NBER WORKING PAPER SERIES

SEASONAL CREDIT CONSTRAINTS AND AGRICULTURAL LABOR SUPPLY: EVIDENCE FROM ZAMBIA

Günther Fink
B. Kelsey Jack
Felix Masiye

Working Paper 20218
http://www.nber.org/papers/w20218

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2014

We thank Seema Jayachandran and audiences at University of Cape Town SALDRU, the University of Washington, the University ofWisconsin-Milwaukee, the University of Zambia, the UC Berkeley development lunch, the Harvard Development faculty retreat, the Gates Foundation, IGC Growth Week and PacDev for comments and suggestions. We are grateful to the International Growth Centre and the Agricultural Technology Adoption Initiative (JPAL/CEGA) for financial support and to Innovations for Poverty Action for logistical support. Austin Land and Chantelle Boudreaux provided excellent assistance with the field work and data, respectively. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Seasonal Credit Constraints and Agricultural Labor Supply: Evidence from Zambia
Günther Fink, B. Kelsey Jack, and Felix Masiye
NBER Working Paper No. 20218
June 2014
JEL No. J22,O16,Q12

ABSTRACT

Small-scale farming remains the primary source of income for a majority of the population in developing countries. While most farmers primarily work on their own fields, off-farm labor is common among small-scale farmers. A growing literature suggests that off-farm labor is not the result of optimal labor allocation, but is instead driven by households' inability to cover short-term consumption needs with savings or credit. We conduct a field experiment in rural Zambia to investigate the relationship between credit availability and rural labor supply. We find that providing households with access to credit during the growing season substantially alters the allocation of household labor, with households in villages randomly selected for a loan program selling on average 25 percent less off-farm labor. We also find that increased credit availability is associated with higher consumption and increases in local farming wages. Our results suggest that a substantial fraction of rural labor supply is driven by short-term constraints, and that access to credit markets may improve the efficiency of labor allocation overall.

Günther Fink
Harvard School of Public Health
Department of Global Health and Population
665 Huntington Ave.
Boston, MA 02115
gfink@hsph.harvard.edu

B. Kelsey Jack
Department of Economics
Tufts University
314 Braker Hall
Medford, MA 02155
and NBER
kelsey.jack@tufts.edu

Felix Masiye
Department of Economics
University of Zambia
Lusaka, Zambia
fmasiye@yahoo.com

1 Introduction

In Zambia, like in much of Sub-Saharan Africa, agriculture employs the vast majority of the rural population, with generally low levels of productivity and farming income.[1] A lack of irrigation infrastructure combined with a long dry season mean that harvest income arrives only once per year, and must cover household needs for the subsequent 10-12 months. If available resources are insufficient to cover consumption needs, households can either finance consumption gaps through credit or engage in alternative consumption smoothing strategies including off-farm labor.[2] From a household production perspective, selling family labor off-farm can be optimal as long as the marginal product of labor off-farm is higher than the marginal product of labor on the farm. However, a growing literature suggests that off-farm labor reflects a coping strategy rather than income maximization (Kerr 2005; Bryceson 2006; Orr et al. 2009; Michaelowa et al. 2010; Cole and Hoon 2013). The logic of the underlying argument is simple: since most small scale farmers have limited ownership of liquid assets and virtually no access to short-term credit, the least costly (or only possible) way to finance short-term consumption may be to generate wage income from local labor markets.

Conceptually, frictions in the credit market may thus lead to two inefficiencies in the farm's production schedule. First, the farm may alter the the total quantity of land used and crop mix chosen to ensure that the available family labor, net of anticipated work required off-farm, is sufficient to maximize yields. In most settings, the production plan chosen in the constrained environment will not be the same as the optimal choice in an unconstrained environment (Fafchamps 1993; Rosenzweig and Binswanger 1993), which may lower the farm's net income. Second, these extensive margin or ex ante inefficiencies can be further compounded by intensive margin inefficiencies if households experience unanticipated income or expenditure shocks during the farming season. To cover short-term consumption and liquidity needs, farming households may deviate from their original (adjusted) production plan by reallocating labor to earn off-farm wages (Kochar 1995, 1999; Rose 2001; Ito and Kurosaki 2009).[3]

[1]A recent study in the region of proposed research shows average gross production value of USD 500 for a family of six (Fink and Masiye 2012). Once capital input, land and labor costs are considered, net profits for many of these families may be negative.

[2]Most rural households own some assets such as livestock, agricultural livestock or cooking ware which they could theoretically liquidate; in practice, finding a buyer for such assets during the most constrained periods may be difficult, and the transaction cost of liquidating assets may be prohibitive.

[3]Common income shocks in the area are the loss of stored food reserves due to pests or theft; expenditure shocks include funerals, school uniforms and medical costs.

To investigate the relationship between credit availability and agricultural labor supply, we conducted a cluster-randomized experiment with small-scale farmers from selected villages in rural Zambia. As shown by the descriptive plots in in Figures 1 and 2 food shortages and off-farm labor (locally referred to as *ganyu*) are common in the study area, with pronounced seasonal patterns.[4] Zambia's agricultural cycle is centered around the rainy season from November to April. Harvest takes place in May and June, and generates income that must cover all consumption and investment needs until the following year's harvest. Household reserves are diminished by the beginning of the rainy season (six months after harvest), and are most scarce from January to March before early crops become available for consumption. The January to March period is generally referred to as the "hungry season" by farmers, and is the period we directly targeted with the intervention. The intervention offered selected small-scale farming households access to interest-free maize loans to be repaid in kind after harvest, approximately 3 months after the last loan installment was received. The maximum loan size was 3 bags of ground maize, a quantity large enough to cover the consumption needs of an average household over a three-month period.[5] Treatment intensity was varied at the village level, with a smaller fraction of households receiving loans in partial (as opposed to full) treatment villages.

Both the demand for and the willingness to repay loans was high, with over 90 percent take up among eligible households and close to 95 percent repayment. To assess the impact of loans on local labor markets and labor allocation decisions, we develop a series of predictions through a simple agricultural labor market model, and test them empirically using our experimental data. Consistent with the model, we find relatively large effects of the loans on off-farm labor supply. In villages where all farmers were eligible for the maize loan, the probability that a household reported off-farm labor over a two week period during the hungry season fell by around 25 percent (11.8 percentage points). In the same villages, the total number of days of *ganyu* reported fell by around one-third, with a corresponding (but statistically imprecise) increase in the number of days worked on-farm. We also test whether the reduction in off-farm labor during the hungry season is offset by an increase in off-farm

[4]Seasonality in food consumption may reflect an optimal response to seasonal price variation. As shown in Appendix figure A.1, seasonal fluctuations in maize prices are less pronounced in Zambia than in other nearby countries (see, for example, Burke (2014)).

[5]Households could receive up to 75 kilograms (in 25 kg bags) of ground "breakfast meal" maize, which they had to repay with 150 kilograms harvested (ungrounded) maize. A 50 kilogram bag of unground maize yields approximately 25 kilograms of ground maize; the market value of the ground maize received and the harvested maize paid back is very similar. Households in the program therefore receive a transfer equal to the value of the interest on the loan. In Section 3, we calibrate the size of the transfer.

labor to generate resources to repay the loan. In the two months preceding loan repayment, reported off-farm casual labor was significantly lower in treated villages, and the overall probability of doing any casual labor at any point during the entire agricultural season was 9.1 percentage points lower in the fully treated villages than in the control villages.

Consistent with the observed reductions in farmers' willingness (or need) to engage in off-farm labor, we find sizable increases in average wages paid locally for off-farm labor. In fully treated villages, the average wage reported among (the smaller number of) workers engaging in *ganyu* increased by almost 50 percent immediately following the loan transfers and by 25 to 35 percent over the treatment period. The wage effects were smaller but remain positive in partial treatment villages, where only 50 percent of farmers were eligible for the loan program. We also find substantial declines in reported hunger in treated villages. Specifically, the intention-to-treat (ITT) estimates show that in villages where all households were eligible for the loan, the number of meals consumed by the respondent in the 24 hours preceding the survey increased by 11 percent (0.2 more meals) and household members were almost 50 percent (16.5 percentage points) less likely to miss meals due to food shortages in the week before the survey.

For both the consumption and the labor allocation results, the observed effect sizes are smaller in partial treatment villages, with relatively large - but imprecisely estimated - spillovers for households whose neighbors receive loans. All of the data collected suggests that these spillovers are not the result of within community sharing or redistribution of resources: we observe no effect of treatment on reported gifts, loans or transfers among eligible or ineligible households. Given the large wage response, some positive spillovers through the labor market seem likely, since higher local wages allow ineligible households to meet their consumption needs with fewer hours of off-farm labor. The relatively small sample size of the study does not allow us to precisely pin down the magnitude of the labor supply elasticity. Larger studies in this area would help to isolate the channels underlying the observed spillover and general equilibrium effects.

The results presented in this paper are linked to a large literature in development economics documenting the often complex and seasonal nature of rural labor supply (for a review, see Rosenzweig 1988). A number of previous studies have investigated households' use of off-farm labor to smooth consumption in response to, or anticipation of, productivity shocks (Kochar 1995, 1999; Rose 2001; Ito and Kurosaki 2009), suggesting that productivity shocks may alter labor supply and affect rural wages. Consistent with the findings in this paper, Jayachandran (2006) shows that incomplete credit markets lower the elasticity of labor supply, further increasing wage volatility. Mobarak and Rosenzweig (2014) show that wage volatility is reduced and sizable welfare improvements are possible when workers can be insured against aggregate shocks. More recently, a growing number of field experiments examine labor supply determinants in developing country settings, some of which shed light on questions of seasonality and consumption constraints. For example, Goldberg (2013) finds very low labor supply elasticities following the harvest period in Malawi. Also in Malawi, Guiteras and Jack (2014) examine labor supply during the hungry season and after harvest and note the seasonal coincidence of high labor demand and liquidity constraints.[6]

The results in this paper are also linked to a larger literature on consumption smoothing. In the face of missing credit markets, households may resort to a range of consumption smoothing strategies, ranging from livestock and asset sales to migration (for summaries of the literature, see Morduch (1995) and Besley (1995)). The substantial increase in consumption that we observe among treated households during the hungry season suggests that farming households in the study adjust to resource scarcity not only through off-farm labor, but also by reducing consumption during the growing season, where caloric needs are likely highest. These reductions in calorie intake during the peak labor season are also consistent with a related nutrition-focused literature which suggests that farmers' constrained nutritional intake leads to suboptimal production (Pitt and Rosenzweig 1986; Strauss 1986; Behrman et al. 1997), an idea also supported by recent evidence from India (Schofield 2013).

The general relationship between income seasonality and consumption smoothing is not well established in the literature. While some studies suggest that precautionary savings are sufficient to smooth consumption even if income is highly seasonal (Paxson 1993; Chaudhuri and Paxson 2002), others have highlighted the pronounced consumption differences between lean or hungry season (Dercon and Krishnan 2000; Khandker 2012), consistent with the patterns shown in Figure 2.

Few studies have examined the impact of explicitly seasonal transfer or loan programs. Burke (2014) offered farmers in Kenya a loan product that allowed them to exploit seasonal variation in maize prices and finds significant effects on total maize revenues and household expenditures. Bryan, Chowdhury and Mobarak (2013) provide credit and grants for short run seasonal labor migration in Bangladesh and argue that credit market failures and highly uncertain returns likely keep long-distance labor supply below welfare-maximizing levels.

[6]Inefficient labor supply does not, however, imply credit or liquidity constraints. In a nonagricultural setting, Dupas and Robinson (2014) rule out consumption constraints as a driver of decisions in their study of reference dependent labor supply among bicycle taxis in Kenya.

Most similar to our study, Basu and Wong (2012) study a seasonal food credit program in Indonesia and find that food loans increase consumption during the lean season, but do not analyze labor supply impacts. Our findings contribute to that literature by providing the first direct evidence that seasonal food shortages drive off-farm labor supply and local wages. We also show that relaxing credit constraints leads to an adjustment of household labor supply and an equilibrium increase in local wages. Both the labor supply and the wage results are the product of multiple different margins adjusting as credit markets are relaxed, including better nutrition, a lower number of credit-constrained households and differential selection into the off-farm labor market. In terms of final labor allocation, the relaxation of credit constraints reallocated labor from relatively better-off farms able to hire *ganyu* to smaller farms generally dependent on doing *ganyu* as a coping mechanism.

The paper proceeds as follows. In the next section, we present a simple model of agricultural labor decisions in the presence of credit constraints. Section 3 describes the study context, experimental design and implementation. Section 4 sets up the identification strategy. Section 5 presents the results, and Section 6 concludes.

2 Conceptual framework

Consider a simple two-period model of agricultural production, where rational farming households maximize their utility over consumption and leisure. In the first period (the growing season), households make their labor and consumption decision. In the second period (the harvesting season), households collect the harvest outcomes and get utility from consuming the harvest outcomes. Land endowments are fixed. There are two types of farms: those with relatively large capital endowments that are net buyers of labor, and small-scale farms, which divide their labor between their own farm and off-farm work at a locally determined wage rate w. For ease of exposition, we refer to the former as large and the latter small, though land size need not differ between the farm types. Large farms, indexed by b, choose their labor inputs optimally to maximize profits such that

$$Y_b^{'}(l_b^*) = w \tag{1}$$

where $Y_b^{'}(l)$ is the marginal product of labor on the large farm, with $Y^{'}(l) > 0$ and $Y^{''}(l) < 0$, and w is the local wage rate. Small-scale farmers have their own land, and maximize their utility $u(.)$, which is an increasing and concave function of consumption c and leisure time

f. We assume that small farmers have a fixed time endowment T, which they can allocate to selling labor off-farm to the large farms (l_o), to work on their own farm (l_s), or for leisure f, so that the time budget constraint is given by

$$l_o + l_s + f \leq T. \tag{2}$$

In the first period, crops are not available yet, so that the farm's consumption constraint is given by

$$c_1 \leq l_o w + k \tag{3}$$

where $l_o w$ is off-farm labor income of households and $k \geq 0$ are savings or other reserves that the farm can consume. Small farms maximize their total utility over the two periods, which is given by

$$u(c_1, f_1) + \rho u(c_2, f_2), \tag{4}$$

where $0 < \rho \leq 1$ is the discount rate applied to second period utility.

Unconstrained Equilibrium

In the absence of credit constraints, small farmers choose their on-farm labor inputs such that the marginal product of on-farm inputs equals the market wage

$$Y_s^{'}(l_s^*) = rw, \tag{5}$$

where $Y_s^{'}$ is the marginal product (harvest) generated by the labor input, and r is the market interest rate paid on capital (k). Given that utility is strictly increasing in consumption, the optimal amount of labor provided off-farm is $l_o^* = L - l_s^* - f^*$.

The equilibrium wage w^* will simply be determined by the intersection of the aggregate labor demand and supply function, i.e. w^* is such that

$$\int l_o^*(w^*)db = \int l_s^*(w^*)ds, \tag{6}$$

where the left-hand side of equation (6) integrates over all large farms b and the right hand side integrates over all small farms s.

Constrained Equilibrium

Let us define a farm as credit constrained if the farm is not able to cover optimal consumption in the unconstrained equilibrium with the first-best (unconstrained) labor allocation choice, i.e.

$$c^* > l_o^* w + k \tag{7}$$

Then, we can state the following propositions:

*Proposition 1*a: The total amount of off-farm labor provided in equilibrium increases with the fraction of small farms that are credit constrained.
*Proposition 1*b: Average consumption levels decrease with the fraction of small farms that are credit constrained.
*Proposition 1*c: The wages observed in equilibrium decline with the fraction of small farms that are credit constrained.

The intuition behind the propositions is straightforward. In the absence of financing through savings or borrowing, consumption in the lean season can only be financed through an increase in *ganyu* labor supply ($l_o > l_o^*$). This increase in *ganyu* (beyond the optimal point) mechanically leads to a decrease in leisure and/or a decrease in consumption, because the total resources available are lower than in the unconstrained equilibrium. This implies a shift in the supply curve to the right; given that the demand curve will stay the same, a higher fraction of credit constrained farmers will lower equilibrium wages and increase the total amount of *ganyu* labor hired.

The stylized model has several implications in terms of increases in credit availability. First, for any credit constrained farmer, as k increases, l_o falls as long as the interest rate on borrowing is sufficiently low, which means that farmers will shift their labor allocation from off-farm to on-farm labor. Second, increasing credit availability will increase consumption. As farmers get closer to their optimal on-farm labor allocation l_s^*, they will be able to increase the net present value of available resources, and thus to increase consumption (and leisure) in both periods. Last, according to Proposition 1c, an increase in credit availability will increase wages.

This stylized model clearly abstracts from a variety of other relevant mechanisms, including the effect of relaxing consumption constraints on labor productivity (Strauss 1986; Pitt et al. 1990; Subramanian and Deaton 1996; Strauss and Thomas 1998). These additional margins would add complexity without substantially affecting the predictions outlined above, which we explore in our empirical analysis. Our model derives equilibrium predictions at the market level, which, in a friction-free labor market, should be observable at the regional or national level. In the local study context, however, mobility is severely limited by poor road infrastructure and a lack of transportation options. As we will show in the empirical analysis below, most *ganyu* labor is provided locally; accordingly, we consider each village as a separate labor market.

3 Context and experimental design

The study was implemented over the course of a year in Chipata District in Eastern Zambia. Chipata District is located at the southeastern border of Zambia, with an estimated population of 456,000 in 2010 (CSO 2010). Approximately 100,000 people live in Chipata town, the district capital; the remaining population lives in rural areas, with small-scale farming as primary source of income. According to the 2010 Living Conditions and Measurement Survey, rural households in Chipata are on average poorer than in the rest of the country, with 47 percent of household classified as "very poor" in the district overall, and 63 percent of households classified as very poor in the rural parts of Chipata. Average monthly expenditure of rural households is about one third of the national average, and access to electricity and piped water close to zero in rural areas (see Appendix table A.1 for a summary of differences between Chipata and the rest of Zambia).

3.1 Study sample

The study sample is drawn from a complete listing of villages in Chipata District obtained from the Zambian Ministry of Agriculture's registry of farmers for 2011. Using the Zambian Ministry of Agriculture definition, we classified farms as small-scale if they farmed on fewer than five hectares of land.[7] Sampling was restricted to communities with 15-25 small-scale farmers; we excluded peri-urban villages as well as a few villages which could not be reached with one day's drive from Chipata town. All remaining eligible villages were visited by study enumerators who conducted a listing of farming households. Based on the listing visits,

[7]We restricted our sample to households with at least one hectare of land to distinguish households with very small scale home gardens from households engaged in crop production, and also to increase the likelihood of sufficient harvest to repay the loan.

further villages were eliminated if they could not be reached during the rainy season or if the number of households or the size of landholdings contradicted the Ministry's records.[8]

3.2 Experimental design and implementation

We study the effects of a short run loan on household labor allocation. The loan was offered to households in January 2013. Under the loan program, eligible households could obtain one 25 kilogram bag of ground maize flour in each of three months during the hungry season (January, February and March), which they were instructed to repay at harvest (June) in the form of (unground) maize. The typical sales price in Zambian Kwacha (KR) of one 25kg bag of ground maize in February was KR 45 (USD 9), while the price of one 50kg bag of unground maize was between KR 50 (USD 10) and KR 65 (USD 13) in June depending on whether maize was sold locally or sold to the government.[9] This means that the implicit interest rate over the 4-6 month loan period was between 11 and 44 percent, which translates to an annual interest rate between 27 and 107 percent. At the same time, the loan saved the farmers the cost of transporting maize to the miller and the cost of grinding, and the cost of transporting harvested maize for sale.

We varied the treatment intensity of the loan program at the village level. In 25 percent of the 40 study villages, all eligible households were offered the chance to take up the loan. In half of the study villages, households that expressed interest in the loan entered in a lottery, through which half were selected for the loan program, resulting in approximately equal numbers of households with access to the loan program in full and in partial treatment villages. The final 25 percent of villages served as the control and received no intervention.

The loan offer and within-village lottery were implemented through a village meeting to which all small-scale farming households were invited in advance. At the meeting, study enumerators explained the loan to the assembled participants. Households were given loan forms and told to return them to their village head person the following day. If the village was assigned to the partial eligibility treatment, households attending the meeting were

[8]By dropping households more than one day's drive from the main town and those inaccessible during the rainy season, we potentially eliminate some of the more vulnerable households. While this potentially detracts from the external validity of the results, randomization is conditional on eligibility so sample selection is orthogonal to treatment assignment.

[9]Zambia has a heavily controlled maize sector, with government intervention through a fixed (above local rate) price for harvested maize (KR 65 in 2013) and a price ceiling on ground maize at the retail level (KR 45 in 2013). Though maize prices tend to be lower at harvest, the government intervention limits the seasonal fluctuations in maize prices that are observed in neighboring countries, though year-to-year variation remains substantial (see Appendix figure A.1).

entered into a lottery during which names of participants were drawn from a basket by a child unaffiliated with the study. Half of the names were drawn, resulting in 50 percent treatment relative to the full eligibility treatment arm.[10] Households drawn in the lottery made their participation decisions by returning their loan form to the head person, as in the full eligibility arm.

Maize flour was delivered to ten different distribution centers to cover the 30 treatment villages during the second half of each of January, February and March. Villages were informed in advance about the time and place of the maize flour delivery and households that had turned in their loan forms to their head person were eligible to borrow during each of the three delivery dates. Because households might have difficulty transporting maize flour, some flexibility was allowed in the pick up procedures. Specifically, the national registration (NRC) card for the eligible borrower was needed to obtain the flour, but the borrower did not have to come in person.[11] In late June and early July, loan repayment was due. Village heads were involved in reminding households to bring their maize for repayment to a central point in the village, and were rewarded with a bag of maize for full repayment. Unlike the distribution of maize flour, collection of maize was done at each village. In a couple of cases, loans were repaid late. Further summary statistics on repayment patterns are described in Section 5.

Households were sampled and village-level treatments were assigned using min-max T randomization (Bruhn and McKenzie 2009). The approach relies on repeated village-level assignment to treatment and selects the draw that results in the smallest maximum t-statistic for any pairwise comparison across treatment arms. Balance was tested for 13 village and household level variables, as well as geographic blocks, with results described in Section 4.

3.3 Data

We rely on both survey and administrative data in our analysis. Survey data comes from three sources: (i) a baseline survey conducted before the start of the agricultural season in October 2012, (ii) a midline survey conducted between the second and third maize delivery in February and March 2013, and (iii) an endline survey conducted after harvest but before repayment in June 2013. We collect data on a host of household and farm character-

[10]If an odd number of farmers participated in the lottery, the number of eligible farmers was rounded up, i.e. $n/2 + 0.5$ farmers selected into the program.

[11]In a small number of cases, households were allowed to use a note from the village head person in lieu of a missing NRC card. While no instances of fraud were reported, any fraud in the loan collection process will bias us against seeing a result.

istics at baseline and endline, and focus the midline survey on labor allocation in the two weeks preceding the survey. An effort was made to keep data collection independent of the intervention, to minimize the role of experimenter demand in influencing survey responses.

3.4 Local credit and labor markets

The conceptual framework builds on several contextual features, namely local labor and credit markets. We provide additional qualitative background on these features of the study setting. As described above, the study sample was limited to small farmers – those with land holdings of 5 hectares (12 acres) or less. The attribute of "small-scale" is somewhat misleading since it suggests that these farmers are unusually small; in fact, small-scale farmers represent the overwhelming majority of households in rural villages in Zambia. In our study villages, over 90 percent of listed households fall into this category.

Credit markets In terms of borrowing opportunities, the study setting is also fairly representative of rural areas in developing countries, where credit markets are typically incomplete. In the baseline survey, 4.5 percent of household respondents report accessing formal loans for cash. Informal borrowing channels are slightly more common: 8.6 percent report taking high interest loans, locally referred to as *kaloba*, with interest rates over 100 percent. Loans between friends and family are most common, with 13.2 percent of baseline respondents reporting borrowing food or cash in the previous year. Some in-kind input loans are available from agricultural companies in the area.

Ganyu labor Local wage earning opportunities for study households are defined largely by casual labor or *ganyu*. In focus groups, small-scale farmers in Chipata described *ganyu* labor as a coping strategy that households would prefer to avoid. Figure 3 shows the age and gender profile of *ganyu* among adults at baseline; on average, men are more likely to engage in *ganyu* than women; the highest *ganyu* rates at baseline were observed for men 25-34, 60 percent of whom reported *ganyu* activity in the prior season. In the baseline survey, the most common response to why an individual in the household worked *ganyu* during the previous agricultural season was hunger. The second most common reason was to access cash for a personal purchase, and the third was to deal with an emergency. Households appear only moderately accurate in their forecast of whether they will have to engage in *ganyu* in a given year. Among control group households that predicted at baseline that they would have to do *ganyu* in the coming year, around 70 percent did; among those that predicted not doing *ganyu*, around 40 percent ended up working off-farm.

The model that we present in Section 2 simplifies a complex rural labor market. In the study setting, road infrastructure is extremely bad, there is no motorized public transport and distances between villages are substantial. Most casual labor takes place in or near the worker's own village. At baseline, around half of individuals who did *ganyu* worked at least some of the time for other farms in the same village. While the conceptual framework distinguishes between larger farms that are buyers of *ganyu* and small farms that are sellers, the distinction is not always as clear in practice. Over 50 percent of respondents report working on other small farms (less than 12 acres) at baseline, and about one third of households selling *ganyu* also report occasional *ganyu* hiring. None of these patterns are inconsistent with our model, which predicts that the (net) off-farm labor supply can be positive or negative, depending on local wages and the relative marginal product of labor; they do, however mean that the boundaries between *ganyu* buyers and sellers are more fluid, and the same farm may sell *ganyu* at one point in the year, and purchase labor at another when more resources are available.

Wages are also seasonal, with highest wages reported during planting (October to early December) and at harvest (May to June), as shown in Appendix figure A.2. In the absence of private vehicles as well as public transportation, wages are mostly determined locally, which either means within a village, or within a small group of villages. *Ganyu* wage rates are typically negotiated on a case-by-case basis, and anecdotally are highly responsive both to demand and supply shocks. In the hungry season (when resources are most constrained), wages are likely to be suppressed both by increased supply from credit-constrained farms and by reduced local demand (inability to hire). As a result, farms with sufficient resources to hire *ganyu* may be able to hire labor at rates below the marginal product of *ganyu* on their land; as such, *ganyu* contracts may constitute a within-village or within-community transfer from smaller or more constrained to larger or less constrained households. Relaxing credit constraints therefore should lead to a reallocation in the opposite direction.

4 Identification

To evaluate the effect of the maize loan we estimate:

$$y_{ivt} = \alpha + \beta_1 full_v + \beta_2 partial_v + X_{iv}\phi + \delta_t + u_{ivt} \tag{8}$$

where y_{ivt} is an outcome of interest for household i located in village v and month t, $full_v$ indicates that the village was assigned to the full eligibility treatment, $partial_v$ indicates that the village was assigned to the partial eligibility treatment, X_{iv} is a vector of controls measured at baseline and δ_t are month dummies to capture seasonal effects. Errors are clustered at the level of the randomization unit, the village v. We include all households in our estimation of equation (8), regardless of whether the household took up the loan or loan size. The coefficient β_1 therefore captures the effect of being a small-scale farmer in a village assigned to the full eligibility treatment, for both takers and non-takers. The coefficient β_2 captures the average effect of being in a partial eligibility village, regardless of lottery outcome and take up decision. In much of the analysis, we pool across months or analyze self-reported behavior in the month of the midline survey only.

Since a household's eligibility in the partial treatment villages depends on the lottery outcome, which was randomly assigned, we can also estimate the reduced form effect of being eligible for a loan on outcomes of interest.[12] We estimate:

$$y_{iv} = \alpha + \delta eligible_i + \phi X_{iv} + u_{iv} \tag{9}$$

where the coefficient δ represents the average effect of eligibility for the loan across treatment intensities. This specification may be invalid in partial eligibility villages if there are spillovers across households. To empirically test the relevance of spillovers, both as an outcome of interest and as a test of the identifying assumptions for equation (9), we interact the $eligible_i$ indicator variable with an indicator for being in a partial treatment village.

4.1 Balance and summary statistics

The coefficients presented in the subsequent analysis identify the causal effect of the loan under the identifying assumption that treatment assignment is orthogonal to u_{iv}. Table 1 presents the means and standard deviations of baseline survey characteristics among study households. Column 1 shows averages for control villages and columns 2 and 3 report the coefficients from a regression of each household characteristic on treatment dummies, with standard errors clustered at the village level. The significance of the coefficient represents the test of the characteristic mean in each treatment against that in the control, while column

[12] For consistency, we define eligibility in both the full and the partial treatment villages conditional on attending the informational meeting.

4 shows the p-values from a test that the treatment coefficients are equal.

A few of the summary statistics in Table 1 are worth noting. Household heads are around 46 years old and just under 25 percent are female. Households contain around 5.5 members, with around one child under 5. In the previous season, households produced 1.3 acres of local maize, 0.97 acres of hybrid maize and 1.7 acres of cotton on average, and used a total of 4.8 acres for agricultural production. Households on average engaged in 12 days of casual labor (*ganyu*) in the last season and hired around 4 days of *ganyu*. The average household stored 25 bags of maize after the previous harvest and had a baseline harvest value of around USD 484. With an average household size of more than 5 members and very limited income beyond agriculture, this means that virtually all households in our sample were well below the USD 1.25 per day per capita poverty threshold. In general, household characteristics are well balanced, with 5 out of 39 pairwise tests significant at $p < 0.10$, which is just above what we would expect from normal sampling with a 10 percent significance threshold. Most notably for our analysis, households in the full treatment villages engaged in marginally significantly fewer days of casual labor in the previous season than did households in the partial treatment villages. Given the relatively small sample size and small imbalances, we control for a full set of baseline characteristics throughout our analysis.

4.2 Attrition and selection

The main identifying assumption of our empirical analysis could be violated if households select into eligibility status or drop out of the loan program differentially across treatments. Households could exit the study both during the midline survey and the endline survey. Overall attrition rates are low as shown in Table 2. The table reports means and standard deviations at baseline in column 1 and the coefficients from a univariate regression of the binary attrition or selection outcome on each explanatory variable with standard errors clustered at the village level. Each coefficient therefore reflects a separate regression. A total of 439 households were enrolled into the study at baseline, 98 percent of whom were in the midline and endline surveys. The probability of being in the midline and endline surveys is not affected by treatment status (columns 2 and 4).

Among households assigned to treatment, attrition could also occur at the stage of the loan meeting, which took place in each partial and full treatment village. Column 3 of Table 2 shows the probability of of attending the loan meeting as a function of each household characteristic (again, each coefficient is its own regression), conditional on residing in a treatment village and being in the baseline sample. Among eligible households, the probability of attending the loan meeting is increasing in the age of the household head and land size, but is not significantly affected by treatment intensity.

5 Results

We begin with descriptive statistics associated with take up and repayment of the loan, and a general manipulation test for changes in consumption, then move on to the analysis of impacts on labor supply. Then we investigate the equilibrium effect on wages, as described in the conceptual framework (Section 2). We conclude our analysis with suggestive results on within-village treatment spillovers.

5.1 Take up and repayment

Table 3 summarizes take up rates by treatment intensity. In the partial treatment villages, a total of 221 households were surveyed at baseline. 208 households participated in the information session, and 104 households were selected for participation in the loan program through the public lottery. All of the selected households decided to enroll in the loan program. Out of 104 households in the fully treated villages, 99 households came to the loan information session, and 90 households (91 percent) decided to sign up for the loan program. On average, households picked up 2.79 bags of maize flour in the partial treatment group (among eligible households), and 2.62 bags of maize flour in the full treatment group. The high take up rate across treatments means that, in practice, our ITT estimates are close approximations of the average treatment effect on the treated. Though take up rates vary slightly with treatment intensity, they are not predicted by any of the household characteristics presented in Table 1 and the intention to treat estimation strategy mitigates remaining concerns about selection into the loan.[13]

The bottom of Table 3 shows repayment rates, which were very high in both groups. Ninety-eight percent of participants repaid at least a part of the loan, and 95.2 and 97.7 percent fully repaid their loans in the partial and full treatment villages, respectively. No differences in repayment rates were found with respect to household size or financial resources (results not shown).

[13]The lottery itself may have affected take up rates by increasing the perceived value of the program by introducing scarcity.

5.2 Consumption

As a first step toward identifying the impact of the loan program, we analyze food consumption. As described above, we conducted a midline survey in the middle of the hungry season, during which detailed information on food intake was collected. Maize is the most commonly grown crop in Chipata District, and also the staple food in Zambia; well-off households consume maize in the form of a porridge (*nshima*) three times per day. To assess changes in nutrition, we examine two summary measures of food consumption: the number of times the respondent reported eating *nshima* during the 24 hours preceding the interview, and whether any family member had to skip a meal in the week preceding the interview because of food shortages. In the control group, the average respondent consumed *nshima* 1.74 times in the 24 hours preceding the interview, and 36 percent report that a family member had to miss a meal.

Table 4 shows the corresponding treatment effects. In full treatment villages, respondents consumed 0.2 more meals, an increase of 11 percent in eating frequency over the control group. The increase in the partial treatment villages was smaller and not statistically different from consumption in the control group. When we analyze the impact of loan eligibility at the household level in column 2, the overall effects appear smaller than the full village treatment effect, but larger than the shifts observed in the partial treatment villages. We find similar impacts for the missed meals among family members. The probability of missing a meal was 16.5 percentage points lower in full treatment villages (a 45 percent reduction relative to the control group); the average treatment effect was 6 percentage points in the partial treatment villages, but not statistically different from the control villages. The individual ITT is once again in between the partial and full treatment village impact.

5.3 Labor supply

We begin our analysis of labor supply effects with the midline survey. The main advantage of the midline survey is that recall was restricted to a two-week period preceding the survey (limiting recall issues), and that the timing of the survey (conducted in the last week of February and first week of March 2013) allows us to directly investigate labor shifts in the middle of the hungry season, i.e. the most constrained part of the agricultural season.

Panel A of Table 5 shows the village level intention-to-treat results for household labor supply at midline. In column (1) we analyze a binary indicator for whether the household engaged in any off-farm *ganyu* labor activities. In the two-week period preceding the survey, 46 percent of households in the control group indicated that they provided *ganyu* labor on other farms.[14] In full treatment villages, this proportion was reduced by 11.8 percentage points (25 percent), while the reduction in partial treatment villages was 7.6 percentage points (16.5 percent) and imprecisely estimated. Days of *ganyu* also decreased following a similar pattern, though the treatment effect in the partial treatment villages is close to zero. While we find positive point estimates for labor supply on the the household farm in both groups (column 3), the confidence intervals are large and the effects are not statistically significant. Finally, we observe no detectable change in hiring of outside labor (column 4).

In Panel B, we show the intention-to-treat results at the household level. The results look similar to Panel A: the estimated coefficient on *ganyu* provision (column 1) suggests an average reduction in *ganyu* labor of 8 percentage points. Since we now also include non-eligible households in partial treatment areas (who may have benefited from the intervention) in the comparison category, the slightly smaller coefficient could in theory be evidence for spillover effects in these villages. At the household level, neither days of *ganyu* nor days on farm (columns 2 and 3) are precisely estimated.

The main advantage of the midline data is its relatively short recall period with a presumably more limited scope for recall bias. The longer recall period at endline allows us to investigate whether households compensated for lower rates of *ganyu* during the hungry season with an increase before loan repayments were due. Table 6 shows results from the endline data, covering the full agricultural calendar, that are qualitatively similar to the midline results in Table 5. The probability of any off-farm labor at any point during the agricultural season falls by 9 percentage points in the fully treated villages relative to the control. Days of *ganyu* fall, days of on-farm labor increase and days of hired labor increase in full treatment villages, all with coefficients that are large in magnitude (39 percent, 7 percent and 46 percent, respectively) but imprecisely estimated. Panel B shows the individual level ITT results, which is consistent with the midline outcomes. Because these measures also include the months preceding the loan offer, we also repeat the analysis in Appendix table A.2, restricting outcomes to the months during which households had loans. This narrower analysis requires that households accurately recall the timing of their labor decisions; per-

[14] The share of households in the control villages who report doing *ganyu* in the two weeks prior to the midline survey is substantially higher than the average *ganyu* rates reported for the previous February during the baseline survey. This can be attributed to two factors. First, the midline and baseline measures refer to different years, and ganyu rates may vary substantially from year to year. Second, the short recall period in the midline survey may allow respondents to remember *ganyu* activities that they had forgotten in the intervening 8 months when asked at baseline about the previous February.

haps due to measurement error, estimates are consistent with but less precise than in Table 6.

In our final set of results on labor supply, we focus on the period immediately preceding loan repayment and the endline survey. Column 4 of Table 6 shows the probability that a household did any *ganyu* in the two months preceding loan repayment. Because the recall period in these questions is only two months, we expect the reports to be reasonably accurate. On average, in control village households, 29 percent of respondents report some *ganyu* in their household in the two months prior to loan repayment. In full treatment villages, that drops by 13.8 percentage points (47.6 percent). The average effect on eligible households is a decrease of 8 percentage points.[15]

Additional analysis by gender using endline data suggests that treatment effects are larger for female household members who have a lower baseline rate of *ganyu* (recall Figure 3). In full treatment villages, the probability that a household supplies female *ganyu* fell by 6.8 percentage points (p-value of 0.052) or by 29.3 percent, while the probability that a household supplies male *ganyu* fell by 3.7 percentage points (p-value of 0.50) or 7.8 percent. Note that these estimates assume equal probabilities of adult males and females able to supply labor across treatment and control households.

5.4 Wages

As outlined in our theoretical framework, individual (or community) changes in *ganyu* labor participation should have only marginal effects on equilibrium wages if labor markets were national or frictionless. As we have argued and shown above, workers' ability to engage in short-term labor opportunities beyond walking distance is somewhat limited, which means that large differentials between local labor market outcomes are likely.

To evaluate the wage impact of the loan programs, we start by analyzing wages reported by households in the midline survey. As part of the midline survey, households were first asked to report the total number of person-days of *ganyu* provided in the 2 weeks prior to the survey, and then asked to report total *ganyu* income for the household in total, as well the days of *ganyu* worked and income earned by each household member. We start

[15]The lack of reported off-farm labor in the treatment villages preceding repayment are unlikely due to a lack of off-farm work opportunities. As shown in Appendix figure A.2, wages are high at harvest relative to other times of the year. Labor supply is scarce, however, because most households are engaged on their own land. At baseline, households reported a total of 64 person-days of *ganyu* during the harvest. For comparison, 232 person-days of *ganyu* were reported during weeding, 57 during planting and 199 during field preparation.

our analysis by looking at the average wage reported at the household level: the total *ganyu* income divided by the total number of person-days reported. Two things are worth highlighting. First, the duration of a typical *ganyu* work day varies substantially; the wage number calculated should thus be interpreted as the wage income a worker receives for a typical day of *ganyu*. Second, a substantial fraction of households do not do any *ganyu*, for which this variable is not defined; wages thus do not represent wage offers, but only wage offers accepted.

Panel A of Table 7 shows that average wage rates reported at midline were around KR 9.65, which corresponds to slightly less than USD 2 at 2013 exchange rates. In column 1 we regress reported wages on the two treatment arms; in column 2 we use the fraction of farmers eligible for the program at the village level. Independent of the chosen specification, the magnitude of the observed effect is large, suggesting an increase of local wages of up to 49 percent. The analysis is repeated in columns 3 and 4, using reported wages (total *ganyu* income divided by total days worked) for each individual in the household who did *ganyu*. The individual-level wage is lower, on average, and the treatment effects are less precisely estimated. We show the full distribution of reported wages at midline, by treatment arm, in Appendix figure A.3.

In the baseline and endline surveys, households were also asked to indicate the total number of person-days worked off-farm over the agricultural season. For this longer-term recall module, households were asked to separately report off-farm labor activities for each of the four main parts of the agricultural season: field preparation, planting, weeding, and harvesting. Respondents indicated the total number of (full-time) person-days for each activity, as well as the total remuneration earned by the household, including food as well as other in-kind services. As Panel B of Table 7 shows, the average wage of KR 18 (USD 3.5) was substantially higher at endline than at midline. The higher rates observed in the endline survey reflect seasonal variation in wages, as well as potentially longer working hours.

In Panel B of Table 7, we show the results of a basic household-level panel model, where each observation corresponds to a wage reported in a given season (2011/12 in baseline, 2012/12 in endline) and for a given activity. A total of 1025 *ganyu* wages were reported across the two seasons. In Panel B of Table 7, we show the results from a basic difference-in-difference estimator with village fixed effects, where all activities in the year of the intervention are considered treated in villages with loan programs (columns 1 and 2). In columns 3 and 4, we report the results of an alternative specification, where we only consider weeding and harvesting related activities in the second year as treated. Given that the vast majority of *ganyu* activities fall in this category, the results from the two models are relatively similar. All of the estimated coefficients go in the expected (positive) direction, but the standard errors are large and coefficients are statistically insignificant. While the magnitudes of the wage treatment effects are (insignificantly) larger at endline than at midline, relative to the respective control group means, they are smaller (27-35 percent versus 49 percent). This difference in the relative shift is due in part to the higher control group average observed at endline, but also perhaps to a dissipation of the treatment effect. Finally, as in the labor analysis, we expect measurement error to be larger at endline than at midline.

The observed wage effects may be driven by adjustment on multiple margins. First, the number of households supplying labor to the market falls as a result of the intervention, as shown above. Second, the composition of individuals in the market also adjusts; in particular, as described in Section 5.3 above, female labor supply falls disproportionately. Finally, to the extent that the consumption gains shown above translate into better nutrition, productivity may also increase, conditional on working in the labor market. We cannot disentangle these different channels, but the presence of all three helps explain the magnitudes of the wage effects that we observe. In addition, our focus on small villages means that a relatively large fraction of the total village population is treated, producing a substantial shock to aggregate labor supply.

5.5 Within-village spillovers

The results shown in Sections 5.2 and 5.3 suggest that ITT effects are smaller in partial treatment villages than in the full treatment villages. To investigate the spillovers associated with the loan program, we separately estimate treatment effects for eligible and ineligible households in partial treatment villages. We investigate three possible spillover channels: first, eligible households may share some of the additional resources by transferring food to ineligible households or family members as gifts or on a loan basis. Second, eligible households may hire more ineligible households to do *ganyu*. Third, general equilibrium wage effects may impact labor allocation decisions of ineligible households in partial treatment villages, as described in Section 2.

Figure 4 and Appendix table A.3 show spillover effects on labor and consumption outcomes and on inter-household transfers, respectively. First, transfers from eligible to ineligible households did not increase (Appendix table A.3). Interestingly, the likelihood of households reporting incoming gifts and loans is substantially higher among eligible households, which is likely a reflection of eligible households reporting the maize loan program under this category. Second, *ganyu* by ineligible households did not increase (Panels c and d of Figure 4), which is consistent with the observations that eligible households in partial treatment villages did not increase hiring (not shown). While we lack the precision to detect spillover effects at conventional levels of significance, Figure 4 shows that ineligible households experience outcomes that move toward the treatment effects and away from the controls for consumption and labor impact. The graphs plot marginal effects from the interaction of household-level eligibility and village-level treatment interactions. A comparison across ineligible households in pure control and partial villages shows the spillover effect on ineligible households of being in a village where some households receive loans. A comparison across eligible households in partial and full treatment households shows the spillover impact on eligible households. In all panels of the figure, the effects are statistically indistinguishable for eligible and ineligible in partial treatment villages. Conceptually, and based on the model developed in Section 2, these results are consistent with income effects: for a given amount of off farm labor, higher wages relax the consumption constraint of even ineligible households. Given the very small number of clusters, all of these estimates are imprecise; larger studies will be needed to more closely identify local labor market spillovers.

6 Discussion and Conclusions

We implemented a simple and potentially scalable maize loan program in rural Zambia to identify the effect of increased access to seasonal credit on off-farm labor supply. The results of the experiment suggest that additional access to short-term credit increases food consumption, decreases off-farm labor and increases local wages in targeted communities. Even though the intervention was relatively simple in its design, increased resource availability during the growing season can potentially affect labor market outcomes through multiple mechanisms. First, short-term credit lowers the degree to which consumption-constrained farmers depend on short-term labor income, reducing local labor supply from particularly resource-limited households. Second, improved nutrition potentially allows household members to work longer and harder (as in Strauss 1986), which may positively affect productivity on their own farm as well as for external labor activities. Both effects will be associated with increases in local wages, which will in return also change the composition of the *ganyu* workforce from individuals engaging in *ganyu* primarily to smooth consumption to individuals engaging in *ganyu* activities to maximize the household's net income.

Even though the credit program was designed to be broadly consistent with local credit market conditions, the extremely high uptake and repayment rates could be seen as evidence that the loan program was considered financially beneficial, and thus provided a net transfer to households. Given the high variability of local short-term lending rates reported, it is hard to quantify how large of a net transfer the loan program offered. With a maximum loan size of USD 27 and a market value of the repaid maize between USD 30 and USD 45 (depending on the buyer), large net transfers seem unlikely. At an annual interest rate of 100 percent, and assuming that the maize would have been sold cheaply in the absence of the program ($30), the estimated net transfer over the four-months loan period would be approximately USD 3, which corresponds to about one day of *ganyu* income. In comparison, the labor responses seem rather large. At the mean, our results suggest that households are forgoing around KR 72 (~ USD 14) of *ganyu* income as a result of their reallocation away from off-farm labor, which makes it seem unlikely that income effects would have driven the observed shifts in labor allocation. In our view, the empirical patterns observed are most consistent with loan programs easing consumption-constraints for the households with the smallest food or financial reserves. This fits well with the positive correlation between missed meals and *ganyu* labor at midline, in both treatment and control villages.

Overall, our findings suggest that credit constraints affect consumption choices made by farming households, and also distort farming households' labor allocation decisions. This implies that expanding seasonal credit to small farmers may allow them to optimize their labor allocation and to benefit from increased local wages. The rise in wages generated by increased access to credit would lead to a redistribution of income from net buyers to net sellers of *ganyu.*

From a policy perspective, short-term loans may potentially be able to improve aggregate productivity; our results suggest that repayment rates can be high enough to make such programs possible and sustainable in good years; in bad years - with large aggregate weather or pest shocks - large drops in repayment rates seem likely; the addition of higher level insurance mechanisms may be necessary to ensure the financial stability and viability of such programs over time. Ongoing data collection at a substantially larger scale will allow us to examine the extent to which the reallocation of labor resources to farmers' own land can increase the productivity of small-scale farmers and will track both sides of the labor market to generate a more complete picture of the overall welfare implications of the intervention.

References

Basu, Karna and Maisy Wong, "Evaluating seasonal food security programs in East Indonesia," *Mimeo*, 2012.

Behrman, Jere R, Andrew D Foster, and Mark R Rosenzweig, "The dynamics of agricultural production and the calorie-income relationship: Evidence from Pakistan," *Journal of Econometrics*, 1997, *77* (1), 187–207.

Besley, Timothy, *Savings, credit and insurance* Handbook of Development Economics, North Holland, 1995.

Bruhn, Miriam and David McKenzie, "In Pursuit of Balance: Randomization in Practice in Development Field Experiments," *American Economic Journal: Applied Economics*, 2009, pp. 200–232.

Bryan, G., S. Chowdhury, and A.M. Mobarak, "Seasonal Migration and Risk Aversion," *Mimeo*, 2013.

Bryceson, Deborah Fahy, "Ganyu casual labour, famine and HIV/AIDS in rural Malawi: causality and casualty," *Journal of Modern African Studies*, 2006, *44* (2), 173.

Burke, Marshall, "Selling low and buying high: An arbitrage puzzle in Kenyan villages," *Mimeo*, 2014.

Chaudhuri, Shubham and Christina Paxson, "Smoothing consumption under income seasonality: Buffer stocks vs. credit markets," *Mimeo*, 2002.

Cole, Steven M and Parakh N Hoon, "Piecework (Ganyu) as an Indicator of Household Vulnerability in Rural Zambia," *Ecology of food and nutrition*, 2013, *52* (5), 407–426.

Dercon, Stefan and Pramila Krishnan, "Vulnerability, seasonality and poverty in Ethiopia," *The Journal of Development Studies*, 2000, *36* (6), 25–53.

Dupas, Pascaline and Jonathan Robinson, "The Daily Grind: Cash Needs, Labor Supply and Self-Control," *Mimeo*, 2014.

Fafchamps, Marcel, "Sequential labor decisions under uncertainty: An estimable household model of West-African farmers," *Econometrica*, 1993, *61* (5), 1173–1197.

Fink, Günther and Felix Masiye, "Assessing the impact of scaling-up bednet coverage through agricultural loan programmes: evidence from a cluster randomised controlled trial in Katete, Zambia," *Transactions of the Royal Society of Tropical Medicine and Hygiene*, 2012, *106* (11), 660–667.

Goldberg, Jessica, "Kwacha gonna do? Experimental evidence about labor supply in rural Malawi," *Mimeo*, 2013.

Guiteras, Raymond and B. Kelsey Jack, "Incentives, Productivity and Selection in Labor Markets: Evidence from Rural Malawi," *Mimeo*, 2014.

Ito, Takahiro and Takashi Kurosaki, "Weather risk, wages in kind, and the off-farm labor supply of agricultural households in a developing country," *American Journal of Agricultural Economics*, 2009, *91* (3), 697–710.

Jayachandran, Seema, "Selling labor low: Wage responses to productivity shocks in developing countries," *Journal of Political Economy*, 2006, *114* (3), 538–575.

Kerr, Rachel Bezner, "Informal Labor and Social Relations in Northern Malawi: The Theoretical Challenges and Implications of Ganyu Labor for Food Security," *Rural sociology*, 2005, *70* (2), 167–187.

Khandker, Shahidur R, "Seasonality of income and poverty in Bangladesh," *Journal of Development Economics*, 2012, *97* (2), 244–256.

Kochar, Anjini, "Explaining household vulnerability to idiosyncratic income shocks," *The American Economic Review*, 1995, pp. 159–164.

_ , "Smoothing consumption by smoothing income: Hours-of-work responses to idiosyncratic agricultural shocks in rural India," *Review of Economics and Statistics*, 1999, *81* (1), 50–61.

Michaelowa, Katharina, Ralitza Dimova, and Anke Weber, "Ganyu Labour in Malawi: Understanding Rural Household's Labour Supply Strategies," *Mimeo*, 2010.

Mobarak, Ahmed Mushfiq and Mark Rosenzweig, "Risk, Insurance and Wages in General Equilibrium," *Mimeo*, 2014.

Morduch, Jonathan, "Income smoothing and consumption smoothing," *Journal of Economic Perspectives*, 1995, *9*, 103–103.

Orr, Alastair, Blessings Mwale, and Donata Saiti-Chitsonga, "Exploring seasonal poverty traps: the 'six-week window' in southern Malawi," *The Journal of Development Studies*, 2009, *45* (2), 227–255.

Paxson, Christina H, "Consumption and income seasonality in Thailand," *Journal of Political Economy*, 1993, *101* (1), 39.

Pitt, M. M. and M.R. Rosenzweig, *Agricultural prices, food consumption and the health and productivity of Indonesian farmers*, Baltimore: Johns Hopkins University Press, 1986.

Pitt, Mark M, Mark R Rosenzweig, and Md Nazmul Hassan, "Productivity, health, and inequality in the intrahousehold distribution of food in low-income countries," *American Economic Review*, 1990, *80* (5), 1139–1156.

Ricker-Gilbert, Jacob, Nicole M Mason, Francis A Darko, and Solomon T Tembo, "What are the effects of input subsidy programs on maize prices? Evidence from Malawi and Zambia," *Agricultural Economics*, 2013, *44* (6), 671–686.

Rose, Elaina, "Ex ante and ex post labor supply response to risk in a low-income area," *Journal of Development Economics*, 2001, *64* (2), 371–388.

Rosenzweig, Mark R, *Handbook of Development Economics*, Vol. 1, Elsevier,

_ **and Hans P Binswanger**, "Wealth, Weather Risk and the Composition and Profitability of Agricultural Investments," *Economic Journal*, 1993, *103* (416), 56–78.

Schofield, Heather, "The Economic Costs of Low Caloric Intake: Evidence from India," *Mimeo*, 2013.

Strauss, John, "Does better nutrition raise farm productivity?," *The Journal of Political Economy*, 1986, pp. 297–320.

_ **and Duncan Thomas**, "Health, nutrition, and economic development," *Journal of Economic Literature*, 1998, *36* (2), 766–817.

Subramanian, Shankar and Angus Deaton, "The demand for food and calories," *Journal of Political Economy*, 1996, pp. 133–162.

Zambian Central Statistics Office (CSO), "2010 Census of Population," *Government of the Republic of Zambia*, 2010.

Table 1: Balance

	Control village	Partial treatment village	Full treatment village	Full = Partial
	Mean [SD]	Coeff [SE]	Coeff [SE]	p-value
	(1)	(2)	(3)	(4)
Age of household head	46.307	-5.115**	-3.160	0.396
	[16.755]	(2.273)	(2.900)	
Female headed household	0.230	-0.022	0.039	0.319
	[0.423]	(0.051)	(0.057)	
Number of household members under 5	0.825	0.207**	0.041	0.190
	[0.833]	(0.098)	(0.123)	
Family size at baseline	5.456	-0.040	-0.100	0.842
	[2.990]	(0.411)	(0.454)	
Last season (2011/2012), acres of Local Maize	1.316	-0.203	0.019	0.140
	[1.172]	(0.143)	(0.177)	
Last season (2011/2012), acres of Hybrid Maize	0.965	-0.083	-0.249	0.151
	[1.505]	(0.209)	(0.205)	
Last season (2011/2012), acres of Cotton	1.680	-0.273	-0.689**	0.123
	[1.523]	(0.262)	(0.314)	
How many mobile phones	0.561	0.172*	0.054	0.306
	[0.799]	(0.098)	(0.122)	
Household-level person days of ganyu last season	11.781	2.319	-2.156	0.058
	[24.463]	(3.562)	(3.441)	
Total acres currently used	4.768	-0.286	-0.320	0.931
	[2.275]	(0.378)	(0.489)	
Total days of ganyu hired	3.982	1.683	-0.213	0.274
	[9.544]	(1.431)	(1.693)	
How much maize stored (50 Kg bags)	25.368	-4.181	-4.493	0.918
	[50.072]	(5.075)	(5.119)	
Baseline harvest value (USD)	484.496	-3.643	-41.402	0.644
	[393.329]	(61.367)	(85.518)	

Notes: Column 1 shows means and standard deviations for the control villages. Columns 2 and 3 show the coefficients from a regression of the household variable on the treatment indicator variables with standard errors clustered at the village level. Column 4 reports the p-value from the test that the two treatment coefficients are equal. * $p < 0.10$ ** $p < 0.05$.

Table 2: Attrition and selection

	Baseline mean	Attrition		
		In midline	Attend loan meeting	In endline
	(1)	(2)	(3)	(4)
Control village	0.26	-0.002		0.007
	[0.439]	(0.015)		(0.014)
Partial treatment village	0.503	-0.013	-0.011	-0.009
	[0.501]	(0.014)	(0.025)	(0.014)
Full treatment village	0.237	0.020	0.011	0.005
	[0.426]	(0.013)	(0.025)	(0.015)
Age of household head	42.991	0.000	0.001*	0.000
	[15.932]	(0.001)	(0.001)	(0.001)
Female headed household	0.228	-0.006	0.019	-0.009
	[0.420]	(0.020)	(0.021)	(0.024)
Number of household members under 5	0.938	0.006	-0.005	0.001
	[0.918]	(0.009)	(0.014)	(0.008)
Family size at baseline	5.412	0.005**	0.004	0.002
	[2.785]	(0.002)	(0.004)	(0.002)
Last season (2011/2012), acres of Local Maize	1.218	-0.000	0.009	0.001
	[1.097]	(0.006)	(0.011)	(0.007)
Last season (2011/2012), acres of Hybrid Maize	0.864	0.001	0.013	0.006
	[1.175]	(0.006)	(0.008)	(0.004)
Last season (2011/2012), acres of Cotton	1.379	0.008**	0.004	0.005
	[1.400]	(0.004)	(0.011)	(0.005)
How many mobile phones	0.661	-0.009	0.006	-0.004
	[0.842]	(0.011)	(0.010)	(0.009)
Household-level person days of ganyu last season	12.437	-0.001	-0.000	0.000
	[24.834]	(0.001)	(0.000)	(0.000)
Total acres currently used	4.548	0.007**	0.010**	0.004*
	[2.195]	(0.003)	(0.004)	(0.002)
Total days of ganyu hired	4.779	0.000	-0.000	0.000**
	[15.252]	(0.000)	(0.001)	(0.000)
How much maize stored (50 Kg bags)	22.199	0.000	0.000	0.000*
	[31.602]	(0.000)	(0.000)	(0.000)
Baseline harvest value (USD)	472.854	0.000	0.000	0.000
	[428.580]	(0.000)	(0.000)	(0.000)
Sample restriction		baseline	treatment	baseline
Dependent var mean		0.975	0.945	0.977
Observations	439	439	325	439

Notes: Column 1 shows means and standard deviations for the full sample at baseline. Columns 2 and 4 show the coefficients from a regression of whether the household was in the midline and endline sample, respectively, on the explanatory variable, with standard errors clustered at the village level. Column 3 shows the same analysis but for whether the household attended the meeting where the loan program was explained, conditional on being in a treatment village. * $p < 0.10$ ** $p < 0.05$.

Table 3: Loan take up and repayment

	Control		Partial Treatment		Full Treatment	
	Number	%	Number	%	Number	%
Panel A: Sample						
Households	114		221		104	
Eligible households	0	0%	104	47%	99	95%
Panel B: Take up						
Any maize loan	0	0%	104	100%	90	91%
Number of bags	0		2.79		2.62	
Panel C: Repayment						
Loan contracts	0		104		90	
Some repayment			102	98.1%	88	97.8%
Full repayment			99	95.2%	87	96.7%
Percentage repaid			96.9%		97.7%	

Notes: Eligibility is defined as attending the loan information meeting in both the partial and full treatment villages. In the partial treatment villages, eligibility is further determined through a public lottery. In panel B, take up rates are shown conditional on eligibility.

Table 4: Consumption impacts - midline survey

	Number of meals in last 24 hours		Any missed meals in past week?	
Control village mean	*1.736*		*0.364*	
	(1)	(2)	(3)	(4)
Partial treatment village	0.0629		-0.0599	
	(0.0609)		(0.0659)	
Full treatment village	0.213***		-0.165**	
	(0.0594)		(0.0639)	
Eligible		0.123**		-0.0766*
		(0.0537)		(0.0435)
R-squared	0.110	0.102	0.092	0.083

Notes: N=424. Regressions of midline hunger outcomes on corresponding treatment specifications. Columns 1 and 3 show village level intention to treat OLS regressions. Columns 2 and 4 show individual level intention to treat OLS regressions. All regressions include controls and cluster standard errors at the village level. * $p < 0.10$ ** $p < 0.05$.

Table 5: Labor supply impacts - midline survey

	Any ganyu?	Days of ganyu	Days on-farm	Days of hired ganyu
Control village mean	*0.464*	*3.055*	*16.17*	*2.009*
	(1)	(2)	(3)	(4)
Panel A: Village-level ITT				
Partial treatment village	-0.0757	-0.00514	2.248	-0.174
	(0.0567)	(0.565)	(1.948)	(0.686)
Full treatment village	-0.118*	-1.028*	1.407	0.158
	(0.0607)	(0.599)	(1.862)	(0.878)
R-squared	0.144	0.160	0.202	0.106
Panel B: Individual-level ITT				
Eligible	-0.0835**	-0.260	1.578	-0.562
	(0.0373)	(0.456)	(1.292)	(0.572)
R-squared	0.144	0.154	0.201	0.107

Notes: N=424 in columns 1-3; N=418 in column 4. Panel A shows village level intention to treat OLS regressions. Panel B shows individual level intention to treat OLS regressions. All columns control for covariates shown in Table 1 and cluster standard errors at the village level. The outcome variable in column 1 is a binary variable that equals one if the family reported any ganyu at midline. Column 2 reports the corresponding person-days of ganyu reported. Column 3 is the person-days of on-farm work and Column 4 is the person-days of hired ganyu. For all questions, the recall period is 2 weeks. * $p < 0.10$ ** $p < 0.05$.

Table 6: Labor supply impacts - endline survey

	Any ganyu?	Days of ganyu	Days on-farm	Days of hired ganyu	Before repayment
Control village mean	*0.595*	*10.96*	*338.3*	*3.126*	*0.294*
	(1)	(2)	(3)	(4)	(5)
Panel A: Village-level ITT					
Partial treatment village	-0.0825*	-0.0746	1.766	0.866	-0.0651
	(0.0485)	(2.833)	(23.60)	(1.020)	(0.0835)
Full treatment village	-0.0908*	-4.340	24.04	1.453	-0.138*
	(0.0454)	(2.792)	(23.91)	(1.380)	(0.0732)
R-squared	0.156	0.214	0.318	0.272	0.124
Panel B: Individual-level ITT					
Eligible	-0.0777*	-0.0852	9.398	0.0273	-0.0802*
	(0.0418)	(2.000)	(16.51)	(0.967)	(0.0419)
R-squared	0.157	0.208	0.316	0.269	0.122

Notes: N=424 in columns 1-4. Panel A shows village level intention to treat regressions. Panel B shows individual level intention to treat OLS regressions. All columns control for covariates shown in Table 1 and cluster standard errors at the village level. The outcome variable in column 1 is a binary variable that equals one if the family reported any ganyu for the agricultural season at endline. Column 2 reports the corresponding person-days of ganyu reported. Column 3 is the person-days of on-farm work and Column 4 is the person-days of hired ganyu. Column 5 is the same outcome as column 1 but for the two months preceding the loan repayment only. For the outcome in columns 1-4, the recall period is the past agricultural season (9 months). For the outcome in column 5, the recall period is the 2 months preceding the endline survey (May and June). * $p < 0.10$ ** $p < 0.05$.

Table 7: Wage impacts

	Daily self-reported wage (among workers)			
	(1)	(2)	(3)	(4)
Panel A: Midline				
	Household level		Individual level	
Partial treatment village	1.396		-0.515	
	(2.116)		(1.045)	
Full treatment village	4.666*		3.040	
	(2.561)		(2.364)	
Fraction eligible		4.775*		3.073
		(2.623)		(2.430)
Control village mean	9.65		7.981	
Observations	154	154	292	292
R-squared	0.096	0.095	0.168	0.160
Panel B: Endline (difference in difference)				
	Activity fixed effects		Season fixed effects	
Partial treatment village	4.664		4.116	
	(3.631)		(3.024)	
Full treatment village	6.437		4.897	
	(4.008)		(4.824)	
Fraction eligible		6.229		4.522
		(4.387)		(5.070)
Control village mean	18.42		18.42	
Observations	1,024	1,024	1,024	1,024
R-squared	0.174	0.173	0.173	0.172

Notes: Panel A shows the results for average ganyu wages reported at the household level (columns 1 and 2) and individual level (columns 3 and 4). The dependent variable in Panel A is the average wage received per day of ganyu in the two weeks preceding the midline survey. Panel B shows difference in difference estimates for ganyu wages reported at baseline and endline, with village fixed effects. Columns 1 and 2 include activity fixed effects for field preparation, planting, weeding and harvesting. Columns 3 and 4 include season fixed effects and restricts treatment effects to the time period after which loans were announced. All specifications include the household level controls shown in Table 1 and cluster standard errors at the village level. * $p < 0.10$ ** $p < 0.05$.

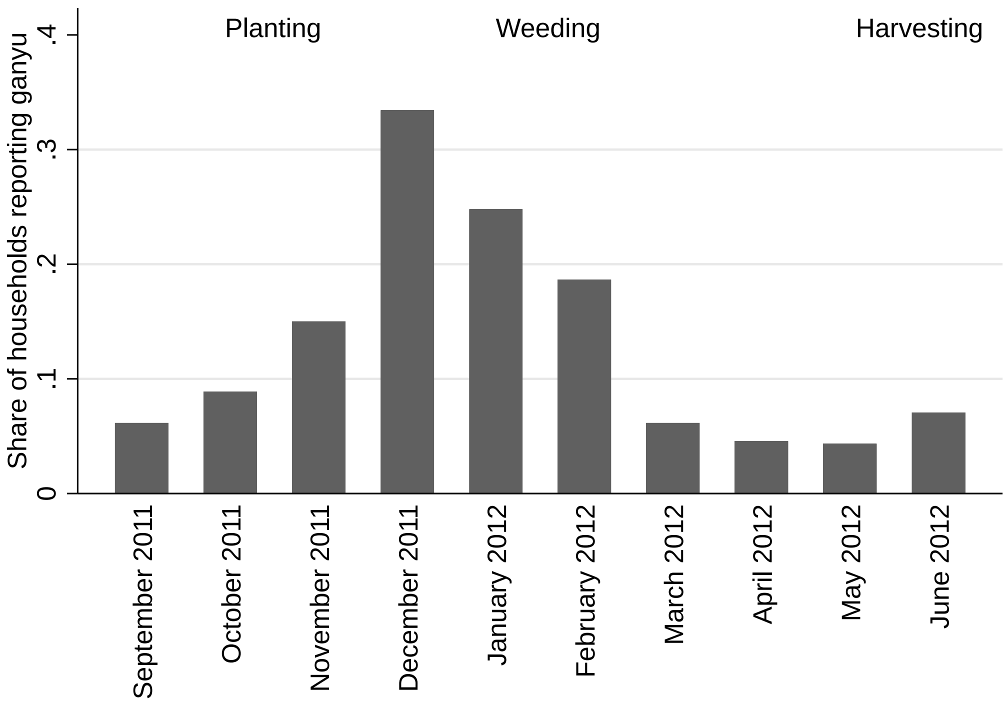

Figure 1: Off-farm labor by month

Notes: Share of households reporting any *ganyu* labor by month in the agricultural season preceding the baseline survey.

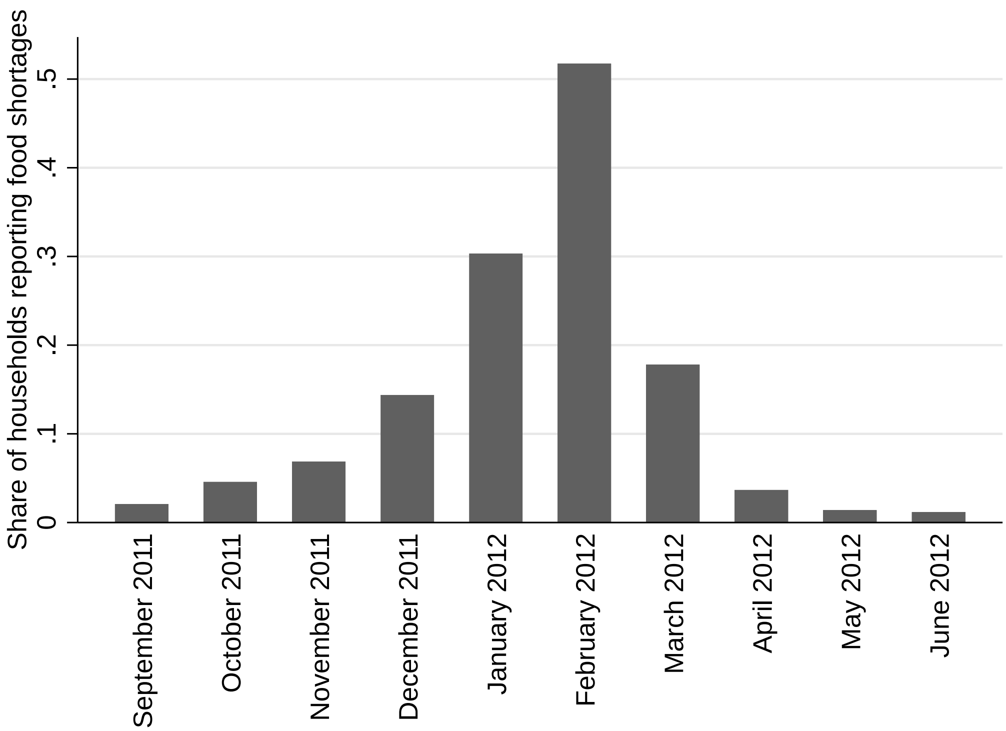

Figure 2: Food shortages by month

Notes: Share of households reporting food shortages by month in the agricultural season preceding the baseline survey.

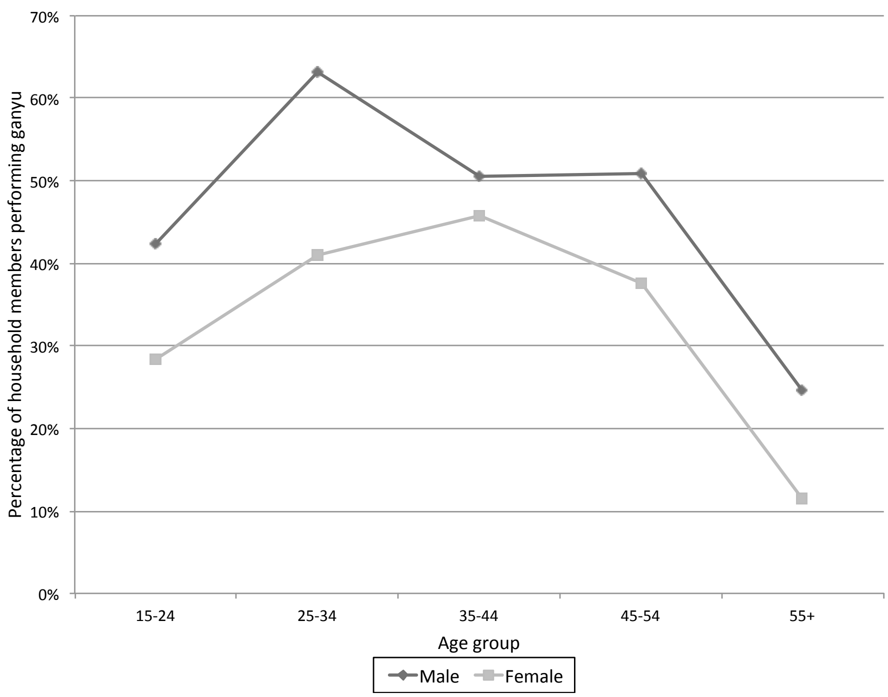

Figure 3: Off-farm labor by gender and age

Notes: Share of respondents by age and gender who report any *ganyu* at baseline.

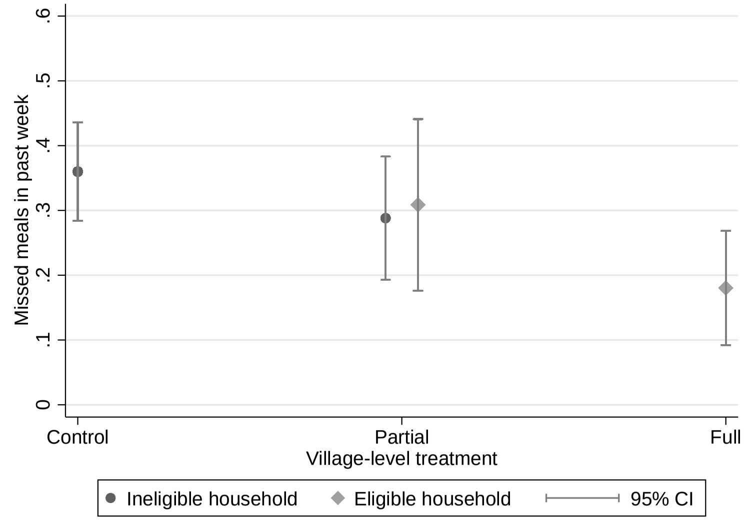

(a) Any missed meals (1 week)

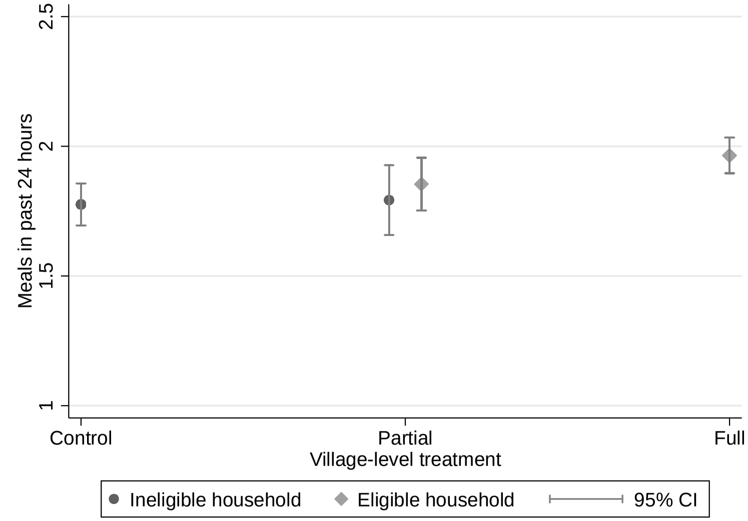

(b) Number of meals (24 hours)

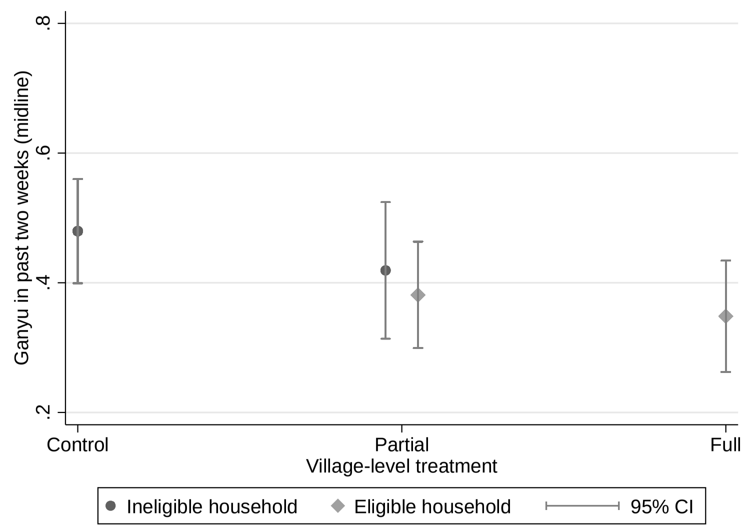

(c) Any *ganyu* (midline)

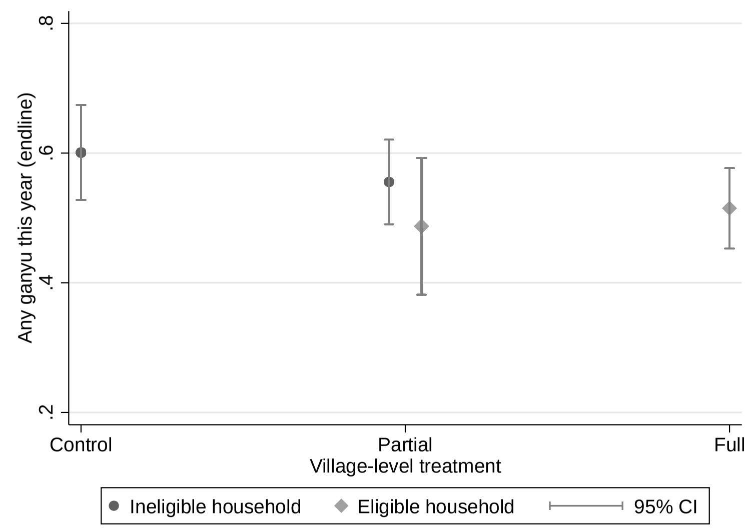

(d) Any *ganyu* (endline)

Figure 4: Within-village spillovers

Notes: Marginal effects by village and household level treatment assignment. A comparison of marginal effects across ineligible households shows the effect of being in a village with partial treatment on ineligible households. A comparison of marginal effects across eligible households shows the effect of being in a village with partial treatment on eligible households. Marginal effects are from regressions in household-level controls and standard errors clustered at the village level.

Appendix

Table A.1: Chipata District household characteristics (LCMS 2010)

	All households			Rural households only		
	Zambia excluding Chipata	Chipata	H0: Equal means (p-value)	Zambia excluding Chipata	Chipata	H0: Equal means (p-value)
Number of household members	5.31	5.11	0.21	5.41	5.37	0.88
Number of rooms	3.29	3.02	0.11	2.93	2.60	0.02
Electricity access	0.33	0.22	0.64	0.07	0.02	0.00
Private water access	0.26	0.23	0.01	0.03	0.00	0.00
Private toilet	0.71	0.59	0.14	0.65	0.52	0.00
Grows crops	0.50	0.60	0.01	0.80	0.89	0.00
Monthly expenditure in USD	333.97	219.21	0.00	219.16	121.94	0.00
Classified as very poor	0.32	0.47	0.00	0.42	0.63	0.00
Observations	18948	449	19397	8243	225	8468

Notes: Tables summarizes average household characteristics as reported in the Zambian Living Conditions Measurement Survey (LCMS) 2010. P-values reported are based on a two-sample mean comparison; standard errors adjusted for 940 survey clusters in the LCMS.

Table A.2: Labor supply impacts - endline survey, loan months only

	Any ganyu?	Days of ganyu	Days on-farm	Days of hired ganyu
Control village mean	*0.505*	*7.216*	*196.3*	*1.919*
	(1)	(2)	(3)	(4)
Panel A: Village-level ITT				
Partial treatment village	-0.0347	0.0900	1.069	0.680
	(0.0553)	(2.087)	(13.73)	(0.660)
Full treatment village	-0.0126	-1.314	18.24	0.845
	(0.0518)	(1.882)	(14.14)	(0.916)
R-squared	0.157	0.165	0.316	0.231
Panel B: Individual-level ITT				
Eligible	-0.0702*	-0.585	8.499	0.266
	(0.0414)	(1.280)	(9.306)	(0.601)
R-squared	0.161	0.164	0.314	0.228

Notes: N=424. Panel A shows village level intention to treat OLS regressions. Panel B shows individual level intention to treat OLS regressions. All columns control for covariates shown in Table 1 and cluster standard errors at the village level. The outcome variable in column 1 is a binary variable that equals one if the family reported any ganyu at endline. Column 2 reports the corresponding person-days of ganyu reported. Column 3 is the person-days of on-farm work and Column 4 is the person-days of hired ganyu. For all questions, the recall period is the months covered by the loan (Jan-June). * $p < 0.10$ ** $p < 0.05$.

Table A.3: Gifts and transfers

	Midline		Endline	
	Outgoing gifts and transfers	Incoming gifts and transfers	Outgoing food gifts	Incoming food gifts
Control village mean	*0.336*	*0.236*	*0.658*	*0.387*
	(1)	(2)	(3)	(4)
Partial village: ineligible	-0.599	-0.0321	-0.113*	-0.0726
	(1.069)	(0.0541)	(0.0592)	(0.0548)
Partial village: eligible	10.43	0.342***	-0.0450	-0.0528
	(10.35)	(0.0768)	(0.0731)	(0.0633)
Full treatment village	0.0563	0.312***	-0.0180	-0.0387
	(0.913)	(0.0774)	(0.0858)	(0.0626)
Observations	424	424	423	423
R-squared	0.018	0.183	0.040	0.068

Notes: Regressions of self-reported gifts and transfers at midline (two week recall window) and endline (year recall window). All outcomes are binary and estimated using a linear probability model. Regressions include controls and cluster standard errors at the village level. * $p < 0.10$ ** $p < 0.05$.

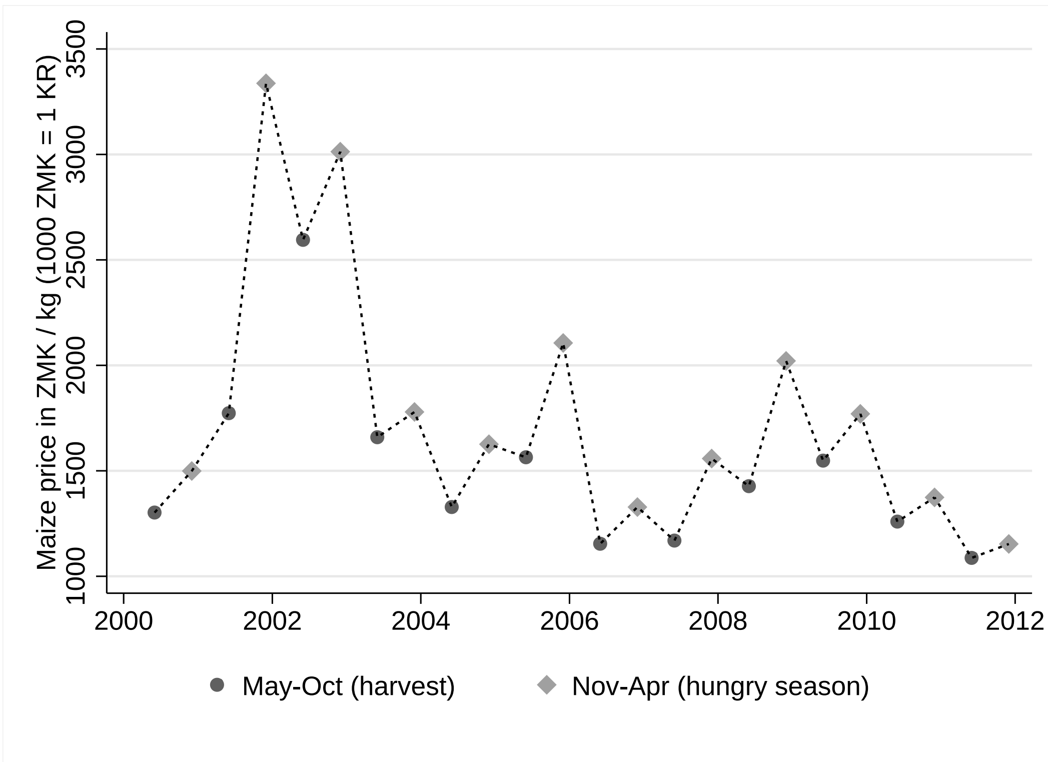

Figure A.1: Seasonal maize prices

Notes: Figure is adapted from data compiled in Ricker-Gilbert (2013), from the Zambian Central Statistics Office and based on district retail prices, averaged across the hungry (lean) and harvest seasons. Prices are in old currency (1000 ZMK = 1 KR) per kilogram.

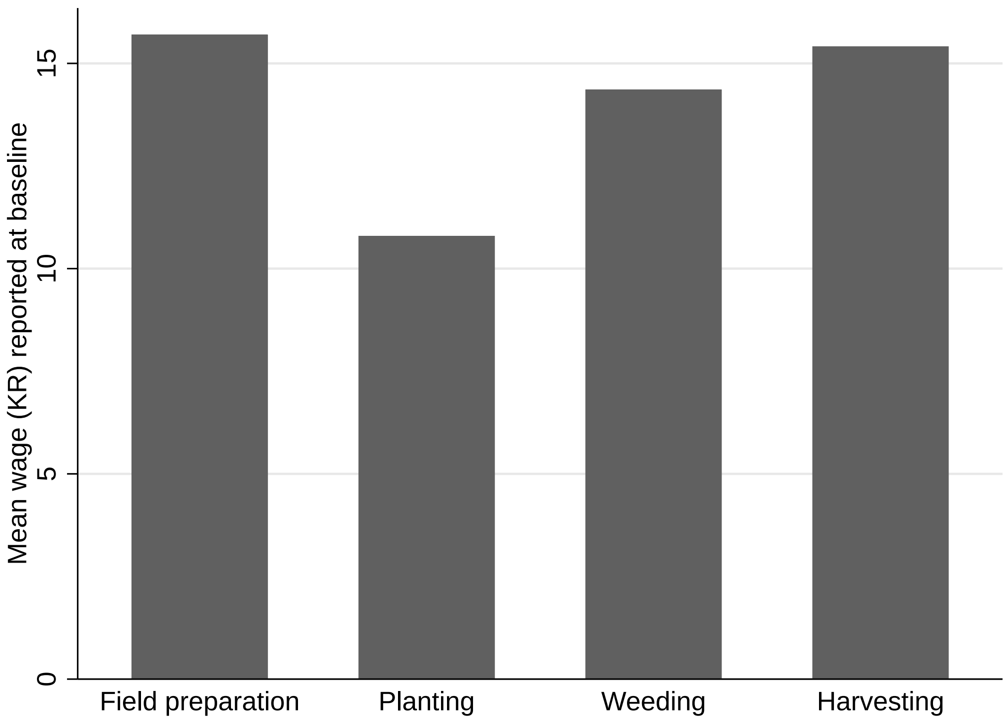

Figure A.2: Average reported wage at baseline by agricultural activity

Notes: Wages reported by baseline respondents by four main agricultural labor activities for the previous season. Wages are in Zambian Kwacha.

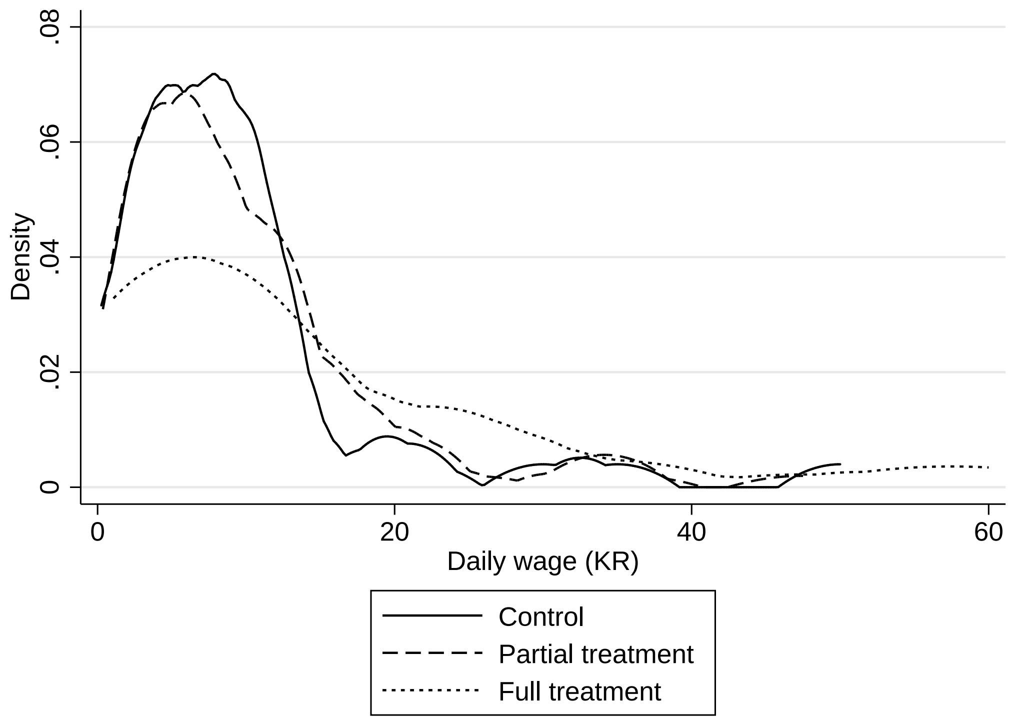

Figure A.3: Wage densities by treatment arm

Notes: Kernel density estimates of wages in the two weeks preceding the midline survey. Figure shows household-level average reported daily wage earnings.