Selling Low and Buying High: Arbitrage and Local Price Effects in Kenyan Markets

Marshall Burke,[1,2,3]* Lauren Falcao Bergquist,[4] Edward Miguel[3,5]

[1]Department of Earth System Science, Stanford University
[2]Center on Food Security and the Environment, Stanford University
[3]National Bureau of Economic Research
[4]Becker Friedman Institute, University of Chicago
[5]Department of Economics, University of California, Berkeley

November 28, 2017

Abstract

Large and regular seasonal price fluctuations in local grain markets appear to offer African farmers substantial inter-temporal arbitrage opportunities, but these opportunities remain largely unexploited: small-scale farmers are commonly observed to "sell low and buy high" rather than the reverse. In a field experiment in Kenya, we show that credit market imperfections limit farmers' abilities to move grain inter-temporally. Providing timely access to credit allows farmers to purchase at lower prices and sell at higher prices, increasing farm profits and generating a return on investment of 28%. To understand general equilibrium effects of these changes in behavior, we vary the density of loan offers across locations. We document significant effects of the credit intervention on seasonal price fluctuations in local grain markets, and show that these GE effects greatly affect our individual level profitability estimates. In contrast to existing experimental work, our results thus indicate a setting in which microcredit can improve firm profitability, and suggest that GE effects can substantially shape estimates of microcredit's effectiveness. Failure to consider these GE effects could lead to substantial misestimation of the social welfare benefits of microcredit interventions.

JEL codes: D21, D51, G21, O13, O16, Q12

Keywords: storage; arbitrage; microcredit; credit constraints; agriculture

*We thank Kyle Emerick, Jeremy Magruder, and Chris Barrett for useful discussions, and thank seminar participants at Berkeley, Stanford, Kellogg, ASSA, and PacDev for useful comments. We also thank Peter LeFrancois, Ben Wekesa, and Innovations for Poverty Action for excellent research assistance in the field, and One Acre Fund for partnering with us in the intervention. We gratefully acknowledge funding from the Agricultural Technology Adoption Initiative and an anonymous donor. All errors are our own.

1 Introduction

Imperfections in credit markets have long been considered to play a central role in underdevelopment (Banerjee and Newman, 1993; Galor and Zeira, 1993; Banerjee and Duflo, 2010), with these imperfections thought to have particularly large consequences for small and informal firms in the developing world and for the hundreds of millions of poor people who own and operate them. This thinking has motivated a large-scale effort to expand credit access to existing or would-be microentrepreneurs around the world, and it has also motivated a subsequent attempt on the part of academics to rigorously evaluate the effects of this expansion on the productivity of these microenterprises and on the livelihoods of their owners.

Findings in this rapidly growing literature have been remarkably heterogenous. Studies that provide cash grants to households and to existing small firms suggest high rates of return to capital in some settings but not in others.[1] Further, experimental evaluations of traditional microcredit products (small loans to poor households) have generally found that individuals randomly provided access to these products are subsequently no more productive on average than those not given access, but that subsets of recipients often appear to benefit.[2]

Here we study a unique microcredit product designed to improve the profitability of small farms – a setting that has been largely outside the focus of the experimental literature on credit constraints. Farmers in our setting in Western Kenya, as well as throughout much of the rest of the developing world, face large and regular seasonal fluctuations in grain prices, with increases of 40-50% between post-harvest lows and pre-harvest peaks common in local markets. Nevertheless, most of these farmers have difficulty using storage to move grain from times of low prices to times of high prices, and this inability appears at least in part due to limited borrowing opportunities: lacking access to credit or savings, farmers report selling their grain at low post-harvest prices to meet urgent cash needs (e.g., to pay school fees). To meet consumption needs later in the year,

[1] Studies finding high returns to cash grants include De Mel et al. (2008); McKenzie and Woodruff (2008); Fafchamps et al. (2013); Blattman et al. (2013). Studies finding much more limited returns include Berge et al. (2011) and Karlan et al. (2012).

[2] Experimental evaluations of microcredit include Attanasio et al. (2011); Crepon et al. (2011); Karlan and Zinman (2011); Banerjee et al. (2013); Angelucci et al. (2013) among others. See Banerjee (2013) and Karlan and Morduch (2009) for nice recent reviews of these literatures.

many then end up buying back grain from the market a few months after selling it, in effect using the maize market as a high-interest lender of last resort (Stephens and Barrett, 2011).

Working with a local agricultural microfinance NGO, we study the role that credit constraints play in farmers' inability to store grain and arbitrage these seasonal price fluctuations. We offer randomly selected smallholder maize farmers a loan at harvest,[3] and study whether access to this loan improves their ability to use storage to arbitrage local price fluctuations relative to a control group. We find that farmers offered this harvest-time loan sell significantly less and purchase significantly more maize in the period immediately following harvest, and this pattern reverses during the period of higher prices 6-9 months later. This change in the marketing behavior results in a statistically significant increase in revenues (net of loan interest) of 545Ksh, suggesting that the loan produces a return on investment of 28%. We replicate the experiment in two back-to-back years to test the robustness of these results and find remarkably similar results on primary outcomes in both years.

Given the high transport costs in our rural African setting, we also study whether storage-related changes in marketing behavior affected local market prices. Did this individual-level intervention have market-level effects? To answer this, we experimentally varied the density of treated farmers across locations and tracked market prices at 52 local market points. We find that the greater storage of grain at the market level (induced by the credit intervention) led to significantly higher prices immediately after harvest and to lower (albeit not significantly so) prices during the lean season. Discernible price effects from such a localized shift in supply imply that agricultural markets in the region are highly fragmented.

We find that these general equilibrium effects greatly alter the profitability of the loan. By dampening the arbitrage opportunity posed by season price fluctuations, treated individuals in high saturated areas show diminished revenue impacts relative to farmers in lower saturation areas. We find that while treated farmers in high-saturation areas store significantly more than their control counterparts, doing so is not significantly more profitable; the reduction in seasonal price

[3] This is unusual - and seemingly counter-intuitive - timing for a loan to agricultural households; our microfinance NGO partner and many other groups offer loans at planting time in order to facilitate farmer adoption of high quality inputs such as fertilizer.

dispersion in these area may reduce the benefits of loan adoption. In contrast, treated farmers in low-density areas have both significantly higher inventories and significantly higher profits relative to control.

These general equilibrium effects — and their impact on loan profitability at the individual level — have lessons for both policy and evaluation. In terms of policy, the general equilibrium effects shape the distribution of the welfare gains of the harvest-time loan: while recipients gain relatively less than they would in the absence of such effects, we find suggestive evidence that non-recipients benefit from smoother prices, even though their storage behavior remains unchanged. Though estimated effects on non-treated individuals are measured with substantial noise, a welfare calculation taking the point estimates at face-value suggests that 70% of overall gains in high-treatment-intensity areas accrued to program non-recipients. These gains to non-recipients, which cannot be readily recouped by private sector lending institutions, may provide some incentive for public provision of such products.

The eroding profitability of arbitrage that we observe also has implications for impact evaluation in contexts of highly fragmented markets. In these settings in which general equilibrium effects are likely to be more pronounced and the SUTVA assumption (Rubin, 1986) more likely to be violated, an evaluation of a simple individually-randomized loan product could have difficulty discerning null effects from large positive effects on social welfare. While this issue may be particularly salient in our context of a loan explicitly designed to enable arbitrage, it is by no means unique to our setting. Any enterprise operating in a small, localized market or in a concentrated industry may face price responses to shifts in own supply, and credit-induced expansion may therefore be less profitable than it would be in more integrated market or in a less concentrated industry. Proper measurement of these impacts requires a study design with exogenous variation in these general equilibrium effects.

We also run a long-run follow-up survey with respondents 1-2 years after harvest-time credit intervention had been discontinued by the NGO, to test whether farmers are able to use the additional revenues earned from this loan product to "save their way out" of credit constraints in future years. While we find no evidence of sustained shifts in the timing of farm sales in subsequent seasons, nor do we see long-run effects on sales or revenues in future years, we find some evidence of heterogeneity by treatment saturation.

Why do we find positive effects on firm profitability when many other experimental studies on microcredit do not? Existing studies have offered a number of explanations for why improved access to capital does not appear beneficial on average. First, many small businesses or potential micro-entrepreneurs simply might not face profitable investment opportunities (Banerjee et al., 2013; Fafchamps et al., 2013; Karlan et al., 2012; Banerjee, 2013).[4] Second, profitable investment opportunities could exist but microentrepreneurs might lack either the skills or ability to channel capital towards these investments - e.g. if they lack managerial skills (Berge et al., 2011; Bruhn et al., 2012), or if they face problems of self-control or external pressure that redirect cash away from investment opportunities (Fafchamps et al., 2013). Third, typical microcredit loan terms require that repayment begin immediately, and this could limit investment in illiquid but high-return business opportunities (Field et al., 2012). Finally, as described above, general equilibrium effects of credit expansion could alter individual-level treatment effect estimates in a number of ways, potentially shaping outcomes for both treated and untreated individuals. This is a recognized but unresolved problem in the experimental literature on credit, and few experimental studies have been explicitly designed to quantify the magnitude of these general equilibrium effects (Acemoglu, 2010; Karlan et al., 2012).[5]

All of these factors likely help explain why our results diverge from existing estimates. Unlike most of the settings examined in the literature, using credit to "free up" storage for price arbitrage does not require starting or growing a business among this population of farmers, is neutral to the scale of farm output, does not appear to depend on entrepreneurial skill (all farmers have stored

[4]For example, many microenterprises might have low efficient scale and thus little immediate use for additional investment capital, with microentrepreneurs then preferring to channel credit toward consumption instead of investment. Relatedly, marginal returns to investment might be high but total returns low, with the entrepreneur making the similar decision that additional investment is just not worth it.

[5]For instance, Karlan et al. (2012) conclude by stating, "Few if any studies have satisfactorily tackled the impact of improving one set of firms' performance on general equilibrium outcomes. . . . This is a gaping hole in the entrepreneurship development literature." Indeed, positive spillovers could explain some of the difference between the experimental findings on credit, which suggest limited effects, and the estimates from larger-scale natural experiments, which tend to find positive effects of credit expansion on productivity – e.g. Kaboski and Townsend (2012). Acemoglu (2010) uses the literature on credit market imperfections to highlight the understudied potential role of GE effects in broad questions of interest to development economists.

before, and all are very familiar with local price movements), and does not require investment in a particularly illiquid asset (inventories are kept in the house and can be easily sold). Farmers do not even have to sell grain to benefit from credit in this context: a net-purchasing farm household facing similar seasonal cash constraints could use credit and storage to move its purchases from times of high prices to times of lower prices.

Furthermore, our results also suggest that – at least in our rural setting – treatment density matters and market-level spillovers can substantially shape individual-level treatment effect estimates. Whether these GE also influenced estimated treatment effects in the more urban settings examined in many previous studies is unknown, although there is some evidence that spillovers do matter for microenterprises who directly compete for a limited supply of inputs to production.[6] In any case, our results suggest that explicit attention to GE effects in future evaluations of credit market interventions is likely warranted.

Beyond contributing to the experimental literature on microcredit, our paper is closest to a number of recent papers that examine the role of borrowing constraints in households' storage decisions and seasonal consumption patterns. Using secondary data from Kenya, Stephens and Barrett (2011) also suggest that credit constraints substantially alter smallholder farmers' marketing and storage decisions, and Basu and Wong (2012) show that allowing farmers to borrow against future harvests can substantially increase lean-season consumption. Similarly, Dillion (2017) finds that an administrative change in the school calendar that moved the timing of school fee payments to earlier in the year in Malawi forced credit constrained households with school-aged children to sell their crops earlier and at a lower price, and Fink et al. (2014) find that agricultural loans aimed at alleviated seasonal labor shortages can improve household welfare in Zambia.

As in these related papers, our results show that when borrowing and saving are difficult, households turn to increasingly costly ways to move consumption around in time. In our particular setting, credit constraints combined with post-harvest cash needs cause farmers to store less than they would in an unconstrained world. In this setting, even a relatively modest expansion of credit affects local market prices, to the apparent benefit of both those with and without access to this

[6]See De Mel et al. (2008) and their discussion of returns to capital for firms in the bamboo sector, all of whom in their setting compete over a limited supply of bamboo.

credit.

Finally, our results speak to an earlier literature showing how credit market imperfections can combine with other features of economies to generate observed broad-scale economic patterns (Banerjee and Newman, 1993; Galor and Zeira, 1993). These earlier papers showed how missing markets for credit, coupled with an unequal underlying wealth distribution, could shape large-scale patterns of occupational choice. We show that missing markets for credit combined with climate-induced seasonality in rural income can help generate widely-observed seasonal price patterns in rural grain markets, patterns that appear to further worsen poor households' abilities to smooth consumption across seasons. Evidence that the expansion of harvest-time credit access helps reduce this price dispersion suggests an under-appreciated but likely substantial additional benefit of credit expansion in rural areas.

The remainder of the paper proceeds as follows. Section 2 describes the setting and the experiment. Section 3 describes our data, estimation strategy, and pre-analysis plan. Section 4 presents baseline estimates ignoring the role of general equilibrium effects. Section 5 presents the market level effects of the intervention. Section 6 shows how these market-level effects shape the individual-level returns to the loan. Section 7 concludes.

2 Setting and experimental design

2.1 Arbitrage opportunities in rural grain markets

Seasonal fluctuations in prices for staple grains appear to offer substantial intertemporal arbitrage opportunities, both in our study region of East Africa as well as in other parts of Africa and elsewhere in the developing world. While long term price data unfortunately do not exist for the small markets in very rural areas where our experiment takes place, price series are available for major markets throughout the region. Average seasonal price fluctuations for maize in available markets are shown in Figure 1. Increases in maize prices in the six to eight months following harvest average roughly 25-40% in these markets, and these increases appear to be a lower bound on seasonal price increases reported elsewhere in Africa.[7]

These increases also appear to be a lower bound on typical increase observed in the smaller markets in our study area, which (relative to these much larger markets) are characterized with much smaller "catchments" and less outside trade. We asked farmers at baseline to estimate average monthly prices of maize at their local market point over the five years prior to our experiment. As shown in Figure 4, they reported a typical doubling in price between September (the main harvest month) and the following June.[8] We also collected monthly price data from local market points in our sample area during the two years of this study's intervention, as well as for a year after the intervention ended (more on this data collection below).[9] Figure 5 presents the price fluctuations observed during this period. Unfortunately, because data collection began in November 2012 (two months after the typical trough in September), we cannot calculate the full price fluctuation for the 2012-2013 season. However, in the 2013-2014 and 2014-2015 seasons we observe prices increasing by 42% and 45% respectively. These are smaller fluctuations than those seen in prior years (as reported by farmers in our sample) and smaller than those seen in subsequent years, which saw increases of 53% and 125% respectively.[10] There is therefore some variability in the precise size of the price fluctuation from season to season. Nevertheless, we see price consistently rise by more than 40% and, in some years, by substantially more.

Farmers do not appear to be taking advantage of these apparent arbitrage opportunities. Figure A.1 shows data from two earlier pilot studies conducted either by our NGO Partner (in 2010/11, with 225 farmers) or in conjunction with our partner (in 2011/12, with a different sample of 700 farmers). These studies tracked maize inventories, purchases, and sales for farmers in our study

[7]For instance, Barrett (2008) reports seasonal rice price variation in Madagascar of 80%, World Bank (2006) reports seasonal maize price variation of about 70% in rural Malawi, and Aker (2012) reports seasonal variation in millet prices in Niger of 40%.

[8]In case farmers were somehow mistaken or overoptimistic, we asked the same question of the local maize traders that can typically be found in these market points. These traders report very similar average price increases: the average reported increase between October and June across traders was 87%. Results available on request.

[9]The study period covers the 2012-2013 and 2013-2014 season. We also collect data for one year after the study period, covering the 2014-2015 season, in order to align with the long-run follow-up data collection on the farmer side.

[10]For the 2015-2016 season, we combine our data with that collected by Bergquist (2017) in the same county in Kenya and estimate that maize prices increased by 53% from November to June. For the 2016-2017 season, we thank Pascaline Dupas for her generosity in sharing maize price data collected in the same county in November 2016 and June 2017, from which we estimate an increase of 125%.

region. In both years, the median farmer exhausted her inventories about 5 months after harvest, and at that point switched from being a net seller of maize to a net purchaser as shown in the right panels of the figure. This was despite the fact that farmer-reported sales prices rose by more than 80% in both of these years in the nine months following harvest.

Why are farmers not using storage to sell grain at higher prices and purchase at lower prices? Our experiment is designed to test the role of credit constraints in shaping storage and marketing decisions. In extensive focus groups with farmers prior to our experiment, credit constraints were the (unprompted) explanation given by the vast majority of these farmers as to why they were not storing and selling maize at higher prices. In particular, because nearly all of these farm households have school aged kids, and a large percentage of a child's school fees are typically due in the few months after harvest in January, given the calendar-year school year schedule, many farmers report selling much of their harvest to pay these fees. Indeed, many schools in the area will accept in-kind payment in maize during this period. Farmers also report having to pay other bills they have accumulated throughout the year during the post-harvest period.

Further, as with poor households throughout much of the world, these farmers appear to have very limited access to formal credit. Only eight percent of households in our sample reported having taking a loan from a bank in the year prior to the baseline survey.[11] Informal credit markets also appear relatively thin, with less than 25% of farmers reporting having given or received a loan from a moneylender, family member, or friend in the 3 months before the baseline.

Absent other means of borrowing, and given these various sources of "non-discretionary" consumption they report facing in the post-harvest period, farmers end up liquidating grain rather than storing. Furthermore, a significant percentage of these households end up buying back maize from the market later in the season to meet consumption needs, and this pattern of "selling low and buying high" directly suggests a liquidity story: farmers are in effect taking a high-interest quasi-loan from the maize market (Stephens and Barrett, 2011). Baseline data indicate that 35% of our sample both bought and sold maize during the previous crop year (September 2011 to August

[11] Note that even at the high interest rates charged by formal banking institutions (typically around 20% annually), storage would remain profitable, given the 40% plus (often much larger) increases in prices that are regularly observed over the 9-month post-harvest period and relatively small storage losses (e.g., due to spoilage), which we estimate to be less than 5%.

2012), and that over half of these sales occurred before January (when prices were low). 40% of our sample reported only purchasing maize over this period, and the median farmer in this group made all of their purchases after January. Stephens and Barrett (2011) report very similar patterns for other households in Western Kenya during an earlier period.

Nevertheless, there could be other reasons beyond credit constraints why farmer are not taking advantage of apparent arbitrage opportunities. The simplest explanations are that farmers do not know about the price increases, or that it is actually not profitable to store – i.e. arbitrage opportunities are actually much smaller than they appear because storage is costly. These costs could come in the form of losses to pests or moisture-related rotting, or they could come in the form of "network losses" to friends and family, since maize is stored in the home and is visible to friends and family, and there is often community pressure to share a surplus. Third, farmers could be highly impatient and thus unwilling to move consumption to future periods in any scenario. Finally, farmers might view storage as too risky an investment.

Evidence from pilot and baseline data, and from elsewhere in the literature, argues against several of these possibilities. We can immediately rule out an information story: farmers are well-aware that prices rise substantially throughout the year. When asked in our baseline survey about expectations for the subsequent season's price trajectory, the average farmer expected prices to increase by 107% in the nine months following the September 2012 harvest (which was actually an over-estimate of the realized price fluctuation that year).[12] Second, pest-related losses appear surprisingly low in our setting, with farmers reporting losses from pests and moisture-related rotting of 2.5% for maize stored for six to nine months. Similarly, the marginal costs associated with storing for these farmers are small (estimates suggest that the cost per bag is about 3.5% of the harvest-time price) and the fixed costs have typically already been paid (all farmers store at least some grain; note the positive initial inventories in Figure A.1), as grain in simply stored in the household or in small sheds previously built for the purpose.[13] Third, while we cannot rule out impatience as a driver of low storage rates, extremely high discount rates would be needed to rationalize this behavior

[12] The 5th, 10th, and 25th percentiles of the distribution are a 33%, 56%, and 85% increase, respectively, suggesting that nearly all farmers in our sample expect substantial price increases.

[13] Though note that Aggarwal et al. (2017) find that offering group-based grain storage can encourage greater storage.

in light of the substantial prices increase seen over a short nine-month period.[14] Furthermore, farm households are observed to make many other investments with payouts far in the future (e.g. school fees), meaning that rates of time preference would also have to differ substantially across investments and goods. Fourth, existing literature shows that for households that are both consumers and producers of grain, aversion to price risk should motivate *more* storage rather than less: the worst state of the world for these households is a huge price spike during the lean season, which should motivate "precautionary" storage (Saha and Stroud, 1994; Park, 2006).

Costs associated with network-related losses appear a more likely explanation for an unwillingness to store substantial amounts of grain. Existing literature suggests that community pressure is one explanation for limited informal savings (Dupas and Robinson, 2013; Brune et al., 2011), and in focus groups farmers often told us something similar about stored grain (itself a form of savings). As described below, our main credit intervention might also provide farmers a way to shield stored maize from their network. To further test this hypothesis, in the first year of the experiment we add an additional treatment arm to determine whether this shielding effect is substantial on its own.

2.2 Experimental design

The study sample is drawn from existing groups of One Acre Fund (OAF) farmers in Webuye and Matete districts in Western Kenya. OAF is a microfinance NGO that makes in-kind, joint-liability loans of fertilizer and seed to groups of farmers, as well as providing training on improved farming techniques. OAF group sizes typically range from 8-12 farmers, and farmer groups are organized into "sublocations" – effectively clusters of villages that can be served by one OAF field officer. OAF typically serves about 30% of farmers in a given sublocation.

The Year 1 sample consists of 240 existing OAF farmer groups drawn from 17 different sublocations in Webuye district, and our total sample size at baseline was 1,589 farmers. The Year 2 sample attempted to follow the same OAF groups as Year 1; however, some groups dissolved such

[14] Given a minimum price increase of 40%, post-harvest losses of 2.5%, and storage costs of 3.5% of price, an individual would have to discount the 9-month future by over 33% to make the decision to sell at harvest rational under no other constraints. Given the distribution of estimated discount rates from a time preference question asked at baseline, this would apply to only 12% of our sample.

that in Year 2 we are left with 171 groups. In addition, some of the groups experienced substantial shifting of the individual members; therefore some Year 1 farmers drop out of our Year 2 sample, and other farmers are new to our Year 2 sample.[15] Ultimately, of the 1,019 individuals in our Year 2 sample, 602 are drawn from the Year 1 sample and 417 are new to the sample.

Figure 2 displays the experimental design. There are two main levels of randomization. First, we randomly divided the 17 sublocations in our sample into 9 "high" intensity" sites and 8 "low intensity" sites. In high intensity sites, we enrolled 80% of OAF groups in the sample (for a sample of 171 groups), while in low intensity sites, we only enrolled 40% of OAF groups in the sample (for a sample of 69 groups). Then, within each sublocation, groups were randomized into treatment or control. In Year 1, two-thirds of individuals in each sublocation were randomized into treatment (more on this below) and one-third into control. In Year 2, half of individuals in each sublocation were randomized into treatment and half into control. As a result of this randomization procedure, high intensity sublocations have double the number of treated individuals as in low intensity sublocations.

The group-level randomization was stratified at the sublocation level (and in Year 1, for which we had administrative data, further stratified based on whether group-average OAF loan size in the previous year was above or below the sample median). In Year 2, we maintained the same saturation treatment status at the sublocation level,[16] but re-randomized groups into treatment and control, stratifying on their treatment status from Year 1.[17] Given the roughly 35% reduction in overall sample size in Year 2, overall treatment saturation rates (the number of treated farmers per sublocation) were effectively 35% lower in Year 2 as compared to Year 1.

In Year 1, there was a third level of randomization pertaining to the timing of the loan offer.

[15]Shifting of group members is a function of several factors, including whether farmers wished to participate in the overall OAF program from year to year. There was some (small) selective attrition based on treatment status in Year 1; treated individuals were 10 percentage points more likely to return to the Year 2 sample than control individuals (significant at 1%). This does slightly alter the composition of the Year 2 sample (see Table K.2 and Section K), but because Year 2 treatment status is stratified by Year 1 treatment status (as will be described below), it does not alter the internal validity of the Year 2 results.

[16]Such that, for example, if a sublocation was a high intensity sublocation in Year 1 it remained a high intensity sublocation in Year 2.

[17]This was intended to result in randomized duration of treatment – either zero years of the loan, one year of the loan, or two years – however, due to selective attrition of the Year 1 sample based on treatment status, duration of loan treatment is no longer entirely random.

In focus groups run prior to the experiment, farmers were split on when credit access would be most useful, with some preferring cash immediately at harvest, and others preferring it a few months later timed to coincide with when school fees were due (the latter preferences suggesting that farmers may be sophisticated about potential difficulties in holding on to cash between the time it was disbursed and the time it needed to be spent). In order to test the importance of loan timing, in Year 1, a random half of the treated group (so a third of the total sample) received the loan in October (immediately following harvest), while the other half received the loan in January (immediately before school fees are due, although still several months before the local lean season). As will be described in Section 4, results from Year 1 suggested that the earlier loan was more effective, and therefore in Year 2 the NGO only offered the earlier timed loan to the full sample (though due to administrative delays, the actual loan was disbursed in November in Year 2).

Although all farmers in each loan treatment group were offered the loan, we follow only a randomly selected 6 farmers in each loan group, and a randomly selected 8 farmers in each of the control groups.

Loan offers were announced in September in both years. To qualify for the loan, farmers had to commit maize as collateral, and the size of the loan they could qualify for was a linear function of the amount they were willing to collateralize (capped at 7 bags in Year 1 and 5 bags in Year 2). In Year 1, to account for the expected price increase, October bags were valued at 1500Ksh, and January bags at 2000Ksh. In Year 2, bags were valued at 2500Ksh. Each loan carried with it a "flat" interest rate of 10%, with full repayment due after nine months.[18][19] These loans were an add-on to the existing in-kind loans that OAF clients received, and OAF allows flexible repayment of both – farmers are not required to repay anything immediately.

Collateralized bags of maize were tagged with a simple laminated tag and zip tie. When we mentioned in focus groups the possibility of OAF running a harvest loan program, and described the details about the collateral and bag tagging, many farmers (unprompted) said that the tags

[18] Annualized, this interest rate is slightly lower than the 16-18% APR charged on loans at Equity Bank, the main rural lender in Kenya.

[19] For example, a farmer who committed 5 bags when offered the October loan in Year 1 would receive 5*1500 = 7500Ksh in cash in October (∼$90 at current exchange rates), and would be required to repay 8250Ksh by the end of July.

alone would prove useful in shielding their maize from network pressure: "branding" the maize as committed to OAF, a well-known lender in the region, would allow them to credibly claim that it could not be given out.[20] Because tags could represent a meaningful treatment in their own right, we wished to separate the effect of the credit from any effect of the tag, and therefore in the Year 1 study offered a separate treatment arm in which groups received only the tags.[21]

Finally, because self- or other-control problems might make it particularly difficult to channel cash toward productive investments in settings where there is a substantial time lag between when the cash is delivered and when the desired investment is made, in Year 1, we also cross-randomized a simple savings technology that had shown promise in a nearby setting (Dupas and Robinson, 2013). In particular, a subset of farmers in each loan treatment group in Year 1 were offered a savings lockbox (a simple metal box with a sturdy lock) which they could use as they pleased. While such a savings device could have other effects on household decision making, our hypothesis was that it would be particularly helpful for loan clients who received cash before it was needed.

The tags and lockbox treatments were randomized at the individual level during Year 1. These treatments were not included in Year 2 due to minimal treatment effects in Year 1 data (discussed below), as well as the somewhat smaller sample size in Year 2. Using the sample of individuals randomly selected to be followed in each group, we stratified individual level treatments by group treatment assignment and by gender. So, for instance, of all of the women who were offered the October loan and who were randomly selected to be surveyed, one third of them were randomly offered the lockbox (and similarly for the men and for the January loan). In the control groups, in which we were following 8 farmers, 25% of the men and 25% of the women were randomly offered the lockbox, with another 25% each being randomly offered the tags. The study design allows identification of the individual and combined effects of the different treatments, and our approach for estimating these effects is described below.

[20]Such behavior is consistent with evidence from elsewhere in Africa that individuals take out loans or use commitment savings accounts mainly as a way to demonstrate that they have little to share (Baland et al., 2011; Brune et al., 2011).

[21]This is not the full factorial research design – there could be an interaction between the tag and the loan – but we did not have the sample size to do the full 2 x 2 design to isolate any interaction effect.

3 Data and estimation

The timing of the study activities is shown in Figure 3. In August/September 2012 (prior to the Year 1 experiment), a baseline survey was conducted with the entire Year 1 sample. The baseline survey collected data on farming practices, on storage costs, on maize storage and marketing over the previous crop year, on price expectations for the coming year, on food and non-food consumption expenditure, on household borrowing, lending, and saving behavior, on household transfers with other family members and neighbors, on sources of non-farm income, on time and risk preferences, and on digit span recall.

We then undertook three follow-up rounds over the ensuing 12 months, spanning the spring 2013 "long rains" planting (the primary growing season) and concluding just prior to the 2013 long rains harvest (which occurs August-September). The multiple follow-up rounds were motivated by three factors. First, a simple inter-temporal model of storage and consumption decisions suggests that while the loan should increase total consumption across all periods, the per-period effects could be ambiguous – meaning that consumption throughout the follow-up period needs to be measured to get at overall effects. Second, because nearly all farmers deplete their inventories before the next harvest, inventories measured at a single follow-up one year after treatment would likely provide very little information on how the loan affected storage and marketing behavior. Finally, as shown in McKenzie (2012), multiple follow-up measurements on noisy outcomes variables (e.g consumption) has the added advantage of increasing power. A similar schedule of three follow-up rounds over 12 months were run in Year 2.[22] The follow-up surveys tracked data on storage inventory, maize marketing behavior, consumption, and other credit and savings behavior. Follow-up surveys also collected information on time preferences and on self-reported happiness.

In order to explore the long-run effects of the loan, we also ran a Long-Run Follow-Up (LRFU) survey from November-December 2015. This was two (one) years following loan repayment for the

[22]Because the Year 2 experiment was meant to follow the sample sample as Year 1, a second baseline was not run prior to Year 2. However, as described in Section 2, due to administrative shifts in farmer group composition, 417 of the 1,019 individuals in the Year 2 sample were new to the study. For these individuals, we do not have baseline data (there was insufficient time between receiving the updated administrative records for Year 2 groups and the disbursal of the loan to allow for a second baseline to be run). Therefore, balance tables can only be run with the sample that was present in Year 1. Because the loan offer was randomized, however, this should not meaningfully affect inference regarding the impacts of the loan.

Year 1 (Year 2) treatment group. This survey followed up on the entire Year 2 sample (1,091 individuals) and a representative subset of the Year 1 only sample (another 481 individuals), for a total sample of 1500 individuals. The survey collected information on maize harvests, sales, purchases, and revenues from 2014-2015 (broken down by harvest and lean season). It also collected data on farm inputs (labor and capital), food consumption and expenditure, household consumption, educational expenditure and attendance among children, non-farm employment and revenues, and a self-reported happiness measure. We were able to track 91.5% of the intended sample. There is no differential attrition based on Year 2 treatment status. While there is some suggestive evidence of differential attrition based on Year 1 treatment status (being treated in Year 1 is associated with 3 percentage point increase in the likelihood of being found in the long-run follow up survey, significant at 10%), this is partially driven by the fact that Year 1 treated individuals were more likely to be in the Year 2 sample (and therefore had been more recently in touch with our survey team). After controlling for whether an individual was present in the Year 2 sample, Year 1 treatment status is no longer significantly correlated with attrition.

In addition to farmer-level surveys, we also collected monthly price surveys at 52 market points in the study area. The markets were identified prior to treatment based on information from local OAF staff about the market points in which client farmers typically buy and sell maize. Data collection for these surveys began in November 2012 and continued through December 2015. Finally, we utilize administrative data on loan repayment that was generously shared by OAF.

Table 1 shows summary statistics for a range of variables at baseline, and shows balance of these variables across the three main loan treatment groups. Groups are well balanced, as would be expected from randomization. Table J.1 shows the analogous table comparing individuals in the high- and low-treatment-density areas; samples appear balanced on observables here as well. Attrition was also relatively low across our survey rounds. In Year 1, overall attrition was 8%, and not significantly different across treatment groups (8% in the treatment group and 7% in the control). In Year 2, overall attrition was 2% (in both treatment and control, with no significant difference). There was some small selective attrition the Year 1 to the Year 2 sample based on Year 1 treatment status, as mentioned above. This does slightly alter the composition of the Year

2 sample (see Table K.2), but because Year 2 treatment status is stratified by Year 1 treatment status, it does not alter the internal validity of the Year 2 results. Appendix K explores this further.

3.1 Pre-analysis plan

To limit both risks and perceptions of data mining and specification search (Casey et al., 2012), we specified and registered a pre-analysis plan (PAP) for Year 1 prior to the analysis of any follow-up data.[23] The Year 2 analysis follows a near identical analysis plan. Both the PAP and the complete set of results are available upon request.

We deviate significantly from the PAP in one instance: the PAP specifies the outcome of interest to be the percent price spread from November to June. However, because in practice the loan was offered at slight different points in time (October and January in Year 1; November in Year 2) and because there is year-to-year variation in when markets hit their peak and trough, this measure may fail to capture the full effect of treatment on prices. Moreover, this measure is vastly underpowered, ignoring 77% of our monthly data by focusing solely on the price gap between two months, rather than exploiting the full nine months of data collected over the season.

Therefore, in our primary specifications, we relax our attachment to this underpowered and perhaps misspecified measure November-June price gap, instead showing the non-parametric effect of treatment on the evolution of monthly prices, as well as a level and time trend effect.[24]. We maintain our original hypothesis that effect of high-density treatment on prices will be initially positive if receipt of the loan allows farmers to pull grain off the market in the post-harvest surplus period and later negative as stored grain is released onto the market.

In two other instances we add to the PAP. First, in addition to the regression results specified in the PAP, we also present graphical results for many of the outcomes. These results are based on non-parametric estimates of the parametric regressions specified in the PAP, and are included because they clearly summarize how treatment effects evolve over time, but since they were not explicitly specified in the PAP we mention them here. Second, we failed to include in the PAP the

[23]The pre-analysis plan is registered here: https://www.socialscienceregistry.org/trials/67, and was registered on September 6th 2013.

[24]Appendix H.3 presents the pre-specified November-June effect.

(rather obvious) regressions in which the individual-level treatment effect is allowed to vary by the sublocation-level treatment intensity, and present these below.

3.2 Estimation of treatment effects

In all analyses, we present results separately by year and pooled across years. Because the Year 2 replication produced results that are quantitatively quite similar to the Year 1 results for most outcomes, we rely on the pooled results as our specification of primary interest. However, for the sake of transparency and comparison, we report both.

There are three main outcomes of interest: inventories, maize net revenues, and consumption. Inventories are the number of 90kg bags of maize the household had in their maize store at the time of the each survey. This amount is visually verified by our enumeration team, and so is likely to be measured with minimal error. We define maize net revenues as the value of all maize sales minus the value of all maize purchases, and minus any additional interest payments made on the loan for individuals in the treatment group. We call this "net revenues" rather than "profits" since we likely do not observe all costs; nevertheless, costs are likely to be very similar across treatment groups (fixed costs of storing at home were already paid, and variable costs of storage are very low). The values of sales and purchases were based on recall data over the period between each survey round. Finally, we define consumption as the log of total per capita household expenditure over the 30 days prior to each survey. For each of these variables we trim the top and bottom 0.5% of observations, as specified in the pre-analysis plan.

Letting T_{jy} be an an indicator for whether group j was assigned to treatment in year y, and y_{ijry} as the outcome of interest for individual i in group j in round $r \in (1, 2, 3)$ in year y. The main specification pools data across follow-up rounds 1-3 (and for the pooled specification, across years):

$$Y_{ijry} = \alpha + \beta T_{jy} + \eta_{ry} + \varepsilon_{ijry} \tag{1}$$

The coefficient β estimates the Intent-to-Treat and, with round-year fixed effects η_{ry}, is identified from within-round variation between treatment and control groups. β can be interpreted as the average effect of being offered the loan product across follow-up rounds, though as we detail below, loan take-up was high. Standard errors are clustered at the loan group level.

To absorb additional variation in the outcomes of interest, we also control for survey date in the regressions. Each follow-up round spanned over three months, meaning that there could be (for instance) substantial within-round drawdown of inventories. Inclusion of this covariate should help to make our estimates more precise without biasing point estimates.

The assumption in (1) is that treatment effects are constant across rounds. In our setting, there are reasons why this might not be the case. In particular, if treatment encourages storage, one might expect maize revenues to be *lower* for the treated group immediately following harvest, as they hold off selling, and *greater* later on during the lean season, when they release their stored grain. To explore whether treatment effects are constant across rounds, we estimate:

$$Y_{ijry} = \sum_{r=1}^{3} \beta_r T_{jy} + \eta_{ry} + \varepsilon_{ijry} \tag{2}$$

and test whether the β_r are the same across rounds (as estimated by interacting the treatment indictor with round dummies). Unless otherwise indicated, we estimate both (1) and (2) for each of the hypotheses below.

To quantify market level effects of the loan intervention, we tracked market prices at 52 market points throughout our study region, and we assign these markets to the nearest sublocation. To estimate price effects we begin by estimating the following linear model:

$$p_{msty} = \alpha + \beta_1 H_s + \beta_2 month_t + \beta_3 (H_s * month_t) + \varepsilon_{mst} \tag{3}$$

where p_{mst} represents the maize sales price at market m in sublocation s in month t in year y.[25] H_s is a binary variable indicating whether sublocation s is a high-intensity sublocation, and $month_t$ is a time trend (in each year, Nov = 0, Dec = 1, etc). If access to the storage loan allowed farmers to shift purchases to earlier in the season or sales to later in the season, and if this shift in marketing behavior was enough to alter supply and demand in local markets, then our prediction is that $\beta_1 > 0$ and $\beta_3 < 0$, i.e. that prices in areas with more treated farmers are higher after harvest

[25] Prices are normalized to 100 among the "low" intensity markets in the first month ($H_s = 0$, $month_t = 0$). Therefore, price effects can be interpreted as a percentage change from control market post-harvest prices.

but lower closer to the lean season.

While H_s is randomly assigned, and thus the number of treated farmers in each sublocation should be orthogonal to other location-specific characteristics that might also affect prices (e.g. the size of each market's catchment), we are only randomizing across 17 sublocations. This relatively small number of clusters could present problems for inference (Cameron et al., 2008). We begin by clustering errors at the sublocation level when estimating Equation 3. We also report standard errors estimated using both the wild bootstrap technique described in Cameron et al. (2008) and the randomization inference technique used in Cohen and Dupas (2010).

To understand how treatment density affects individual-level treatment effects, we estimate Equations 1 and 2, interacting the individual-level treatment indicator with the treatment density dummy. The pooled equation is thus:

$$Y_{ijsry} = \alpha + \beta_1 T_{jy} + \beta_2 H_s + \beta_3 (T_{jy} * H_s) + \eta_{ry} + \varepsilon_{ijsry} \tag{4}$$

If the intervention produces sufficient individual level behavior to generate market-level effects, we predict that $\beta_3 < 0$ and perhaps that $\beta_2 > 0$ - i.e. treated individual in high-density areas do worse than in low density areas, and control individuals in high density areas do better than control individuals in low density areas. As in Equation 3, we report results with errors clustered at the sublocation level.

For long-run effects, we first estimate the following regression for each year separately:

$$Y_{ij} = \alpha + \beta T_{jy} + \varepsilon_{ij} \tag{5}$$

in which Y_{ij} is the outcome of interest for individual i in group j. The sample is restricted to those who were in the Year y study.

We further also estimate the following specification:

$$Y_{ij} = \alpha + \beta_1 T_{j1} + \beta_2 T_{j2} + \beta_3 T_{j1} * T_{j2} + \varepsilon_{ij} \tag{6}$$

in which T_{j1} is an indicator for being an in treated group in year 1, T_{j12} is an indicator for being in a treated group in year 2, and $T_{j1} * T_{j2}$ is an interaction term for being in a group that was treated in both years. The sample is restricted to those who were in the study for both years. Because of this sample restriction, and because attrition from the Year 1 to Year 2 study was differential based on treatment status (see Appendix K), this last specification is open to endogeneity concerns and therefore should not be interpreted causally. For the sake of transparency, we present it regardless, but with the aforementioned caveat.

4 Individual level results

4.1 Harvest loan take up

Take-up of the loan treatments was quite high. Of the 954 individuals in the Year 1 treatment group, 610 (64%) applied and qualified for the loan. In Year 2, 324 out of the 522 treated individuals (62%) qualified for and took up the loan. Unconditional loan sizes in the two treatment groups were 4,817 Ksh and 6,679 Ksh, or about \$57 and \$79 USD, respectively. The average loan sizes conditional on take-up were 7,533 Ksh (or about \$89 USD) for Year 1 and 10,548 Ksh (or \$124) for Year 2.[26]

Relative to many other credit-market interventions in low-income settings in which documented take-up rates range from 2-55% of the surveyed population (Karlan et al., 2010), the 60-65% take-up rates of our loan product were very high. This is perhaps not surprising given that our loan product was offered as a top-up for individuals who were already clients of an MFI. Nevertheless, OAF estimates that about 30% of farmers in a given village in our study area enroll in OAF, which implies that even if *no* non-OAF farmers were to adopt the loan if offered it, population-wide take-up rates of our loan product would still exceed 15%.

[26]Recall in Year 1 there were two versions of the loan, one offered in October and the other in January. Of the 474 individuals in the 77 groups assigned to the October loan treatment (T1), 329 (69%) applied and qualified for the loan. For the January loan treatment (T2), 281 out of the 480 (59%) qualified for and took up the loan. Unconditional loan sizes in the two treatment groups were 5,294 Ksh and 4,345 Ksh (or about \$62 and \$51 USD) for T1 and T2, respectively, and we can reject at 99% confidence that the loan sizes were the same between groups. The average loan sizes conditional on take-up were 7,627Ksh (or about \$90 USD) for T1 and 7,423Ksh (or \$87) for T2, and in this case we cannot reject that conditional loan sizes were the same between groups.

Default rates were extremely low, at less than 2%.

4.2 Primary effects of the loan offer

We begin by estimating treatment effects in the standard fashion, assuming that there could be within-randomization-unit spillovers (in our case, the group), but that there are no cross-group spillovers. In all tables and figures, we report results broken down by each year and pooled. As explained in Section 3, the Year 2 replication produced results that are quantitatively quite similar to the Year 1 results for most outcomes, and as such, we report in the text the pooled results, unless otherwise noted.

Tables 2-4 and Figure 6 present the results of estimating Equations 1 and 2 on the pooled treatment indicator, either parametrically (in the table) or non-parametrically (in the figure). The top panels in Figure 6 show the means in the treatment group (broken down by year and then pooled, in the final panel) over time for our three main outcomes of interest (as estimated with Fan regressions). The bottom panels present the difference in treatment minus control over time, with the 95% confidence interval calculated by bootstrapping the Fan regression 1000 times.

Farmers responded to the intervention as anticipated. They held significantly more inventories for much of the year, on average about 25% more than the control group mean (Column 6 in Table 2). Inventory effects are remarkably similar across both years of the experiment.

Net revenues[27] are significantly lower immediately post harvest and significantly higher later in the year (Column 6 in Table 3 and middle panel of Figure 6). The net effect on revenues averaged across the year is positive in both years of the experiment, and is significant in the Year 2 and the pooled data (see Columns 1, 3, and 5 in Table 3). Breaking down Year 1 results by the timing of loan suggest that the reason results in Year 1 are not significant is that the later loan, offered in January to half of the treatment group, was less effective than the October loan. Table C.1 presents results for the Year 1 loan, broken down by loan timing. We see in Column 5 that the October loan (T1) produced revenue effects that are more similar in magnitude (and now significant, at 5%) to those of the Year 2 loan (which was offered almost at the same time). The January loan (T2) had

[27] From which loan interest rates were subtracted for those who took out a loan.

no significant effect on revenues. Appendix Section C explores the effects of loan timing in greater detail. The total effect across the year can be calculated by adding up the coefficients in Column 6 of Table 3,which yields an estimate of 1548 Ksh, or about $18 at the prevailing exchange rate at the time of the study. Given the unconditional average loan size of 5,476 Ksh in the pooled data, this is equivalent to a 28% return (net of loan and interest repayment), which we consider large.

The final panel of Figure 6 and Table 4 present the consumption effects (as measured by logged total household consumption). While point estimates are positive in both years, they are not significant at traditional confidence levels when pooled (in Year 2, treatment is associated with a 7 percentage point increase in consumption, significant at 10%, but in Year 1, estimated effects are only slightly greater than zero and are not significant).

Tables B.1-B.3 in Appendix **??** present the pre-specified dimensions of heterogeneity in treatment effects on inventories, revenues and log consumption. We see that greater measures of impatience are associated with significantly lower inventories and revenues among the control group, but greater treatment effects of the loan (in fact, almost all of the increase in revenues appear concentrated among the impatient). Present-biasedness shows similar expected patterns of heterogeneity, though these effects are not significant (Table B.1. We see no heterogeneity by the number of school-aged children. We see somewhat surprising effects by wealth status: while we see no heterogeneity in inventories, we see large heterogeneity in revenue effects, with almost all of the revenue increases concentrated among the wealthy (Table B.2). Finally, while we see no heterogeneity by the seasonal price increase expected at baseline, we do see strong heterogeneity by the percent of early sales in the season prior to Year 1; treatment effects are significantly larger for those who sold a larger percentage early in the previous year. It may be that these households have the greatest room for movement in storage behavior and/or that these households were most constrained at baseline (Table B.3).

Table 5 presents effects on purchase quantities, prices, and values. We observe that purchases are significantly higher in Round 1 and lower in Round 3 among treated individuals, suggesting a shift forward in time of purchase to the post-harvest period, when prices are lower. Consistent with this, we observe that treated individuals pay a lower price for their purchases.[28] As a result, purchase values are initially higher and later in the season lower among treated individuals.

As expected, sales behavior shows the opposite effect. Table 6 displays these effects. Sales quantities are lower post-harvest (though not significantly so) and are higher later in the season. Price received is also higher.[29] As a result, quantity values are initially lower and later in the season greater among treated individuals.

Finally, Table 7 presents the full effect on net sales (quantity sold - quantity purchased). We see net sales are initially lower in the treated group, and later in the season are higher; both results are highly significantly. Similarly, the effective purchase prices paid is significantly higher in the treated group due to the shift in purchases earlier in the season. Sales show the opposite effect, also highly statistically significant.[30]. The total impact on net sales is a weakly positive effect, which – off of a negative average net sales amount – means that gap between how much households must purchase to cover their needs, compared to the amount the household sells, is slightly less negative among treated households.[31]

4.3 Secondary effects of the loan offer

Appendix Section D presents outcomes on potential secondary outcomes of interest. We find no significant effects on profits earned from and hours worked at non-farm household-run businesses (Tables D.1, nor on D.2), wages earned from and hours worked in salaried employment (Tables

[28] The "overall" purchase price effect shown in in Column 3 is likely an underestimate of the full effective reduction in prices paid by treated individuals on average, because this specification includes round fixed effects, and so only captures shifts in timing of sales *within* round. It also weights average prices across rounds evenly, rather than weighting by quantities. A more accurate estimate of the impact on price paid is displayed in Column 3 of Table 7, which regresses the "effective purchase price," constructed by the dividing the total value of all purchases over the full year (summed across rounds) by the total quantity of all purchases over the full year, on treatment. This regression includes only one observation per individual and does not include round fixed effects, as the price paid changes *because* of shifts in the timing of transactions. Here we see the effective price paid for purchase is 3.4% lower among treated individuals.

[29] Again, Column 3 is likely an underestimate of the full shift in sales price received (see footnote 28 above). Column 4 of Table 7 presents the full effect on effective sales price received, which is 4.7% higher among treated individuals.

[30] See footnote 28 above.

[31] Unlike the impact on net sales *per round*, on which we have strong theoretical predictions, the impact on total net sales was ex-ante ambiguous, from a theoretical perspective. In practice, the total effect on net sales will be a combined response of the increase in purchases in response to lower effective purchase prices and increases in sales in response to higher effective sales prices. The point estimate on the sales price change is slightly larger than that on the purchase price increase (though note this variable is endogenous, since it is affected by the timing and quantities sold/purchased). In addition, the response will demand on household's demand and supply elasticities for maize.

D.3 and D.4). We also find no significant effects on schools fees paid (the primary expenditure that households say constrain them to sell their maize stocks early; see Table D.7), nor do we find significant effects on food expenditure (Table D.5). We do in Year 1 find a significant 0.07 point increase on a happiness index (an index for the following question: "Taking everything together, would you say you are very happy (3), somewhat happy (2), or not happy (1)"). However, we find no significant increase in this measure in Year 2.

4.4 Long-run effects

Appendix Section E presents the long-run follow-up effects of the loan, as measured in the Long-Run Follow-Up (LRFU) survey conducted November-December 2015, which measures outcomes one to two years after the completion of the intervention (for the Year 2 and Year 1 loan respectively). In this section, we primarily focus on the effects of each year of the study as estimated separately, because these results can be interpreted causally. For the main results on timing of sale and revenues, we also present pooled results on the effect of being ever treated, though these must be interpreted with caution given the slight differential attrition from the Year 1 sample to the Year 2 sample.

We first explore outcomes for the 2014 long-rains harvest, the season immediately following the completion of the Year 2 study. If farmers are able to use revenues from the one- (sometimes two-) time loan to "save their way" out of this credit constraint, we should expect to see sustained shifts in the timing of sales, as well as long-run revenue effects. Table E.1 presents these results.[32] We see no significant change in net sales in 2014-2015 in Columns 1-3. We also see no evidence of a sustained shift in the timing of sales. We break up sales and purchases into those that occurred before January 1 (a period of relatively low price, entitled "harvest") and those that occurred after January 1 (a period of relatively high price, entitled "lean'). If the loan drove sustained shifts into improved arbitrage, we would expect to see long-run increases in the percent of total sales made in the "lean" period and in the percent of total consumption purchased in the "harvest" period.

[32]Note we find no long-run treatment effects on 2014 harvest levels. The lack of effect on subsequent harvests is interesting in its own right and will be discussed further below, but for now note that all effects on sales and revenues are off of similar base harvest levels.

However, as can be seen in Columns 4-6 and 7-9, we see no meaningful shifts either in the timing of sales or consumption. While we see no significant changes in long-run annual revenue (Columns 10-12), this effect is measured with substantial noise and we cannot rule out large effects on revenues (in fact, point estimates, if taken seriously, would suggest a doubling of net revenues).[33]

In Table E.2, we further break down sales and purchase behavior, exploring long-run treatment effects on amount and value sold/purchased. However, we find no significant effects on any of these outcomes. We find no evidence of significant long-run treatment impacts when breaking down these outcomes separately by season (Tables E.3 and E.4), though again estimates are relatively noisy.

We now turn to effects on 2015 long-rains input use and harvest levels. Specifically, we test the hypothesis that loan access produced long-run increases in on-farm investment. This could occur if revenues from the loan relaxed credit constraints that previously restricted farmers' ability to invest in inputs. Alternatively, if the loan led to long-run improvements in the price farmers receive for their crops, this increased output price could increase incentives to invest in production-enhancing inputs; the marginal value product of a given amount of input use is now higher.[34]. However, Table E.5 suggests little movement on this margin. We estimate fairly precise null effects on labor inputs, non-labor inputs, and 2015 long-rains harvest levels. We therefore find no evidence that a one-time increase in storage and revenues crowds in other inputs and increases harvests in future years.

We also explore other outcomes for the 2015 year. Table E.6 explores long-run effects on maize eaten, food expenditures, and overall household (log) consumption. We find no significant effects. Table E.6 also explores the long-run effects on the happiness index (an index for the following

[33]While we therefore see no significant changes in sales timing or revenue in among the pooled treatment group, we see when breaking these results down by treatment status some interesting heterogeneity (see Tables E.10 and E.11. Point estimates suggest (and are significant in Year 2) that the percent sold in the lean season and the percent purchased in the harvest season are higher in low-saturation areas. In high saturation areas, the negative interaction terms cancels this effect out (see Table E.11). This is consistent with the idea that in low intensity areas, the lack of effect on prices means storage is highly profitable, encouraging individuals to purchase more in the post-harvest period and sell more in the lean season. In contrast, in high intensity areas, price effects dampen the returns to arbitrage, and there is lower incentive to store. However, we see that control individuals in high intensity areas may be storing more, buying more (significant among Year 2 individuals) in the harvest period, when prices are low. As a result, we see cannot rule out sizable increases in revenues for control individuals in high-intensity areas; though this effect is measured with considerable noise, it is consistent with the idea that control individuals may benefit from the loan. See Appendix E for greater discussion of this heterogeneity.

[34]An improved price could be attained either in the lean season, if the farmer in question himself stores, or at harvest time, if other farmers are arbitraging and producing lower overall season price fluctuations (though note in Tables E.1 and E.9 we see no evidence of such long-run shifts in either sales timing or prices).

question: "Taking everything together, would you say you are very happy (3), somewhat happy (2), or not happy (1)"). We find an increase of 0.1 points on the index from the Year 1 treatment, which is around the same size as (and in fact, a larger point estimate than) the immediate effects on happiness. However, we find no effect of the Year 2 treatment on the long-run happiness index (consistent with the lack of an immediate effect for this study year). We also find no significant long-run effects on educational expenditure or school attendance (Table E.7).

Table E.8 displays long-run effects on the hours spent on non-farm businesses owned by the household (Columns 1-3) and profits from these businesses (Columns 4-6). We see no significant effects. The same table also presents long-run impacts on hours work and wages at salaried employment positions. We find no effects on hours worked. The point estimate on wages are positive, but is only significant in Year 2.

Finally, consistent with the lack of long-run effects on the timing of sales at the individual level, in Table E.9 we observe no long-run effects of treatment density on price trends the year after the loan was removed (and, in fact, point estimates go in the opposite direction from that expected).

In summary, while we cannot rule out potentially large long-run effects on revenues, we find no significant evidence that the loan permanently alters farmers' timing of sales or a variety of other household-level economic outcomes. Consistent with this, we find no long-run effects on local market prices (though effects are in the same direction as the short-run effects, but are much muted). We therefore find little evidence that a one-time injection of credit can permanently ameliorate the underlying constraints limiting grain arbitrage in a rural Kenyan setting.

4.5 Temptation and kin tax

To test whether self-control issues or social pressure to share with others limits storage, we test the impact of laminated tags that brand the maize as committed to OAF. Estimates are shown in Table F.1. We find no significant difference in inventories, revenues, or consumption for individuals who receive only the tags (without the loan), and point estimates are small. Therefore, the tags do not appear to have any effect on storage behavior. However, this may simply be because tags are a weak form of commitment, either to one's self or to others.

4.6 Savings one's way put of the credit constraint

How long might it take for a farmer to "save his way out" of this credit constraint? In Appendix G, we present various estimates suggesting that it would take the farmer 3-6 years to self-finance the loan, if he were to save the full returns from his investment, but 34 years if he saved at a more standard rate for Kenyan households of 10%. Therefore, low savings rates are important to understanding why credit constraints persist in the presence of high return, divisible investment opportunities.

In order to test the importance of savings constraints, we examine the impact of the lockbox, as well as its interaction with the loan. Table G.1 presents these results. We observe no significant effects of the lockbox on inventories, revenues, or consumption in the overall sample. Interestingly, when interacted with the loan, we see that receiving the lockbox alone is associated with significantly *lower* inventories; perhaps the lockbox serves as a substitute savings mechanism, rather than grain. However, receiving both the lockbox and the loan is associated with a reversal of this pattern. We see no such heterogeneity on revenues. Interestingly, the point estimates on consumption are negative (though not significant) for the lockbox and loan when received separately; however, the interaction of the two is large and positive (and significant, at 95%), canceling out this effect. In Appendix G we also explore the long-run effects of the loan, finding similar heterogeneous patterns by interaction with the loan.

5 General equilibrium effects

Because the loan resulted in greater storage, shifting supply across time, and given the high transport costs common in the region, we might expect this intervention to affect the trajectory of local market prices. By shifting sales out of a relative period of abundance, we would expect the loan to result in higher prices immediately following harvest. Conversely, by shifting sales into a period of relative scarcity, we would expect the loan to result in lower prices later in the lean season. These effects will of course only be discernible if the treatment affects a sufficiently large portion of the available grain supply on the market. This requires (1) that a substantial percentage of local farmers are treated, such that local maize supply faces a sizable shock, and (2) that markets are somewhat isolated, such that local prices are at least partially determined by local supply.

On the first count, the percent of local farmers affected by this treatment was considerable. In "mature" areas where OAF has been working for a number of years (such as Webuye district where our experiment took place), approximately 30% of all farmers sign up for OAF. This means that in high treatment density sublocations, where 80% of OAF groups were enrolled in the study (a little more than half of whom were in the treatment group), 14% of all farmers in the area were offered the loan (compared to 7% in the low saturation areas).[35]

There is also evidence that rural agricultural markets in the region are not well-integrated. Transport costs and search costs have been shown to generate substantial transaction costs between markets (Teravaninthorn and Raballand, 2009; Aker, 2012), and high mark-ups charged by intermediaries appear to drive wedges between producers and consumers (Bergquist, 2017). As a result, markets may remain isolated and quite strongly affected by local shifts in supply and demand. This is shown empirically in other papers, such as Cunha et al. (2011), where local food supply shocks have substantial on local prices even in settings (in their case, Mexico) where markets are likely much less isolated than ours.

In this section, we explore whether the loan offer – and the resulting shifts in storage behavior at the micro-level – produced price movements at the market-level. We then consider how such general equilibrium effects shape individual-level results, and discuss the implications these spillover effects have for the distribution of overall welfare benefits driven by this intervention.

5.1 Market level effects

To understand the effect of our loan intervention on local maize prices, we identified 52 local market points spread throughout our study area that OAF staff indicated were where their clients typically bought and sold maize, and our enumerators tracked monthly maize prices at these market points. We then match these market points to the OAF sublocation in which they fall. "Sublocations"

[35]Given an average take-up rate of 63%, this means about 9% of farmers in high saturation areas and 4.5% in low saturation areas actually received the loan. More details on the percent of treated populations in high vs. low treatment sublocations are provided below.

here are simply OAF administrative units that are well defined in terms of client composition (i.e. which OAF groups are in which sublocation), but less well defined in terms of their exact geographic boundaries. Given this, we use GPS data on both the market location and the location of farmers in our study sample to calculate the "most likely" sublocation, based on the designated sublocation to which the modal study farmer falling within a 3km radius belongs, as pre-specified. As also pre-specified, we test robustness to alternative radii of 1km and 5km.[36]

We then utilize the sublocation-level randomization in treatment intensity to identify market-level effects of our intervention, estimating Equation 3 and clustering standard errors at the sublocation level. Regression results are shown in Table 8 and plotted non-parametrically in Figure 7. In each year, we explore the price changes from the period following loan disbursal (November in Year 1, December in Year 2) until the beginning of the subsequent harvest (August in both years). In Figure 7, which presents the pooled data, we see prices in high-intensity markets on average start out almost 4% higher in the immediate post-harvest months. As the season goes on, prices in high-density markets begin to converge and then dip below those low-density markets, ending almost 2% lower in high-density areas compared to low-density. Table 8 presents these results according to the empirical specifically outlined in Section 3. In line with the graphic results visible in Figure 7, here we see the interaction term on "Hi" treatment intensity is positive (and significant at 5%), while the interaction term between the monthly time trend and the high intensity dummy is negative (though not significant). Columns 4-5 display robustness to alternative radii; we find similar point estimates.

The overall picture painted by the market price data is consistent with the individual-level results presented above. Price effects are most pronounced (and statistically significant) early on in the season. This is when we observe the largest and most concentrated shock to the supply on the market (note in Table 2 that the greatest shift in inventories is seen in Round 1). Sensibly, treatment effects are most concentrated around the time of the loan disbursal, which represents a common shock affecting all those taking out the loan; this produces a simultaneous inward shift in

[36]Because we draw twice the sample from high-intensity areas compared to low (in accordance with our randomized intensity), we weight the low-intensity observations by two to generate a pool reflective of the true underlying OAF population. From this pool, we identify the modal farmer sublocation.

supply in the post-harvest period. In contrast, the release of this grain onto the market in the lean period appears to happen with more diffuse timing among those the treatment group (as can be seen in Figure 6, in which we note a gradual reduction in the treatment-control gap in inventories, rather than the sharp drop we would expect if all treated individuals sold at the same time). Anecdotally, farmers report that the timing of sales is often driven by idiosyncratic shocks to the household's need for cash, such as the illness of a family member, which may explain the observed heterogeneity in timing in which the treatment group releases its stores. Perhaps as a result of these more diffuse treatment effects in the lean season, price effects are smaller and measured with larger standard errors in the second half of the year. Finally, prices across high and low intensity areas appear to equalize around the time treated individuals switch from being net sellers to net buyers, as one would expect if treatment is producing a contraction in supply while treated individuals are net sellers and later an expansion in supply once treated individuals become net buyers.

Are the size of these observed price effects plausible? A back-the-envelope calibration exercise suggests yes. One Acre Fund works with about 30% of farmers in the region. Of these farmers, 80% were enrolled in the study in high density areas, while 40% were enrolled in low-density areas. About 58% of those enrolled received the loan offer [37] Together, this implies that about 14% of the population was offered treatment in high-intensity sublocations and 7% in low-intensity areas, such that the treatment was offered to 7 percentage points more of the population in high-density areas. Table 2 suggests that treated individuals experienced average increases in inventory (i.e. inward supply shifts) of 24.5%. Taken together, this suggests a contraction in total quantity available in the high-density markets by 1.7%. Experiments conducted in the same region in Kenya suggest an average demand elasticity of -1.1 (Bergquist, 2017). This would imply that we should expect to see an overall price increase of 1.5%. In the period immediately following harvest, when the inventory effects are most concentrated – during which time inventories are 47.7% higher among treatment individuals – we should expect to see a 3.0% increase in price. This is quite close to what we observe in Figure 7. We see a jump in price of about 2.5% during this period,[38] which

[37]In Year 1, 66% of the sample received the loan offer (1/3 received the offer in October, 1/3 received the loan offer in January, and 1/3 served as control). In Year 2, 50% of the sample received the loan offer (1/2 received the offer in November and 1/2 served as control). In this calibration exercise, we use the average of the two years' rates.

[38]We measure shifts in post-harvest inventories in Round 1 of the survey, which conducted roughly January-

then peters out to a slightly negative (though not significant) effect towards the end of the season.

5.2 Robustness checks

We check the robustness of the regression results to functional form assumptions. Table H.1 presents a binary version of Equation 3, replacing $month_t$ with an indicator $lean_t$ for being in the lean season (defined as April-August) and the interaction term with $lean_t * H_s$. Results suggest similar significant increases in price post-harvest in high-intensity markets. The lean season interaction term suggests that prices in high-intensity markets are lower overall in the lean season, although the point estimate on the interaction term is only slightly larger in absolute value than the the main H_s treatment coefficient, such that the combined effect of treatment in the lean season is to lower prices in high-intensity markets only slightly below those in low-intensity overall. Comparing these effects to Figure 8, we observe this is because at the beginning of the lean season prices are still higher in high intensity markets, with a cross-over mid-lean season as prices in high-intensity markets drop below those low-intensity markets. However, the 1km and 5km specifications shown in the right panel in Figure 8 shows suggest that this crossover occurs closer to the transition from the harvest to lean season; therefore the 1km and 5km specification of the binary specification, shown in Columns 4-5 of Table H.1, estimate a more substantial decrease in price for the full lean season.

We also check the robustness of these results to a more continuous measure of treatment at the market-level, following the technique described in Miguel and Kremer (2004). We construct an estimate of the ratio of total treated farmers to the total farmers in our sample within a 3km radius around each market.[39] We re-estimate an equation identical to Equation 3 with H_s replaced with $ratio_m$, the aforementioned ratio. Results are presented in Table H.2. We also present non-parametric estimates of this specification in Figure H.1, displaying average prices in markets with above- vs. below-median ratios. While results are slightly noisier in this specification, the broad

February for the average respondent. We therefore estimate the change in price change in January-February from Table 8 to be $3.97 + 2.5 * (-0.57) = 2.5$.

[39] Because we draw twice the sample from high-intensity areas compared to low (in accordance with our randomized intensity), for the total farmer count, we weight the low-intensity observations by two to generate a count reflective of the true underlying OAF population.

patterns remain consistent: prices are higher in the post-harvest period and lower in the lean period in markets with a greater proportion of treated individuals in the area.

We also check robustness to small cluster standard error adjustments. These market-level price results rely on the treatment saturation randomization being conducted at the sublocation level (a higher level than the group-level randomization employed in the individual-level results). While we cluster standard errors at the sublocation level,[40] one might be concerned due to the small number of sublocations – of which we have 17 – that asymptotic properties may not apply to our market-level analyses and that our standard errors may therefore be understated. We run several robustness checks to address these small sample concerns. In Appendix H, we use a nonparametric randomization inference approach employed by Bloom et al. (2013) and Cohen and Dupas (2010) to draw causal inferences in the presence of small samples. Results using these alternative approaches are broadly consistent with those from the primary specifications (see Appendix H for further details). We also check the robustness of our results by conducting the wild bootstrap procedure proposed by Cameron et al. (2008). While we do see some decrease in statistical precision, these adjustments are small. Finally, to ensure that results are not sensitive to a single outlier sublocation, we drop each sublocation one-by-one and re-run our analysis; the pattern observed in the full data is generally robust to this outlier analysis. We also See Appendix H for further details.

5.3 Related Outcomes

We also check whether treatment intensity affected other outcomes of interest related to market price. First, we check whether treatment effects can be seen in farmgate prices (see Table H.5). We see similar patterns in these prices as well. We also explore whether trader movement responds to treatment. We see some evidence that fewer traders enter high-intensity treated markets in the immediate post-harvest period in Year 2 (see Table H.6), which may be a sensible demand respond to the increase in price observed during a time when traders are typically purchasing. This, along with the overall weaker treatment intensity in Year 2, may contribute to the weaker price effects

[40] For all analyses in this paper, we cluster our standard errors at the level of randomization. For the individual results shown in Section 4, this is at the group level. For the results presented in this section, which relying on the sublocation-level randomized saturation, we cluster at the sublocation level.

observed in Year 2.[41]

6 Individual results with spillovers

Mass storage appears to raise prices at harvest time and lower price in the lean season, thereby smoothing out seasonal price fluctuations. What effect does this have on the individual profitability of the loan, which is designed to help farmers to take advantage of these price variations? That is, how do the individual-level returns to arbitrage vary with the stock of arbitrageurs?[42]

To answer this question, we revisit the individual results, re-estimating them to account for the variation in treatment density across sublocations. Tables 9 - 11 and Figure 8 display how our main outcomes respond in high versus low density areas for treated and control individuals.

We find that inventory treatment effects do not significantly differ as a function of treatment intensity for the pooled treatment.

Effects on net revenues, however, paint a different picture. Treatment effects in low-intensity areas are much larger – roughly double — those estimated in the pooled specification. This is because most of the revenue effects seen in the pooled specification are concentrated among treated individuals in low-intensity sublocations. In contrast, revenue effects for treated individuals in high-intensity sublocations are substantially lower (and, in fact, are statistically indistinguishable from zero in the pooled results presented Column 3 of Table 10).[43,44]. Therefore, while individuals in both high and low-intensity sublocations store significantly more as a result of treatment, only

[41]In terms of weaker treatment intensity, note that the sample size in Year 2 is only about 65% that of Year 1. As a result, the intensity in Year 2 is only about 65% what it was in Year 1. Note that the point estimate on “Hi” in column 2 (Y2) of Table 8 is almost exactly 65% of the coefficient on column 1 (Y1). The coefficient on “Hi Intensity * Month” in column 2 (Y2) is close to (a bit more than) 65% of the coefficient on column 1 (Y1).

[42]Shifts in local market prices may not be the only channel through which treatment density affected individual-level results. For example, sharing of maize or informal lending between households could also be affected by the density of loan recipients. Appendix J explores these alternative channels and presents evidence suggesting that the individual-level spillover results are most consistent with spillovers through market prices. However, we do not rule out that such additional mechanisms could also be at play.

[43]Tables 9-11 display “p-val T+TH=0,” which indicates the joint significance of $\beta_1 + \beta_3$ from Equation 4; this represents the full effect of treatment for individuals in high-intensity sublocations.

[44]While the interaction term “Treat*Hi” is only significant at traditional levels in Year 1, we attribute at least some of the weakened Year 2 interaction term to the lower treatment intensity in Year 2. Recall that the sample size in Year 2 is only about 65% that of Year 1. As a result, the intensity in Year 2 is only about 65% what it was in Year 1. If we scale the coefficient on “Treat*Hi” in Year 2 (column 2) to account for this difference (i.e. divide by 0.65), we get an estimate much closer to the Year 1 estimate. In addition, any trader movement that dampened Year 2 market-level effects may have further contributed to this weaker Year 2 effect.

treated individuals in low-intensity sublocations earn significantly higher revenues.

Table 11 presents effects on consumption. As with earlier estimates, they remain relatively imprecisely estimated.[45,46]

Why might loan profitability be lower in high treatment density areas? It appears that as more farmers store, producing the smoother prices documented in the above section, the (direct) benefits to arbitrage fall. Sensibly, arbitrage – the exploitation of price differentials – is most profitable to an individual when he is the only one arbitraging. As others begin to arbitrage as well, general equilibrium effects drive down these differentials and therefore diminish the direct returns to arbitrage.

Conversely, for those who do *not* engage in arbitrage, these spillovers may be positive. Though the timing of their sales will not change, they may benefit from relatively higher sale prices at harvest-time and relatively lower purchase prices during the lean season. We see some evidence of these positive spillovers to control group revenues in high-intensity treatment areas (see middle panel of Figure 8 and the estimate on the Hi dummy in Column 3 of Column 3 of Table 10). However, it should be noted that this effect is measured with considerable noise and and thus remains more speculative.[47] Given the diffuse nature of spillover effects, it should perhaps not be surprising that identifying these effects with great statistical precision is challenging. However, they are suggestive of important distributional dynamics for welfare, which we explore below.

6.1 Discussion

The randomized saturation design allows us to capture how both direct and indirect treatment effects vary with saturation level. Table 12 breaks down the distribution of welfare gains from the loan, based on saturation rate and revenue effects drawn from the pooled results. While this

[45] Interestingly, they are strongly positive for treated individuals in the high-intensity areas in Year 2. However, because there is no clear pattern across years, we avoid speculating or over-interpreting this figure.

[46] In light of this heterogeneity, and given that the sample is drawn twice as heavily from high-intensity areas, Appendix I presents the main treatment effects from Section 4 weighted by the inverse probability that a given observation is included in our sample, with the probability defined as the percentage of our sample drawn from high and low intensity areas respectively.

[47] And even goes in the opposite direction in the Year 2 results alone; see Column 2 of Table 10.

exercise takes all point estimates as given, note that some are less precisely measured than others.[48] As a result, there are likely large standard errors around some of the figures presented in Table 12. This exercise should therefore be interpreted as an illustration of how general equilibrium effects can shape the distribution of welfare gains in isolated markets, rather than precise quantitative estimates.

In the first row, we present that the direct gains per person, representing the increase in revenues driven by treatment for those who are treated (specifically calculated as the coefficient on the "Treat" dummy in low saturation areas and as the coefficient on the "Treat" dummy plus the coefficient on the "Treat*Hi" interaction term in high saturation areas). We see, as discussed before, that the direct treatment effects are much greater for those in low saturation sublocations, where treated individuals are closer to "being the only one arbitraging" than in high saturation areas.

The second row presents the indirect gains per person. This is estimated as zero in low saturation areas and as the coefficient on "Hi" in high saturation areas.[49] We see in row 3 that, in the high saturation areas, the indirect gains are 58% the size of the direct gains. When we account for the much larger size of the total population relative to that of just the direct beneficiaries (presented in rows 5 and 4 respectively), we find that the total size of the indirect gains would swamp that of the direct gains in high saturation areas (rows 7 and 6 respectively).

These findings have two implications. First, the total gains from the intervention (presented in row 7) are much higher in high saturation areas than they are in low saturation areas. Although the direct gains to the treatment group are lower in areas of high saturation, the small per-person

[48]For example, the point estimate on "Treat*Hi" is not quite significant at traditional levels, while the point estimate on "Hi" is measured with large noise.

[49]Because the coefficient on "Treat*Hi" captures the differential value of direct treatment in high saturation areas, we include this in the calculation of direct benefits, since we view the general equilibrium effects observed in the high saturation areas as mitigating the direct treatment effect. However, an alternative formulation could view this as a negative spillover for the treatment group and include this as a indirect (negative) gain, restricting the direct gains only to be the coefficient as estimated on the "Treat" dummy (though, as will be discussed, even this "pure" direct effect could be affected by spillovers, as it is estimated by comparing the outcomes of the treated in the low and high saturation areas, neither of which is truly a pure zero saturation area). In this alternative formulation, the indirect gains per person, which would be a weighted average of the negative gains for the treated group and the positive gains for the control group, would be much smaller 51Ksh/person. The indirect gains would then only account for 25% of the total gains, rather than the 81% estimated under the current formulation. Regardless, the private gains, which will be subsequently discussed, are unambiguously defined.

indirect gains observed in these areas accrue to a large number of untreated individuals, resulting in an overall increase in total gains.[50,51] Second, the distribution of gains shifts in the presence of general equilibrium effects. While in low saturation areas all of the gains appear to come from direct gains, in high saturation areas, 81% of the total gains are indirect gains (row 9).[52] General equilibrium effects therefore more evenly distribute gains across the entire population, reducing the proportion of the gains that direct beneficiaries exclusively receive and increasing the share enjoyed by the full population.[53]

This redistribution of gains has implications for private sector investment in arbitrage. Row 10 presents the per-person private gains accruing to arbitragers, as estimated by the coefficient on the "Treat" indicator in low saturation areas and by the coefficient on the "Treat" dummy plus the coefficient on the "Treat*Hi" interaction term plus the coefficient on the "Hi" interaction term in high saturation areas. This represents the per-person gains accruing to treated farmers in our sample, under each level of saturation. It also represents the most that private sector banks or other financial institutions could hope to extract from each farmer to whom they might provide loans for storage. Row 11 presents the total private gains, multiplying the per-person gains by the number of treated individuals. Despite the fact that high saturation areas have two times the number of treated farmers, the total private gains are still lower in these areas compared to low saturation areas.

These calculations suggest that private sector financial institutions may face incentives that result in the under-provision of finance for arbitrage. Although overall social gains are higher at greater levels of saturation (row 8), because much of these gains are indirect, private sector institutions will not be able to capture them. For private sector institutions, the available gains

[50] Also contributing is the fact that although the direct benefits/person are only a quarter of the size in high areas, there are twice the number of beneficiaries, which makes up some of the gap in terms of total direct gains.

[51] Note that even if the indirect gains were only 38Ksh/individual (substantially less than $1 USD), the total gains would still be larger under high saturation than low saturation.

[52] It is possible that there are general equilibrium effects – and therefore indirect gains – occurring in the low saturation areas that we simply cannot detect in the absence of a pure control group. If this is the case, it would mean that our current estimates underestimate the total gains, as well as the percentage of gains coming from indirect gains, in low saturation areas. However, it would also mean that we are underestimating these figures in the high intensity areas as well.

[53] Note that even if the indirect gains were only 40Ksh/individual, the indirect gains would still be larger than the direct gains under high saturation.

for capture are actually lower at high levels of saturation (row 11). Row 12 attempts to quantify this disincentive. At low levels of saturation, private sector institutions could fully internalize all gains, capturing up to 100% of the total revenue increases generated by the product (under our assumption of no indirect gains in the low saturation case). However, at high saturation rates, only 31% of the total gains are private. Financial institutions therefore will fail to internalize 69% of the gains at these higher saturation levels, which will likely result in under-provision of financial products, compared to the socially optimal amount. Socially oriented NGOs, such as our partner organization in this project, or public sector entities may be better positioned to internalize these benefits and offer such credit products.

7 Conclusion

We study the effect of offering Kenyan maize farmers a cash loan at harvest. This is unusual timing for an agricultural loan in low income settings, where such credit is typically offered before planting. The timing of this loan is motivated by two facts: the large observed average increase in maize prices between the post harvest season and the lean season six to nine months later, and the inability of most poor farmers appear to successfully arbitrage these prices due to a range of "non-discretionary" consumption expenditures they must make immediately after harvest. Instead of putting maize in storage and selling when the price is higher, farmers are observed to sell much of it immediately, sacrificing potential profits.

We show that access to credit at harvest "frees up" farmers to use storage to arbitrage these prices. Farmers offered the loan shift maize purchases into the period of low prices, put more maize in storage, and sell maize at higher prices later in the season, increasing farm profits. Using experimentally-induced variation in the density of treatment farmers across locations, we document that this change in storage and marketing behavior aggregated across treatment farmers also affects local maize prices: post harvest prices are significantly higher in high-density areas, consistent with more supply having been taken off the market in that period, and are lower later in the season (but not significantly so). These general equilibrium effects feed back to our profitability estimates, with farmers in low-density areas – where price differentials were higher and thus arbitrage opportunities greater – differentially benefiting.

The findings make a number of contributions. First, our results are some of the first experimental results to find a positive and significant effect of microcredit on the profits of microenterprises (farms in our case), and the first experimental study to directly account for general equilibrium effects in this literature. At least in our particular setting, failing to account for these GE effects substantially alters the conclusions drawn about the average benefits of improved credit access. This suggests that explicit attention to GE effects in future evaluations of credit market interventions could be warranted.

Second, we show how the absence of financial intermediation can be doubly painful for poor households in rural areas. Lack of access to formal credit causes households to turn to much more expensive ways of moving consumption around in time, and aggregated across households this behavior generates a large scale price phenomenon that further lowers farm income and increases what these households must pay for food. Our results suggest that in this setting, expanding access to affordable credit could reduce this price variability and thus have benefits for recipient and non-recipient households alike. Welfare estimates suggest that a large portion of the benefits of expanded loan access could accrue indirectly to non-borrowers. Under such a distribution of welfare gains, private sector financial institutions may be less willing to offer products in this sector, and thus that these socially beneficial credit products could more realistically be offered by the public sector or socially minded non-profits.

What our results do not address is why larger actors – e.g. large-scale private traders – have not stepped in to bid away these arbitrage opportunities. Traders do exist in the area and can commonly be found in local markets. In a panel survey of local traders in the area, we record data on the timing of their marketing activities and storage behavior. But we find little evidence of long-run storage. When asked to explain this limited storage, trader report being able to make even higher total profits by engaging in spatial arbitrage across markets (relative to temporal arbitrage). Nevertheless, this does not explain why the scale or number of traders engaging in spatial arbitrage have not expanded; imperfect competition among traders may play a role (Bergquist, 2017).

References

Acemoglu, Daron, "Theory, general equilibrium and political economy in development economics," *Journal of Economic Perspectives*, 2010, *24* (3), 17–32.

Aggarwal, Shilpa, Eilin Francis, and Jonathan Robinson, "Grain Today, Gain Tomorrow: Evidence from a Storage Experiment with Savings Clubs in Kenya," *Working Paper*, 2017.

Aker, Jenny C, "Rainfall shocks, markets and food crises: the effect of drought on grain markets in Niger," *Center for Global Development, working paper*, 2012.

Angelucci, Manuela, Dean Karlan, and Jonathan Zinman, "Win some lose some? Evidence from a randomized microcredit program placement experiment by Compartamos Banco," Technical Report, National Bureau of Economic Research 2013.

Attanasio, Orazio, Britta Augsburg, Ralph De Haas, Emla Fitzsimons, and Heike Harmgart, "Group lending or individual lending? Evidence from a randomised field experiment in Mongolia," 2011.

Baland, Jean-Marie, Catherine Guirkinger, and Charlotte Mali, "Pretending to be poor: Borrowing to escape forced solidarity in Cameroon," *Economic Development and Cultural Change*, 2011, *60* (1), 1–16.

Banerjee, Abhijit V and Andrew F Newman, "Occupational choice and the process of development," *Journal of political economy*, 1993, pp. 274–298.

_ **and Esther Duflo**, "Giving credit where it is due," *The Journal of Economic Perspectives*, 2010, *24* (3), 61–79.

Banerjee, Abhijit Vinayak, "Microcredit Under the Microscope: What Have We Learned in the Past Two Decades, and What Do We Need to Know?," *Annual Review of Economics*, 2013, (0).

Banerjee, A.V., E. Duflo, R. Glennerster, and C. Kinnan, "The Miracle of Microfinance?: Evidence from a Randomized Evaluation," *working paper, MIT*, 2013.

Barrett, C., "Displaced distortions: Financial market failures and seemingly inefficient resource allocation in low-income rural communities," *working paper, Cornell.*, 2008.

Basu, Karna and Maisy Wong, "Evaluating Seasonal Food Security Programs in East Indonesia," *working paper*, 2012.

Berge, Lars Ivar, Kjetil Bjorvatn, and Bertil Tungodden, "Human and financial capital for microenterprise development: Evidence from a field and lab experiment," *NHH Dept. of Economics Discussion Paper*, 2011, (1).

Bergquist, Lauren Falcao, "Pass-through, Competition, and Entry in Agricultural Markets: Experimental Evidence from Kenya," *Working Paper*, 2017.

Blattman, Christopher, Nathan Fiala, and Sebastian Martinez, "Credit Constraints, Occupational Choice, and the Process of Development: Long Run Evidence from Cash Transfers in Uganda," *working paper*, 2013.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts, "Does management matter? Evidence from India," *The Quarterly Journal of Economics*, 2013, *128* (1), 1–51.

Bruhn, Miriam, Dean S Karlan, and Antoinette Schoar, "The impact of consulting services on small and medium enterprises: Evidence from a randomized trial in mexico," *Yale University Economic Growth Center Discussion Paper*, 2012, (1010).

Brune, L., X. Giné, J. Goldberg, and D. Yang, "Commitments to save: A field experiment in rural malawi," *University of Michigan, May (mimeograph)*, 2011.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-based improvements for inference with clustered errors," *The Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Casey, Katherine, Rachel Glennerster, and Edward Miguel, "Reshaping Institutions: Evidence on Aid Impacts Using a Preanalysis Plan*," *The Quarterly Journal of Economics*, 2012, *127* (4), 1755–1812.

Cohen, Jessica and Pascaline Dupas, "Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment," *Quarterly Journal of Economics*, 2010.

Crepon, B., F. Devoto, E. Duflo, and W. Pariente, "Impact of microcredit in rural areas of Morocco: Evidence from a Randomized Evaluation," *working paper, MIT*, 2011.

Cunha, Jesse M, Giacomo De Giorgi, and Seema Jayachandran, "The price effects of cash versus in-kind transfers," Technical Report, National Bureau of Economic Research 2011.

Dillion, Brian, "Selling Crops Early to Pay for School: A Large-scale Natural Experiment in Malawi," *Working Paper*, 2017.

Dupas, P. and J. Robinson, "Why Don't the Poor Save More? Evidence from Health Savings Experiments," *American Economic Review, forthcoming*, 2013.

Fafchamps, Marcel, David McKenzie, Simon Quinn, and Christopher Woodruff, "Microenterprise Growth and the Flypaper Effect: Evidence from a Randomized Experiment in Ghana," *Journal of Development Economics*, 2013.

Field, Erica, Rohini Pande, John Papp, and Natalia Rigol, "Does the Classic Microfinance Model Discourage Entrepreneurship Among the Poor? Experimental Evidence from India," *American Economic Review*, 2012.

Fink, Gunther, Kelsey Jack, and Felix Masiye, "Seasonal credit constraints and agricultural labor supply: Evidence from Zambia," *NBER Working Paper*, 2014, (20218).

Galor, Oded and Joseph Zeira, "Income distribution and macroeconomics," *The review of economic studies*, 1993, *60* (1), 35–52.

Kaboski, Joseph P and Robert M Townsend, "The impact of credit on village economies," *American economic journal. Applied economics*, 2012, *4* (2), 98.

Karlan, D., J. Morduch, and S. Mullainathan, "Take up: Why microfinance take-up rates are low and why it matters," Technical Report, Financial Access Initiative 2010.

Karlan, Dean and Jonathan Morduch, "Access to Finance," *Handbook of Development Economics, Volume 5*, 2009, (Chapter 2).

_ **and Jonathan Zinman**, "Microcredit in theory and practice: using randomized credit scoring for impact evaluation," *Science*, 2011, *332* (6035), 1278–1284.

_ **, Ryan Knight, and Christopher Udry**, "Hoping to win, expected to lose: Theory and lessons on micro enterprise development," Technical Report, National Bureau of Economic Research 2012.

McKenzie, D., "Beyond baseline and follow-up: the case for more T in experiments," *Journal of Development Economics*, 2012.

McKenzie, David and Christopher Woodruff, "Experimental evidence on returns to capital and access to finance in Mexico," *The World Bank Economic Review*, 2008, *22* (3), 457–482.

Mel, Suresh De, David McKenzie, and Christopher Woodruff, "Returns to capital in microenterprises: evidence from a field experiment," *The Quarterly Journal of Economics*, 2008, *123* (4), 1329–1372.

Miguel, E. and M. Kremer, "Worms: identifying impacts on education and health in the presence of treatment externalities," *Econometrica*, 2004, *72* (1), 159–217.

Park, A., "Risk and household grain management in developing countries," *The Economic Journal*, 2006, *116* (514), 1088–1115.

Rubin, Donald, "Which Ifs Have Causal Answers? Discussion of Holland's Statistics and Causal Inference.," *Journal of the American Statistical Association*, 1986, *81*, 961–962.

Saha, A. and J. Stroud, "A household model of on-farm storage under price risk," *American Journal of Agricultural Economics*, 1994, *76* (3), 522–534.

Stephens, E.C. and C.B. Barrett, "Incomplete credit markets and commodity marketing behaviour," *Journal of Agricultural Economics*, 2011, *62* (1), 1–24.

Teravaninthorn, Supee and Gael Raballand, "Transport Prices and Costs in Africa," *World Bank*, 2009.

World Bank, "Malawi Poverty and Vulnerability Assessment: Investing in our Future," 2006.

Tables and Figures

Figure 1: **Monthly average maize prices**, shown at East African sites for which long-term data exist, 1994-2011. Data are from the Regional Agricultural Trade Intelligence Network, and prices are normalized such that the minimum monthly price = 100. Our study site in western Kenya is shown in green, and the blue squares represent an independent estimate of the months of the main harvest season in the given location. Price fluctuations for maize (corn) in the US are shown in the lower left for comparison

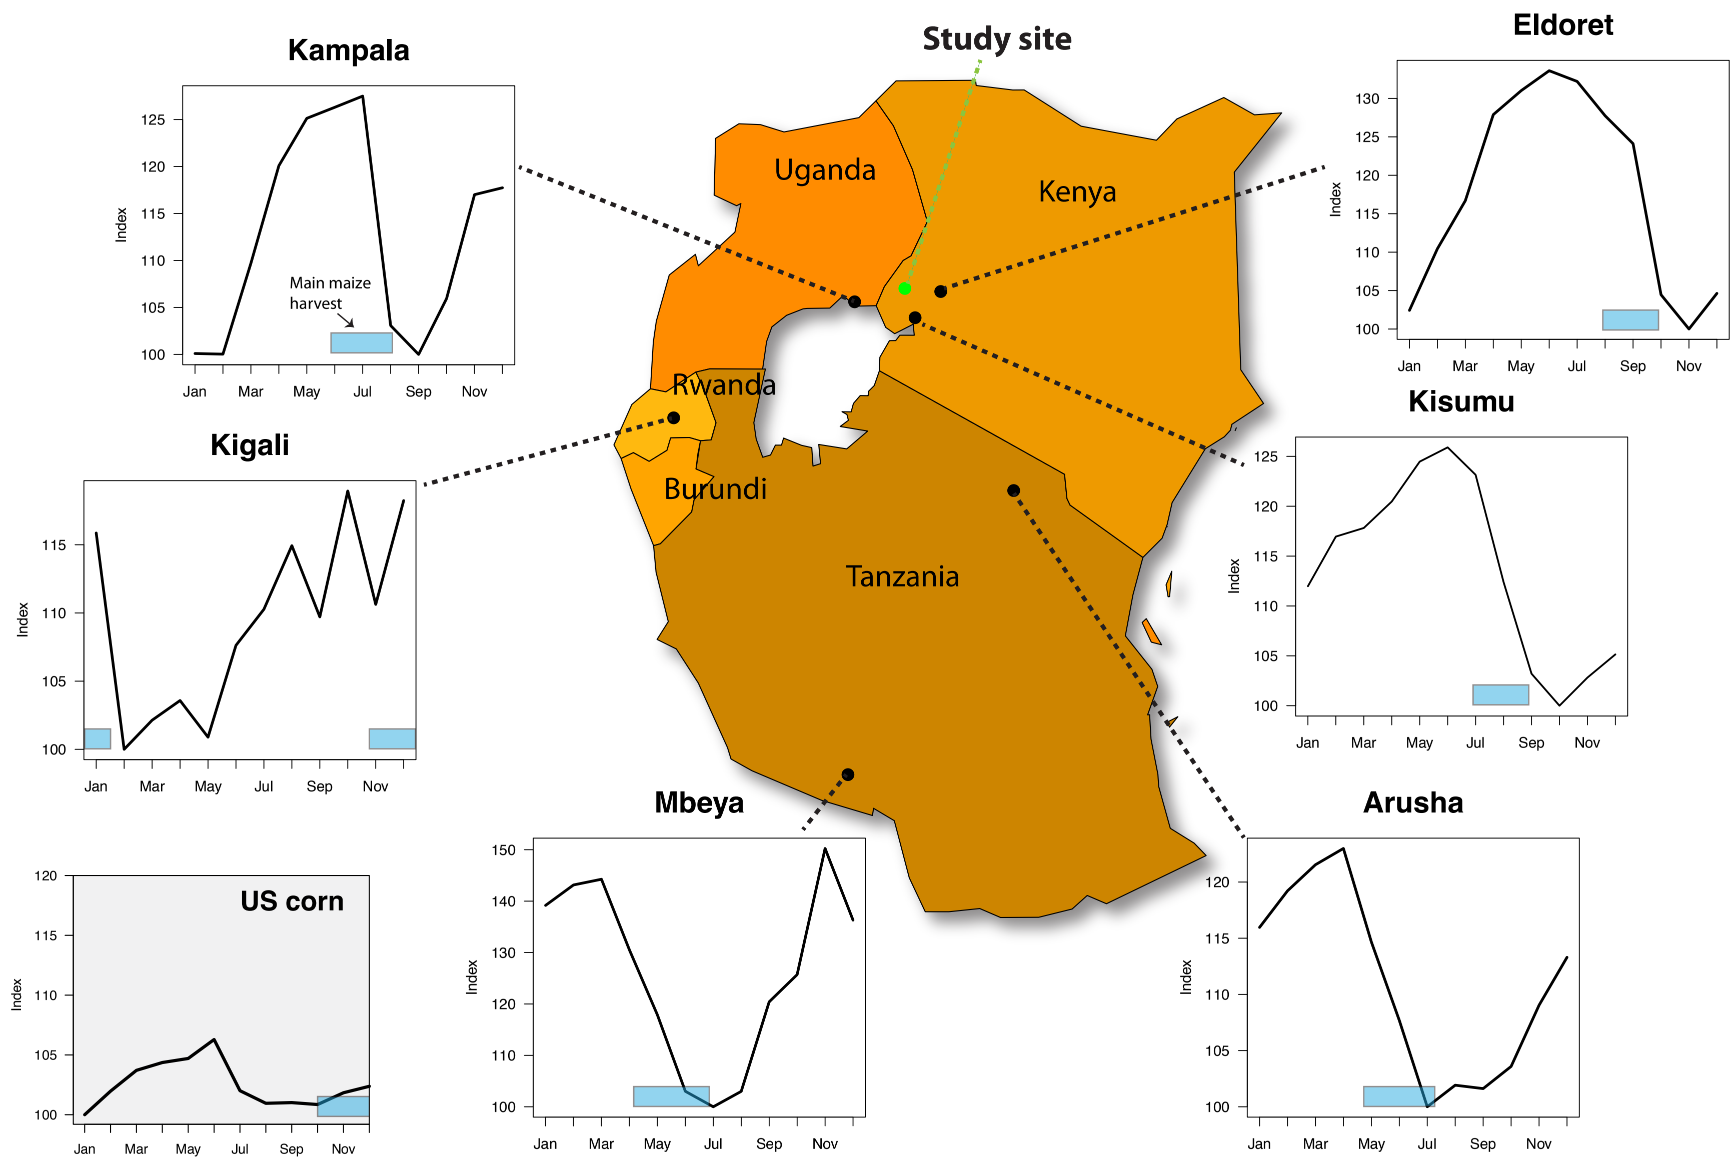

Figure 2: **Study design**. Randomization occurs at three levels. First, treatment intensity was randomized across 17 sublocations (top level, each box represents a sublocation). This randomization was held constant across the two years. Second, treatment was randomized at the group level within sublocations (second level, each box representing a group in a given sublocation). In Year 1, treatment groups were further divided into October and January loans. In Year 2, only one timing of the loan was offered (in November). Finally, in Year 1, there was a third level of randomization at the individual level, in which the tags and lockbox were cross-randomized (bottom level). In Year 2, no individual level treatments were offered. Total numbers of randomized units in each bin are given on the left.

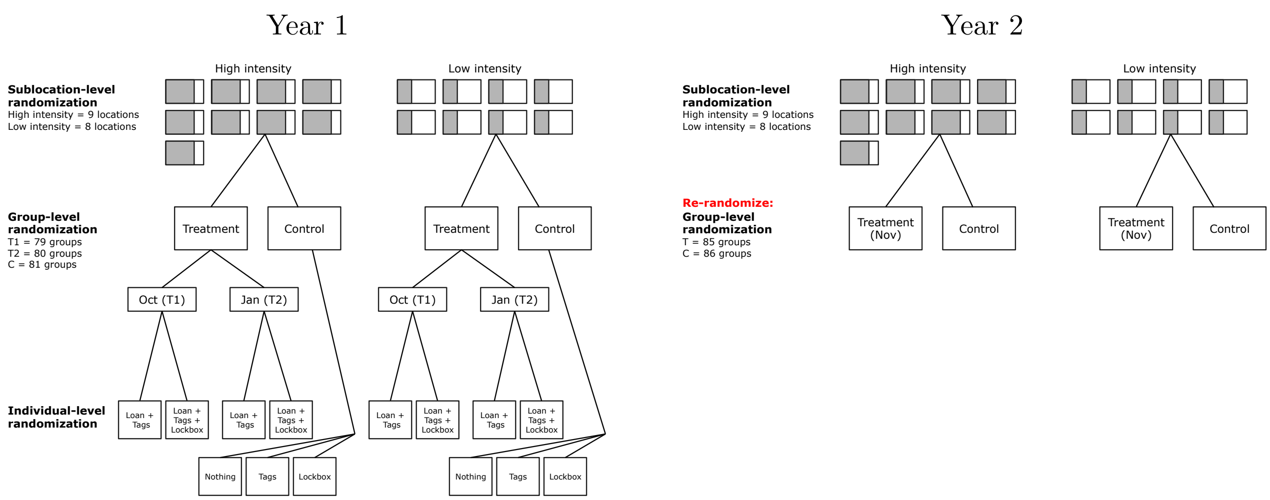

Figure 3: **Study timeline**. The timing of the interventions and data collection are show. Year 1 spanned 2012-2013, Year 2 spanned 2013-2014, and the long-run follow-up data collection occurred 2014-2015. The market survey, shown in red, ran from November 2012-December 2015. The baseline was run August-September 2012. Three rounds of household data were collected on a rolling basis in each year of the main study. A long-run follow-up survey was run September-December 2015. Light blue arrows show the timing of the loan announcement (immediately as harvest was ending), while dark blue arrows display the date of loan disbursal (October and January in Year 1, November in Year 2). Harvest time is highlighted in grey and occurs in September-October each year.

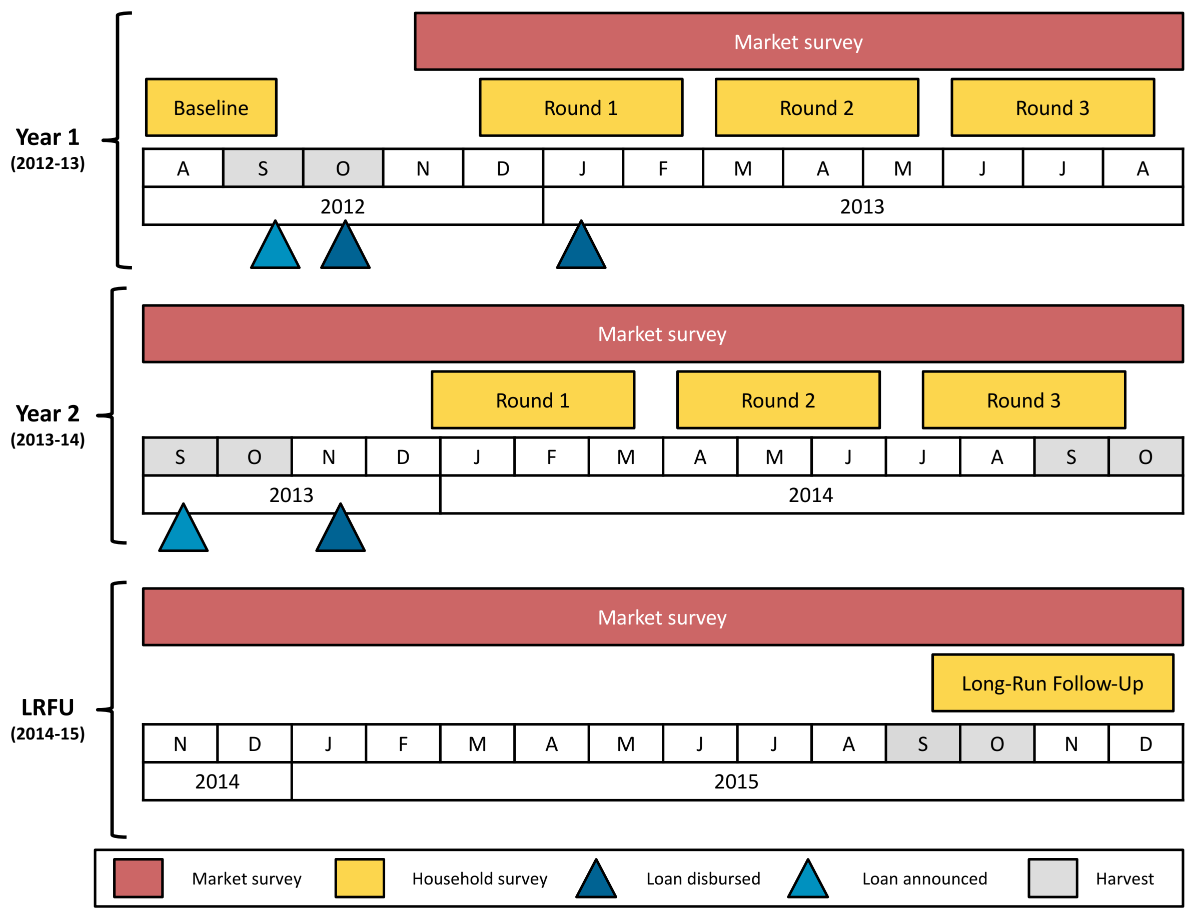

Figure 4: **Maize price trends (pre-study period)**. Farmer-reported average monthly maize prices for the period 2007-2012, averaged over all farmers in our sample. Prices are in Kenyan shillings per goro (2.2kg).

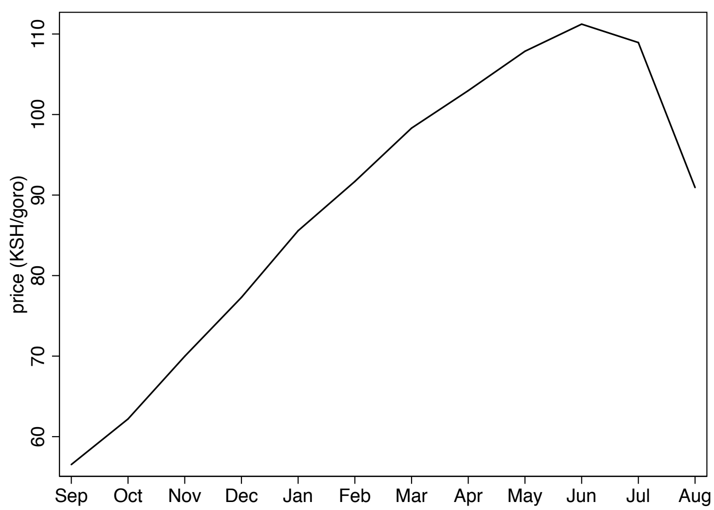

Figure 5: **Maize price trends (study period & post-study period)**. Author-collected average monthly maize prices for the period 2012-2014 (study period) and 2014-2015 (post study period), averaged over all markets in our sample. Prices are in Kenyan shillings per goro (2.2kg).

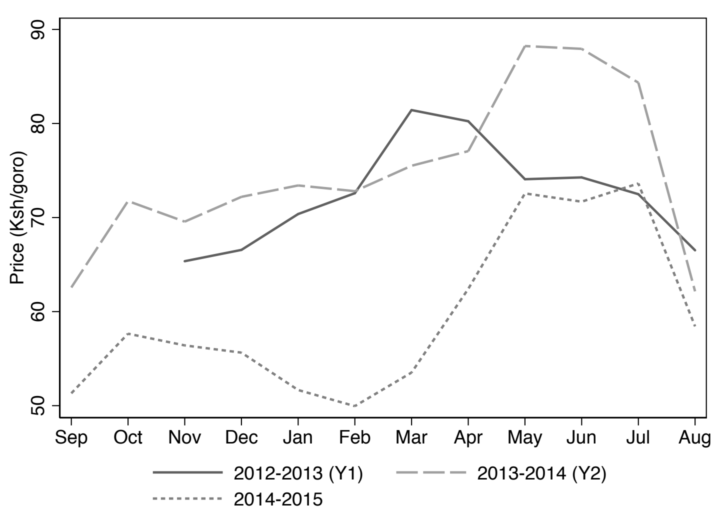

Figure 6: **Pooled Treatment effects**. The top row of plots shows how average inventories, net revenues, and log household consumption evolve from December to August in Y1 and Y2 (pooled) in the treatment group versus the control group, as estimated with fan regressions. The bottom row shows the difference between the treatment and control, with the bootstrapped 95% confidence interval shown in grey (100 replications drawing groups with replacement).

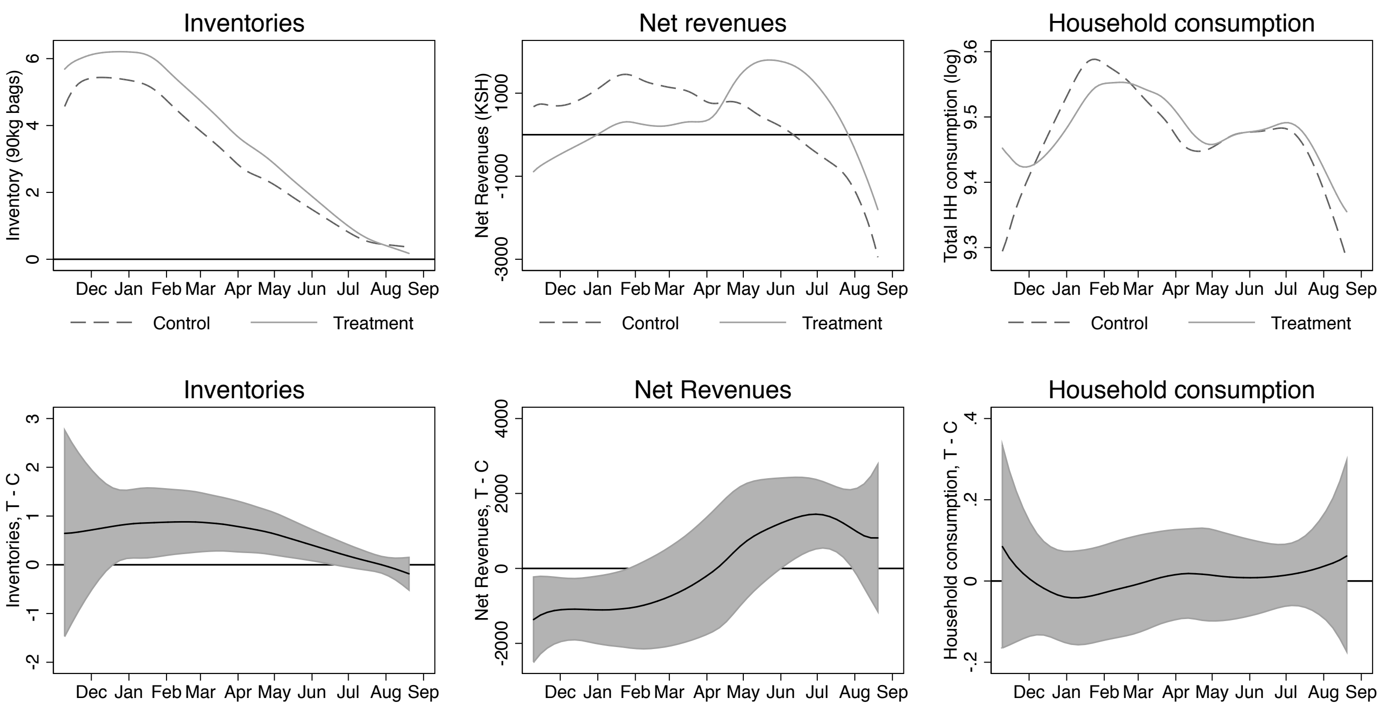

Figure 7: **Pooled market prices for maize as a function of local treatment intensity.** Markets matched to treatment intensity using sublocation of the modal farmer within 3km of each market. The left panel shows the average sales price in markets in high-intensity areas (solid line) versus in low-intensity areas (dashed line) over the study period. The middle panel shows the average difference in log price between high- and low-intensity areas over time, with the bootstrapped 95% confidence interval shown in light grey and the 90% confidence interval shown in dark grey. The right panel shows the robustness of results to alternative radii (1km, 3km, and 5km)

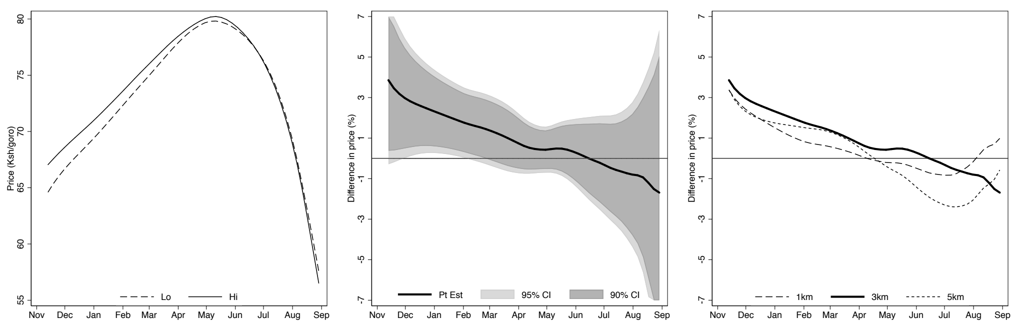

Figure 8: **Pooled Treatment effects by treatment intensity**. Average inventories, net revenues, and log HH consumption over the study period in the treatment group versus the control group, split apart by high intensity areas (orange lines) and low-intensity areas (black lines).

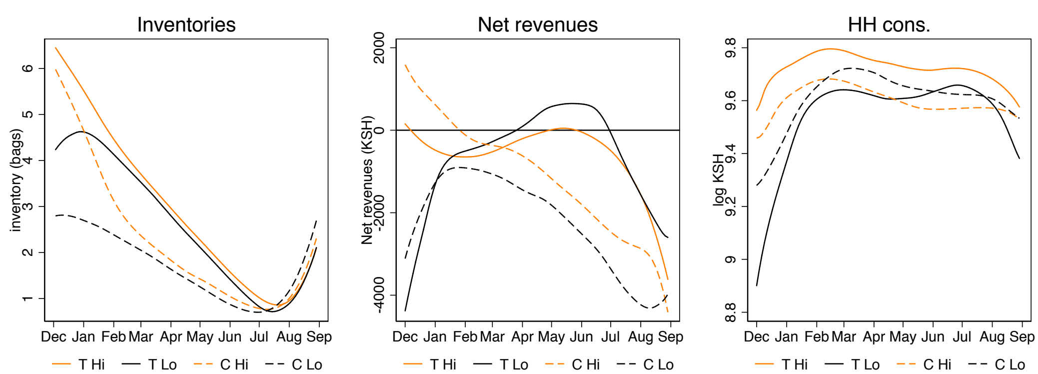

Table 1: **Summary statistics and balance among baseline covariates.** Balance table for the Y1 study (restricted to the Y1 sample, for which we have baseline characteristics. The first two columns give the means in each treatment arm. The 3rd column gives the total number of observations across the two groups. The last two columns give differences in means normalized by the Control sd, with the corresponding p-value on the test of equality.

Baseline characteristic	Control	Treat	Obs	C - T	
				sd	*p-val*
Male	0.33	0.30	1,589	0.08	0.11
Number of adults	3.20	3.00	1,510	0.09	0.06
Kids in school	3.07	3.00	1,589	0.04	0.46
Finished primary	0.77	0.72	1,490	0.13	0.02
Finished secondary	0.27	0.25	1,490	0.04	0.46
Total cropland (acres)	2.40	2.44	1,512	-0.01	0.79
Number of rooms in hhold	3.25	3.07	1,511	0.05	0.17
Total school fees (1000 Ksh)	29.81	27.24	1,589	0.06	0.18
Average monthly cons (Ksh)	15,371.38	14,970.86	1,437	0.03	0.55
Avg monthly cons./cap (log Ksh)	7.96	7.97	1,434	-0.02	0.72
Total cash savings (KSH)	8,021.50	5,157.40	1,572	0.09	0.01
Total cash savings (trim)	5,389.84	4,731.62	1,572	0.05	0.33
Has bank savings acct	0.43	0.42	1,589	0.01	0.82
Taken bank loan	0.08	0.08	1,589	0.02	0.73
Taken informal loan	0.25	0.24	1,589	0.01	0.84
Liquid wealth	97,280.92	93,878.93	1,491	0.03	0.55
Off-farm wages (Ksh)	3,797.48	3,916.82	1,589	-0.01	0.85
Business profit (Ksh)	1,801.69	2,302.59	1,589	-0.08	0.32
Avg %Δ price Sep-Jun	133.18	133.49	1,504	-0.00	0.94
Expect 2011 LR harvest (bags)	9.03	9.36	1,511	-0.02	0.67
Net revenue 2011	-4,088.62	-3,303.69	1,428	-0.03	0.75
Net seller 2011	0.30	0.32	1,428	-0.05	0.39
Autarkic 2011	0.06	0.07	1,589	-0.03	0.51
% maize lost 2011	0.01	0.02	1,428	-0.03	0.57
2012 LR harvest (bags)	11.03	11.18	1,484	-0.02	0.74
Calculated interest correctly	0.73	0.71	1,580	0.03	0.50
Digit span recall	4.58	4.57	1,504	0.01	0.89
Maize giver	0.26	0.26	1,589	0.00	0.99

"Liquid wealth" is the sum of cash savings and assets that could be easily sold (e.g. livestock). Off-farm wages and business profit refer to values over the previous month. Net revenue, net seller, and autarkic refer to the household's maize marketing position. "Maize giver" is whether the household reported giving away more maize in gifts than it received over the previous 3 months.

Table 2: **Inventory Effects, Individual Level.** Regressions include round-year fixed effects, with errors clustered at the group level.

	Y1		Y2		Pooled	
	(1) Overall	(2) By rd	(3) Overall	(4) By rd	(5) Overall	(6) By rd
Treat	0.52***		0.50***		0.53***	
	(0.16)		(0.14)		(0.12)	
Treat - R1		0.82***		1.21***		1.03***
		(0.31)		(0.24)		(0.20)
Treat - R2		0.71***		0.24		0.52***
		(0.19)		(0.15)		(0.12)
Treat - R3		0.06		0.04		0.07
		(0.07)		(0.37)		(0.19)
Observations	3836	3836	2944	2944	6780	6780
Mean DV	2.67	2.67	1.68	1.68	2.16	2.16
R squared	0.35	0.35	0.18	0.19	0.29	0.30

Table 3: **Net Revenue Effects, Individual Level.** Regressions include round-year fixed effects, with errors clustered at the group level.

	Y1		Y2		Pooled	
	(1) Overall	(2) By rd	(3) Overall	(4) By rd	(5) Overall	(6) By rd
Treat	279.78		800.24**		524.66**	
	(292.16)		(330.63)		(220.25)	
Treat - R1		-1146.56***		-23.71		-608.68**
		(325.13)		(478.41)		(285.70)
Treat - R2		534.85		1917.28***		1170.71***
		(485.80)		(532.81)		(359.84)
Treat - R3		1371.95***		520.76		985.79***
		(436.12)		(403.27)		(302.09)
Observations	3795	3795	2935	2935	6730	6730
Mean DV	334.41	334.41	-3434.38	-3434.38	-1616.12	-1616.12
R squared	0.01	0.01	0.04	0.05	0.09	0.09

Table 4: **HH Consumption (log) Effects, Individual Level.** Regressions include round-year fixed effects, with errors clustered at the group level.

	Y1		Y2		Pooled	
	(1) Overall	(2) By rd	(3) Overall	(4) By rd	(5) Overall	(6) By rd
Treat	0.00		0.07*		0.04	
	(0.03)		(0.04)		(0.03)	
Treat - R1		-0.04		0.07		0.01
		(0.05)		(0.05)		(0.03)
Treat - R2		0.02		0.08*		0.05
		(0.04)		(0.05)		(0.03)
Treat - R3		0.03		0.06		0.04
		(0.05)		(0.05)		(0.03)
Observations	3792	3792	2944	2944	6736	6736
Mean DV	9.48	9.48	9.61	9.61	9.55	9.55
R squared	0.00	0.00	0.01	0.01	0.02	0.02

Table 5: **Purchase Effects, Individual Level.** Regressions include round-year fixed effects, with errors clustered at the group level.

	Purchase Qty		Purchase Price		Purchase Val	
	Overall	By rd	Overall	By rd	Overall	By rd
Treat	0.01		-26.12**		-141.88	
	(0.04)		(12.44)		(125.87)	
Treat - R1		0.17***		-22.35		458.29***
		(0.06)		(25.14)		(158.47)
Treat - R2		-0.01		-16.89		-383.12**
		(0.06)		(19.50)		(192.49)
Treat - R3		-0.10*		-36.07**		-489.55***
		(0.06)		(17.99)		(188.89)
Observations	6145	6145	4190	4190	6745	6745
Mean DV	1.18	1.18	3193.11	3193.11	3248.86	3248.86
R squared	0.18	0.18	0.45	0.45	0.12	0.12

Table 6: **Sales Effects, Individual Level.** Regressions include round-year fixed effects, with errors clustered at the group level.

	Sales Qty		Sales Price		Sales Val	
	Overall	By rd	Overall	By rd	Overall	By rd
Treat	0.17***		24.53		535.30***	
	(0.04)		(18.19)		(132.51)	
Treat - R1		-0.02		8.35		-29.98
		(0.07)		(29.19)		(196.07)
Treat - R2		0.28***		46.70*		896.21***
		(0.07)		(25.59)		(213.60)
Treat - R3		0.25***		13.96		723.41***
		(0.06)		(35.03)		(186.98)
Observations	6777	6777	2060	2060	6761	6761
Mean DV	0.57	0.57	2899.93	2899.93	1640.18	1640.18
R squared	0.04	0.04	0.41	0.41	0.04	0.04

Table 7: **Net Sales and Effective Prices, Individual Level.** Columns 1-2 regressions on net sales (quantity sole minus quantity purchased) include round-year fixed effects, with errors clustered at the group level. Columns 3-4 include only one observation per individual (per year). Round fixed effects are omitted in these specifications in order to estimate the effect of treatment on prices paid and received, which change *because* of shifts in the timing of transactions; therefore round controls are not appropriate. "Effective purchase price" is constructed by the dividing the total value of all purchases over the full year (summed across rounds) by the total quantity of all purchases over the full year. "Effective sales price" is constructed similarly.

	Net Sales		Effective Price	
	Overall	By rd	Purchase	Sales
Treat	0.12*		-104.94***	131.70***
	(0.06)		(31.57)	(40.85)
Treat - R1		-0.26**		
		(0.10)		
Treat - R2		0.27***		
		(0.10)		
Treat - R3		0.29***		
		(0.09)		
Observations	6108	6108	2014	1428
Mean DV	-0.62	-0.62	3084.78	2809.76
R squared	0.16	0.16	0.01	0.01

Table 8: **Market prices for maize as a function of local treatment intensity.** "Hi" intensity is a dummy for a sublocation randomly assigned a high number of treatment groups. "Month" is a linear month time trend (beginning in Nov at 0 in each year). Standard errors are clustered at the sublocation level. Prices measured monthly following loan disbursal (Nov-Aug in Y1; Dec-Aug in Y2). Price normalized to 100 in Nov in low-intensity sublocations.

	Main Specification (3km)			Robustness (Pooled)	
	Y1	Y2	Pooled	1km	5km
Hi	4.41*	2.85	3.97**	2.79	3.77*
	(2.09)	(1.99)	(1.82)	(1.72)	(1.82)
Month	1.19***	1.22***	1.36***	1.33***	1.54***
	(0.36)	(0.38)	(0.35)	(0.34)	(0.29)
Hi Intensity * Month	-0.57	-0.48	-0.57	-0.52	-0.83**
	(0.42)	(0.46)	(0.39)	(0.39)	(0.37)
Observations	491	381	872	872	872
R squared	0.08	0.03	0.06	0.06	0.06

Table 9: **Inventory Effects, Accounting for Treatment Intensity.** Regressions include round-year fixed effects with errors clustered at the sublocation level. P-values on the test that the sum of the treated and treated*hi equal zero are provided in the bottom rows of the table.

	(1) Y1	(2) Y2	(3) Pooled
Treat	0.76***	0.55***	0.74***
	(0.19)	(0.18)	(0.15)
Hi	0.12	-0.03	0.02
	(0.36)	(0.22)	(0.24)
Treat*Hi	-0.33	-0.07	-0.29
	(0.23)	(0.25)	(0.19)
Observations	3836	2944	6780
Mean DV	2.74	1.38	2.04
R squared	0.35	0.18	0.29
p-val T+TH=0	0.01	0.02	0.01

Table 10: **Net Revenue Effects, Accounting for Treatment Intensity.** Regressions include round-year fixed effects with errors clustered at the sublocation level. P-values on the test that the sum of the treated and treated*hi equal zero are provided in the bottom rows of the table.

	(1) Y1	(2) Y2	(3) Pooled
Treat	1059.60**	1193.77	1101.39**
	(437.73)	(685.05)	(430.09)
Hi	533.90	-152.60	164.94
	(551.18)	(558.95)	(479.68)
Treat*Hi	-1114.63*	-555.21	-816.77
	(535.59)	(804.86)	(520.04)
Observations	3795	2935	6730
Mean DV	-253.51	-3620.40	-1980.02
R squared	0.01	0.04	0.09
p-val T+TH=0	0.86	0.15	0.41

Table 11: **HH Consumption (log), Accounting for Treatment Intensity.** Regressions include round-year fixed effects with errors clustered at the sublocation level. P-values on the test that the sum of the treated and treated*hi equal zero are provided in the bottom rows of the table.

	(1) Y1	(2) Y2	(3) Pooled
Treat	0.01	-0.05	-0.01
	(0.04)	(0.04)	(0.02)
Hi	-0.00	-0.08	-0.05
	(0.05)	(0.05)	(0.04)
Treat*Hi	-0.01	0.17***	0.07*
	(0.05)	(0.06)	(0.04)
Observations	3792	2944	6736
Mean DV	9.47	9.65	9.56
R squared	0.00	0.02	0.03
p-val T+TH=0	0.97	0.01	0.08

Table 12: **Distribution of gains in the presence of general equilibrium effects** Calculations employ per-round point estimate on revenues β_1, β_2, and β_2 (estimated in Ksh) from Column 3 of Table 10 (multiplied by three to get the annual revenue gains). They also include the following assumptions: (A_1) Total population in the study area is 7,105 households (HH) (this figure is an approximation, as the sublocations used in this study are One Acre Fund (OAF) administrative districts and therefore do not directly correspond to the Kenyan census administrative districts. OAF estimates that it works with 30% of all farmers in the area. While this figure affects the total gains estimates, it does not affect any estimates of per-HH gains, ratios, or fractions in the table, nor does it affect any comparisons between low and high saturation areas); (A_2) 50% of the study population resides in low saturation sublocations (this is roughly accurate; moreover, it allows a comparison of the size of the benefits across low and high saturation rates that is unconfounded by differences in underlying population sizes); (A_2) 30% of HH in the region are One Acre Fund (OAF) members, a figure provided by OAF administrative records; (A_{4a}) 40% of all OAF members were enrolled in the study in low saturation sublocations and (A_{4b}) 80% were enrolled in high saturation sublocation (A_5) In each sublocation, 58% of individuals in the sample were randomly assigned to receive treatment (average across the pooled data from Year 1 and Year 2).

	Low Saturation	**High Saturation**
1. Direct gains/HH	3,304[a]	854[b]
2. Indirect gains/HH	0	495[c]
3. Ratio of indirect: direct gains[d]	0.00	0.58
4. Direct beneficiary population (HH)	247[e]	495[f]
5. Total local population (HH)	3,553[g]	3,553[h]
6. Total direct gains[i]	816,984	422,248
7. Total indirect gains[j]	0	1,757,880
8. Total gains (direct + indirect)[k]	816,984	2,180,128
9. Fraction of gains indirect[l]	0.00	0.81
10. Private gains/HH	3,304[m]	1,349[n]
11. Total private gains[o]	816,984	666,945
12. Fraction of gains private[p]	1.00	0.31

[a] $3 * \beta_1$
[b] $3 * (\beta_1 + \beta_3)$
[c] $3 * \beta_2$
[d] *Row 2/Row 3*
[e] $A_1 * A_2 * A_3 * A_{4a} * A_5 = 7,105 * 0.5 * 0.3 * 0.4 * 0.58$
[f] $A_1 * (1 - A_2) * A_3 * A_{4b} * A_5 = 7,105 * 0.5 * 0.3 * 0.8 * 0.58$
[g] $A_1 * A_2 = 7,105 * 0.5$
[h] $A_1 * (1 - A_2) = 7,105 * 0.5$
[i] *Row 1*Row 4*
[j] *Row 2*Row 5*
[k] *Row 6+Row 7*
[l] *Row 7/Row 8*
[m] $3 * \beta_1$
[n] $3 * (\beta_1 + \beta_2 + \beta_3)$
[o] *Row 10*Row 4*
[p] *Row 11*Row 8*

Supplementary Appendix

A Pilot Results

Figure A.1: **Pilot data on maize inventories and marketing decisions over time**, using data from two earlier pilot studies conducted with One Acre Fund in 2010/11 with 225 farmers (top row) and 2011/12 with 700 different farmers (bottom row). *Left panels*: inventories (measured in 90kg bags) as a function of weeks past harvest. The dotted line is the sample median, the solid line the mean (with 95% CI in grey). *Right panels*: average net sales position across farmers over the same period, with quantities shown for 2010/11 (quantity sold minus purchased) and values shown for 2011/12 (value of all sales minus purchases).

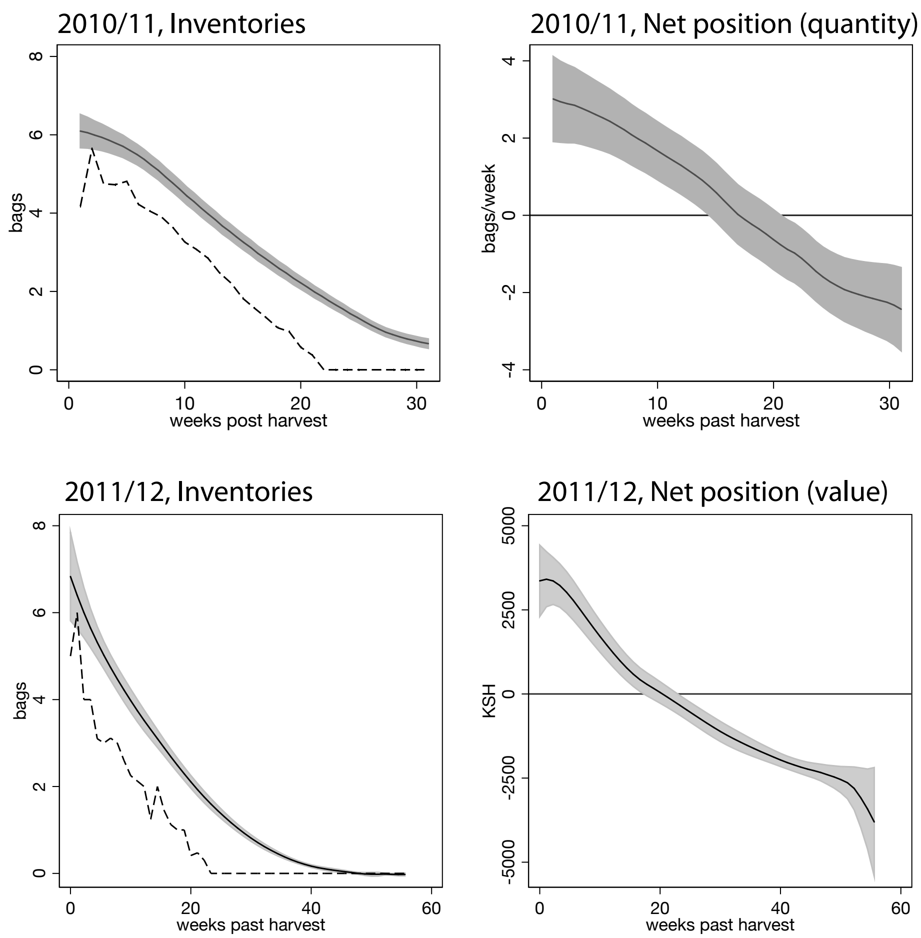

B Treatment Heterogeneity

Table B.1: **Heterogeneity in Treatment Effects.** Heterogeneity in treatment effects, as pre-specified in pre-analysis plan. All variables are from the baseline run in 2012 prior to Year 1. Because those who are new to the sample in Year 2 are missing baseline variables, the specification presented below only presents Year 1 results, for which we have full baseline data. The outcome variable is regressed on a treatment indicator, an indicator for whether the household is above the median on the dimension of heterogeneity, and an interaction term (except for the "Present Biased" specification, when the outcome variable is regressed on a treatment indicator, an indicator for whether the household expresses preferences consistent with being present biased, and an interaction term. "Delta" is the percent allocated to the early period (versus later period) in the time preference questions, such that a higher delta represents less patience. "Present Biased" represents whether the individual allocated more to the earlier period when the earlier period was "today," versus 6 months from now.

	Delta (higher = less patient)			Present Biased		
	Inv	Rev	Cons	Inv	Rev	Cons
Treat	0.395**	71.387	0.027	0.380*	156.749	-0.003
	(0.184)	(326.051)	(0.037)	(0.204)	(356.279)	(0.037)
Abv Med	-0.665**	-1099.797***	0.100	-0.302	-350.559	-0.012
	(0.267)	(397.635)	(0.065)	(0.230)	(361.804)	(0.046)
Treat*Abv Med	0.551*	837.300	-0.093	0.477	396.518	0.017
	(0.330)	(512.171)	(0.073)	(0.296)	(474.266)	(0.060)
Observations	3819	3779	3775	3836	3795	3792
R squared	0.35	0.01	0.00	0.35	0.01	0.00
Median Het. Val.	0.14	0.14	0.14	0.00	0.00	0.00
Mean DV	2.82	629.74	9.46	2.75	459.78	9.48

Table B.2: **Heterogeneity in Treatment Effects.** Heterogeneity in treatment effects, as pre-specified in pre-analysis plan. All variables are from the baseline run in 2012 prior to Year 1. Because those who are new to the sample in Year 2 are missing baseline variables, the specification presented below only presents Year 1 results, for which we have full baseline data. The outcome variable is regressed on a treatment indicator, an indicator for whether the household is above the median on the dimension of heterogeneity, and an interaction term. "Children" is number of school-aged children in the household. "Liquid Wealth" is the combined value of total assets, livestock, and cash savings.

	Children			Liquid Wealth		
	Inv	Rev	Cons	Inv	Rev	Cons
Treat	0.616***	355.280	0.032	0.592***	-73.080	0.018
	(0.174)	(323.270)	(0.038)	(0.179)	(315.373)	(0.039)
Abv Med	0.479**	-201.694	0.241***	1.424***	1008.839***	0.382***
	(0.214)	(333.351)	(0.041)	(0.215)	(306.526)	(0.046)
Treat*Abv Med	-0.186	-198.174	-0.054	-0.139	683.969*	-0.033
	(0.287)	(436.112)	(0.053)	(0.265)	(405.509)	(0.055)
Observations	3836	3795	3792	3761	3722	3719
R squared	0.35	0.01	0.03	0.38	0.02	0.09
Median Het. Val.	3.00	3.00	3.00	63490.00	63490.00	63490.00
Mean DV	2.46	439.48	9.37	2.02	-179.59	9.30

Table B.3: **Heterogeneity in Treatment Effects.** Heterogeneity in treatment effects, as pre-specified in pre-analysis plan. All variables are from the baseline run in 2012 prior to Year 1. Because those who are new to the sample in Year 2 are missing baseline variables, the specification presented below only presents Year 1 results, for which we have full baseline data. The outcome variable is regressed on a treatment indicator, an indicator for whether the household is above the median on the dimension of heterogeneity, and an interaction term. "Early Sales" is the percentage of 2011-2012 total season sales that occurred prior to January 1, 2012. "Price Expect" is the percentage expected change in price expected between September 2012 and June 2013.

	Price Expect			Early Sales		
	Inv	Rev	Cons	Inv	Rev	Cons
Treat	0.562***	131.984	0.007	0.236	-300.802	-0.031
	(0.209)	(335.816)	(0.039)	(0.311)	(550.069)	(0.056)
Abv Med	-0.072	-556.603	-0.014	-1.457***	-2056.657***	-0.164**
	(0.269)	(460.098)	(0.046)	(0.339)	(603.377)	(0.068)
Treat*Abv Med	-0.127	312.645	-0.006	0.774*	1388.021*	0.076
	(0.315)	(531.907)	(0.057)	(0.424)	(723.363)	(0.079)
Observations	3811	3771	3768	1884	1871	1874
R squared	0.35	0.01	0.00	0.38	0.02	0.02
Median Het. Val.	120.00	120.00	120.00	0.23	0.23	0.23
Mean DV	2.75	494.86	9.48	3.69	2236.07	9.59

C Effects of Loan Timing

In Year 1, the loan was (randomly) offered at two different times: one in October, immediately following harvest (T1) and the other in January, immediately before school fees are due (T2). Splitting apart the two loan treatment arms in Year 1, results provide some evidence that the timing of the loan affects the returns to capital in this setting. As shown in Figure C.1 and Table C.1, point estimates suggest that those offered the October loan held more in inventories, reaped more in net revenues, and had higher overall consumption. Overall effects on net revenues are about twice as high as pooled estimates, and are now significant at the 5% level (Column 5 of Table C.1), and we can reject that treatment effects are equal for T1 and T2 ($p = 0.04$). Figure C.2 shows non-parametric estimates of differences in net revenues over time among the different treatment groups. Seasonal differences are again strong, and particularly strong for T1 versus control.

Why might the October loan have been more effective than the January loan? Note that while we are estimating the intent-to-treat (ITT) and thus that differences in point estimates could in principle be driven by differences in take-up, these latter differences are probably not large enough to explain the differential effects. For instance, "naive" average treatment effect estimates that rescale the ITT coefficients by the take-up rates (70% versus 60%) still suggest substantial differences in effects between T1 and T2. A more likely explanation is that the January loan came too late to be as useful: farmers in the T2 group were forced to liquidate some of their inventories before the arrival of the loan, and thus had less to sell in the months when prices rose. This would explain why inventories began lower, and why T2 farmers appear to be selling more during the immediate post-harvest months than T1 farmers. Nevertheless, they sell less than control farmers during this period and store more, likely because qualifying for the January loan meant carrying sufficient inventory until that point.

Figure C.1: **Year 1 Treatment effects by loan timing**. Plots shows how average inventories, net revenues, and log per capita consumption evolve over the study period for farmers assigned to T1 (blue line), T2 (red line), and C (black dashed line), as estimated with fan regressions.

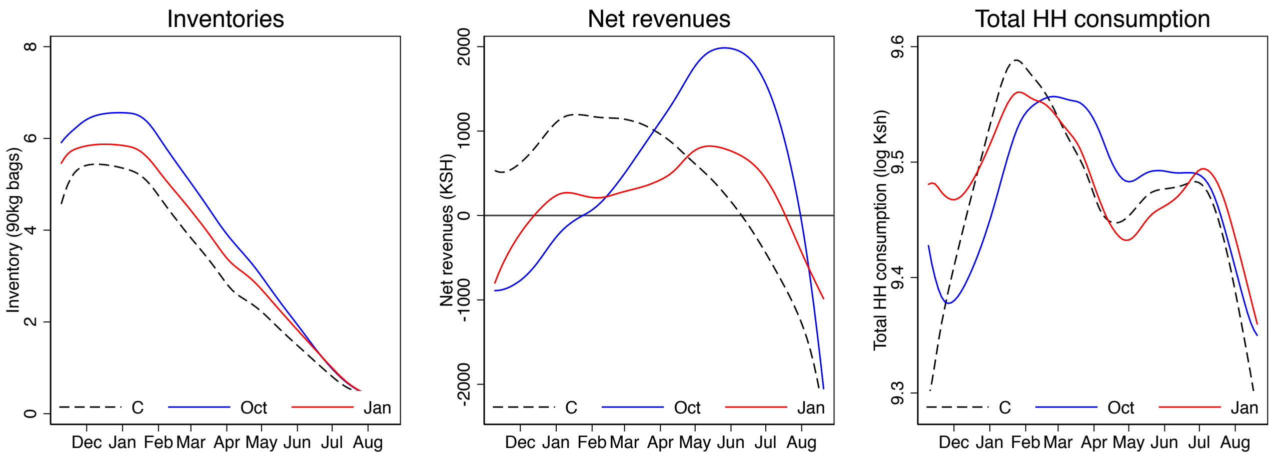

Figure C.2: **Year 1 Revenue treatment effects by loan timing**. Plots show the difference in net revenues over time for T1 versus C (left), T2 versus C (center), and T1 versus T2 (right), with bootstrapped 95% confidence intervals shown in grey.

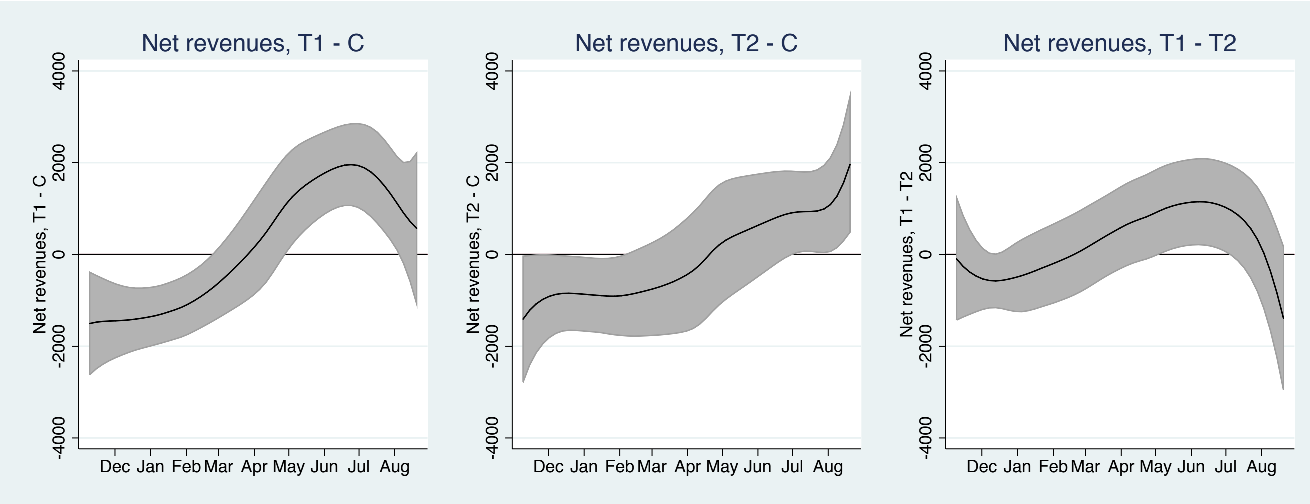

Table C.1: **Year 1 Results by Loan Timing.** Regressions include round-year fixed effects and strata fixed effects, with errors clustered at the group level.

	Inventories		Prices		Revenues		Consumption	
	(1) Pooled	(2) By round	(3) Purchase price	(4) Sales prices	(5) Pooled	(6) By round	(7) Pooled	(8) By round
T1	0.77***		-47.81**	10.51	541.95**		0.04	
	(0.13)		(23.20)	(129.67)	(248.78)		(0.03)	
T2	0.46***		2.47	-34.93	36.03		0.01	
	(0.13)		(22.47)	(114.55)	(248.15)		(0.03)	
T1 - Round 1		1.25***				-1218.96***		-0.00
		(0.27)				(353.43)		(0.05)
T1 - Round 2		0.91***				924.50*		0.08*
		(0.19)				(512.50)		(0.05)
T1 - Round 3		0.18				1840.70***		0.04
		(0.13)				(483.92)		(0.04)
T2 - Round 1		0.54**				-951.27***		-0.01
		(0.27)				(347.35)		(0.05)
T2 - Round 2		0.65***				156.58		0.01
		(0.16)				(503.66)		(0.05)
T2 - Round 3		0.18				851.70**		0.02
		(0.12)				(410.53)		(0.04)
Observations	3816	3816	1914	1429	3776	3776	3596	3596
Mean of Dep Variable	3.03	3.03	2936.14	2991.23	501.64	501.64	8.02	8.02
SD of Dep Variable	3.73	3.73	425.20	2007.53	6217.09	6217.09	0.66	0.66
R squared	0.49	0.50	0.30	0.07	0.13	0.14	0.21	0.21
T1 = T2 (pval)	0.02		0.04		0.04		0.19	

D Secondary Outcomes

Table D.1: **Pooled Non-Farm Profit** Non-farm Profit is the household's profit from non-farm activities in the last month (Ksh).

	Y1		Y2		Pool	
	(1) Overall	(2) By Intensity	(3) Overall	(4) By Intensity	(5) Overall	(6) By Intensity
Treat	197.30	-150.81	-127.45	-309.72	-35.28	-264.58
	(170.57)	(333.30)	(164.75)	(299.21)	(127.06)	(231.41)
Hi		-145.48		-28.99		-55.22
		(308.27)		(256.84)		(208.13)
Treat * Hi		489.84		256.78		323.31
		(385.60)		(357.42)		(275.02)
Observations	1305	1305	2938	2938	4243	4243
Mean DV	984.02	1056.54	1359.52	1337.37	1270.51	1269.33
R squared	0.00	0.00	0.00	0.00	0.00	0.00

Table D.2: **Pooled Non-Farm Hours** Hours Non-Farm is the number of hours worked by the household in a non-farm businesses run by the household in the last 7 days.

	Y1		Y2		Pool	
	(1) Overall	(2) By Intensity	(3) Overall	(4) By Intensity	(5) Overall	(6) By Intensity
Treat	1.40	0.73	0.77	-0.67	0.96	-0.25
	(1.59)	(2.70)	(1.23)	(2.08)	(0.99)	(1.66)
Hi		2.40		1.14		1.41
		(2.76)		(1.69)		(1.44)
Treat * Hi		0.84		2.04		1.69
		(3.32)		(2.56)		(2.02)
Observations	1305	1305	2942	2942	4247	4247
Mean DV	11.90	10.27	13.60	12.49	13.20	11.95
R squared	0.00	0.00	0.00	0.01	0.00	0.01

Table D.3: **Salaried Employment.** Hours Salary is the total number of hours worked by household members in a salaried position.

	Y1		Y2		Pool	
	(1) Overall	(2) By Intensity	(3) Overall	(4) By Intensity	(5) Overall	(6) By Intensity
Treat	0.47	0.86	0.18	-2.07	0.30	-0.96
	(1.42)	(2.43)	(1.16)	(2.18)	(0.90)	(1.64)
Hi		0.17		-1.71		-1.16
		(2.52)		(1.87)		(1.51)
Treat * Hi		-0.56		3.29		1.82
		(2.99)		(2.55)		(1.94)
Observations	1295	1295	2012	2012	3307	3307
Mean DV	11.16	10.70	6.74	7.33	8.12	8.35
R squared	0.00	0.00	0.01	0.01	0.01	0.01

Table D.4: **Average Wage** Avg Wage is the average monthly wage for those household members who are salaried.

	Y1		Y2		Pool	
	(1) Overall	(2) By Intensity	(3) Overall	(4) By Intensity	(5) Overall	(6) By Intensity
Treat	2293.22	-908.20	-333.47	1822.23	1296.43	-743.50
	(1720.62)	(3043.72)	(1620.91)	(4063.57)	(1243.91)	(2251.83)
Hi		-1843.78		-1092.62		-1476.21
		(2710.81)		(2678.38)		(1939.28)
Treat * Hi		4556.76		-2495.62		2933.25
		(3640.89)		(4689.26)		(2759.66)
Observations	284	284	135	135	419	419
Mean DV	11486.64	12087.50	5232.03	5682.00	8984.80	9278.07
R squared	0.02	0.02	0.02	0.04	0.10	0.10

Table D.5: **Food Expenditure** Food Expenditure is the household's expenditure on food purchases in the last month (Ksh).

	Y1		Y2		Pool	
	(1) Overall	(2) By Intensity	(3) Overall	(4) By Intensity	(5) Overall	(6) By Intensity
Treat	-94.37	-205.03	40.18	-359.47	-33.21	-285.49
	(152.11)	(281.91)	(167.47)	(286.88)	(112.34)	(198.39)
Hi		182.75		-197.90		-15.19
		(295.11)		(259.26)		(196.23)
Treat * Hi		147.21		566.21		356.35
		(333.67)		(352.04)		(236.63)
Observations	3817	3817	2919	2919	6736	6736
Mean DV	6665.50	6611.09	7430.94	7617.81	7057.83	7120.57
R squared	0.01	0.01	0.00	0.01	0.03	0.03

Table D.6: **Maize Eaten** Maize Eatern is the household's consumption of maize (in goros, 2.2kg tins) over the past 7 days.

	Y1		Y2		Pool	
	(1) Overall	(2) By Intensity	(3) Overall	(4) By Intensity	(5) Overall	(6) By Intensity
Treat	-0.07	-0.32	-0.02	-0.41	-0.05	-0.37*
	(0.14)	(0.26)	(0.17)	(0.27)	(0.11)	(0.19)
Hi		-0.07		-0.10		-0.09
		(0.28)		(0.24)		(0.18)
Treat * Hi		0.36		0.55		0.45**
		(0.31)		(0.34)		(0.23)
Observations	3844	3844	2947	2947	6791	6791
Mean DV	5.48	5.55	5.55	5.75	5.52	5.65
R squared	0.01	0.01	0.00	0.01	0.00	0.01

Table D.7: **Pooled School Fees Paid.** School Fees Paid are the expenditure on school fees over the past month (Ksh).

	Y1		Y2		Pool	
	(1) Overall	(2) By Intensity	(3) Overall	(4) By Intensity	(5) Overall	(6) By Intensity
Treat	150.82	31.71	213.27	-329.82	191.55	-94.21
	(118.32)	(227.21)	(377.33)	(693.94)	(186.63)	(335.90)
Hi		-272.68		-662.03		-485.39
		(207.59)		(573.79)		(312.27)
Treat * Hi		178.21		773.26		414.02
		(264.46)		(830.63)		(396.15)
Observations	3867	3867	2905	2905	6772	6772
Mean DV	1217.27	1369.71	3851.29	4077.54	2560.84	2740.01
R squared	0.05	0.05	0.03	0.03	0.09	0.09

Table D.8: **Pooled Happiness Index.** Happy is an index for the following question: "Taking everything together, would you say you are very happy (3), somewhat happy (2), or not happy (1)?"

	Y1		Y2		Pool	
	(1) Overall	(2) By Intensity	(3) Overall	(4) By Intensity	(5) Overall	(6) By Intensity
Treat	0.07**	0.04	0.01	0.03	0.04*	0.03
	(0.03)	(0.06)	(0.03)	(0.05)	(0.02)	(0.04)
Hi		-0.03		-0.02		-0.02
		(0.06)		(0.04)		(0.04)
Treat * Hi		0.04		-0.03		0.01
		(0.07)		(0.06)		(0.05)
Observations	3870	3870	2969	2969	6839	6839
Mean DV	2.57	2.58	2.68	2.68	2.63	2.63
R squared	0.01	0.01	0.00	0.00	0.01	0.01

E Long-Run Follow-up (LRFU) Survey Results

The Long-Run Follow-Up (LRFU) survey was run Nov-Dec 2015. Results presented in this appendix show the limited effects of the loan on long-run outcomes.

E.1 Long-Run Main Effects

Table E.1: **LRFU 2014-2015 Outcomes:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) outcomes. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* "Y1" refers to treatment in Year 1, while "Y2" refers to treatment in Year 2. "2014 Harvest" is the size of harvest in 90kg bags. "Net Sales" is the total number of 90kg bags sold - the total number of 90kg bags purchased between the 2014 long-rains harvest and 2015 long-rains harvest. "% Sold Lean" is the percentage of total sales completed from January onward. "% Purch Harvest" is the percentage of total purchases completed prior to January. "Revenues" are the net revenues from all maize sales and purchases from the 2014 long-rains harvest to the 2015 long-rains harvest.

	Net Sales			% Sold Lean			% Purch Harvest			Revenues		
	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both
Treat Y1	0.31		-0.01	0.04		-0.02	-0.02		0.09	350.50		-763.60
	(0.35)		(0.59)	(0.05)		(0.09)	(0.03)		(0.07)	(950.10)		(1854.40)
Treat Y2		0.29	0.29		-0.03	-0.05		-0.03	0.01		1286.62	1330.40
		(0.35)	(0.61)		(0.04)	(0.10)		(0.04)	(0.07)		(1094.42)	(1777.33)
Treat Y1 * Y2			0.21			0.10			-0.10			1126.71
			(0.80)			(0.12)			(0.09)			(2510.70)
Observations	979	937	557	532	534	327	724	665	399	979	938	558
R squared	0.00	0.00	0.00	0.00	0.01	0.00	0.02	0.00	0.05	0.00	0.00	0.01
Mean DV Control	-0.10	0.35	0.46	0.60	0.64	0.64	0.26	0.24	0.20	397.23	1052.01	1422.30

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table E.2: **LRFU 2014-2015 Sales and Purchases:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) sales. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* "Y1" refers to treatment in Year 1, while "Y2" refers to treatment in Year 2. Amounts are in 90 kg bag units and values are in Ksh.

	Tot Amt Sold			Tot Val Sold			Tot Amt Purch			Tot Val Purch		
	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both
Treat Y1	0.10		0.01	557.96		252.96	0.08		0.20	298.39		407.96
	(0.23)		(0.47)	(645.31)		(1363.64)	(0.15)		(0.25)	(452.26)		(726.89)
Treat Y2		0.17	-0.12		338.96	-236.18		-0.23	-0.33		-811.94	-1274.11
		(0.22)	(0.55)		(670.48)	(1534.45)		(0.17)	(0.28)		(531.18)	(792.22)
Treat Y1 * Y2			0.29			773.24			0.13			829.11
			(0.67)			(1893.17)			(0.35)			(1010.21)
Observations	979	935	555	979	936	556	978	938	557	978	938	557
R squared	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.01	0.01	0.01	0.01
Mean DV Control	2.01	2.13	2.26	5646.07	6342.74	6387.60	1.90	1.86	1.72	5560.79	5590.23	5220.76

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table E.3: **LRFU 2014-2015 Sales by Season:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) sales. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* "Y1" refers to treatment in Year 1, while "Y2" refers to treatment in Year 2. Amounts are in 90 kg bag units and values are in Ksh.

	Harv Amt Sold			Harv Val Sold			Lean Amt Sold			Lean Val Sold		
	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both
Treat Y1	0.03		0.25	77.46		530.06	0.22		0.16	679.47		392.38
	(0.09)		(0.16)	(243.35)		(481.72)	(0.20)		(0.41)	(574.93)		(1155.90)
Treat Y2		0.18**	0.22		334.68	600.49		0.04	0.06		303.79	115.41
		(0.08)	(0.21)		(221.93)	(603.64)		(0.20)	(0.49)		(568.03)	(1307.93)
Treat Y1 * Y2			-0.22			-572.62			0.05			513.65
			(0.24)			(707.79)			(0.60)			(1676.81)
Observations	980	937	555	980	935	556	981	937	557	981	935	557
R squared	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Mean DV Control	0.52	0.46	0.36	1346.28	1267.63	1079.90	1.34	1.53	1.49	3974.15	4383.35	4354.60

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table E.4: **LRFU Purchases by Season:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) purchases. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* "Y1" refers to treatment in Year 1, while "Y2" refers to treatment in Year 2. Amounts are in 90 kg bag units and values are in Ksh.

	Harv Amt Purch			Harv Val Purch			Lean Amt Purch			Lean Val Purch		
	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both
Treat Y1	-0.04		0.17	-149.68		347.10	0.10		-0.03	370.11		-294.98
	(0.09)		(0.15)	(233.61)		(375.77)	(0.12)		(0.20)	(356.84)		(628.83)
Treat Y2		-0.08	-0.01		-298.29	-146.51		-0.09	-0.31		-279.60	-1092.92
		(0.08)	(0.17)		(215.31)	(406.71)		(0.13)	(0.21)		(416.36)	(668.77)
Treat Y1 * Y2			-0.19			-370.52			0.34			1432.54
			(0.20)			(494.05)			(0.27)			(869.14)
Observations	977	941	557	977	940	557	982	939	559	979	938	558
R squared	0.01	0.00	0.02	0.01	0.00	0.02	0.00	0.01	0.01	0.00	0.01	0.01
Mean DV Control	0.58	0.52	0.44	1484.23	1317.58	1144.34	1.29	1.25	1.27	3922.78	3926.80	4040.25

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table E.5: **LRFU 2015 Harvest and Input Use:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on 2015 LR harvest and input usage. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* "Y1" refers to treatment in Year 1, while "Y2" refers to treatment in Year 2. Harvests are in 90kg bag units. Non-labor input exp are the amount spent in Ksh on all fertilizers, hybrid seeds, DAP, CAN, and other physical inputs excluding labor. Labor person-days record the number of person-days of labor applied. All results are for maize plots only.

	Labor Person-Days			Non-Labor Input Exp			2015 Harvest		
	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both
Treat Y1	-4.76		-13.76	18.46		315.04	-0.22		-1.53*
	(5.98)		(9.85)	(213.39)		(393.59)	(0.56)		(0.92)
Treat Y2		-9.66	-16.38		122.23	-153.46		0.92	-0.42
		(7.04)	(13.00)		(194.98)	(404.36)		(0.59)	(0.94)
Treat Y1 * Y2			14.63			402.65			2.39*
			(15.84)			(526.04)			(1.27)
Observations	979	940	560	978	940	559	987	946	561
R squared	0.01	0.00	0.06	0.01	0.00	0.01	0.00	0.00	0.02
Mean DV Control	126.15	131.48	142.58	2620.61	2271.07	2001.67	9.78	9.97	10.95

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table E.6: **LRFU 2015 Food Consumption, Food Expenditure, Total Consumption, and Happiness:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on food consumption, expenditure, total consumption, and happiness. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* "Y1" refers to treatment in Year 1, while "Y2" refers to treatment in Year 2. Maize Eaten in the past week in 2kg "goros." Food expenditure is the value of maize purchases, own production consumed, and gifts given to others over the past 30 days. HH consumption is the total household consumption (logged) over the past 30 days. Happy is an index for the following question: "Taking everything together, would you say you are very happy (3), somewhat happy (2), or not happy (1)?"

	Maize Eaten			Food Exp			HH Cons			Happy		
	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both
Treat Y1	-0.11		0.43	40.82		-124.26	-0.03		-0.00	0.10**		0.05
	(0.19)		(0.38)	(247.76)		(492.87)	(0.05)		(0.10)	(0.05)		(0.08)
Treat Y2		-0.26	-0.13		99.58	-97.26		0.04	0.08		0.01	0.00
		(0.22)	(0.41)		(251.35)	(556.87)		(0.05)	(0.11)		(0.04)	(0.10)
Treat Y1*Treat Y2			-0.47			254.32			-0.09			-0.03
			(0.54)			(658.28)			(0.13)			(0.12)
Observations	976	937	554	977	939	557	976	939	556	985	945	560
R squared	0.00	0.00	0.01	0.02	0.00	0.02	0.01	0.00	0.01	0.01	0.00	0.01
Mean DV Control	5.68	5.74	5.51	6840.11	6786.12	6928.43	9.50	9.47	9.49	2.40	2.47	2.48

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table E.7: **LRFU 2015 Education:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment education and non-farm profit. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* "Y1" refers to treatment in Year 1, while "Y2" refers to treatment in Year 2. Education attendance is the proportion of days the children in the household attended school in the last 5 days. Educational expenditure is the total household expenditure on children's education (in Ksh) over the past 12 months.

	Edu Exp			Edu Attend		
	Y1	Y2	Both	Y1	Y2	Both
Treat Y1	-3654.14		-6576.46	0.00		0.02
	(3854.68)		(6998.49)	(0.01)		(0.02)
Treat Y2		-1168.61	-4367.33		-0.01	0.02
		(2917.71)	(8041.06)		(0.01)	(0.02)
Treat Y1*Treat Y2			2391.45			-0.04
			(9231.27)			(0.03)
Observations	979	936	556	927	876	528
R squared	0.00	0.00	0.01	0.00	0.00	0.01
Mean DV Control	38371.63	37452.55	43373.16	0.94	0.95	0.93

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

Table E.8: **LRFU 2015 Non-Farm Business and Salaried Employment:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on non-farm business and salaried employment. The "Year 1" column contains observations that were in the sample in the Year 1 study, the "Year 2" column contains observations that were in the sample in the Year 2 study, and the "Both" column contains the (select) subset of respondents who were in both samples.* "Y1" refers to treatment in Year 1, while "Y2" refers to treatment in Year 2. Hours Non-Farm is the number of hours worked by the household in a non-farm businesses run by the household in the last 7 days. Non-farm profit is the household's profit from non-farm activities in the last month (Ksh). Hours Salary is the total number of hours worked by household members in a salaried position. Avg Wage is the average monthly wage for those household members who are salaried.

	Hours Non-Farm			Non-Farm Profit			Hours Salary			Avg Wage		
	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both	Y1	Y2	Both
Treat Y1	0.94		1.41	-186.29		48.03	-2.28		1.47	1892.96		884.26
	(1.75)		(2.71)	(285.72)		(528.13)	(1.77)		(3.57)	(1697.63)		(3231.62)
Treat Y2		0.22	0.63		-244.86	-47.72		-0.98	-1.74		3651.39**	528.77
		(1.87)	(3.43)		(315.71)	(607.26)		(1.98)	(4.49)		(1700.71)	(3525.65)
Treat Y1*Treat Y2			4.05			-47.91			-4.57			3027.24
			(4.25)			(744.40)			(5.19)			(4752.24)
Observations	979	937	556	975	933	552	982	939	559	292	274	155
R squared	0.01	0.00	0.02	0.00	0.00	0.01	0.00	0.00	0.01	0.00	0.02	0.02
Mean DV Control	15.97	14.87	13.32	2138.25	2019.84	1966.83	15.03	14.30	15.50	13014.88	12646.63	12714.71

*Note that differential attrition from Year 1 to Year 2 means that those in the "Both" column are a select subsample. "Y1" should be interpreted with particular caution in this column, given the possibility that treatment in year 1 affected selection into this sample and may therefore no longer represent a causal effect. While "Y2" was re-randomized among the remaining sample and therefore represents a causal effect, it should be remembered that this the causal effect among a specific subset of respondents.

E.2 Long-Run Price Effects

Table E.9 displays the long-run effects on market prices for our main specification (3km radius), as well as the robustness to 1km and 5km radii. While point estimates are substantially weaker (at least, on the "Hi" treatment coefficient) and are measured with noise, point estimates go in the same direction as the main effects. There may therefore be some small lingering effects on market price following the intervention.

Table E.9: **LRFU Market prices for maize as a function of local treatment intensity.** "Hi" intensity is a dummy for a sublocation randomly assigned a high number of treatment groups. "Month" is a linear month time trend (beginning in Nov at 0 in each year). Standard errors are clustered at the sublocation level. Prices measured during the long-run follow-up year (Nov-Aug in the year following Y2 (2014-2015)). Price normalized to 100 in Nov 2014 in "low" sublocations.

	3km	1km	5km
Hi	1.87	0.90	0.93
	(2.73)	(2.80)	(2.50)
Month	3.34***	3.22***	3.06***
	(0.29)	(0.32)	(0.29)
Hi Intensity * Month	-0.67	-0.45	-0.04
	(0.75)	(0.76)	(0.71)
Observations	253	253	253
R squared	0.25	0.25	0.25

E.3 Long-Run Effects Interacted with Treatment Intensity

Point estimates suggest (and are significant in Year 2) that the percent sold in the high price period (lean season) and the percent purchased in the low price season (post-harvest) are higher in low-saturation areas. In high saturation areas, the negative interaction terms cancels this effect out (see Table E.11). This is consistent with the idea that in low intensity areas, the lack of effect on prices means storage is highly profitable, encouraging individuals to purchase more in the post-harvest period and sell more in the lean season. In contrast, in high intensity areas, price effects dampen the returns to arbitrage, and there is less incentive to store.

The point estimates on revenue effects, although not significant, reflect this heterogeneity. Point estimates suggest there may be large gains from treatment in low intensity areas, but that these gains are partially or fully washed out in high intensity areas. Point estimates on the effect of living in a high intensity areas for control individuals are also positive, albeit far from significant.

Table E.10: **LRFU 2014-2015 Outcomes:** Effect of being ever treated (either in Year 1 or Year 2) on Year 3 (2014-2015) outcomes. "Net Sales" is the total number of 90kg bags sold - the total number of 90kg bags purchased between the 2014 long-rains harvest and 2015 long-rains harvest. "% Sold Lean" is the percentage of total sales completed from January onward. "% Purch Harvest" is the percentage of total purchases completed prior to January. "Revenues" are the net revenues from all maize sales and purchases from the 2014 long-rains harvest to the 2015 long-rains harvest.

	% Lean Sales		% Harvest Purch		Revenues	
	Overall	By Int	Overall	By Int	Overall	By Int
Ever Treated	-0.02	0.03	-0.01	0.07	477.35	1624.37
	(0.04)	(0.08)	(0.03)	(0.04)	(935.32)	(1667.68)
Hi		-0.03		0.13***		1855.76
		(0.08)		(0.05)		(1469.26)
Ever Treated*Hi		-0.07		-0.11*		-1636.83
		(0.09)		(0.06)		(1928.35)
Observations	739	739	990	989	1359	1358
R squared	0.00	0.01	0.00	0.01	0.01	0.01

Table E.11: **LRFU 2014-2015 Outcomes:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) outcomes. The "Year 1" column contains observations that were in the sample in the Year 1 study, while the "Year 2" column contains observations that were in the sample in the Year 2 study. "Y1" refers to treatment in Year 1, while "Y2" refers to treatment in Year 2. "Net Sales" is the total number of 90kg bags sold - the total number of 90kg bags purchased between the 2014 long-rains harvest and 2015 long-rains harvest. "% Sold Lean" is the percentage of total sales completed from January onward. "% Purch Harvest" is the percentage of total purchases completed prior to January. "Revenues" are the net revenues from all maize sales and purchases from the 2014 long-rains harvest to the 2015 long-rains harvest.

	Net Sales		% Hi Sales		% Lo Purch		Revenues	
	Y1	Y2	Y1	Y2	Y1	Y2	Y1	Y2
Treat Y1	0.32		0.04		0.02		1089.62	
	(0.59)		(0.08)		(0.05)		(1701.46)	
Treat Y1*Hi	-0.01		-0.00		-0.06		-1052.68	
	(0.73)		(0.10)		(0.07)		(2080.75)	
Treat Y2		0.65		0.08		0.10**		2156.20
		(0.74)		(0.07)		(0.05)		(2269.60)
Treat Y2*Hi		-0.49		-0.16*		-0.19***		-1204.33
		(0.81)		(0.08)		(0.06)		(2477.80)
Hi	-0.01	0.26	-0.10	-0.01	0.08	0.18***	1007.50	648.42
	(0.59)	(0.46)	(0.08)	(0.06)	(0.05)	(0.05)	(1585.07)	(1520.18)
Observations	979	936	532	534	724	664	979	937
R squared	0.00	0.00	0.01	0.02	0.02	0.03	0.00	0.00
Mean DV Control	-0.12	0.14	0.67	0.64	0.22	0.12	-428.76	501.34

F Effects of Tags

Table F.1: **Effects of tags.** Regressions include round fixed effects, with errors clustered at the group level.

	(1)	(2)	(3)	(4)	(5)	(6)
	Inventories	Inventories	Revenues	Revenues	Consumption	Consumption
Year 1 - Treat	0.52***		276.29		0.00	
	(0.16)		(291.42)		(0.03)	
T1 (Oct Loan)		0.69***		520.96		-0.00
		(0.19)		(321.38)		(0.04)
T2 (Jan Loan)		0.36*		31.78		0.00
		(0.19)		(346.02)		(0.04)
Tags	0.06	0.06	71.00	71.30	-0.01	-0.01
	(0.23)	(0.23)	(411.42)	(411.44)	(0.05)	(0.05)
Observations	4273	4273	4229	4229	4223	4223
R squared	0.34	0.34	0.01	0.01	0.00	0.00
Year 1 Treat-tags p-val	0.06		0.63		0.89	
T1-tags p-val		0.02		0.31		0.93
T2-tags p-val		0.26		0.93		0.86

G Savings Constraints and Effect of Lockboxes

How long might it take for a farmer to "save his way out" of this credit constraint? While the amount he would need to be fully released from this credit constraint is an ill-defined concept, one useful threshold is the point at which the farmer would be able to self-finance the loan.

We consider a few scenarios as benchmarks. If he receives the loan continuously each year and saves all of the additional revenue generated by the loan (1,548Ksh each year, according to our pooled estimate) under his mattress, he should be able to save the full average amount of the loan in 3.5 years. If instead the farmer reinvested this additional revenue, such that it compounds, he could save the full amount of the loan in a little less than 3 years. If the loan is only offered once, it would take more than 6 years of reinvesting his returns to save the full amount of the loan.

These may seem like fairly short time periods required for the farmer to save his way out of his credit constraint. However, the above estimates assume the the farmer saves 100% of the return from the loan. This may not be empirically accurate, nor optional, given that the farmer has urgent competing needs for current consumption. As an example, take the case in which the farmer instead saves only 10% of his return under her mattress. It would then take him 34 years to save the the full amount of the loan, even if it were continually offered during that period. Therefore, low savings rates are important to understanding why credit constraints persist in the presence of high return, divisible investment opportunities.

G.1 Short-Run Effects of the Lockbox

In order to test the importance of savings constraints, we examine the impact of the lockbox, as well as its interaction with the loan. First, in Table G.1, we explore the immediate effects of the lockbox for outcomes in Year 1 (recall the lockbox was only offered in Year 1, and was crosscut with the loan treatment). We observe no primary significant effects of the lockbox on inventories, revenues, or consumption (Columns 1, 3, and 5). Interestingly, when interacted with the loan, we see that receiving the lockbox alone is associated with significantly *lower* inventories; perhaps the lockbox serves as a substitute savings mechanism, rather than grain (see Column 2). However, receiving both the lockbox and the loan is associated with a reversal of this pattern. We see no such heterogeneity on revenues (Column 4). Interestingly, the point estimates on consumption are negative (though not significant) for the lockbox and loan when received separately; however, the interaction of the two is large and positive (and significant, at 95%), canceling out this effect.

Table G.1: **Effects of lockboxes.** Regressions include round fixed effects, with errors clustered at the group level.

	(1) Inventories	(2) Inventories	(3) Revenues	(4) Revenues	(5) Consumption	(6) Consumption
Lockbox	-0.08	-0.45*	96.63	-1.35	0.02	-0.07
	(0.15)	(0.25)	(234.65)	(419.79)	(0.03)	(0.05)
Treat		0.35*		232.93		-0.04
		(0.21)		(356.31)		(0.04)
Lockbox*Treat		0.54*		141.94		0.13**
		(0.31)		(509.35)		(0.06)
Observations	3836	3836	3795	3795	3792	3792
Mean DV	3.06	2.82	432.03	317.41	9.47	9.50
R squared	0.34	0.35	0.01	0.01	0.00	0.00

G.2 Long-Run Effects of the Lockbox

Because the hypothesized link between an inability to save one's way out of a credit constraint is an inherently intertemporal one, we also explore the long-run effects of the lockbox on storage behavior. Table G.4 explores these effects using data from our long-run follow-up survey, which explores outcomes two years after the lockboxes were distributed. The amount and value sold in the lean season are significantly lower among those who received the lockbox (but there appears to be no net effect among those who received the lockbox and the loan). There are no significant effects on sales in the harvest season or purchases in either period.

We also look for broader long-run effects on the lockbox on farming outcomes, which might compound over time. Table G.3 presents the long-run effects of the lockbox (and its interaction with the Year 1 loan) on harvest levels, sales, and revenues. While we observe no long-run shifts in the size or timing of sales (Columns 2-4), interestingly we do see large differences in harvests as a function of the lockbox and loan. Receiving the lockbox alone or the loan alone is associated with a significant and fairly large *reduction* in harvest levels in 2014 (which is also seen again in 2015; see Column 3 of Table G.4). However, receiving both cancels these effects out, producing a net effect of the combined treatment of roughly zero. These effects are reflected in the revenues (Column 5), and are replicated in the 2015 Harvest levels (Table G.4 Column 3). We see no significant difference in labor-days or input expenditure used that can explain these results (Table G.4 Column 1-2), though estimates on input expenditures are noisy.

Table G.2: **LRFU Lockbox 2014-2015 Sales and Purchases by Season:** Effect of lockbox and Year 1 (2012-2013) loan treatment on Year 3 (2014-2015) outcomes. The sample consists of those in the Year 1 study. Amounts are in 90 kg bag units and values are in Ksh. The “Lo” season refers to the period from harvest to January 1, a time of typically low prices. “Hi” refers to the period following January 1 to harvest, a time of typically high prices.

	Sales				Purchases			
	Amt Lo	Val Lo	Amt Hi	Val Hi	Amt Lo	Val Lo	Amt Hi	Val Hi
Lockbox	-0.11	-338.34	-0.54*	-1512.38*	-0.02	-42.08	-0.10	-174.98
	(0.16)	(422.15)	(0.29)	(824.08)	(0.16)	(385.78)	(0.18)	(545.83)
Loan (Y1)	-0.01	-43.81	0.05	209.82	-0.01	-100.42	0.07	223.68
	(0.11)	(291.29)	(0.24)	(688.00)	(0.10)	(264.05)	(0.13)	(364.42)
Lockbox*Loan (Y1)	0.15	470.95	0.67*	1897.95*	-0.06	-135.24	0.12	495.37
	(0.21)	(556.39)	(0.37)	(1069.13)	(0.18)	(444.76)	(0.22)	(696.86)
Observations	980	980	981	981	977	977	982	979
R squared	0.00	0.00	0.00	0.00	0.01	0.01	0.00	0.00
Mean DV Control	0.54	1423.58	1.46	4313.49	0.59	1505.99	1.32	3972.83

Table G.3: **LRFU Lockbox 2014-2015 Outcomes:** Effect of lockbox and Year 1 (2012-2013) loan treatment on Year 3 (2014-2015) outcomes. The sample consists of those in the Year 1 study. "2014 Harvest" is the size of harvest in 90kg bags. "Net Sales" is the total number of 90kg bags sold - the total number of 90kg bags purchased between the 2014 long-rains harvest and 2015 long-rains harvest. "% Hi Sales" is the percentage of total sales completed from January onward. "% Lo Purch" is the percentage of total purchases completed prior to January. "Revenues" are the net revenues from all maize sales and purchases from the 2014 long-rains harvest to the 2015 long-rains harvest.

	Harvest 2014	Net Sales	% Sold Lean	% Purch Harvest	Revenue
Lockbox	-1.57***	-0.57	-0.02	-0.00	-2398.15*
	(0.60)	(0.51)	(0.10)	(0.06)	(1332.45)
Loan (Y1)	-1.00*	0.17	0.02	-0.02	-174.04
	(0.51)	(0.42)	(0.05)	(0.04)	(1078.48)
Lockbox*Loan (Y1)	1.96**	0.64	0.07	0.01	2372.39
	(0.80)	(0.63)	(0.11)	(0.07)	(1987.18)
Observations	973	979	532	724	979
R squared	0.01	0.01	0.00	0.02	0.01
Mean DV Control	9.37	0.01	0.60	0.26	878.43

Table G.4: **LRFU Lockbox 2015 Harvest and Input Use:** Effect of lockbox and Year 1 (2012-2013) loan treatment on Year 3 (2014-2015) outcomes. The sample consists of those in the Year 1 study. Harvests are in 90kg bag units. Non-labor input exp are the amount spent in Ksh on all fertilizers, hybrid seeds, DAP, CAN, and other physical inputs excluding labor. Labor person-days record the number of person-days of labor applied. All results are for maize plots only.

	Labor	Inputs	Harvest 2015
Lockbox	-2.97	-469.77	-1.36*
	(8.75)	(371.68)	(0.79)
Loan (Y1)	-7.14	-62.24	-0.61
	(6.79)	(261.76)	(0.65)
Lockbox*Loan (Y1)	8.06	397.06	1.63
	(11.19)	(440.10)	(1.04)
Observations	979	978	987
R squared	0.01	0.01	0.00
Mean DV Control	127.17	2734.48	10.08

H Price Effects Robustness

H.1 Binary and Ratio Treatment Estimates

In this subsection, we test the robustness of price effects to functional form assumptions. Table H.1 presents a binary version of Equation 3, replacing $month_t$ with an indicator $lean_t$ for being in the lean season (defined as April-August) and the interaction term with $lean_t * H_s$. Results suggest similar significant increases in price post-harvest in high-intensity markets.

Table H.1: **Market prices for maize as a function of local treatment intensity (binary).** "Lean" is a binary variable for being in the lean season (Apr-Aug). "Month" is a linear month time trend (beginning in Nov at 0 in each year). Standard errors are clustered at the sublocation level. Prices measured monthly following loan disbursal (Nov-Aug in Y1; Dec-Aug in Y2). Price normalized to 100 in Nov control ("low") sublocations.

	Main Specification (3km)			Robustness (Pooled)	
	Y1	Y2	Pooled	1km	5km
Hi	3.69**	1.24	2.75**	1.61	2.12
	(1.46)	(1.17)	(1.19)	(1.13)	(1.23)
Lean	5.89***	11.01***	8.70***	8.44***	9.65***
	(1.84)	(1.29)	(1.58)	(1.54)	(1.26)
Hi Intensity * Lean	-3.74*	-1.25	-2.80	-2.39	-4.37**
	(2.00)	(1.60)	(1.66)	(1.61)	(1.51)
Observations	491	381	872	872	872
R squared	0.06	0.12	0.09	0.08	0.09

We also check the robustness of these results to a more continuous measure of treatment at the market-level, following the technique described in Miguel and Kremer (2004). We construct an estimate of the ratio of total treated farmers to the total farmers in our sample within a 3km radius around each market.[54]. We re-estimate an equation identical to Equation 3 with H_s replaced with $ratio_m$, the aforementioned ratio. Results are presented in Table H.2.

We also present non-parametric estimates of this specification in Figure H.1, displaying average prices in markets with above- vs. below-median ratios. While results are slightly noisier in this specification, the broad patterns remain consistent: prices are higher in the post-harvest period and lower in the lean period in markets with a greater proportion of treated individuals in the area.

[54] Because we draw twice the sample from high-intensity areas compared to low (in accordance with our randomized intensity), for the total farmer count, we weight the low-intensity observations by two to generate a count reflective of the true underlying OAF population.

Table H.2: **Market prices for maize as a function of local treatment intensity (ratio).** "Ratio" is the number of treated farmers within a given radius around the market/the total number of farmers (weighted) in our sample within the same radius. "Month" is a linear month time trend (beginning in Nov at 0 in each year). "Lean" is a binary variable for being in the lean season (Apr-Aug). Standard errors are clustered at the sublocation level. Prices measured monthly following loan disbursal (Nov-Aug in Y1; Dec-Aug in Y2). Price normalized to 100 in Nov control ("low") sublocations.

	Main Specification (3km)			Robustness (Pooled)	
	Y1	Y2	Pooled	1km	5km
Ratio	9.52*	7.19	4.33	2.23	4.78
	(5.27)	(4.11)	(4.12)	(2.45)	(4.88)
Month	1.27**	1.01**	1.33***	1.29***	1.34**
	(0.55)	(0.40)	(0.41)	(0.33)	(0.49)
Ratio * Month	-0.83	0.03	-0.59	-0.57	-0.59
	(0.95)	(0.91)	(0.69)	(0.60)	(0.87)
Observations	491	381	872	872	872
R squared	0.07	0.04	0.05	0.05	0.05

Figure H.1: **Pooled market prices for maize as a function of local treatment intensity (ratio)**. Market prices for maize as a function of the Miguel-Kremer treatment intensity ratio. The ratio is the total number of treated farmers/total OAF population within 3km radius. The left panel shows the average sales price in markets whose treatment ratio is above the median (solid line) versus below the median (dashed line) over the study period. The middle panel shows the average difference in log price between above- and below-median-ratio markets over time, with the bootstrapped 95% confidence interval shown in light grey and the 90% confidence interval shown in dark grey. The right panel shows prices over time in markets binned by the quarter of this ratio.

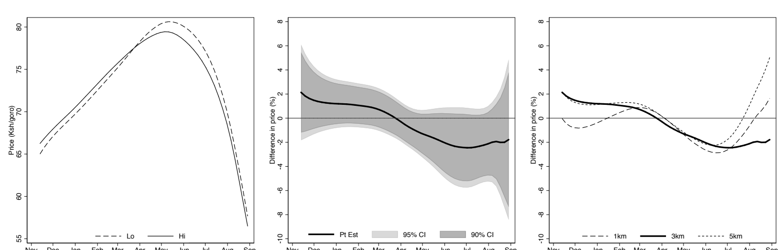

H.2 Randomization Inference, Wild Bootstrap, and Outlier Robustness

These market-level price results rely on the treatment saturation randomization being conducted at the sublocation level, a higher level than the group-level randomization employed in the individual-level results. While we cluster standard errors at the sublocation level, one might be concerned due to the small number of sublocations – of which we have 17 – that asymptotic properties may not apply to our market-level analyses and that our standard errors may therefore be understated. We run several robustness checks to address these small sample concerns.

First, building on other experimental work with small numbers of randomization units (Bloom et al., 2013; Cohen and Dupas, 2010), we use nonparametric randomization inference to confirm our results. We generate 1000 placebo treatment assignments and compare the estimated price effects under the "true" (original) treatment assignment to estimated effects under each of the placebo assignments.[55] Results are shown in Figure H.2. The left-hand panel of each figure shows price differences under the actual treatment assignment in black, and the placebo treatment assignments in grey. "Exact" p-values on the test that the price difference is zero are then calculated by summing up, at each point in the support, the number of placebo treatment estimates that exceed the actual treatment estimate (in absolute value) and dividing by the total number of placebo treatments (1000 in this case); these are shown in the right-hand panel of each figure.

Figure H.2 suggests that prices differences observed in the pooled data are significant at conventional levels from December to mid-February. This is roughly consistent with the results shown in Figure 7.

Figure H.2: **Nonparametric Randomization Inference** *Left panel*: price effects under the "true" treatment assignment (black line) and 1000 placebo treatment assignments (grey lines). *Right panel*: randomization-inference based p-values, as derived from the left panel.

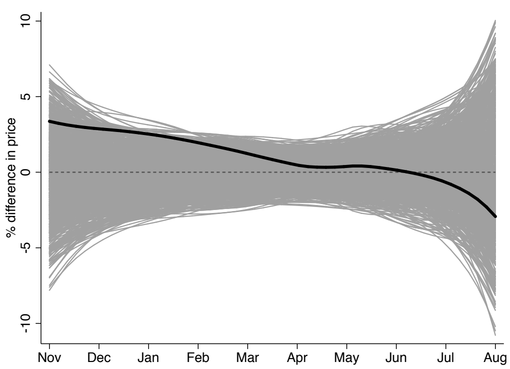

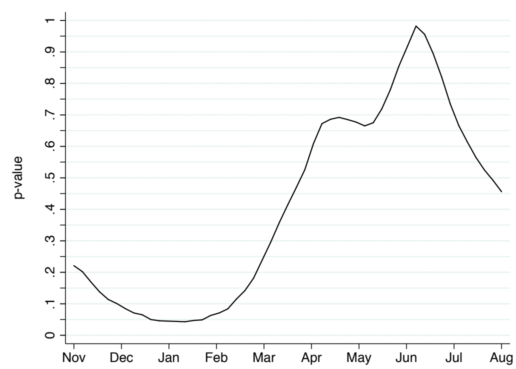

As an alternative method of accounting for the small number of clusters, we implement the wild bootstrap procedure proposed by Cameron et al. (2008). As a point of reference, Columns 1, 3, and 5 of Table H.3 present the results from the primary specification (that presented in Table 8) with p-values presented in the notes. Columns 2, 4, and 6 present the results from the wild bootstrapping exercise, with the empirical p-values in the notes (empirical p-values represent twice

[55] With 17 sublocations, 9 of which are "treated" with a high number of treatment farmers, we have 17 choose 9 possible treatment assignments (24,310). We compute treatment effects for a random 1,000 of these possible placebo assignments.

the fraction of t-statistics from the bootstrap samples that are above (below) the initial t-statistic for positive (negative) t-statistics). Comparing columns of Table H.3, we see only a small decrease in statistical precision.

Table H.3: **Wild bootstrap** Specifications as presented in Table 8, but with empirical p-values assessed using the wild bootstrap procedure proposed by Cameron et al. (2008), clustering at the sublocation level.

	Y1		Y2		Pooled	
	(1)	(2)	(3)	(4)	(5)	(6)
Hi	4.41	4.41	2.85	2.85	3.97	3.97
Month	1.19	1.19	1.22	1.22	1.36	1.36
Hi Intensity * Month	-0.57	-0.57	-0.48	-0.48	-0.57	-0.57
Observations	491	491	381	381	872	872
Mean of Dep Var	62.15	62.15	62.15	62.15	62.15	62.15
R squared	0.08	0.08	0.03	0.03	0.06	0.06
Wild Bootstrap	No	Yes	No	Yes	No	Yes
P-val Hi	0.05	0.10	0.17	0.20	0.04	0.08
P-val Month	0.01	0.04	0.01	0.00	0.00	0.03
P-val Hi*Month	0.19	0.18	0.32	0.32	0.16	0.17

Finally, to ensure that the trends observed are not driven by a single sublocation, we drop sublocations one-by-one and re-estimating prices differences. The results of this exercise are presented in Figure H.3. Differential trends over time in the two areas do not appear to be driven by particular sublocations.

Figure H.3: **Robustness to dropping each sublocation** Difference in prices between high and low-density markets over time for the full sample (black line) and for the sample with each sublocation dropped in turn (grey lines).

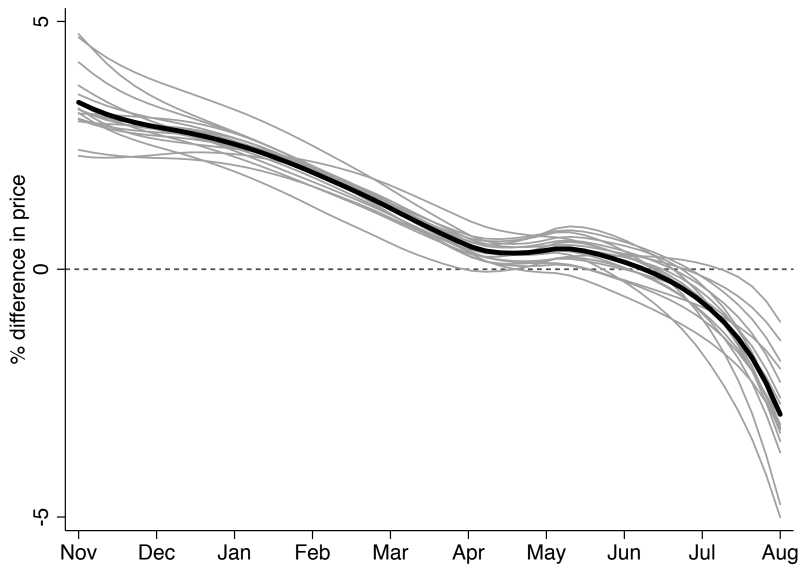

H.3 Pre-Specified Measures of Price Effects

As noted in Section 3, the pre-analysis plan (PAP) specifies the outcome of interest to be the percent price spread from November to June. We selected these dates to roughly match (i) the trough and peak price periods, respectively; and (ii) the period during which the loan was disbursed. However, there is variation in timing of both periods. For example, in Year 1 prices peaked in April (the exact trough is unknown, as price data collection only began in November of that year) and in Year 2 prices reached their trough in September and peaked in June. As for the loan disbursal period, loans were offered in October and January in Year 1 and in November in Year 2. Therefore, the impact of the loan may not map exclusively to the November-to-June price change. To allow for greater flexibility in the timing of these effects, the primary specification employed in the main text presents the non-parametric effect of treatment on the evolution of monthly prices, as well as a level and time trend effect. This also allows greater use of the full data. While we have 872 monthly observations of price across these markets over the pooled study period, because the pre-specified metric only allows for a single outcome per market per year, our observations fall to 95 in this specification.

However, for completeness, here we present the pre-specified effect of treatment saturation on the percentage change in prices from November to June. We hypothesized that the treatment would cause a reduction in this gap in treated areas, representing smoother prices across the season. We observe no effect of the treatment on the percent price increase from November to June. Looking at Figure 8, we observe a sizable increase in prices in the immediate post-harvest period in November, a gap which slow tappers off until June, when prices equalize in high and low treatment density markets. The simple comparison of November to June, which bookends this period, ignores data from the interim period, during which we also observe differences in prices between high and low treatment intensity markets. It also ignores the subsequent fall in prices in high markets relative to low in the following period. This analysis is therefore vastly underpowered relative to the analysis conducted in the main text.

Table H.4: **Pooled Price Gap Nov - June** Percent increase in price from November to June regressed on indicator for being in a high saturation sublocation.

	(1) Y1	(2) Y1	(3) Pooled
Hi	-0.02	0.02	0.00
	(0.04)	(0.02)	(0.03)
Observations	52	43	95
Mean DV	0.14	0.25	0.19
R squared	0.01	0.01	0.00

H.4 Effect on Related Outcomes

We explore whether treatment intensity had effects on related outcomes. First we check whether treatment effects can be seen in farmgate prices (see Table H.5). Using individual-level sales prices as reported in the household survey, we estimate a specification identical to Equation 1. We normalize prices in the low-intensity households in round 1 to 100, such that estimates can be interpreted as percentage changes relative to this baseline. We see similar patterns to those presented in Table 8. Point estimates suggest that prices are 3.32% higher in round 1 (significant at 5%), 2.92% higher in rough 2 (significant at 10%), and 0.72% lower in round 3 (not significant).

Note that these results should be interpreted with caution, as farmgate sales price is only observed for farmers who sell maize during the round in question. Any extensive margin response to treatment may bias these estimates. However, it is reassuring that they roughly aline with the main estimates using the market data (which does not suffer from such selection biases).

We also explore whether trader movement responds to treatment intensity. In Table H.6, we see some evidence that fewer traders enter high-intensity treated markets in the immediate post-harvest period in Year 2, which may be a sensible demand respond to the increase in price observed during a time when traders are typically purchasing. This may also contribute to the weaker price effects observed in Year 2.

Table H.6 presents effects of treatment intensity on the number of traders present in the market. We see these local markets are quite small; there are only 0.55 traders in a given market on average.

Table H.5: **Farmgate prices for maize as a function of local treatment intensity.** "Hi" intensity is a dummy for a sublocation randomly assigned a high number of treatment groups. "Round" represents the round of the survey (1, 2, or 3). Standard errors are clustered at the sublocation level. Regression includes round-year fixed effects and a control for the interview date. Price normalized to 100 in round 1"low" sublocations.

	(1) Y1	(2) Y2	(3) Pooled
Hi - R1	4.66**	1.52	3.32**
	(2.03)	(1.27)	(1.40)
Hi - R2	3.16*	2.21	2.95*
	(1.59)	(2.86)	(1.47)
Hi - R3	-0.35	-3.51	-0.72
	(1.27)	(5.31)	(1.56)
Observations	1582	636	2218
R squared	0.45	0.20	0.42

Table H.6: **Number Traders**

	Y1		Y2		Pooled	
Hi	-0.13	-0.07	-0.34	-0.37**	-0.22	-0.17*
	(0.11)	(0.09)	(0.24)	(0.16)	(0.15)	(0.09)
Month		0.02		0.03		0.04*
		(0.02)		(0.03)		(0.02)
Hi Intensity * Month		-0.02		0.01		-0.01
		(0.02)		(0.04)		(0.02)
Observations	451	451	419	419	870	870
Mean of Dep Var	0.32	0.32	0.82	0.82	0.55	0.55
R squared	0.01	0.01	0.02	0.03	0.01	0.02

I Main Results Weighted by Sample Weights

Because the sample consists of 80% of OAF farmers in high intensity areas and 40% of OAF farmers in low intensity areas, the main effects displayed in this paper over-represent the effects experienced by the farmers in high intensity areas. Given the heterogeneity identify in Section 5, these estimates are likely to be lower than a sample that was evenly pulled from high and low intensity areas.

We therefore attempt to identify the average effects that would have prevailed for such a sample here, by weighting observations by the inverse probability that they are in our sample, with the probability defined as the percentage of our sample drawn from high and low intensity areas respectively (essentially weighting observations from low-intensity areas by two). Table I.1 presents the results for our main pooled specification.

Table I.1: **Main Treatment Effects, Weighted by Inverse Sampling Probabilities.** Observations are weighted by the inverse probability that they are in our sample, with the probability defined as the percentage of our sample drawn from high and low intensity areas respectively.

	Inventory		Net Revenues		Consumption	
	Overall	By rd	Overall	By rd	Overall	By rd
Treat	0.57***		700.69***		0.02	
	(0.14)		(247.64)		(0.03)	
Treat - R1		1.06***		-455.42		0.01
		(0.22)		(298.82)		(0.04)
Treat - R2		0.57***		1318.27***		0.02
		(0.14)		(385.29)		(0.03)
Treat - R3		0.10		1223.68***		0.02
		(0.21)		(353.84)		(0.04)
Observations	6780	6780	6730	6730	6736	6736
Mean DV	2.16	2.16	-1616.12	-1616.12	9.55	9.55
R squared	0.30	0.30	0.09	0.10	0.02	0.02

Comparing Table I.1 to the main pooled effects on inventories (Table 2), we see that inventory effects are unaffected by this weighting. This is because there is little evidence of heterogeneity in inventory effects by treatment saturation, as seen in Table 9.

However, we do see a meaningful increase in the estimates of the effect on revenues (see previous estimates on revenues in Table 3). This is because the previous estimate drew more observations from on individuals in high-intensity areas, and these individuals experience lower revenue gains (see Table 10). The weighted specification, which corrects for the lower number of observations drawn from low intensity areas, finds higher effects.

Consumption effects are a bit lower with these weights (see previous estimates in Table 4), as consumption gains are higher in high intensity areas (see Table **??**); however, in neither case (unweighted or weighted) are they significant.

J Alternative Explanations for Individual Results Varying with Treatment Intensity

While our experiment affected local market prices differentially in high- and low-treatment density areas, changes in treatment density could precipitate other spillovers beyond output price effects. In Section 5, we attempt to clarify the sign and magnitude of these potential spillovers, as well as document one possible channel: price effects.

However, here we explore some alternative channels through which the differential net revenue effects could have occurred. Covariates were balanced at baseline in Year 1 between high- and low-intensity areas (Table J.1), as expected given the random assignment, so we can rule out simple concerns of imbalance.

Are these effects driven by differences in take-up? Among the pooled data, we see no differences in the (unconditional) loan size across the low and high intensity groups. We do, however, find some imbalances in loan take-up by intensity (see Table J.2). In high intensity areas, loan take-up is 5 percentage points lower than in low areas (significant at 5%) overall (Row 3). Interestingly, though, this pattern reverses from Year 1 (when loan take-up is 13 percentage points lower in high intensity areas) to Year 2 (when loan take-up is 8 percentage points higher in high intensity areas).[56] This differential take-up could matter for our treatment effects because we estimate the Intent-to-treat, and given a constant treatment-effect-on-the-treated, ITT estimates should be mechanically closer to zero in cases where take-up is lower. One might worry that, in particular in Year 1 when take-up is lower in the high intensity areas, this explains why revenue effects are also lower in high intensity areas. Two factors argue against this concern. First, the difference appears too small to explain our results fully. If there were no other spillovers, and treatment-on-treated effects were the same in high and low intensity areas, then ITT estimates in the high intensity ares should be 83% as large (0.61/0.74). However, point estimates on revenue treatment effects in Year 1 are roughly *zero* in the high-intensity areas (compared to 1,060 in low-intensity areas), a much bigger gap that could be explained by differential take-up. Second, and moreover, in Year 2, the differential take-up pattern switches; in this year, take-up is *higher* in high-intensity areas. If take-up were driving these results, we should see that a switch in the take-up patterns by intensity results in a switch in the revenue effects by intensity. However, we consistently across Years 1 and 2 see that revenue effects are greater among low-intensity areas. Take-up is therefore unlikely to be driving results.

We do additionally see some differences in loan size by intensity in Year 2. In this year of the experiment, loans were larger in high intensity areas.[57] However, this should have driven *greater* revenue effects in high intensity areas, rather than the lower effects that we find. We therefore believe it is unlikely that differential take-up or loan size are driving these results.

Finally, given the importance of social safety nets in rural communities, it is possible that informal lending between households could also be differentially affected by having a locally higher

[56] The Year 1 results may be the result of repayment incentives faced by OAF field staff: our loan intervention represented a substantial increase in the total OAF credit outlay in high-intensity areas, and given contract incentives for OAF field staff that reward a high repayment rate for clients in their purview, these field officers might have more carefully screened potential adopters. We are still exploring why the Year 2 results would have switched; given that the returns are more concentrated among low-intensity individuals, we would expect if anything higher take-up in Year 2 among the low-intensity individuals.

[57] Again, we are exploring why this might be the case, given that we would have expected, if anything, the lower returns in Year 1 in the high-intensity areas to lead to *smaller* rather than larger loans. It may be that given the price effects, a larger loan is necessary to arbitrage (e.g. if prices are higher at harvest, farmers would require a greater infusion of cash to supplement their outside option of sale at harvest and/or or fund purchases of maize at harvest).

density of loan recipients; as an untreated household, one's chance of knowing someone who received the loan is higher if one lives in a high-treatment-density areas. Perhaps high-intensity households have lower revenue effects because they share more with neighbors and others in their social network. Table J.3 explores this possibility, testing the impact of treatment on maize given away (as a gift or loan) and cash given away (as a loan). We find that the amount of transfers other households does not appear to respond to either treatment or to treatment intensity.

Overall, then, the individual-level spillover results are perhaps most consistent with spillovers through market prices.

Table J.1: **Balance among baseline covariates, high versus low treatment intensity areas.** The first two columns give the means in the low or high treatment intensity areas, the 3rd column the total number of observations across the two groups, and the last two columns the differences in means normalized by the standard deviation in the low intensity areas, with the corresponding p-value on the test of equality.

	Lo	Hi	Obs	Lo - Hi	
				sd	*p-val*
Male	0.32	0.31	1,589	0.02	0.72
Number of adults	3.11	3.07	1,510	0.02	0.74
Kids in school	3.15	2.98	1,589	0.09	0.11
Finished primary	0.71	0.75	1,490	-0.08	0.13
Finished secondary	0.27	0.25	1,490	0.04	0.51
Total cropland (acres)	2.60	2.35	1,512	0.08	0.15
Number of rooms in hhold	3.31	3.08	1,511	0.08	0.10
Total school fees (1000 Ksh)	29.23	27.88	1,589	0.04	0.51
Average monthly cons (Ksh)	15,586.03	14,943.57	1,437	0.05	0.38
Avg monthly cons./cap (log Ksh)	7.98	7.97	1,434	0.02	0.77
Total cash savings (KSH)	5,776.38	6,516.09	1,572	-0.04	0.56
Total cash savings (trim)	5,112.65	4,947.51	1,572	0.01	0.82
Has bank savings acct	0.42	0.42	1,589	-0.01	0.91
Taken bank loan	0.07	0.09	1,589	-0.06	0.30
Taken informal loan	0.25	0.24	1,589	0.02	0.72
Liquid wealth	87,076.12	98,542.58	1,491	-0.12	0.06
Off-farm wages (Ksh)	3,965.65	3,829.80	1,589	0.01	0.84
businessprofitmonth	1,859.63	2,201.34	1,589	-0.04	0.53
Avg %Δ price Sep-Jun	121.58	138.18	1,504	-0.21	0.00
Expect 2011 LR harvest (bags)	10.52	8.70	1,511	0.08	0.03
Net revenue 2011	-2,175.44	-4,200.36	1,428	0.03	0.45
Net seller 2011	0.34	0.30	1,428	0.08	0.16
Autarkic 2011	0.06	0.07	1,589	-0.04	0.53
% maize lost 2011	0.01	0.01	1,428	0.00	0.95
2012 LR harvest (bags)	11.57	10.94	1,484	0.07	0.19
Calculated interest correctly	0.68	0.74	1,580	-0.12	0.03
Digit span recall	4.49	4.60	1,504	-0.10	0.08
Maize giver	0.25	0.27	1,589	-0.05	0.37
delta	0.14	0.13	1,512	0.07	0.28

See Table 1 and the text for additional details on the variables.

Table J.2: **Loan Take-up and Size by Treatment Intensity.**

	Loan Take-up					**Loan Size (Cond)**					**Loan Size (Uncond)**				
	Low Mean	High Mean	N Obs	Diff SD	Diff p-val	Low Mean	High Mean	N Obs	Diff SD	Diff p-val	Low Mean	High Mean	N Obs	Diff SD	Diff p-val
Year 1	0.74	0.61	954	0.30	0.00	7,457.50	7,573.14	617	-0.05	0.60	5,524.07	4,616.96	954	0.23	0.00
Year 2	0.56	0.64	525	-0.17	0.07	9,434.52	11,281.25	324	-0.53	0.00	5,248.34	7,239.30	525	-0.37	0.00
Pooled	0.67	0.62	1,479	0.11	0.05	8,042.25	8,927.70	941	-0.30	0.00	5,425.18	5,543.95	1,479	-0.03	0.68

Table J.3: **Effect of Treatment on Transfers.** "Maize Given" represents the amount of maize (in terms of 90kg bags) given away to others outside the household, either as a gift or loan, in the past round (∼3 months). "Cash Given" represents the amount of cash (in Ksh) given to others outside the household as a loan in the past round.

	Maize Given		Cash Given	
	(1)	(2)	(3)	(4)
Treat	0.44	1.43	-31.12	-1.41
	(0.78)	(1.94)	(93.64)	(183.97)
Hi		-0.77		52.16
		(0.95)		(178.97)
Treat*Hi		-1.37		-42.92
		(2.07)		(224.83)
Observations	6850	6850	5987	5987
Mean DV	3.96	4.44	541.97	460.80
R squared	0.03	0.03	0.03	0.03

K Attrition

Attrition was relatively low in both years. In Year 1, overall attrition was 8%, and not significantly different across treatment groups (8% in the treatment group and 7% in the control). In Year 2, overall attrition was 2% (in both treatment and control, with no significant difference).

However, there was some non-random selection of the Year 2 study sample. Recall that the Year 1 sample consists of 240 existing One Acre Fund (OAF) farmer groups drawn from 17 different sublocations in Webuye district, and our total sample size at baseline was 1589 farmers. The Year 2 sample attempted to follow the same OAF groups as Year 1. However, a prerequisite for inclusion in the study sample is membership in OAF. Each year, farmers must opt into renewed engagement with OAF's services. There is some natural churn in this membership from year-to-year, with some existing members dropping out while new members join. Treatment in Year 1 had the effect of increasing farmers' interest in renewed engagement with OAF (a sensible result, given that the maize storage loan offer appears to be beneficial for farmers and therefore likely increased the perceived value of OAF's services).

As a result, the Year 2 sample, which was designed to include all farmers from Year 1 of the study, in practice includes a disproportionate number of farmers from the Year 1 treatment group.[58] Treated individuals were 10 percentage points more likely to return to the Year 2 sample than control individuals (significant at 1%).

Because Year 2 treatment status is stratified by Year 1 treatment status, the sample composition does not alter the internal validity of the Year 2 results. However, it may still have implications for our results, which we explore in this Appendix.

K.1 Year 2 Sample Composition

First, because this effect slightly alters the composition of the Year 2 sample, we may be interested in exploring how this affects the external validity, or generalizability, of our results. This is particularly relevant in the presence of heterogeneous treatment effects. For example, it may be that those for whom treatment was more beneficial were more likely to return to OAF, such that the Year 2 results are estimated on a sample for whom treatment effects are particularly strong. This would not affect the internal validity of the results for the sample in question, but it may affect our ability to generalize these results to other populations.

Table K.1 presents several key Year 1 outcome variables regressed on a dummy for Year 1 treatment status, a dummy for whether the individual stayed in the sample in Year 2, and an interaction term. In Column 1, for example, we see that those who stayed in the sample were famers with larger inventories. However, the insignificant interaction term suggests no evidence of a differential treatment effect on inventories (at least in Year 1) for those who stayed. In Column 2, we observe that stayers, on average, are those farmers who face higher purchase prices (perhaps for these farmers, the loan is more useful because they are facing high consumer prices). The interaction term is significant and negative, suggesting that treatment results in a particularly low purchase prices for stayers. This is consistent with the idea that those who stayed were those for whom the loan was most beneficial. We see similar patterns for sales prices (but with opposite

[58] Note that a second, broader result of this churn was a mix in the composition of the Year 2 sample between those drawn from the Year 1 sample (those who stayed from Year 1, comprising 602 individuals) and those who were new to the sample (417 individuals).

signs, as expected), though these results are not significant. We see no significant interaction for revenues or consumption.

Table K.1: **Selective attrition** Year 1 outcome variables regressed on dummy for whether treated in Y1, dummy for whether stayed in the sample in Y2, and interaction term. Sample is all Y1 subjects. Treatment effects at the individual level, all rounds. Regressions include round-year fixed effects and strata fixed effects, with errors clustered at the group level.

	(1)	(2)	(3)	(4)	(5)
	Invent	Purchase price	Sales prices	Rev	Log HH Cons
Treat Y1	0.53***	9.88	-19.52	315.01	-0.01
	(0.17)	(23.90)	(27.55)	(302.74)	(0.04)
Stayed Y2	0.68***	78.91**	-44.41	380.62	0.01
	(0.24)	(31.38)	(39.10)	(338.42)	(0.05)
Treat Y1 * Stayed Y2	-0.06	-100.28**	43.30	-158.30	0.05
	(0.29)	(38.84)	(45.41)	(408.97)	(0.06)
Observations	3836	1914	1425	3776	3792
Mean of Dep Variable	2.67	2982.02	2827.58	334.41	8.00
R squared	0.37	0.30	0.47	0.13	0.03
Controls	Yes	Yes	Yes	Yes	Yes

Table K.2 presents additional results on how stayers may differ from attriters. Stayers have significantly more children in school and pay more in school fees. This is consistent with focus groups that stated that farmers are often forced to sell maize early to pay for school fees; this group may get the most benefit from the loans and therefore be more eager to return to OAF with the hopes of taking up the loan. Stayers also had significantly larger harvests in 2011 and 2012, and were more likely to be net sellers in 2011. This is consistent with the idea that those with the most to sell have the most to gain from properly timing their sales. It could also reflect some underlying correlation between wealth and staying behavior. Consistent with this later interpretation, stayers are more likely to have a bank savings account. They also have greater liquid wealth, higher average monthly consumption, and more rooms in their household. Interestingly, despite being more likely to have completed primary school, stayers have significantly lower digit span recall. Sensible, stayers have higher values of δ, representing greater patience.

The stayer characteristics of greatest interest are those that are relevant in terms of treatment heterogeneity. Table K.3 presents this heterogeneity. It appears the treatment effect for men has a bigger effect on revenue, though a smaller effect on consumption. More kids are associated with smaller treatment effects on inventories. More rooms is associated with larger treatment effects on revenues. Having more cash savings is associated with larger effects for inventories, revenues, and consumption, though oddly baseline revenues are associated with smaller treatment effects on revenues. Having greater digit recall is associated with lower consumption effects.

With the exception of liquid wealth, which appears to be associated with both staying and larger treatment effects, there are not clear patterns that are suggest stayers are selected according to characteristics that are important for treatment heterogeneity.

K.2 Impacts of Two Years of Treatment

A second issue of note is how the selective attrition between Year 1 and Year 2 of the study may affect the interpretation of the long-run follow-up results. Results presented Appendix E include specifications that explore the long-run effects of the intervention separately by year (Equation 5) and specifications that explore the interaction of the two years' treatment statuses (Equation 6). Specifications from Equation 5 are well-identified, because the treatment was re-randomized within the sample each year. It is these effects on which we focus in the main text.

However, the estimates produced by Equation 6, which attempt to explore the impact of receiving treatment for two years in a row, do face potential selection bias. This specification includes a dummy for treatment in Year 1, a dummy for treatment in Year 2, and an interaction of the two. Because these variables are only defined for subjects present in both years of the study, the sample for this specification is restricted to those individuals. However, this is a sample selected endogenously based on the value of one of the regressors included in the specification (treatment in Year 1) and therefore this particular specification may not produce unbiased treatment estimates. For example, imagine that receiving treatment in Year 1 encourages poorer farmers in the treatment group to stick with OAF in Year 2, while these poorer farmers in the control group drop out. Because these poorer farmers from the control group will not be included in the specification defined by Equation 6, β_1 will produce an underestimate of the effect of the treatment on wealth, as it compares the full distribution of the treatment group to the upper distribution of the control group.

An alternative would be to consider the full Year 1 sample in Equation 6. However, T_2 and $T_1 * T_2$ is undefined for those individuals who dropped out of the sample between Year 1 and Year 2. Because T_2 would have been randomly assigned, had these attriters continued in the sample, one option is to randomly assign them a placebo treatment status for T_2, and simply consider those assigned to treatment in Year 2 to be "non-compliers" who were assigned but did not receive treatment. As a robustness test, we can also consider two alternate specifications that assign all attriters to treatment or all attriters to control, respectively, which allows us to bound these estimates at their extreme.

Tables K.4 - K.11 present these results. For each outcome variable, the first column "Actual" presents the results with the actual treatment status. As a result, attriters drop from the sample, as they are missing a T_2 treatment sample (these are identical results to those presented in Appendix E, but are displayed again here for comparison). The second column "Rand" presents results in which attriters are assigned a random T_2 treatment status. The third column "Treat" and the fourth column "Control" present results in which attriters are all signed the the treatment or control groups in Year 2, respectively.

In most cases, random assignment of treatment diminishes the estimated treatment effect in Y2 (sensibly, given that it essentially involves more non-compliance). Because most results are already insignificant, this does little to change the overall finding of little to no long-run effects of the intervention.

Table K.2: **Attrition and Sample Selection.** "Attrit" is an indicator for having exited the sample between Year 1 (2012-13) and Year 2 (2013-14). "Stay" is an indicator for being in the Year 1 and Year 2 samples

Baseline characteristic	Attrit	Stay	Obs	Attrit - Stay	
				sd	*p-val*
Treatment 2012	0.56	0.66	1,589	-0.20	0.00
Male	0.28	0.25	1,816	0.07	0.13
Number of adults	3.01	3.12	1,737	-0.05	0.30
Kids in school	2.89	3.23	1,816	-0.17	0.00
Finished primary	0.73	0.77	1,716	-0.08	0.10
Finished secondary	0.25	0.25	1,716	-0.01	0.81
Total cropland (acres)	2.26	2.50	1,737	-0.08	0.12
Number of rooms in hhold	2.94	3.34	1,738	-0.16	0.00
Total school fees (1000 Ksh)	25.93	30.08	1,816	-0.11	0.02
Average monthly cons (Ksh)	14,344.56	15,410.58	1,652	-0.09	0.10
Avg monthly cons./cap (log Ksh)	7.94	7.96	1,649	-0.04	0.49
Total cash savings (KSH)	5,355.05	6,966.35	1,797	-0.09	0.13
Total cash savings (trim)	4,675.61	4,918.86	1,797	-0.02	0.70
Has bank savings acct	0.38	0.46	1,816	-0.15	0.00
Taken bank loan	0.07	0.08	1,816	-0.04	0.46
Taken informal loan	0.23	0.24	1,816	-0.01	0.86
Liquid wealth	89,564.21	100,021.77	1,716	-0.10	0.05
Off-farm wages (Ksh)	3,508.17	4,103.66	1,816	-0.05	0.31
Business profit (Ksh)	2,069.13	2,159.55	1,816	-0.01	0.86
Avg %Δ price Sep-Jun	130.30	141.63	1,728	-0.15	0.00
Expect 2011 LR harvest (bags)	8.13	9.55	1,732	-0.09	0.05
Net revenue 2011	-4,983.94	-4,156.75	1,633	-0.02	0.72
Net seller 2011	0.26	0.35	1,633	-0.19	0.00
Autarkic 2011	0.06	0.07	1,816	-0.03	0.53
% maize lost 2011	0.01	0.01	1,609	0.00	0.98
2012 LR harvest (bags)	9.26	11.94	1,708	-0.31	0.00
Calculated interest correctly	0.72	0.72	1,806	-0.01	0.91
Digit span recall	4.61	4.50	1,731	0.09	0.06
Maize giver	0.26	0.26	1,816	0.00	0.98
Delta	0.86	0.87	1,738	-0.08	0.09

Table K.3: **Heterogeneity in Y1 Results.**

	Inventories	Revenues	Log Cons
Male	0.20	1,364.60	-0.08
	(0.32)	(599.24)	(0.06)
Number of adults	0.01	-140.41	0.00
	(0.07)	(128.51)	(0.02)
Kids in school	-0.19	-147.93	-0.01
	(0.08)**	(141.58)	(0.02)
Finished primary	0.40	384.44	-0.03
	(0.31)	(495.19)	(0.06)
Finished secondary	0.06	153.66	-0.02
	(0.36)	(625.11)	(0.06)
Total cropland (acres)	-0.05	15.17	-0.02
	(0.06)	(105.31)	(0.01)
Number of rooms in hhold	0.14	107.76	0.01
	(0.09)	(118.84)	(0.02)
Total school fees (1000 Ksh)	0.01	6.52	-0.00
	(0.00)	(7.03)	(0.00)
Average monthly cons (Ksh)	-0.00	-0.01	-0.00
	(0.00)	(0.02)	(0.00)
Avg monthly cons./cap (log Ksh)	0.38	363.43	-0.03
	(0.27)	(418.44)	(0.05)
Total cash savings (1000 KSH)	0.03	36.88	0.00
	(0.01)***	(17.35)	(0.00)
Total cash savings (1000 KSH, trim)	0.03	47.18	0.00
	(0.01)**	(21.09)	(0.00)
Has bank savings acct	0.47	229.86	0.01
	(0.31)	(474.00)	(0.06)
Taken bank loan	-0.49	-1,245.20	0.02
	(0.58)	(1,153.71)	(0.09)
Taken informal loan	0.18	-151.30	0.05
	(0.30)	(532.21)	(0.07)
Liquid wealth	-0.00	0.00	-0.00
	(0.00)	(0.00)	(0.00)
Off-farm wages (Ksh)	0.00	0.01	-0.00
	(0.00)	(0.02)	(0.00)
Business profit (Ksh)	-0.00	-0.04	0.00
	(0.00)	(0.06)	(0.00)
Avg %Δ price Sep-Jun	0.00	-2.85	0.00
	(0.00)	(2.72)	(0.00)
Expect	(0.00)	(3.93)	(0.00)
2011 LR harvest (bags)	0.02	30.54	0.00
	(0.02)	(31.19)	(0.00)
Net revenue 2011 (1000 KSH)	-0.02	-36.47	-0.00
	(0.01)*	(17.05)	(0.00)
Net seller 2011	0.14	16.71	0.04
	(0.33)	(592.08)	(0.06)
Autarkic 2011	-0.85	-594.17	0.04
	(0.69)	(1,141.05)	(0.13)
% maize lost 2011	1.51	3,284.26	0.13
	(1.80)	(3,093.89)	(0.39)
2012 LR harvest (bags)	0.00	35.39	-0.00
	(0.04)	(53.62)	(0.00)
Calculated interest correctly	0.30	864.63	0.06
	(0.33)	(505.61)	(0.07)
Digit span recall	-0.04	267.19	-0.05
	(0.13)	(206.28)	(0.03)
Maize giver	0.02	-364.06	0.03
	(0.32)	(564.18)	(0.06)
Delta	0.57	-326.76	-0.12
	(1.58)	(1,843.51)	(0.23)

Table K.4: **LRFU 2014-2015 Outcomes:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) outcomes. "Y1" is treatment status in Y1. In columns labeled "Actual", "Y2" is the actual treatment status in Year 2 (as a result, individuals not present in both years drop from the sample). In all other columns, the sample is those present in Year 1 (whether or not they were present in Y2 as well). For these columns, "Y2" treatment status is the actual treatment status for those in the Year 2 sample. For the subset of subjects not actually present in the Year 2 sample, a "Y2" placebo treatment status is randomly assigned for the columns labeled "Rand," assigned to be treatment for all in columns labeled "Treat," and assigned to be control for all in columns labeled "Control." Harvests, amount eaten, and net sales are in 90kg bag units. Revenues are in Ksh.

	Net Sales				% Sold Lean				% Purch Harvest				Revenues			
	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control
Treat Y1	-1.06	-0.81	-1.07	-0.75	-0.01	-0.37	-0.01	0.21	-0.02	0.03	-0.02	0.01	-763.60	-1296.03	-777.38	-108.65
	(0.86)	(0.57)	(0.84)	(0.49)	(0.59)	(0.40)	(0.59)	(0.40)	(0.09)	(0.07)	(0.09)	(0.05)	(1854.40)	(1246.72)	(1835.89)	(1089.72)
Treat Y2	0.45	0.25	-0.80	1.39*	0.29	-0.76	-0.70	1.06*	-0.05	0.02	-0.05	-0.02	1330.40	-1119.03	-1164.96	2889.69*
	(0.96)	(0.61)	(0.79)	(0.73)	(0.61)	(0.57)	(0.59)	(0.55)	(0.10)	(0.07)	(0.08)	(0.08)	(1777.33)	(1372.95)	(1485.96)	(1555.69)
Treat Y1*Treat Y2	0.64	0.53	0.78	0.29	0.21	1.35*	0.46	-0.01	0.10	0.01	0.08	0.08	1126.71	3181.94*	1552.74	486.47
	(1.15)	(0.77)	(0.94)	(0.93)	(0.80)	(0.69)	(0.70)	(0.68)	(0.12)	(0.09)	(0.10)	(0.10)	(2510.70)	(1906.18)	(2087.17)	(2067.51)
Observations	556	973	973	973	557	979	979	979	327	532	532	532	558	979	979	979
R squared	0.01	0.01	0.00	0.02	0.00	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.01	0.01	0.00	0.01
Mean DV Control	9.61	8.90	9.61	8.70	0.46	0.26	0.46	-0.35	0.64	0.59	0.64	0.60	1422.30	918.02	1422.30	-306.72

Table K.5: **LRFU 2014-2015 Sales and Purchases:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) outcomes. "Y1" is treatment status in Y1. In columns labeled "Actual", "Y2" is the actual treatment status in Year 2 (as a result, individuals not present in both years drop from the sample). In all other columns, the sample is those present in Year 1 (whether or not they were present in Y2 as well). For these columns, "Y2" treatment status is the actual treatment status for those in the Year 2 sample. For the subset of subjects not actually present in the Year 2 sample, a "Y2" placebo treatment status is randomly assigned for the columns labeled "Rand," assigned to be treatment for all in columns labeled "Treat," and assigned to be control for all in columns labeled "Control." Amounts are in 90 kg bag units and values are in Ksh.

	Tot Amt Sold				Tot Val Sold				Tot Amt Purch				Tot Val Purch			
	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control
Treat Y1	0.01	-0.28	0.00	-0.01	252.96	-654.13	237.73	253.61	0.20	0.16	0.21	0.04	407.96	386.18	416.01	125.50
	(0.47)	(0.30)	(0.46)	(0.25)	(1363.64)	(889.11)	(1336.95)	(732.92)	(0.25)	(0.21)	(0.25)	(0.18)	(726.89)	(604.24)	(724.56)	(548.02)
Treat Y2	-0.12	-0.43	-0.33	0.17	-236.18	-1492.48*	-974.59	632.42	-0.33	-0.08	0.19	-0.62**	-1274.11	-294.66	339.74	-2001.35***
	(0.55)	(0.33)	(0.42)	(0.46)	(1534.45)	(892.90)	(1184.98)	(1280.59)	(0.28)	(0.26)	(0.25)	(0.24)	(792.22)	(734.83)	(688.45)	(702.48)
Treat Y1* Treat Y2	0.29	0.77*	0.14	0.29	773.24	2442.31**	452.27	724.49	0.13	-0.15	-0.18	0.30	829.11	-123.92	-166.60	1118.34
	(0.67)	(0.42)	(0.54)	(0.53)	(1893.17)	(1195.80)	(1545.34)	(1504.27)	(0.35)	(0.31)	(0.30)	(0.30)	(1010.21)	(918.44)	(865.72)	(882.77)
Observations	555	979	979	979	556	979	979	979	557	978	978	978	557	978	978	978
R squared	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.01	0.01	0.01	0.02	0.01	0.01	0.01	0.01
Mean DV Control	2.26	2.22	2.26	1.97	6387.60	6341.52	6387.60	5494.64	1.72	1.94	1.72	2.05	5220.76	5700.84	5220.76	6045.99

Table K.6: **LRFU 2014-2015 Sales by Season:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) outcomes. "Y1" is treatment status in Y1. In columns labeled "Actual", "Y2" is the actual treatment status in Year 2 (as a result, individuals not present in both years drop from the sample). In all other columns, the sample is those present in Year 1 (whether or not they were present in Y2 as well). For these columns, "Y2" treatment status is the actual treatment status for those in the Year 2 sample. For the subset of subjects not actually present in the Year 2 sample, a "Y2" placebo treatment status is randomly assigned for the columns labeled "Rand," assigned to be treatment for all in columns labeled "Treat," and assigned to be control for all in columns labeled "Control." Amounts are in 90 kg bag units and values are in Ksh.

	Harv Amt Sold				Harv Val Sold				Lean Amt Sold				Lean Val Sold			
	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control
Treat Y1	0.25	-0.02	0.24	0.01	530.06	-113.39	519.07	86.03	0.16	0.03	0.15	0.18	392.38	47.32	383.38	473.64
	(0.16)	(0.13)	(0.16)	(0.11)	(481.72)	(351.04)	(474.29)	(282.97)	(0.41)	(0.26)	(0.41)	(0.22)	(1155.90)	(757.15)	(1145.46)	(632.65)
Treat Y2	0.22	-0.00	0.20	0.11	600.49	-116.83	359.33	448.04	0.06	-0.15	-0.21	0.28	115.41	-486.27	-535.15	672.18
	(0.21)	(0.15)	(0.14)	(0.19)	(603.64)	(396.21)	(414.90)	(530.76)	(0.49)	(0.28)	(0.35)	(0.41)	(1307.93)	(797.65)	(960.34)	(1141.01)
Treat Y1*Treat Y2	-0.22	0.09	-0.30	0.00	-572.62	368.12	-604.37	-160.71	0.05	0.36	0.09	0.03	513.65	1231.33	411.34	416.16
	(0.24)	(0.18)	(0.18)	(0.22)	(707.79)	(481.09)	(532.50)	(603.79)	(0.60)	(0.38)	(0.48)	(0.49)	(1676.81)	(1091.22)	(1347.98)	(1363.44)
Observations	555	980	980	980	556	980	980	980	557	981	981	981	557	981	981	981
R squared	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Mean DV Control	0.36	0.52	0.36	0.49	1079.90	1401.79	1079.90	1245.66	1.49	1.41	1.49	1.27	4354.60	4204.56	4354.60	3826.17

Table K.7: **LRFU Purchases by Season:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) outcomes. “Y1” is treatment status in Y1. In columns labeled “Actual”, “Y2” is the actual treatment status in Year 2 (as a result, individuals not present in both years drop from the sample). In all other columns, the sample is those present in Year 1 (whether or not they were present in Y2 as well). For these columns, “Y2” treatment status is the actual treatment status for those in the Year 2 sample. For the subset of subjects not actually present in the Year 2 sample, a “Y2” placebo treatment status is randomly assigned for the columns labeled “Rand,” assigned to be treatment for all in columns labeled “Treat,” and assigned to be control for all in columns labeled “Control.” Amounts are in 90 kg bag units and values are in Ksh.

	Harv Amt Purch				Harv Val Purch				Lean Amt Purch				Lean Val Purch			
	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control
Treat Y1	0.17	0.08	0.17	-0.02	347.10	227.08	340.53	-110.12	-0.03	0.02	-0.03	0.05	-294.98	-55.17	-288.59	160.40
	(0.15)	(0.12)	(0.15)	(0.11)	(375.77)	(305.89)	(373.97)	(292.22)	(0.20)	(0.17)	(0.19)	(0.14)	(628.83)	(538.01)	(626.54)	(430.56)
Treat Y2	-0.01	0.05	0.16	-0.17	-146.51	192.61	399.03	-574.68	-0.31	-0.19	-0.00	-0.41**	-1092.92	-791.16	-228.58	-1212.97**
	(0.17)	(0.13)	(0.14)	(0.15)	(406.71)	(328.37)	(355.19)	(371.99)	(0.21)	(0.18)	(0.18)	(0.18)	(668.77)	(515.93)	(559.30)	(563.74)
Treat Y1*Treat Y2	-0.19	-0.23	-0.29*	-0.01	-370.52	-720.25*	-671.33	56.40	0.34	0.17	0.17	0.27	1432.54	896.08	897.14	997.15
	(0.20)	(0.16)	(0.17)	(0.18)	(494.05)	(410.54)	(434.01)	(439.11)	(0.27)	(0.23)	(0.22)	(0.24)	(869.14)	(709.52)	(694.52)	(742.60)
Observations	557	977	977	977	557	977	977	977	559	982	982	982	558	979	979	979
R squared	0.02	0.02	0.01	0.02	0.02	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Mean DV Control	0.44	0.56	0.44	0.62	1144.34	1394.71	1144.34	1628.81	1.27	1.38	1.27	1.39	4040.25	4292.15	4040.25	4213.79

Table K.8: **LRFU 2015 Harvest and Input Use:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) outcomes. "Y1" is treatment status in Y1. In columns labeled "Actual", "Y2" is the actual treatment status in Year 2 (as a result, individuals not present in both years drop from the sample). In all other columns, the sample is those present in Year 1 (whether or not they were present in Y2 as well). For these columns, "Y2" treatment status is the actual treatment status for those in the Year 2 sample. For the subset of subjects not actually present in the Year 2 sample, a "Y2" placebo treatment status is randomly assigned for the columns labeled "Rand," assigned to be treatment for all in columns labeled "Treat," and assigned to be control for all in columns labeled "Control." Harvests are in 90kg bag units. Non-labor input exp are the amount spent in Ksh on all fertilizers, hybrid seeds, DAP, CAN, and other physical inputs excluding labor. Labor person-days record the number of person-days of labor applied. All results are for maize plots only.

	Labor Person-Days				Non-Labor Input Exp				2015 Harvest			
	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control
Treat Y1	-13.76	-8.79	-14.41	-6.65	315.04	-17.83	316.72	-191.32	-1.53*	-1.13	-1.56*	-0.88
	(9.85)	(7.07)	(9.64)	(6.49)	(393.59)	(312.97)	(393.24)	(255.00)	(0.92)	(0.75)	(0.91)	(0.64)
Treat Y2	-16.38	-8.48	-24.84***	2.22	-153.46	131.40	801.41**	-951.18***	-0.42	-0.53	-1.60*	1.01
	(13.00)	(8.65)	(9.48)	(11.61)	(404.36)	(317.45)	(363.26)	(339.00)	(0.94)	(0.69)	(0.81)	(0.78)
Treat Y1*Treat Y2	14.63	8.55	13.60	4.92	402.65	49.98	-422.70	917.65**	2.39*	1.75*	1.85*	1.66
	(15.84)	(11.11)	(11.69)	(14.76)	(526.04)	(429.78)	(455.81)	(440.10)	(1.27)	(0.96)	(1.05)	(1.11)
Observations	560	979	979	979	559	978	978	978	561	987	987	987
R squared	0.06	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.02	0.01	0.00	0.02
Mean DV Control	142.58	130.20	142.58	126.05	2001.67	2559.89	2001.67	2850.73	10.95	10.03	10.95	9.55

Table K.9: **LRFU 2015 Food Consumption, Food Expenditure, Total Consumption, and Happiness:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) outcomes. “Y1” is treatment status in Y1. In columns labeled “Actual”, “Y2” is the actual treatment status in Year 2 (as a result, individuals not present in both years drop from the sample). In all other columns, the sample is those present in Year 1 (whether or not they were present in Y2 as well). For these columns, “Y2” treatment status is the actual treatment status for those in the Year 2 sample. For the subset of subjects not actually present in the Year 2 sample, a “Y2” placebo treatment status is randomly assigned for the columns labeled “Rand,” assigned to be treatment for all in columns labeled “Treat,” and assigned to be control for all in columns labeled “Control.” Maize Eaten in the past week in 2kg “goros.” Food expenditure is the value of maize purchases, own production consumed, and gifts given to others over the past 30 days. HH consumption is the total household consumption (logged) over the past 30 days. Happy is an index for the following question: “Taking everything together, would you say you are very happy (3), somewhat happy (2), or not happy (1)?”

	Maize Eaten				Food Exp				HH Cons				Happy			
	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control
Treat Y1	0.43	-0.10	0.44	-0.09	-124.26	-36.86	-127.53	-13.50	-0.00	-0.01	-0.00	-0.02	0.05	0.10	0.05	0.12^{**}
	(0.38)	(0.27)	(0.38)	(0.22)	(492.87)	(353.95)	(489.83)	(300.31)	(0.10)	(0.06)	(0.10)	(0.06)	(0.08)	(0.07)	(0.08)	(0.05)
Y2	-0.13	-0.46	0.25	-0.42	-97.26	-135.65	-236.22	99.61	0.08	0.05	0.00	0.11	0.00	-0.01	-0.11	0.10
	(0.41)	(0.29)	(0.34)	(0.33)	(556.87)	(371.56)	(454.95)	(430.77)	(0.11)	(0.07)	(0.10)	(0.08)	(0.10)	(0.08)	(0.08)	(0.09)
Y1*Y2	-0.47	0.03	-0.74	0.08	254.32	161.05	233.28	133.53	-0.09	-0.05	-0.05	-0.07	-0.03	0.00	0.06	-0.09
	(0.54)	(0.38)	(0.46)	(0.43)	(658.28)	(459.99)	(550.80)	(534.88)	(0.13)	(0.08)	(0.11)	(0.10)	(0.12)	(0.09)	(0.09)	(0.10)
Observations	554	976	976	976	557	977	977	977	556	976	976	976	560	985	985	985
R squared	0.01	0.01	0.01	0.00	0.02	0.02	0.02	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Mean DV Control	5.51	5.90	5.51	5.77	6928.43	6908.15	6928.43	6840.65	9.49	9.47	9.49	9.48	2.48	2.41	2.48	2.38

Table K.10: **LRFU 2015 Education:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) outcomes. “Y1” is treatment status in Y1. In columns labeled “Actual”, “Y2” is the actual treatment status in Year 2 (as a result, individuals not present in both years drop from the sample). In all other columns, the sample is those present in Year 1 (whether or not they were present in Y2 as well). For these columns, “Y2” treatment status is the actual treatment status for those in the Year 2 sample. For the subset of subjects not actually present in the Year 2 sample, a “Y2” placebo treatment status is randomly assigned for the columns labeled “Rand,” assigned to be treatment for all in columns labeled “Treat,” and assigned to be control for all in columns labeled “Control.” Education attendance is the proportion of days the children in the household attended school in the last 5 days. Educational expenditure is the total household expenditure on children’s education (in Ksh) over the past 12 months.

	Edu Exp				Edu Attend			
	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control
Treat Y1	-6576.46	-4030.81	-6567.57	-3548.56	0.02	0.02	0.02	0.01
	(6998.49)	(5219.91)	(6969.11)	(4494.20)	(0.02)	(0.02)	(0.02)	(0.02)
Y2	-4367.33	-2493.22	-7206.06	1051.14	0.02	0.03	0.01	0.02
	(8041.06)	(5553.62)	(6582.92)	(6447.60)	(0.02)	(0.02)	(0.02)	(0.02)
Y1*Y2	2391.45	1037.39	4119.35	-656.67	-0.04	-0.03	-0.02	-0.03
	(9231.27)	(6503.79)	(7682.01)	(7587.93)	(0.03)	(0.02)	(0.03)	(0.03)
Observations	556	979	979	979	528	927	927	927
R squared	0.01	0.00	0.00	0.00	0.01	0.01	0.00	0.00
Mean DV Control	43373.16	39540.19	43373.16	38180.17	0.93	0.92	0.93	0.93

Table K.11: **LRFU 2015 Non-Farm Business and Salaried Employment:** Effect of Year 1 (2012-2013) and Year 2 (2013-2014) treatment on Year 3 (2014-2015) outcomes. "Y1" is treatment status in Y1. In columns labeled "Actual", "Y2" is the actual treatment status in Year 2 (as a result, individuals not present in both years drop from the sample). In all other columns, the sample is those present in Year 1 (whether or not they were present in Y2 as well). For these columns, "Y2" treatment status is the actual treatment status for those in the Year 2 sample. For the subset of subjects not actually present in the Year 2 sample, a "Y2" placebo treatment status is randomly assigned for the columns labeled "Rand," assigned to be treatment for all in columns labeled "Treat," and assigned to be control for all in columns labeled "Control." Hours Non-Farm is the number of hours worked by the household in a non-farm businesses run by the household in the last 7 days. Non-farm profit is the household's profit from non-farm activities in the last month (Ksh). Hours Salary is the total number of hours worked by household members in a salaried position. Avg Wage is the average monthly wage for those household members who are salaried.

	Hours Non-Farm				Non-Farm Profit				Hours Salary				Avg Wage			
	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control	Actual	Rand	Treat	Control
Treat Y1	1.41	0.22	1.32	-0.89	48.03	-131.45	40.31	-228.25	1.47	-0.50	1.50	-1.59	884.26	-0.89	953.70	1055.37
	(2.71)	(2.35)	(2.73)	(1.95)	(528.13)	(405.82)	(524.02)	(328.83)	(3.57)	(2.49)	(3.56)	(2.01)	(3231.62)	(2295.78)	(3189.99)	(1861.07)
Y2	0.63	-0.68	3.11	-1.98	-47.72	-215.99	168.88	-215.69	-1.74	-0.17	-0.53	-1.86	528.77	-523.66	343.68	207.24
	(3.43)	(2.48)	(2.55)	(3.15)	(607.26)	(411.13)	(459.69)	(503.22)	(4.49)	(2.84)	(3.60)	(3.59)	(3525.65)	(2371.79)	(2995.19)	(2690.15)
Y1*Y2	4.05	1.42	-0.57	6.09	-47.91	-73.42	-311.23	194.64	-4.57	-3.26	-5.09	-1.51	3027.24	3953.47	1338.23	3252.25
	(4.25)	(3.38)	(3.28)	(3.81)	(744.40)	(522.26)	(588.43)	(602.61)	(5.19)	(3.55)	(4.16)	(4.22)	(4752.24)	(3242.78)	(3823.14)	(3917.34)
Observations	556	979	979	979	552	975	975	975	559	982	982	982	155	292	292	292
R squared	0.02	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.01	0.00	0.01	0.00	0.02	0.01	0.01	0.01
Mean DV Control	13.32	16.31	13.32	16.54	1966.83	2240.21	1966.83	2198.80	15.50	15.11	15.50	15.41	12714.71	13249.44	12714.71	12978.62