trees is controlled by a parameter (*min.node.size*) in the algorithm and is tuned jointly with the other parameters listed above.

Overfitting

The grf algorithm uses honest estimation and the use of out-of-bag predictions to minimize the risk of overfitting. The goal is to avoid overfitting and allow for generalizability without giving up part of the sample when training the forest. Honesty is defined by Wager and Athey (2018) as “A tree is honest if, for each training example i , it only uses the response Y_i to estimate the within-leaf treatment effect or to decide where to place the splits, but not both.” Nevertheless, Davis and Heller (2017) demonstrate that overfitting can occur even with honest estimation. They propose out-of-bag predictions in addition to honest estimation to reduce the overfitting risk. In practice, this means that the prediction for a given observation is calculated using only trees that were not trained with that observation (or cluster, when using cluster-robust estimation as in our case, which we discuss below).

The grf package, released after Davis and Heller (2017), uses both out-of-bag predictions and honest estimation by default.

Clustered RCT design

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 CATEs relate with the actual treatment effects. First, we group the observations in each sample into 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.

Appendix A3: Unobservable versus observable predictors of marginal returns

i. 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, some which are likely to be 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 loans 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 Appendix Table 5 show that there is limited evidence of heterogeneity using the variables that we saw to be important in Table 1, including baseline gross profits, baseline land size, and baseline value of livestock. However, the estimates of the interaction terms with observable characteristics are very imprecise, and noise in the data may limit our power to detect heterogeneity. The exercise still demonstrates that it would be difficult for local NGOs or other policymakers to predict returns using easy-to-collect data.

Instead of relying on our intuition for choosing baseline characteristics, we also exploit a machine learning algorithm to estimate heterogeneity in treatment effects (Athey & Imbens, 2016; Athey et al., 2019; Wager & Athey, 2018). 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. Appendix A2 above provided details on the implementation of the causal forests algorithm.

In column (3) of Appendix Table 5, 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 that 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. Appendix Table 5, column (4) looks at this loan villages sub-sample. 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). Baseline characteristics, among a selected sample of nonborrowers, can predict heterogeneity in the returns to capital but we can only detect this heterogeneity with the assistance of the two-stage experiment.

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

Appendix Table 5 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.

Appendix Table 6 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 (-US\$33, $se=22$ compared to -US\$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. It is also possible that we lack precision, either due to sample size or too much measurement error.

In section 4 and table 5, we show that a given Z characteristic—for example, gross profits—has a very different relationship with predicted treatment effects (CATEs) depending on whether the algorithm was trained on data from no-loan villages or from loan villages. Those with baseline higher gross profits had higher predicted treatment effects in no-loan villages, but lower predicted treatment effects in loan villages. We may not observe strong evidence of heterogeneous returns in the random sample in Table Appendix Table 5 because 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.

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.

Appendix A4: Randomization inference

We follow Young (2019) to implement the Randomization Inference (RI) procedure.⁴³ First, we generated 10,000 simulations of the assignment of grants. In each simulation, we reproduced the re-randomization routine described in Section 2.1 to ensure that the grant assignments are drawn from the same distribution as the original experiment. We took the villages type (loan village / no-loan village), as well as the selection of households in loan villages into taking the loan, as given. Therefore, the sample of eligible recipients of the grant (i.e., all households in no-loan villages and non-borrowers in loan villages) was pre-determined and identical across all iterations. In each iteration, we reproduced the main analysis using the synthetic treatment assignment and stored the coefficients for all the relevant tests. That is, we re-estimated the effect of receiving a grant and its interaction with village type on all the agricultural outcomes of interest, for each year of the experiment. We then used the results to approximate the covariance matrix of the estimated coefficients of interest across the universe of potential treatment assignments. This allowed us to calculate the randomization-c p-values from a two-tailed test of significance for each treatment effect, as in Young (2019). We also implement randomization-based joint testing procedures to test the null hypothesis that all relevant treatment effects in an equation family are zero. To avoid grouping together aggregate outcomes of interest with their individual components, we divide the agricultural variables into three independent families: (i) agricultural inputs and crop choice, (ii) total input expenses and value of output, and (iii) gross profit. We report RI p-values for joint Wald tests of significance of the

⁴³ We use an adapted version of the Stata command “randcmd” (Young, 2020) which allows for more flexibility in the randomization routine.

treatment effects of the grant and its interaction with village type on all the outcomes in a given family (i.e., an omnibus test of overall experimental significance for that equation group).

Table 1: Comparison of baseline characteristics of clients vs. non-clients in loan treatment villages

	Tookup	Did Not Takeup	Difference ³
	(1)	(2)	(3)
A. Agriculture, Livestock & Business			
Land size (ha)	2.64 (2.71)	2.21 (2.64)	0.59 (0.13)
Total input expenses	205.82 (300.42)	151.87 (285.75)	46.37 (14.22)
Value of output	709.04 (752.17)	596.10 (827.66)	132.60 (39.79)
Gross profit	503.22 (555.12)	444.23 (642.11)	86.23 (30.84)
Total value of livestock	1871.22 (3037.9)	1294.65 (2549.9)	504.65 (135.2)
Predicted grant treatment effects (CATEs)	54.50 (25.39)	53.33 (25.24)	4.45 (1.28)
B. Household Demographics			
Age of female respondent	36.58 (10.29)	34.92 (11.68)	2.46 (0.58)
Married (0/1)	0.98 (0.13)	0.92 (0.27)	0.07 (0.01)
Not first wife (0/1)	0.33 (0.47)	0.19 (0.39)	0.13 (0.02)
Number of children	4.86 (2.34)	4.34 (2.40)	0.70 (0.12)
Risk aversion: safe lottery	0.46 (0.50)	0.50 (0.50)	-0.03 (0.02)
Index of intra-household decision making power	0.08 (0.97)	-0.03 (1.05)	0.14 (0.05)
Index of community action	0.28 (1.03)	-0.03 (0.99)	0.26 (0.05)
Social integration index	0.23 (1.04)	-0.09 (0.98)	0.18 (0.05)
D. Consumption			
Food consumption EQ (past 7 days, USD)	6.89 (4.17)	6.70 (4.22)	0.40 (0.21)
Monthly non-food exp (USD)	48.09 (45.38)	39.77 (38.44)	10.04 (2.03)

Notes

1 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: how frequently she speaks with different village leaders, and different types of participation in village meetings and activities. The social capital index includes questions about 7 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). All three of these variables are indices, normalized by the no-grant households in loan-unavailable villages.

2 Clients are defined by households who took out a loan in the 2010 agricultural season.

Table 2: Agriculture - Year 1

	Land cultivated (ha)	Land planted with rice and groundnut (ha)	Used Plough (0/1)	Quantity Seeds (Kg)	Family labor (days)	Hired labor (days)	Fertilizer and chemical expenses (USD)	Total input expenses (USD)	Value output (USD)	Gross Profit (USD)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Grant β_1	0.19 (0.07) [0.003]	0.08 (0.02) [0.000]	0.06 (0.01) [0.000]	5.86 (2.09) [0.021]	6.50 (4.16) [0.110]	2.86 (0.80) [0.001]	22.60 (5.96) [0.001]	31.38 (8.02) [0.000]	68.85 (18.66) [0.000]	40.72 (15.32) [0.005]
Grant * loan village β_2	-0.17 (0.10) [0.080]	0.02 (0.03) [0.651]	-0.002 (0.02) [0.936]	1.62 (3.41) [0.667]	-6.92 (6.30) [0.264]	1.17 (1.42) [0.374]	-15.87 (8.62) [0.099]	-10.35 (12.04) [0.402]	-44.17 (28.03) [0.129]	-39.21 (22.35) [0.082]
p-value for $\beta_1 + \beta_2 = 0$	0.783	0.000	0.001	0.007	0.930	0.001	0.282	0.020	0.240	0.927
N	5343	5386	5393	5339	5342	5340	5387	5341	5339	5286
Mean of control (year 1)	2.07	0.87	0.80	87.93	134.16	17.03	117.55	186.84	501.91	316.45
SD of control (year 1)	2.22	0.72	0.40	76.57	128.02	23.24	199.27	251.75	595.30	428.12
Per \$100 impact for loan takers	0.59 (0.29)	0.01 (0.10)	0.05 (0.07)	-0.17 (10.14)	23.23 (19.02)	-1.11 (4.16)	58.78 (26.28)	50.24 (36.42)	167.88 (84.77)	134.46 (67.98)

Notes

- 1 Size of grant was \$140. Loan recipients are excluded from the analysis sample.
- 2 Rows showing Grant + Grant * loan village = 0 shows the p value of the test of whether the total effect of grants in loan villages is statistically different from zero.
- 3 Standard errors are in parentheses and clustered at the village level in all specifications.
- 4 In brackets are randomization inference p values following Young (2019). They are the randomization-c p -values from a two-tailed test of significance for each treatment effect. There are three independent families of outcomes: (i) agricultural inputs and crop choice in columns (1)-(7), (ii) total input expenses and value of output in columns (8)-(9), and (iii) gross profit in column (10). The RI p -values for joint Wald tests of significance of the treatment effects of the grant and its interaction with village type on each outcome individually are in brackets. The p values for the omnibus test of the overall experimental significance for each family is as follows: $p < 0.001$; $p < 0.001$; and $p = .029$. Appendix A4 discusses implementation details.
- 5 Total input expenses includes fertilizer, manure, herbicide, insecticide, rental and maintenance costs of farming equipment, purchased seeds, and hired labor but excludes the value of family labor. Gross profit is revenue minus most, but not all, expenses. Specifically, the formula includes value of harvest (whether sold, stored or consumed) minus 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). Thus this does not subtract value of own labor, value of family (i.e., any unpaid) labor, and the implicit rental value of land used.
- 6 Additional controls include: village fixed effects; the baseline value of the dependent variable; an indicator for whether the baseline value is missing; an indicator for the HH being administered the input survey in 2011, and household stratification controls (from baseline: whether the household was part of an extended family; was polygamous; an index of the household's agricultural assets; an index of household's other assets; per capita food consumption; and for the primary female respondent her baseline: land size, 0/1 on whether she used fertilizer in the previous agricultural season, and whether she had access to a plough). Village-level stratification controls are not included since there are village fixed effects.
- 7 Mean of control is the mean of the dependent variable in the column heading among households that received no grants in no-loan villages in year 1.
- 8 The per dollar return for loan takers is calculated as: $(\beta_1 - .79 * (\beta_1 + \beta_2)) / (.21 * 140)$ where .21 is the loan take up rate and 140 is the value of the grant.

Table 3: Additional Outcomes of Grants in Year 1

	Own any livestock (0/1)	Total value of livestock (USD)	HH has a business (0/1)	Food consumption EQ (past 7 days, USD)	Monthly non- food exp (USD)	HH has any financial savings (0/1)	Educ expenses (USD)	Medical expenses (USD)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Grant β_1	0.12 (0.01)	179.98 (69.59)	0.04 (0.02)	0.34 (0.14)	2.80 (1.38)	0.03 (0.02)	2.34 (3.11)	-2.44 (1.82)
Grant * loan village β_2	-0.04 (0.02)	-58.58 (100.24)	0.00 (0.02)	0.07 (0.21)	2.33 (2.06)	0.03 (0.03)	-0.29 (5.58)	4.91 (2.55)
p-value for $\beta_1 + \beta_2 = 0$	0.000	0.093	0.032	0.014	0.001	0.013	0.660	0.164
N	5264	5212	5263	5091	5055	5204	3573	5219
Mean of control (year 1)	0.78	1213.08	0.83	5.96	43.81	0.63	69.87	33.66
SD (year 1)	(0.42)	(2048.50)	(0.37)	(3.16)	(37.31)	(0.48)	(81.20)	(45.92)
Per \$100 impact for loan takers	0.19 (0.07)	285.97 (306.05)	0.04 (0.07)	0.06 (0.64)	-4.25 (6.25)	-0.07 (0.08)	2.46 (16.33)	-14.95 (7.83)

Notes

1 See the notes of Table 2 for details on specification.

Table 4: Heterogeneity in Borrowing Frictions

	Gross Profit				
	(1)	(2)	(3)	(4)	(5)
Grant	15.16 (21.92)	33.26 * (16.90)	-29.06 (27.47)	10.55 (19.16)	40.92 *** (15.40)
Grant * Loan village	36.26 (28.48)	-15.14 (24.31)	109.20 *** (39.39)	32.95 (30.31)	-37.65 (23.05)
Grant * Baseline gross profit (γ_1)	0.07 (0.06)				
Grant * Baseline gross profit * Loan village (γ_2)	-0.18 *** (0.07)				
Grant * Baseline livestock (γ_1)		0.01 (0.01)			
Grant * Baseline livestock * Loan village (γ_2)		-0.02 (0.01)			
Grant * Baseline food consumption (γ_1)			11.05 ** (4.51)		
Grant * Baseline food cons * Loan village (γ_2)			-22.78 *** (5.96)		
Grant * Baseline non-food expenditure (γ_1)				0.75 ** (0.38)	
Grant * Baseline non-food exp * Loan village (γ_2)				-1.78 *** (0.60)	
Grant * Baseline social integration index (γ_1)					-13.96 (14.76)
Grant * Baseline social index * Loan village (γ_2)					17.50 (22.71)
($\gamma_1+\gamma_2$) coefficient	-0.12 (0.04)	-0.01 (0.01)	-11.72 (3.92)	-1.04 (0.46)	3.54 (17.16)
($\gamma_1+\gamma_2$) SE					
N	5286	5285	5189	5121	5285

Notes

1 See the notes of Table 2 for details on additional controls.

Table 5: Correlation of Causal Forest Predicted Treatment Effects with Baseline Characteristics

	(1)	(2)
	No loan villages model CATE	Non-borrowers in loan villages model CATE
Gross profit	0.016 (0.001)	-0.066 (0.002)
Food consumption EQ (past 7 days, USD)	3.875 (0.153)	-3.175 (0.273)
Monthly non-food exp (USD)	0.152 (0.014)	-0.298 (0.035)
Total value of livestock (USD)	0.001 (0.000)	-0.005 (0.001)
Social capital index	-5.153 (0.559)	-4.737 (1.009)
Land cultivated (ha)	3.193 (0.304)	-13.938 (0.783)
Value of agricultural assets owned	0.003 (0.001)	-0.011 (0.009)
Total labor (days)	0.046 (0.004)	-0.199 (0.009)

Notes

- 1 Each row reports the coefficients from two separate regressions of the predicted treatment effect generated by a causal forest algorithm on the sub-sample indicated in the column heading (and predicted only for the households in that sub-sample), on the baseline value of the covariate indicated in the row heading and village fixed effects.
- 2 Standard errors are in parantheses and clustered at the village level in all specifications.

Table 6: Agriculture ITT estimates from Loans

	Land cultivated (ha)	Land planted with rice and groundnut (ha)	Used Plough	Quantity Seeds (Kg)	Family labor (days)	Hired labor (days)	Fertilizer and chemical expenses	Total input expenses	Value output	Gross profit (USD)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Loan village - year 1	0.02 (0.06)	0.01 (0.03)	0.02 (0.02)	-1.07 (2.84)	13.88 (4.50)	0.19 (0.93)	-12.62 (7.30)	-10.46 (9.62)	34.36 (19.65)	50.17 (15.68)
Loan Village - year 2	0.129 (0.07)	0.02 (0.03)	0.02 (0.02)	0.11 (3.06)	-6.65 (4.63)	-1.75 (1.02)	28.02 (7.90)	41.48 (10.75)	23.29 (20.86)	-15.10 (14.28)
N	8725	8871	8848	8763	8770	8769	8879	8768	8767	8687
Mean of control (year 1)	2.07	0.88	0.80	87.93	134.16	17.07	117.04	186.25	500.49	315.43
SD (year 1)	(2.22)	(0.73)	(0.40)	(76.57)	(128.02)	(23.35)	(197.76)	(250.17)	(591.41)	(425.37)
Per \$100 impact, TOT, year 1	0.08 (0.25)	0.03 (0.12)	0.10 (0.07)	-4.50 (11.96)	58.48 (18.96)	0.79 (3.93)	-53.19 (30.76)	-44.10 (40.53)	144.79 (82.81)	211.41 (66.06)
Diff in per \$100 impact: Grants - Loans	0.51	-0.02	-0.06	4.33	-35.25	-1.90	111.97	94.34	23.09	-76.95
SE from Bootstrap on Difference	(0.31)	(0.11)	(0.06)	(11.90)	(20.32)	(4.43)	(31.12)	(42.29)	(93.38)	(71.12)

Notes

- Grant recipients in both loan and no-loan villages are removed from the analysis sample. Probability weights are applied to account for the differences in the sampling probabilities in loan villages, which are a function of loan take-up. The probability weights of nonborrowers in loan villages are calculated as $[(\# \text{ of non-borrowers in sample in a loan village}) / (\# \text{ of these households who did not receive grant})]$, and are 1 for all other households in the sample.
- Total input expenses is the same variable as defined in Table 2.
- Additional controls include: cercle fixed effects; the baseline value of the dependent variable, along with a dummy when missing, interacted with year of survey indicators; and village-level stratification controls: population size, distance to nearest road, distance to nearest paved road, whether the community is all bambara (dominant ethnic group), distance to the nearest market, percentage of households with a plough, percentage of women with access to plough in village, percentage of women in village using fertilizer and the fraction of children enrolled in school. The specification uses probability weights to reflect sampling design. All grant-recipients households are removed from the analysis in both loan and no-loan villages.
- Standard errors are in parantheses and clustered at the village level in all specifications.
- Mean of control is the mean of the dependent variable in the column heading among households in no-loan villages.
- The per dollar return, TOT, year 1 is: the coefficient on Loan village - year 1 $/ (.21 * 113)$ since the average value of the loan was \$113. The standard error on the difference in per dollar impact is the result of a bootstrap of 1000 draws comparing the per dollar impact of the grant vs the loan using re-sampling of households. Probably weights were calculated in each bootstrap sample and used in the estimate of the loan impact.

Table 7: Agriculture - Year 2 & Long-term follow up

	Land cultivated (ha)	Land planted with rice and groundnut (ha)	Used Plough	Quantity Seeds (Kg)	Family labor (days)	Hired labor (days)	Fertilizer and chemical expenses	Total input expenses	Value output	Gross profit (USD)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
A. Impact of grants in Year 2										
Grant β_1	0.10 (0.08) [0.202]	0.07 (0.03) [0.006]	*** 0.04 (0.01) [0.003]	*** 6.91 (2.56) [0.012]	*** -4.63 (3.79) [0.26]	1.17 (0.81) [0.17]	-3.56 (8.88) [0.72]	2.69 (10.13) [0.82]	54.05 (22.68) [0.03]	** 49.85 (17.07) [0.002]
Grant * loan village β_2	0.06 (0.11) [0.641]	0.05 (0.04) [0.17]	-0.01 (0.02) [0.68]	1.09 (3.52) [0.78]	9.54 (5.91) [0.12]	1.52 (1.16) [0.21]	23.08 (14.34) [0.12]	28.34 (17.29) [0.11]	-18.99 (32.85) [0.61]	-41.57 (23.59) [0.10]
Grant + Grant * loan village = 0	0.06	0.00	0.03	0.00	0.28	0.00	0.08	0.03	0.14	0.61
N	5241	5386	5353	5300	5300	5300	5384	5300	5300	5247
Mean of control	2.25	0.92	0.81	90.53	122.99	15.39	170.94	251.20	511.73	257.22
SD of control	(2.39)	(0.74)	(0.39)	(76.89)	(121.30)	(22.53)	(286.85)	(343.16)	(704.24)	(435.18)
B. Impact of grants in Long-term follow up										
Grant β_1	0.13 (0.11) [0.233]	0.03 (0.03) [0.318]	0.03 (0.02) [0.131]	6.50 (3.88) [0.130]	2.94 (4.98) [0.581]	2.01 (1.59) [0.265]	5.92 (11.09) [0.659]	22.94 (16.28) [0.219]	23.65 (42.67) [0.598]	-10.56 (28.31) [0.750]
Grant * loan village β_2	0.08 (0.16) [0.635]	0.03 (0.05) [0.618]	-0.01 (0.03) [0.605]	1.01 (6.08) [0.866]	1.01 (7.69) [0.898]	-1.97 (2.61) [0.475]	8.40 (18.91) [0.686]	-4.56 (25.84) [0.872]	43.16 (61.95) [0.528]	32.29 (44.25) [0.512]
Grant + Grant * loan village = 0	0.081	0.073	0.510	0.109	0.499	0.985	0.348	0.359	0.137	0.524
N	4959	5166	5007	4958	4958	4957	5156	4957	4948	4898
Mean of control	2.12	0.89	0.72	100.80	120.48	23.39	178.01	289.26	694.34	408.91
SD of control	2.57	0.88	0.45	105.20	130.77	42.08	325.44	432.52	1075.91	783.87

Notes

1 See the notes of Table 2 for additional details on the specification.

2 Rows showing Grant + Grant * loan village = 0 shows the p value of the test of whether the total effect of grants in loan villages is statistically different from zero.

3 Total input expenses is the same variable as defined in Table 2.

4 Standard errors are in parentheses and clustered at the village level in all specifications.

5 Mean of control is the mean of the dependent variable in the column heading among households in no-loan villages.

Figure 1: Experimental Design

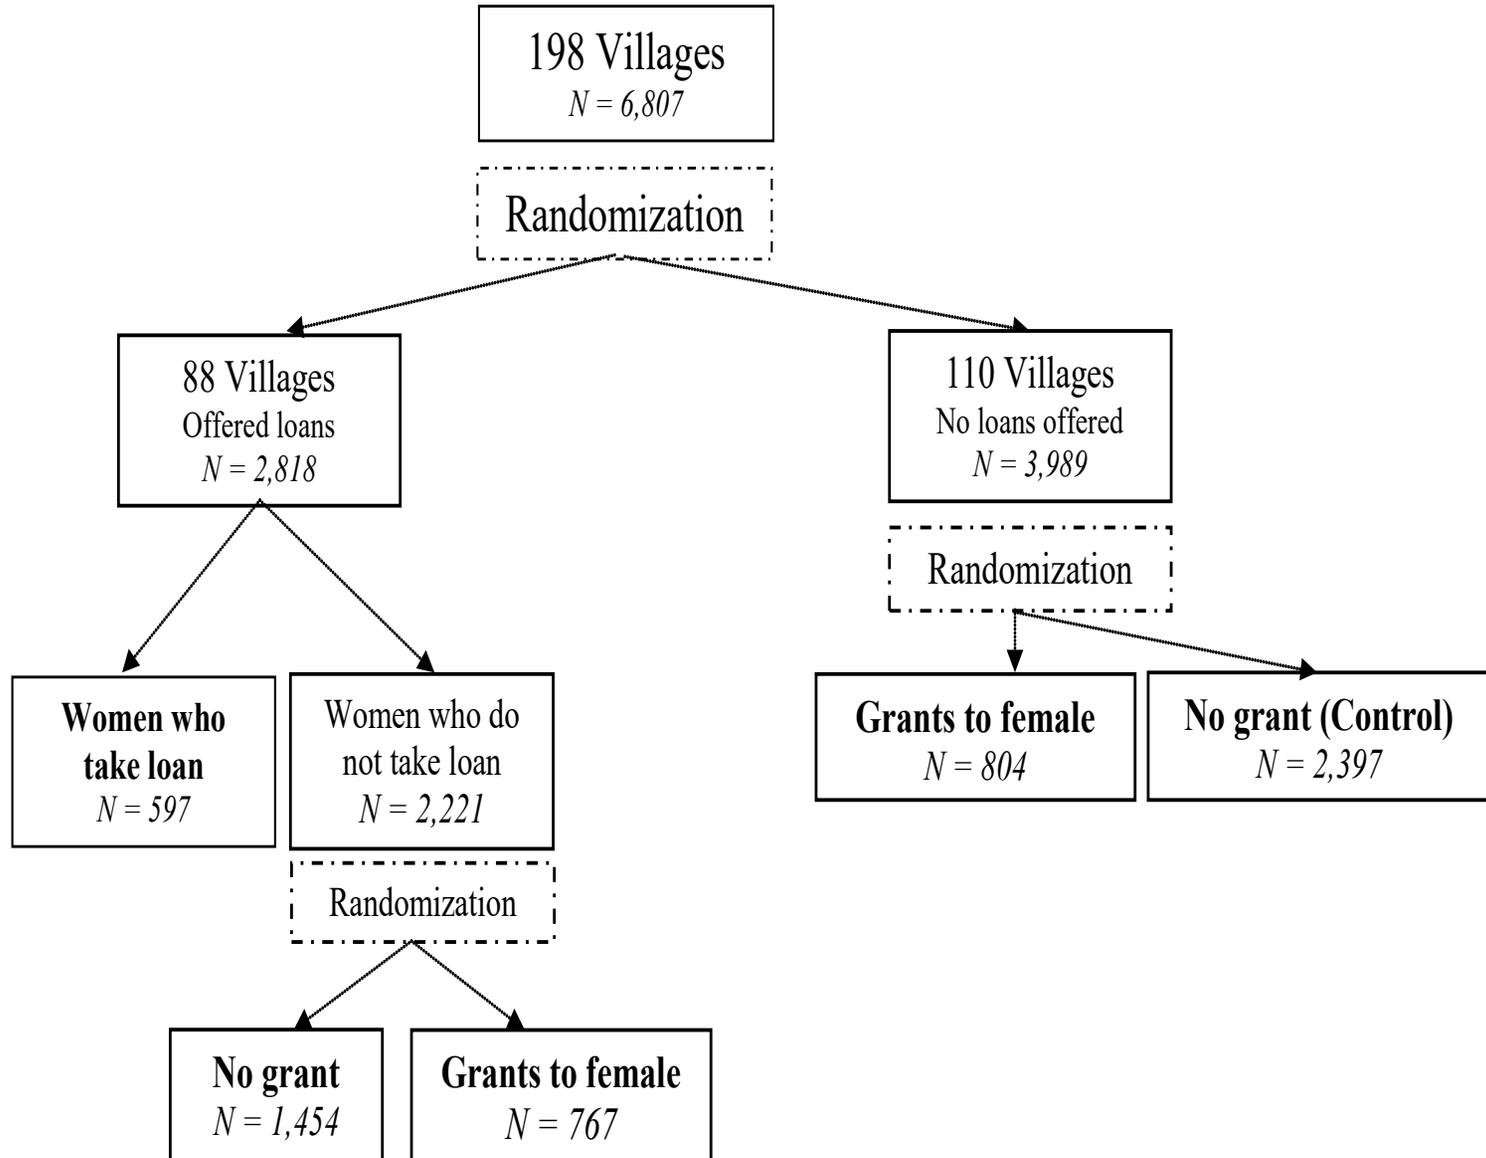


Figure 2: Baseline characteristics of borrowers vs. non borrowers in loan villages

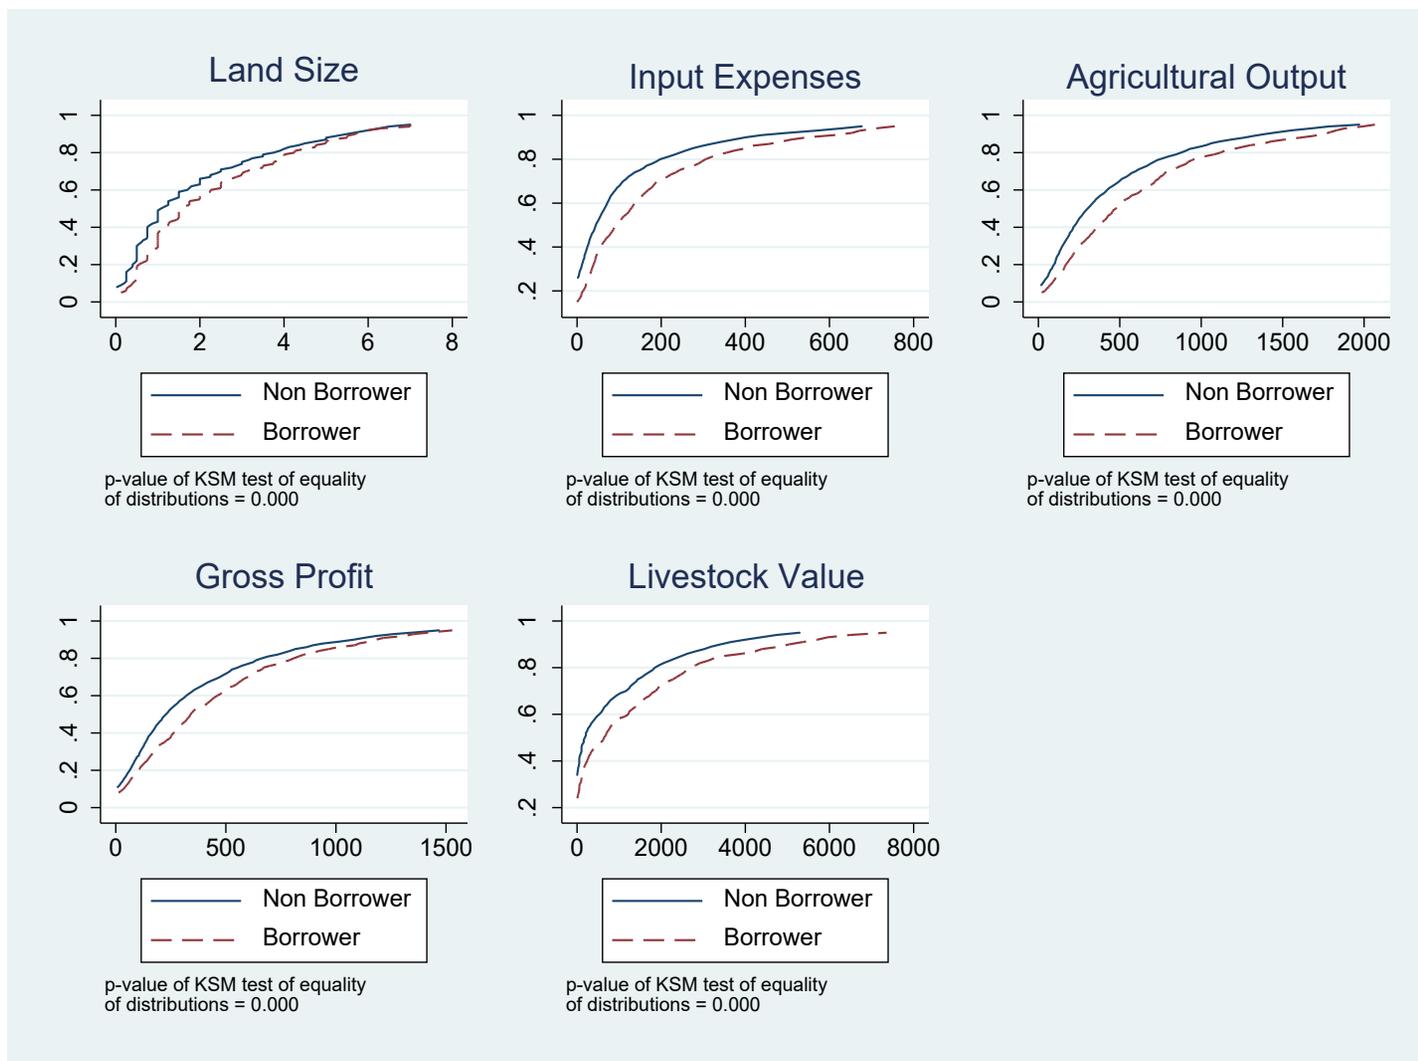
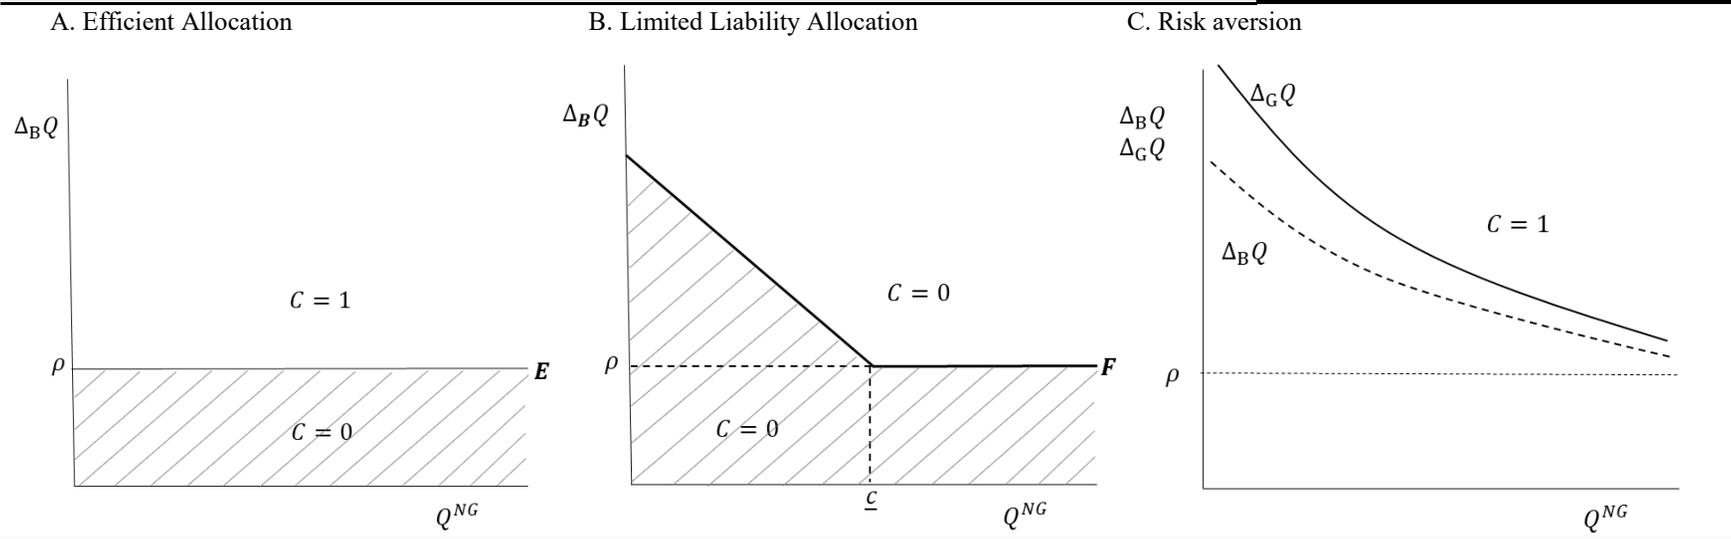


Figure 3: Selection into borrowing



Notes

- 1 The y axis is the change in gross profits in response to receiving a loan. ρ is the lender's gross cost of funds.
- 2 The x axis represents gross profits in the absence of a grant or loan. \underline{c} is the minimum consumption required below which the limited liability constraint binds.

Figure 4: CDF of Gross Profit

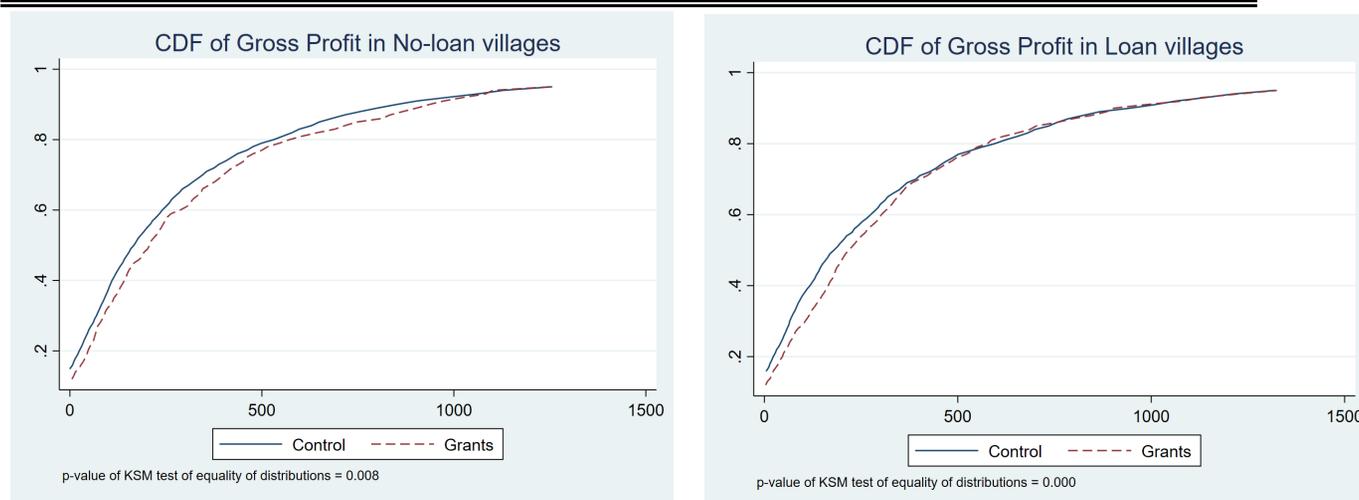
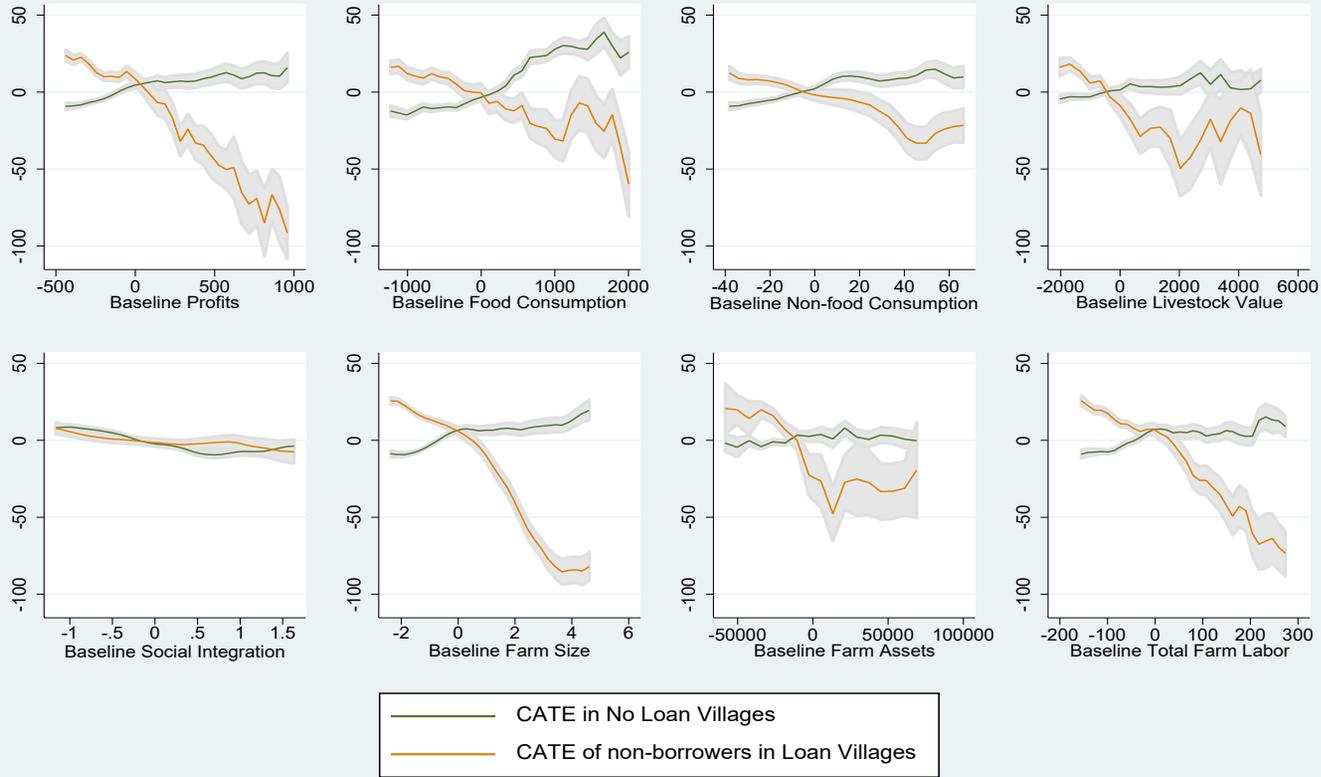


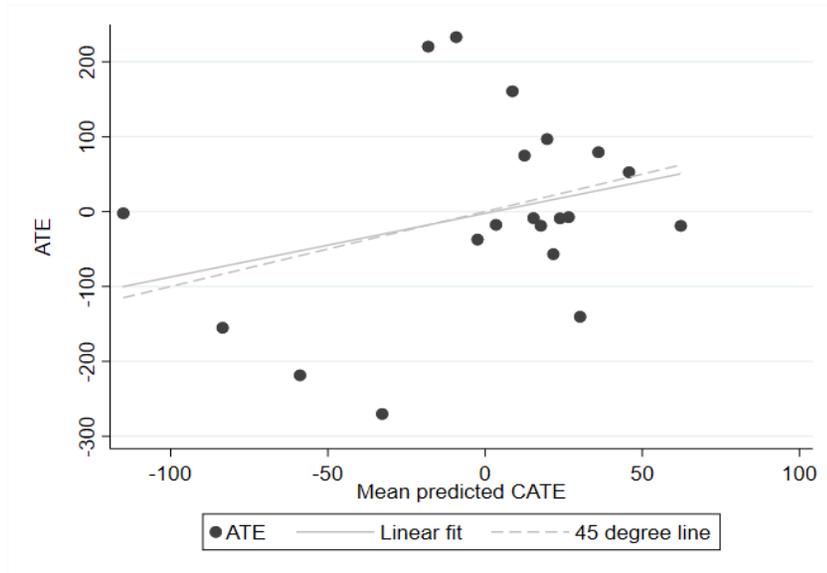
Figure 5: Grant Treatment Effects by Baseline Characteristic



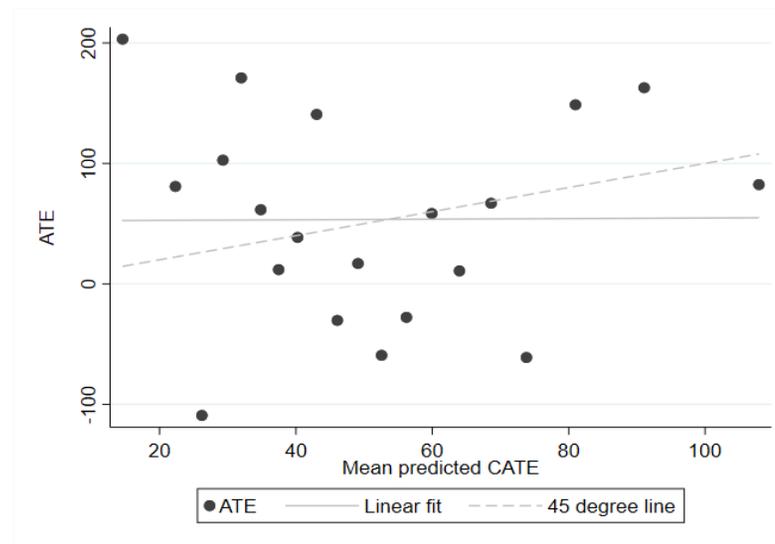
Note: Local linear regression smooth estimates using Epanechnikov kernel. Shaded region is pointwise 95 percent confidence interval.

Appendix figure A1: Predicted versus actual treatment effects on gross profits

A. Loan sample



B. No-loan sample



Notes

- 1 Each figure shows a comparison between the ATE and the mean predicted CATE for each 5-percent bin of the distribution of predicted CATE.
- 2 Predicted CATE are out-of-bag predictions from a causal forest trained on no-loan villages in Panel A and loan villages in Panel B.
- 3 ATE are obtained through regressions that control for baseline net revenue, village fixed effects, and additional controls specified in the notes to Table 2.
- 4 The solid line shows the linear fit between the group ATE and the mean predicted CATE.
- 5 The dashed line is the 45-degree line.

Appendix Table 1: Balance check

	Loan vs no-loan villages				Grants vs no-grants in no-loan villages				Grants vs no-grants in loan villages			
	Mean of control group	Difference between T and C	p-value	N	Mean of control group	Difference between T and C	p-value	N	Mean of control group	Difference between T and C	p-value	N
Household size	7.41	0.03	0.764	6,828	7.43	-0.06	0.617	3,151	7.37	-0.05	0.746	2,415
Land (ha)	1.92	0.22	0.034	6,856	1.92	0.04	0.676	3,174	2.09	-0.00	0.961	2,422
Days of family labor	139.41	-0.13	0.984	6,858	139.61	2.91	0.599	3,165	133.69	4.94	0.292	2,426
Days of hired labor	10.60	1.02	0.317	6,856	10.38	0.08	0.913	3,170	11.30	-0.56	0.451	2,419
Input expenses (USD)	126.95	17.68	0.128	6,856	127.49	9.80	0.253	3,172	138.55	0.55	0.952	2,422
Agricultural output (USD)	522.22	37.48	0.226	6,855	523.74	5.07	0.836	3,176	537.61	11.06	0.657	2,415
Livestock value (USD)	1,520.29	-120.52	0.285	6,924	1,515.83	2.63	0.980	3,199	1,389.71	-36.17	0.785	2,448
Has a Business	0.54	0.01	0.667	6,924	0.53	0.02	0.348	3,200	0.54	0.01	0.610	2,447
Monthly non-food expenses	39.48	0.18	0.917	6,568	39.75	-0.83	0.520	3,041	38.82	0.58	0.677	2,322
Male age	46.57	0.19	0.661	6,427	46.67	-0.35	0.497	2,947	45.93	0.53	0.307	2,272
Male is illiterate	0.77	-0.01	0.446	6,562	0.78	-0.00	0.824	3,015	0.77	0.01	0.583	2,321
F- test for joint significance				0.258				0.911				0.667

Appendix Table 2: Attrition

	Loan vs no-loan villages				Grants vs no-grants in no-loan villages				Grants vs no-grants in loan villages			
	Year 1		Year 2		Year 1		Year 2		Year 1		Year 2	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Treatment	0.001 (0.004)	-0.011 (0.016)	0.010 (0.007)	0.028 (0.023)	0.006 (0.005)	0.019 (0.023)	0.000 (0.005)	0.006 (0.023)	0.000 (0.004)	-0.032 (0.027)	-0.004 (0.006)	-0.041 (0.031)
Interaction of treatment and:												
Household size		0.000 (0.001)		0.000 (0.001)		0.000 (0.002)		0.000 (0.001)		-0.002 (0.002)		0.001 (0.002)
Land (ha)		0.002 (0.003)		-0.002 (0.004)		0.001 (0.005)		0.006 (0.006)		0.001 (0.005)		0.001 (0.005)
Days of family labor†		0.000 (0.001)		-0.001 (0.001)		-0.001 (0.001)		-0.001 (0.001)		-0.001 (0.001)		-0.002 (0.001)
Days of hired labor†		0.001 (0.002)		-0.001 (0.002)		0.002 (0.002)		0.002 (0.002)		-0.004 (0.003)		-0.001 (0.003)
Input expenses*		0.001 (0.004)		0.000 (0.004)		0.003 (0.005)		0.005 (0.004)		0.004 (0.008)		-0.002 (0.007)
Agricultural output *		0.000 (0.001)		0.004 (0.002)		-0.001 (0.002)		-0.002 (0.002)		-0.001 (0.002)		-0.001 (0.003)
Livestock value*		0.000 (0.000)		0.000 (0.000)		0.000 (0.000)		0.000 (0.000)		0.000 (0.000)		0.000 (0.000)
Has a Business		0.012 (0.006)		0.023 (0.008)		0.006 (0.012)		-0.002 (0.010)		0.011 (0.011)		0.021 (0.012)
Monthly non-food expenses*		-0.023 (0.013)		-0.010 (0.016)		-0.011 (0.012)		-0.005 (0.015)		0.042 (0.023)		0.014 (0.022)
Male age		0.000 (0.000)		0.000 (0.000)		0.000 (0.000)		0.000 (0.000)		0.001 (0.000)		0.001 (0.000)
Male is illiterate		-0.002 (0.011)		-0.018 (0.013)		-0.004 (0.015)		-0.010 (0.017)		-0.007 (0.020)		-0.010 (0.021)
N	5649	4757	5649	4757	3201	2702	3201	2702	2448	2055	2448	2055
Mean attrition control	0.014		0.015		0.012		0.015		0.015		0.026	
F- test for joint significance of coefficients of treatment and interaction terms		0.27		0.38		0.33		0.63		0.21		0.11

Notes

* Variables divided by 100 for ease of exposition.

† Variable divided by 10 for ease of exposition.

Appendix Table 3: Timing robustness (No-loan villages)

	Index			Land Size		
	(1)	(2)	(3)	(4)	(5)	(6)
Date (linear)	0.001 (0.004)	0.003 (0.008)		0.002 (0.011)	0.005 (0.023)	
Date squared		0.000 (0.000)			0.000 (0.001)	
1 if before June 1st			-0.045 (0.140)			-0.176 (0.407)
Revisit to Village	-0.022 (0.106)	-0.007 (0.119)	-0.034 (0.121)	0.124 (0.307)	0.147 (0.344)	0.051 (0.351)
Observations	787	787	787	774	774	774
Fixed effects	None	None	None	None	None	None

Notes

- 1 Index includes: land area, number of family labor days, number of hired labor days, an indicator for whether fertilizer was used, value of fertilizer expenses, value of other chemical expenses, value of all input expenses, value of harvest, and profits.
- 2 Sample includes only grant recipients in no-loan villages.

Appendix Table 4: Spillovers in No-loan Villages

	Land cultivated (ha)	Land planted with rice and groundnut (ha)	Family labor (days)	Hired labor (days)	Fertilizer and chemical expenses (USD)	Total input expenses	Value output	Gross Profits	Price Index	Wage index
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
No-loan village	-0.09 (0.13)	0.04 (0.06)	5.71 (7.73)	3.55 *** (1.36)	-12.65 (12.09)	-19.49 (16.53)	-34.90 (45.53)	-15.76 (30.30)	-0.052 (0.170)	0.041 (0.146)
N	3649	3681	3643	3649	3678	3652	3649	3619	175	170
Mean of excluded group	2.08	0.88	122.83	12.73	123.65	181.17	572.51	393.47	0.08	0.09
SD of excluded group	2.44	0.85	133.87	19.75	231.97	277.51	757.53	543.35	0.97	0.97

Notes

- 1 The sample includes households in (i) no-intervention villages and (ii) households in no-loan villages who did not receive a grant. The analysis uses only data from follow-up year 1.
- 2 The excluded group are households in no-intervention villages.
- 3 Additional controls for columns (1)-(8) include: *cercle* fixed effects; the baseline value of the dependent variable, along with a dummy when missing; the baseline value of the dependent variable interacted with the no-intervention village dummy; an indicator for the HH being administered the input survey in 2011; village-level stratification controls: population size, distance to nearest road, distance to nearest paved road, whether the community is all bambara (dominant ethnic group), distance to the nearest market, percentage of households with a plough, percentage of women with access to plough in village, percentage of women in village using fertilizer and the fraction of children enrolled in school; and individual-level stratification controls: whether the household was part of an extended family; was polygamous; an index of the household's agricultural assets and other assets; per capita food consumption; and for the primary female respondent her baseline: land size, fertilizer use, and whether she had access to a plough. Standard errors are clustered at the village level.
- 4 Columns (9) and (10) are village-level regressions. Additional controls include *cercle* fixed effects and the village-level stratification controls listed above.
- 5 Also included are the following individual controls: the number of adult household members, the number of children in the household, the average age of adults in the household and the share of adults with primary school education level.
- 6 Price index is a normalized average of prices of: main grain prices, livestock, and wage labor.

Appendix Table 5: Are Returns Predicted by Baseline Characteristics?

	(1)	(2)	(3)
Grant	15.09 (54.75)	14.55 (56.15)	23.60 (30.00)
Predicted Causal Effects			-0.40 (0.43)
Grant * Predicted Causal Effects			0.33 (0.58)
Grant * Baseline gross profit	0.07 (0.08)	0.07 (0.08)	
Grant * Baseline land	-7.30 (15.32)	-7.73 (15.30)	
Grant * Baseline value of livestock	0.00 (0.01)	0.00 (0.01)	
Grant * Large HH at baseline	57.68 (55.03)	59.54 (55.89)	
Grant * Baseline social index		-25.99 (15.19)	
Grant * Baseline intra-household bargaining index		-15.02 (16.29)	
N	3100	3099	3065
Year	1	1	1
Sample	No loan vill	No loan vill	No loan vill
Additional HH structure controls interacted with grant & year	Yes	Yes	No
HH decision-making/community action interacted with grant & year	No	Yes	No
Mean of Baseline gross profit	395.79		
SD of Baseline gross profit	488.88		
Mean of Baseline land	2.03		
SD of Baseline land	2.43		

Notes

- 1 See the notes 1, 3 and 5 of Table 2 for details on specification.
- 2 Large household is 6 or more adults in the household.
- 3 Other household structure controls include: an indicator for the presence of an extended family and the number of children in the household.
- 4 Predicted causal effects in column (3) are generated by a causal forest algorithm on no-loan village data and then extrapolated to all no-loan village households.
- 5 Predicted causal effects in column (4) are generated by a causal forest algorithm on loan village data and then extrapolated to all loan village households.

Appendix Table 6: Can Heterogeneous Treatment Effects be Predicted by Baseline Characteristics?

	Gross Profits			
	(1)	(2)	(3)	(4)
Grant	40.72 (15.32)	64.76 (40.34)	67.61 (41.40)	83.24 (26.62)
Grant * Loan village	-39.21 (22.35)	-33.67 (22.24)	-34.04 (22.12)	-31.73 (22.74)
Predicted Causal Effects				0.47 (0.34)
Grant * Predicted Causal Effects				-0.83 (0.47)
Grant * Baseline gross profit		0.03 (0.05)	0.03 (0.05)	
Grant * Baseline land		-16.26 (9.65)	-16.10 (9.63)	
Grant * Baseline value of livestock		0.00 (0.01)	0.00 (0.01)	
Grant * Large HH at baseline		84.81 (43.64)	83.47 (43.46)	
Grant * Baseline social index			-10.361 (12.22)	
Grant * Baseline intra-household bargaining index			-25.96 (11.48)	
Grant + Grant * loan village = 0	0.927	0.432	0.416	0.170
N	5286	5285	5283	5207
Year	1	1	1	1
Additional HH structure controls interacted with grant & year	No	Yes	Yes	No
HH decision-making/community action interacted with grant & year	No	No	Yes	No

Notes

- 1 See the notes 1, 3 and 5 of Table 2 for details on specification.
- 2 Col (4): Predicted treatment effects is from Causal Forest model trained on no-loan villages and predicted for entire analysis sample.

Appendix Table 7: Returns to Grant for Bottom Tercile of Baseline Characteristics

	Gross Profits			
	(1)	(2)	(3)	(4)
Grant	44.47 (21.39)	36.48 (19.45)	47.41 (20.33)	58.57 (22.26)
Grant * Loan village	-47.85 (29.97)	-46.19 (29.91)	-79.75 (31.05)	-77.38 (31.38)
Grant * T1 Baseline gross profit	-9.88 (28.20)			
Grant * T1 Baseline gross profit * Loan village	27.30 (38.08)			
Grant * T1 Baseline livestock		14.26 (29.35)		
Grant * T1 Baseline livestock * Loan village		18.53 (43.74)		
Grant * T1 Baseline food consumption			-25.95 (29.17)	
Grant * T1 Baseline food consumption * Loan village			133.61 (48.06)	
Grant * T1 Baseline non-food expenditure				-49.70 (35.15)
Grant * T1 Baseline non-food exp * Loan village				110.89 (49.36)
N	5286	5285	5189	5121
Grant impact for bottom tercile of baseline Z	14.03	23.08	75.33	42.37
SE	(18.38)	(22.00)	(26.72)	(26.39)

Notes

- 1 The covariates T1 Baseline gross profit, T1 Baseline livestock, T1 Baseline food consumption and T1 Baseline non-food consumption are all indicator variables which are 1 if the household was in the bottom tercile of the baseline distribution of a that variable and 0 otherwise.
- 2 See the notes of Table 2 for details on additional controls.

Appendix Table 8: Additional Outcomes for Loan Intent to Treat

	Own any livestock	Total value of livestock	HH has a business	Food consumption EQ (past 7 days)	Monthly non-food exp	HH has any financial savings	Educ expenses	Medical expenses	Intra HH Decision- making Index	Community Action Index	Social Capital Index
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Loan village - year 1	0.01 (0.01)	-35.54 (76.72)	-0.03 (0.02)	-0.46 (0.17)	-0.19 (2.10)	0.02 (0.02)	2.73 (4.01)	-5.03 (1.64)	-0.0005 (0.043)	0.047 (0.049)	-0.001 (0.048)
Loan village - year 2	-0.01 (0.02)	189.48 (94.68)	0.02 (0.01)	0.75 (0.20)	-0.60 (2.50)	0.00 (0.03)	1.91 (3.44)	-1.36 (1.81)	0.038 (0.054)	0.062 (0.048)	0.043 (0.043)
N	8634	8558	8634	8323	8297	8533	6021	8550	7900	7769	7808
Mean of control (year 1)	0.78	1219.43	0.83	5.96	43.97	0.63	69.87	33.46	0.035	-0.115	-0.063
SD (year 1)	(0.42)	(2070.58)	(0.37)	(3.16)	(37.67)	(0.48)	(81.20)	(45.44)	0.958	0.881	0.933
Per \$100 impact, TOT, year 1	0.04 (0.06)	475.88 (315.69)	-0.03 (0.10)	0.89 (0.74)	-0.82 (8.85)	0.07 (0.10)	11.49 (16.89)	-21.18 (6.90)	-0.002 (0.18)	0.20 (0.21)	-0.004 (0.20)

Notes

- 1 See the notes of Table 16 for details on specification.
- 2 Standard errors are in parentheses and clustered at the village level in all specifications.
- 3 Mean of control is the mean of the dependent variable in the column heading among households in no-loan villages.
- 4 The per dollar return, TOT, year 1 is: the coefficient on Loan village - year 1 / (.21*113) since the average value of the loan was \$113. The standard error on the

Appendix Table 9: Additional Outcomes - Impact of Grants in Year 2

	Own any livestock (0/1)	Total value of livestock (USD)	HH has a business (0/1)	Food consumption EQ (past 7 days, USD)	Monthly non- food exp (USD)	HH has any financial savings (0/1)	Educ expenses (USD)	Medical expenses (USD)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Grant	0.093 (0.015)	193.70 (99.51)	0.03 (0.01)	0.24 (0.19)	4.16 (2.10)	0.032 (0.019)	0.61 (3.68)	-0.61 (1.81)
Grant * loan village	0.006 (0.024)	-157.80 (135.20)	-0.02 (0.02)	0.21 (0.24)	-1.52 (2.72)	0.038 (0.026)	1.25 (5.24)	1.41 (2.74)
Grant + Grant * loan village = 0	0.000	0.696	0.704	0.005	0.126	0.000	0.619	0.698
N	5198	5146	5201	5003	5009	5143	3621	5151
Mean of control	0.77	1474.87	0.87	7.29	47.64	0.66	75.55	33.46
SD	(0.42)	(2509.91)	(0.33)	(3.93)	(43.22)	(0.47)	(85.11)	(47.30)

Notes

- 1 Rows showing Grant + Grant * loan village = 0 (year 1) shows the p value of the test of whether the total effect of grants in loan villages is statistically different from zero.
- 2 See the notes of Table 2 for details on specifications.