

Seasonal Liquidity, Rural Labor Markets, and Agricultural Production[†]

By GÜNTHER FINK, B. KELSEY JACK, AND FELIX MASIYE*

Rural economies in many developing countries are characterized by a lean season in the months preceding harvest, when farmers have depleted their cash and grain savings from the previous year. To identify the impacts of liquidity during the lean season, we offered subsidized loans in randomly selected villages in rural Zambia. Ninety-eight percent of households took up the loan. Loan eligibility led to increases in on-farm labor and agricultural output, driving up wages in local labor markets. Larger effects for poorer households suggest that liquidity constraints contribute to inequality in rural economies. (JEL O13, O15, O18, Q11, Q12, R23)

In many agricultural settings, low returns to saving and high costs of borrowing raise the cost of smoothing consumption from one harvest to the next, resulting in a distinct “lean season” or “hungry season” in the months leading up to harvest. In this paper, we show that liquidity shortages during the hungry season affect not only consumption, but also local labor markets and agricultural production. In a two-year randomized controlled trial conducted with 3,139 small-scale farmers in 175 villages in rural Zambia, we test the impacts of small loans designed to cover basic consumption needs during the hungry season. We find almost universal take-up of the loans, high repayment rates, and positive impacts on agricultural production. Most of the increase in agricultural output comes from a reallocation of family labor from the market, where it provides an immediate source of wage income during the hungry season, to the family farm, where it leads to additional harvest income several months later.

Our setting is representative of many parts of sub-Saharan Africa. Agriculture is rain-fed, resulting in a single harvest each year. Access to formal saving opportunities is limited, informal alternatives are low-return, and borrowing opportunities are

*Fink: Swiss Tropical Public Health Institute and University of Basel (email: guenther.fink@swisstph.ch); Jack: University of California–Santa Barbara and NBER (email: kelseyjack@ucsb.edu); Masiye: University of Zambia (email: fmasiye@yahoo.com). Esther Duflo was the coeditor for this article. We thank the two anonymous referees along with participants at numerous seminars and conferences for comments and suggestions. Financial support was provided by the Growth and Labor Markets in Low Income Countries (GLM-LIC), the International Growth Centre, the Agricultural Technology Adoption Initiative (JPAL/CEGA), and an anonymous donor. We thank Innovations for Poverty Action for logistical support, Rachel Levenson for careful oversight of the field work, Daniel Velez Lopez, Chantelle Boudreaux, Simon Jean, and Carlos Riumallo Herl for research assistance, and Jack Cavanagh for data support. The project received IRB approval from the Harvard T. H Chan School of Public Health and the University of Zambia, and is registered on the AEA trial registry as AEARCTR-0000130.

[†]Go to <https://doi.org/10.1257/aer.20180607> to visit the article page for additional materials and author disclosure statements.

accompanied by extremely high interest rates. As a result, food reserves and consumption are highly seasonal, peaking after harvest, and reaching their nadir during the hungry season.¹ When asked how they will cover short-term needs (in addition to restricting consumption as the name “hungry season” suggests), a majority of households in our sample say they will sell family labor in local labor markets. These labor sales, locally referred to as *ganyu*, typically occur within a given village, with better-off farmers hiring labor from relatively poor farmers at an individually negotiated rate. In our setting, poorer households also report higher baseline interest rates on borrowing. This implies a higher marginal product of labor on poorer households’ farms during the hungry season, even if the discounted return to labor is equalized across farms.

In the first year of the study, selected households in two-thirds of villages were eligible for cash or food loans, worth around US\$30, during the hungry season; in the second year, 50 percent of study villages received the program, with rotation of treatment status between years (i.e., some villages received two years of the program, some one year and some zero years). Our sample covers around one-half of the households in a village on average, both for data collection and treatment eligibility.² Despite an implicit interest rate of 4.5 percent per month, more than 98 percent of eligible households took up the offer. These take-up rates are considerably above what is typically reported in the microfinance literature,³ highlighting the high demand for credit during the hungry season, as well as the high cost of alternative consumption-smoothing strategies, including *ganyu* labor sales. Nearly 95 percent repayment in the first year suggests that high take-up was not driven by anticipated default.

We find that the intervention led to adjustments in labor allocation and increases in agricultural output. In year 1 of the program, the likelihood that a family sold any *ganyu* labor during the hungry season fell by 4.8 percentage points (14 percent) in response to treatment, with a 25 percent reduction in hours sold, on average. The likelihood of hiring *ganyu* labor increased by around 3.9 percentage points (43 percent). As a result, average village-level wages, which we measure as daily earnings, increased by about 2.5 Kwacha, or 17 percent relative to control villages.⁴ We also find increases in family labor supply: on average, treated households reported 4.9 additional hours of family labor per week on their own farms during the hungry season in year 1, an increase of 10.6 percent relative to the control group.⁵

¹Other papers highlight the contribution of seasonal variation in grain prices to consumption seasonality (see, for example, Kaminski, Christiaensen, and Gilbert 2014; Devereux, Sabates-Wheeler, and Longhurst 2013). In our setting, grain prices also peak during the hungry season but appear insufficient to fully explain consumption fluctuations.

²Treating only a subset of households in treatment villages results in a “partial treatment, general equilibrium” setup, in which (unmeasured) labor supply responses by untreated households in treatment villages dampen wage effects. We discuss this feature of our design both in presenting our model and interpreting our results.

³The average take-up rate across the six randomized evaluations of microfinance published in a special issue of the *American Economic Journal: Applied Economics* in 2015 was 39 percent (Banerjee, Karlan, and Zinman 2015). The only RCT in this issue with a take-up rate close to 98 percent had a default rate of 46 percent; we observe an average default rate of 11 percent.

⁴Some of the average wage effect may be driven by changes in the composition of who selects into the labor market rather than changes in equilibrium wages. We conduct a bounding exercise that shows wage increases even under conservative assumptions about worker selection.

⁵This result suggests that resource scarcity may constrain labor supply during the hungry season. We examine other outcomes to interpret this result, and find suggestive evidence for psychological or behavioral channels

Agricultural output increased by 9 percent in response to treatment. The value of the increase in agricultural revenue is similar to the repayment due under the loan, yet this comparison is likely to understate the welfare gains, since hungry season consumption also increased. The strongest evidence for a positive net welfare impact comes from repeat take-up across years: out of 937 households offered the loan for a second time in year 2, only 18 (<2 percent) declined to take up a second time.

Households in villages treated for the first time in year 2 show similar treatment effects to those in year 1, though effects are slightly less precisely estimated. Households in villages that receive treatment in both years similarly show reductions in family labor sales and improvements in consumption, but they do not increase labor investments on-farm nor do they increase agricultural output relative to villages never eligible for treatment.

We calculate that, on average, around 65 percent of the loan value was used for additional labor inputs on-farm, i.e., for reallocating family labor from the labor market to the family farm and for hiring additional labor, while the remainder went to other household expenditures, including hungry season consumption. The measured increase in labor inputs almost fully explains the impact on agricultural output at a marginal product of labor equal to the casual labor daily wage during the hungry season. We observe the largest impacts on both consumption and agricultural output among the households with the fewest liquid resources and highest interest rates at baseline. As a result, inequality in consumption and output declines among households eligible for the loan.

Our paper is closely related to an extensive literature that highlights the links between credit market frictions, agricultural labor markets, and aggregate output. Our conceptual framework builds on Jayachandran (2006), who shows that a lack of credit access leads to increased labor supply and lower wages among landless rural laborers when the economy is exposed to aggregate productivity shocks. More directly related to our study, Pitt and Khandker (2002) show a link between seasonal hunger, demand for microcredit, and male labor supply in Bangladesh. The critical role of family labor sales for smoothing consumption has been documented by others (Kochar 1995, 1999; Rose 2001; Ito and Kurosaki 2009).⁶ We extend this literature in two ways. First, we show that family labor sales are not only important in the presence of unanticipated shocks, but also to cover anticipated liquidity shortages. Second, we show that liquidity-induced labor sales lower future (agricultural) income and may result in an inefficient allocation of labor across farms, lowering aggregate output and increasing within-village inequality.

Our findings run counter to some of the recent evidence on the impacts of microfinance, which has found modest take-up and mixed impacts on both consumption and income (Cull and Morduch 2018). We observe high take-up (>95 percent) and high returns on investment (about 30 percent over a six-month period), consistent with the insight from Field et al. (2013) that tailoring loan products to clients' financial flows and consumption smoothing needs can improve both take-up and impacts.

(consistent with the findings in Kaur et al. 2019 and Banerjee et al. 2020) rather than the physical channels that have been the focus of prior literature (e.g., Dasgupta and Ray 1986, 1987).

⁶This consumption smoothing role of local labor markets is also tied to the substantial literature on informal smoothing strategies (see, for example, Morduch 1995 for a review), some of which, like labor sales, may carry long-run costs (e.g., Rosenzweig and Wolpin 1993).

Prior evidence suggests high returns from synchronizing borrowing or investment opportunities with financial flows in rural agricultural settings, where incomes and prices are highly seasonal (Duflo, Kremer, and Robinson 2011; Bryan, Chowdhury, and Mobarak 2014; Basu and Wong 2015; Burke, Bergquist, and Miguel 2019; Casaburi and Willis 2018; Aggarwal, Francis, and Robinson 2018).⁷ The loan product we study targets a particularly constrained time of year for agricultural households, but we cannot rule out similar impacts from offering the loan at other times. Directly testing whether the returns to liquidity are different in the hungry season would have required offering the same product at different times of the year, similar to Casaburi and Willis (2018). Given that our intervention was timed to coincide with the period of peak agricultural labor demand, similar effects on labor allocation and agricultural production at other times of the year seem unlikely.

From a policy perspective, our findings suggest large potential welfare gains from relaxing seasonal liquidity constraints for selected households in a village. We use our model to engage in “structured speculation” (Banerjee, Chassang, and Snowberg 2017) about (general equilibrium) effects outside of our study sample and design. We simulate our model to show that scaling up access to lower interest rate loans leads to larger wage adjustments and more homogeneous returns to labor within the village, resulting in greater reductions in income inequality. In spite of these benefits, the potential for scaling up seasonal consumption loans is hindered by the high transaction costs associated with delivering and collecting loans in remote areas with poor road infrastructure. Bundling seasonal loans with other technologies, such as digital borrowing platforms, or piggybacking on existing rural networks may help bring down costs.⁸ Other strategies for lowering the cost of seasonal consumption smoothing, such as more secure savings, may also decrease reliance on family labor as a costly smoothing strategy.

I. An Agrarian Economy with Capital Market Frictions

We build on the agrarian labor market model introduced in Jayachandran (2006).⁹ Each village economy has a finite number N of farming households that maximize utility over two periods ($t = 1, 2$).¹⁰ Each household i some has initial liquid resources S_{i0} .¹¹ All households allocate their labor endowment h between

⁷Our paper is most closely related to Basu and Wong (2015), who evaluate a seasonal food credit and improved storage program in Indonesia. Similar to the results presented here, they find that food loans increase non-staple food consumption during the hungry season and income from crop sales at harvest, but do not analyze impacts on labor allocation or production.

⁸We test this approach in an ongoing trial in collaboration with a large outgrower cotton company, registered as AEARCTR-0003561.

⁹We modify Jayachandran’s model in two important ways to more closely match our setting: first, we assume that all farmers own land and can thus create income both from their own farms and from selling labor to others. Second, we assume that farming income is earned in the second period rather than the first to highlight the trade-off between financing hungry season consumption and receiving greater output in the future.

¹⁰We use the terms farmer and household interchangeably.

¹¹This initial distribution of liquid resources can be thought of as the result of a stochastic process where all households start with an initial endowment of zero, and accumulate resources over time based on the farm’s (land and labor) productivity and idiosyncratic shocks such as weather or pests. We assume that initial resources are pre-determined and positively correlated with farm productivity A_i (and verify this assumption empirically, see online Appendix Section A.2). In our model we abstract from the stochastic element in the production process. To achieve a stable distribution of baseline reserves in a recursive model, a substantial degree of stochasticity or other form

sales to the market and work on their own farms, which have heterogeneous productivity A_i .¹² Farming output is a function of A_i as well as total labor input (farm-level labor demand) d_i , which includes both own (family) labor on farm and hired labor, and is given by

$$(1) \quad y(d_i) = A_i d_i^\beta,$$

where $\beta \in (0, 1)$ defines the returns to labor.¹³

Households maximize their utility from consumption. Utility is additive and separable across the two periods; second period utility is discounted by a subjective discount factor $\rho < 1$:

$$(2) \quad u(c_{i1}, c_{i2}) = \log(c_{i1}) + \rho \log(c_{i2}).$$

Households can borrow at a rate r_i , which decreases with the farm's initial resources, S_{i0} , i.e., $\partial r_i / \partial S_{i0} < 0$.¹⁴ All borrowing needs to be repaid by the end of the second period.¹⁵

Labor is traded in local markets, which clear at the endogenous wage w such that total farm labor input equals aggregate labor supply:

$$(3) \quad \sum_{i=1}^N d_i(w) = \sum_{i=1}^N (h_i).$$

A. Household Utility Maximization

Rational households maximize utility from consumption over two periods:

$$(4) \quad \max_{c,d} \log(c_{i1}) + \rho \log(c_{i2}),$$

subject to

$$c_{i1} \leq S_{i0} + (h(c_{i1}) - d_i)w + B_i,$$

$$c_{i2} \leq y_i(d_i) - B_i r_i^e,$$

of income redistribution would be necessary, which is consistent with our data: across years, the within-household correlation in hungry season liquid resource rankings is 0.49 and the correlation in agricultural output is 0.63.

¹²We interpret the productivity term A_i as a general measure of a farm's potential output, capturing variation in farming skill, farm size, and land quality.

¹³This production function corresponds to a standard Cobb-Douglas production function $y(d_i) = A_i d_i^\beta k^{1-\beta}$ with the second input factor k normalized to 1. Like Rosenzweig (1980) and others, we assume that labor markets are well functioning and that land owning households (all of our sample) both buy and sell labor on local markets.

¹⁴This assumption is consistent with any model where the expected ability to repay increases with collateral (which is proxied by S_{i0}), ignoring limits on borrower liability.

¹⁵We do not model saving technologies explicitly in our model. Empirically, returns to savings are low in the study setting, explaining at least partially the low levels of reserves during the hungry season. In addition, savings decisions are most relevant following harvest, which precedes period 1 in our model. Our model reflects the sequential nature of agricultural production, which may be subject to period-specific constraints (Behrman, Foster, and Rosenzweig 1997; Skoufias 1996).

where B_i is net resources borrowed during the first period, and $r_i^e = 1 + r_i$ is the effective interest rate. We normalize the price of consumption goods to 1. We allow households' effective labor supply to be constrained by first period consumption, such that $\partial h_i / \partial c_{1i} \geq 0$. This constraint on labor supply can reflect physical constraints in the spirit of Dasgupta and Ray (1986, 1987), but can also represent scarcity-induced reductions in cognition and decision-making or increases in stress and affect that influence the decision or motivation to work (see Kremer, Rao, and Schilbach 2019 for a review of recent related literature).

In period 1, households choose labor inputs d_i on the farm and consumption. In period 2, households receive harvest income y_i . Period 2 net income (and consumption) is given by harvest income minus period 1 borrowing plus interest. Period 1 consumption can be financed through initial resources (S_{i0}), labor income (effective labor supply, h_i , minus on-farm labor inputs, d_i , times the wage rate, w), as well as through borrowing (B_i). For any given wage and interest rate, households will always choose labor inputs such that the discounted marginal product of labor earned in period 2 equals the wage, which implies

$$(5) \quad d_i^* = \left(\frac{\beta A_i}{w r_i^e} \right)^{\frac{1}{1-\beta}}.$$

Optimal labor inputs increase with farm productivity, A_i , and decrease with wages and interest rates. Optimal consumption patterns imply

$$(6) \quad \frac{c_{i2}}{c_{i1}} = \rho r_i^e (1 - h'_{c_{i1}} w).$$

If $h'_{c_{i1}} = \partial h_i / \partial c_{1i} = 0$, this simplifies to $c_{i2} / c_{i1} = \rho r_i^e$, highlighting the basic relationship between interest rates, subjective discount rates, and consumption seasonality. If, instead, effective labor supply increases with first period consumption ($\partial h_i / \partial c_{1i} > 0$), the ratio of second period to first period consumption falls (consumption seasonality declines). Online Appendix Section A.1 provides a full solution for the model.

The Effect of Lowering Hungry Season Interest Rates.—Our experimental intervention subsidized credit access for a subset of small-scale farmers in randomly selected villages by offering hungry season (period 1) loans at a specific interest rate, $\hat{r} < r_i^e$.¹⁶ We derive the impacts on labor allocation, wages, agricultural output, and consumption among the subset of treated households. We discuss the effects on untreated (without access to \hat{r}) households in the same labor market, and the effects of offering \hat{r} to all households in the market, below and in Section VI.

¹⁶This inequality is consistent with the 98 percent loan take-up that we observe in our experiment.

Prediction 1: Labor demand increases among farmers who borrow at $\hat{r} < r_i^e$, causing equilibrium wages to increase.

From equation (5), lower interest rates always increase demand for labor inputs, holding wages constant, i.e., $\partial d_i^*/\partial r_i^e < 0$. If effective labor supply is not constrained by consumption ($\partial h_i/\partial c_{1i} = 0$), this will mechanically increase net demand for labor among treated farmers and result in an increase in equilibrium wages for markets to clear. If effective labor supply depends on consumption ($\partial h_i/\partial c_{1i} \geq 0$), treatment effects on net labor demand and wages stay positive as long as labor demand effects dominate consumption-driven labor supply effects, i.e., if $\partial d_i/\partial r_i^e > (\partial h_i/\partial c_{1i})(\partial c_{1i}/\partial r_i^e)$.

Prediction 2: Agricultural output increases among farmers who borrow at $\hat{r} < r_i^e$.

An increase in on-farm labor demand, d_i^* , from lower interest rates mechanically increases agricultural output ($\partial y_i^*/\partial d_i^* > 0$), holding wages fixed. This output effect is moderated by an increase in wages if labor supply is fixed: the larger the wage response, the smaller the increase in output among treated farmers. If labor supply increases in response to higher first period consumption, wage responses are smaller and output increases are larger.

Prediction 3: Period 1 consumption increases and consumption seasonality decreases among farmers who borrow at $\hat{r} < r_i^e$.

By optimality condition (6), it must always be true that lower interest rates increase the share of resources allocated to the first period, and thus also decrease consumption seasonality (c_2/c_1). Given that treatment increases overall resources, absolute levels of first period consumption must always increase. For second period consumption, the positive income effects from lower interest rates are partially offset by negative substitution effects toward first-period consumption.¹⁷

Prediction 4: The impacts of borrowing at $\hat{r} < r_i^e$ are increasing in r_i^e and decreasing in S_{i0} .

Since effective interest rates are, by assumption, highest for the households with the lowest initial resources (S_{i0}), the change in interest rates induced by the intervention will decrease with S_{i0} . As long as the correlation between productivity A_i and S_{i0} is sufficiently low (see online Appendix Section A.1 for details), labor and output effects will be largest for the farmers with the lowest initial endowment. This also implies a decline in consumption and agricultural income inequality among farmers borrowing at \hat{r} .

¹⁷Higher food prices in the hungry season may contribute to consumption seasonality. Note that this model normalizes the price of consumption to one in all periods, and so suppresses the effect of grain price fluctuations, which may arise due to storage costs, for example, on consumption seasonality. We test for treatment effects on grain prices in Section IV.

Credit Access and General Equilibrium Effects.—Our predictions focus on treated farmers in a “partial treatment, general equilibrium” framework, in keeping with our intervention and data collection that covered a subset of farmers in each village. Quantifying the overall welfare impacts of programs to address credit market frictions, which, like our intervention, often target only a subset of households in a village, requires an assessment of spillovers to untreated farmers. In relative terms, the expected changes in labor allocation and agricultural production are greatest for treated households in a partial treatment setting, because some of their increased labor demand will be met by untreated households in the same village, and the wage response will be muted relative to a full treatment scenario.¹⁸

In Section VI, we engage in “structured speculation” (Banerjee, Chassang, and Snowberg 2017) to explore the broader welfare implications of our experimental results. We calibrate the model above to our study setting, which allows us to (i) assess program impact on untreated farmers in our partial treatment setup, and (ii) to simulate outcomes under a scaled up policy version of the intervention that lowers credit market interest rates for all households in a village.

II. Experimental Design and Implementation

We turn now to our experimental setting, design, and implementation. We offer further detail on local markets in Section IIIB.

A. Study Setting

The study was implemented between October 2013 and September 2015 (with survey data covering three agricultural cycles/years) in Chipata District, Zambia. Chipata District is located at the southeastern border of Zambia, with an estimated population of 456,000 in 2010 (CSO 2010). Three-quarters of the population live in rural areas, with small-scale farming as the primary source of income. The 2010 Living Conditions Monitoring Survey (CSO 2010) estimates an average monthly expenditure by rural households in Chipata of US\$122 (US\$0.8 per person-day), or about one-third of the national average (US\$389).

The study implementation targeted small-scale farmers, i.e., households growing crops on less than 5 hectares (12 acres). The label “small-scale” is somewhat misleading since it suggests that these farmers are unusually small; in fact, small-scale farmers represent the vast majority of rural households in Zambia. In our study villages, we document that over 95 percent of households meet this definition.

Study Sample.—We randomly sampled 5 villages from 50 of the 53 administrative blocks in the district, using village lists from the Ministry of Agriculture’s farm registry and omitting villages with less than 20 or more than 100 farms and those

¹⁸This partial treatment set up increases our power to detect impacts on labor and agricultural output. While we could have also sampled untreated households in intervention villages, this would have come at the cost of a smaller number of villages (or of treated farmers per village) in the sample.

containing Chipata town. Study enumerators visited the 250 sampled villages to screen for eligibility.¹⁹ Our final sample covers 175 villages.

Within each eligible village, households were sampled from the village rosters collected during the initial screening visits. Only small farms, less than 5 hectares according to the Zambian Ministry of Agriculture, were eligible for the program.²⁰ Eligible households were randomly sorted and the first 22 selected for the baseline survey. This resulted in 53 percent of households on average being selected for the project; across all villages, the share of households enrolled in the study ranged from 15 to 100 percent. A total of 3,701 households were sampled for the baseline and 3,139 were surveyed at baseline (85 percent).²¹

B. Experimental Design

The study took place over two years and was designed to coincide with the agricultural cycle (see Figure 1), which starts with field preparation in September, followed by planting activities around the time of the first rains in November. Planting is followed by weeding between January and April, which is also the time locally referred to as the “hungry season” or “lean season.” In April, early crops start to become available; harvest begins in earnest in May. Between August and October, few agricultural activities take place. Our study covered two years: the 2013–2014 agricultural cycle (study year 1), and the 2014–2015 agricultural cycle (study year 2). The study design is summarized in Figure 2.

The study included two main loan treatment arms: a *cash loan treatment* and a *maize loan treatment*, both offered at the start of the hungry season (January). Repayment was due at harvest (July), and loans could be repaid in either cash or maize (or both). The two treatment arms present trade-offs. On the one hand, providing the staple food offers a direct way of targeting food shortages. On the other hand, cash offers more flexibility to address non-food consumption needs, though it may be more prone to wasteful consumption than maize. In year 1, both treatment arms were rolled out in January. Of the 175 study villages, 58 (1,033 farms) were assigned to a control group, which received no intervention, 58 (1,092 farms) were assigned to the cash loan treatment, and 59 (1,095 farms) were assigned to a maize loan treatment. In the second year of the program, the treatment groups were rotated: 20 villages that were in the control group in year 1 were rotated to either the maize loan or cash loan treatment arms (10 each), and 29 cash loan villages and 28 maize loan villages were rotated to the control group. Treatment rotation was designed to investigate the persistence of the results for villages phased out after one

¹⁹ Villages were ineligible if (i) other projects had been conducted there in the recent past, (ii) the village bordered a village that was in the study pilot, (iii) the village bordered a village already listed, (iv) the village had fewer than 17 households, or (v) it was impossible to get a four-wheel drive vehicle within a 5km radius of the village during rainy season. Out of an initial list of 201 eligible villages, 25 were eliminated for a failure to meet one or more of the eligibility criteria that had been overlooked during the screening process. In addition, one village refused to participate in the baseline survey.

²⁰ We also restricted our sample to households with at least 2 acres 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. Together, the land size restrictions excluded less than 0.5 percent of households.

²¹ The majority of households sampled but not interviewed either had moved away from the village ($N = 219$) or turned out to be ineligible because their plots were too small or too large to meet our inclusion criteria ($N = 146$).

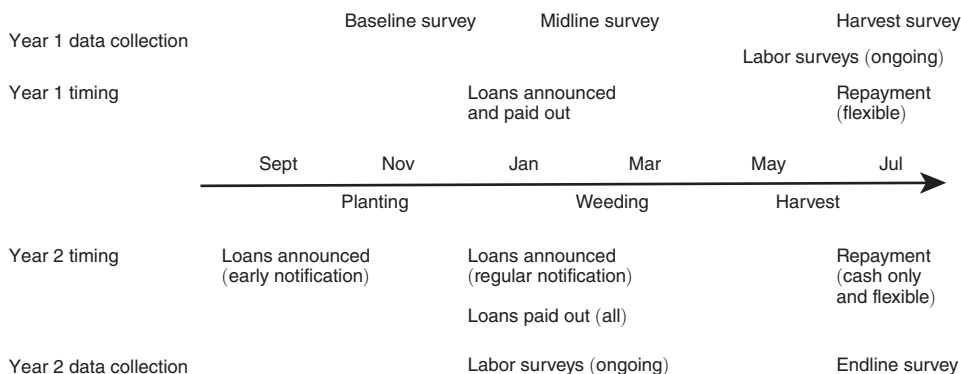


FIGURE 1. STUDY IMPLEMENTATION

Note: Time line of study implementation, including data collection, over the agricultural cropping calendar.

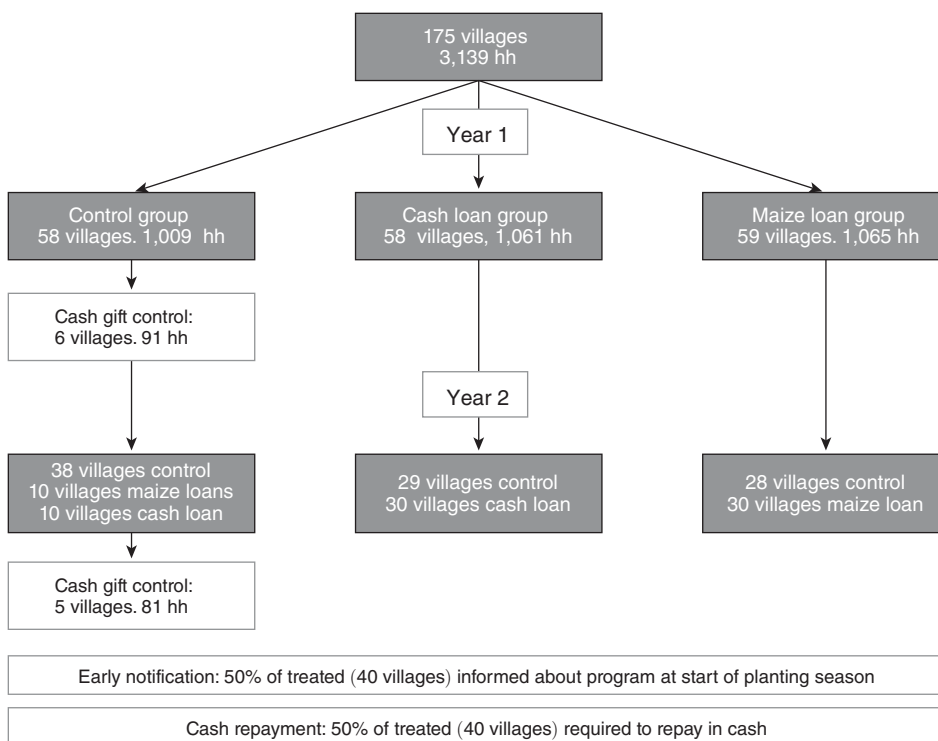


FIGURE 2. STUDY DESIGN

Note: Year 1 and year 2 treatments, all randomized at the village level.

year, and to separate the impact of repeated treatment from first time treatment. We also introduced additional variation in the control group to rule out income effects (a small cash grant in 6 and 5 villages in years 1 and 2, respectively), and announced the loan at planting time in selected treatment villages in year 2.

Details of the Cash and Maize Loans.—In the *maize loan treatment* arm, households were offered three 50 kilogram bags of unpounded maize. Maize is the staple crop in Zambia and 150 kilograms provides enough grain for a family of five to cover its basic consumption needs for at least two months during the peak hungry season. In the *cash loan treatment* arm, households were offered 200 Kwacha (US\$33), which corresponded approximately to the value of the three maize bags at official government prices (65 Kwacha per bag) at baseline.²² In both treatment arms, repayment was due in July when most harvest activities were completed. In the first year of the program, households could repay either 4 bags of maize or 260 Kwacha in cash (or a mix at K65 per bag). Villages randomly selected from both the maize and cash treatment arms for “cash only” repayment in the second year of the study had to repay 260 Kwacha.²³ While both treatment arms were designed to reflect an interest rate of about 30 percent over 5 months (or a roughly 4.5 percent monthly interest rate), actual interest rates are hard to calculate *ex post* for the maize loan due to substantial regional and seasonal fluctuations in grain prices, and limited information on the transaction costs associated with buying and selling maize locally. As shown in online Appendix Table B.1, interest rates in the maize arm vary between –11 and 33 percent depending on which maize price is used in the calculation.

C. Implementation

Both loan treatments were administered by Innovations for Poverty Action (IPA) under the Chipata Loan Project (CLP) to distinguish loan operations from the survey visits, also conducted by IPA.²⁴ This distinction between the CLP and IPA brands and staffing was intended to assure participants that survey responses would not affect loan eligibility. All study households in villages randomly selected for treatment (one-half of the village population on average) were eligible for loans in the first year. In year 2, the same rules applied. In villages treated in both years, eligibility was also restricted to households who fully repaid in year 1. In both years, the loan amount was fixed for logistical reasons; larger or smaller loan amounts were not available.

The loan intervention was announced to households during a village meeting to which eligible households were invited.²⁵ At the meeting, project staff described the terms of the loan and logistics surrounding distribution and repayment.

Loans were distributed between three days and one week after the village meeting at a location convenient for transportation. Project staff confirmed recipient

²²To make the two loan programs as comparable as possible, we conducted a series of hypothetical choice experiments in villages outside of the study sample in November 2013. Responses to these questions determined final loan contract details. Further detail on the implementation of the choice experiments is provided in online Appendix Section C.2.

²³Requiring cash repayment was tested in the second year for programmatic reasons, to see if administration costs could be reduced without affecting program impacts. We observe no effect from this variant in repayment requirements on take-up, repayment, or any of our main outcomes. It is, however, cheaper to administer.

²⁴Repayment risk for the loan program overall was fully borne by the project: substantial default would have led to cancellation of year 2 of implementation and/or the collection of the endline survey.

²⁵Eligible households could send an adult representative if the household head was not available to attend. All village headmen were eligible for the loan, even if they were not sampled for the baseline survey (and are therefore not in our study sample).

identities using national registration cards,²⁶ collected signed enrollment forms, and handed over the cash or maize. Before finalizing the transaction, project staff confirmed that the participant understood the terms of the loan. Repayment was due in early July. Villages were notified in advance about the date of repayment as well as the locations at which repayment would be collected, which were either within the village or at the closest point accessible by a 10-ton truck. Two attempts at collecting repayment were made. Throughout the project, households were told that the program might or might not continue in future years, which accurately represented the study team's knowledge.

Randomization.—All treatments were randomly assigned at the village (cluster) level ($N = 175$). In year 1, villages were divided into three equally sized arms: a control arm, a cash loan arm, and a maize loan arm. Treatments were assigned at the village level using min-max T randomization (Bruhn and McKenzie 2009), checking balance on both household and village characteristics. 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 household level baseline variables, village size, and geographic block dummies, with results described in Section IVA. The smallest p -value for the pairwise comparisons observed in the final draw was $p = 0.213$. In year 2, we randomly selected 50 percent of villages treated in year 1 for program continuation, and phased out loan treatments in the other 50 percent. At the same time, 35 percent of villages in the control group (20 villages) in year 1 were randomly selected for the program in year 2. Treatment assignment in year 2 was balanced on the same variables plus harvest output from year 1, and stratified by year 1 treatment. In other words, year 2 treatment assignment was carried out within each year 1 treatment arm. Within each treatment arm, villages were randomly assigned to both the main treatment arms (control, cash loan, and maize loan) and the subtreatments (income effect control, early notification, and cash repayment).

Attrition and Selection.—Online Appendix Table B.2 reports the number of households sampled in each survey round, and the probability of being surveyed as a function of treatment. The coefficients and standard errors are from OLS regressions for each survey round, with errors clustered at the village level. Overall, attrition rates are low: 3,030 out of the 3,139 households (96.5 percent) enrolled at baseline completed the endline survey. We do not find any differences in attrition overall or in the probability of participating in specific survey rounds across treatment arms.

We also examine whether household self-selection into the program varied by treatment. Online Appendix Table B.3 shows the stages of program implementation. First, households were invited to participate in the village meeting based on random sampling (year 1). To be eligible for borrowing, households had to both attend the meeting and hand in a consent form. The latter step was completed after learning treatment status and so is the most susceptible to nonrandom attrition (column 3). In year 1, there was no selection into meeting attendance or eligibility. In year 2, there

²⁶In select cases, a household representative picked up the loan. In these cases, the representative needed to carry the loan-holder's NRC card with him or her.

was some modest selection into meeting attendance (over 90 percent attendance in all treatments and subtreatments), and no further selection into eligibility. Column 4 of online Appendix Table B.3 also previews our take-up results, which we describe in Section IV.

III. Data and Descriptive Statistics

We start this section with further description of the data and our main outcome variables. Then we turn to a set of descriptive results that provide contextual information and validate some of our design choices.

A. Data and Measurement

We rely on both household survey and administrative loan data in our analysis. Comprehensive surveys of all study households were conducted at baseline (November–December 2013), harvest of year 1 (July–August 2014), and harvest of year 2 (July–August 2015). We refer to these as long recall surveys since they ask questions about the preceding agricultural cycle. Surveys on labor activities, consumption, and farming practices were administered to an ongoing rolling sample. For these surveys, smaller survey teams visited a random sample of villages each week. We refer to these as short recall surveys since they primarily ask about activities in the past two days to two weeks. Sample sizes for each survey round are provided in Table 1, including the number of observations recorded during the hungry season, which is the focus of much of our analysis. During the hungry season of both years, the short recall surveys deliver a total sample size of 3,732 household interviews. Sample sizes for some outcomes vary because of selective refusals (<5 observations in most cases) and because of occasional differences in measurement across short recall survey rounds. Across all survey rounds, a total of 15,044 surveys were conducted with the 3,139 study households. Online Appendix Section C.1 provides further detail on the content of each survey.

Outcome Measures.—We focus on three main outcome types, based on the predictions in our conceptual framework:²⁷ (i) labor allocation and daily earnings, (ii) agricultural output, and (iii) consumption. In many cases, we focus on data collected during the hungry season (January–March) of each year.

We rely on the short recall survey rounds during the hungry season to construct labor allocation measures over the week prior to the survey. The casual labor market transactions that we observe are referred to locally as *ganyu*, as described in greater detail in Section IIIB. Labor allocation outcomes include (i) family labor sold to other farms (*ganyu* sold), (ii) labor purchased (*ganyu* hired), and (iii) family labor invested on-farm. We construct both extensive margin measures and continuous measures at the household level, and winsorize the continuous variables at the ninety-ninth percentile to address unrealistic outliers. To match the predictions of our model, we also construct measures of total household labor supply (family labor

²⁷Our data collection and analysis follows a pre-analysis plan available at <https://www.socialscisceregistry.org/trials/130/>. Data and code are deposited at ICPSR (Fink, Jack, and Masiye 2020).

TABLE 1—SURVEY OVERVIEW

Survey round	Dates	Observations	Hungry season observations
Baseline	Nov 2013–Dec 2013	3,139	0
Harvest	Jul 2014–Sep 2014	3,028	0
Endline	Jul 2015–Sep 2015	3,005	0
Midline	Feb 2014–Apr 2014	1,193	1,190
Labor R1	Jan 2014–Jul 2014	1,276	778
Labor R2	Jul 2014–Jan 2015	1,333	376
Labor R3	Jan 2015–Mar 2015	1,388	1,388
Labor R4	Apr 2015–Jun 2015	680	0
Total		15,042	4,412

Notes: Sample sizes by survey round. Hungry season observations are recorded during January, February, and March. See online Appendix Section C.1 for further description of survey content and implementation.

invested on-farm plus ganyu sold) and total household labor demand (ganyu hired plus family labor invested on-farm).

To measure wages, we construct a measure of daily earnings during the hungry season. We again use the short recall surveys, which ask respondents about earnings from ganyu sold by each household member over the past week, and calculate average daily earnings at the household level based on days worked and total earnings. We winsorize the top 1 or 5 percent of household-level average daily earnings by year and treatment and analyze daily earnings, as our proxy for local wages, first at the household and then at the village level, where we further reduce noise by focusing on mean earnings within the village. Households or villages with no ganyu activities reported receive a missing value.

To measure agricultural output, farmers were asked to report, by crop, output in kilograms as well as the total value of the harvest, including early consumption and crops still on the field at the time of the interview. We aggregate the total value across all crops, and calculate a constant price series to remove fluctuations in crop value across survey rounds in our main specification. We also construct a measure based on own reported prices to allow for the possibility that treatment could affect output prices through increased effort to market output.

Our main consumption outcome is the number of meals consumed in a day by adult members of the household, measured during our short recall data collection rounds. While this is a coarse measure of consumption, reductions in the number of meals per day point to severe food shortages, and are relatively easy to measure. We collected this outcome over a two-day recall period. We supplement the measure of meals consumed with a count of the months during which the household reports having enough food to cover consumption needs, measured during the long-recall surveys. In addition, we collect data on households' perceived food security and construct an index of z -scores based on responses in the control group.

In addition to these main variables, we analyze a few other auxiliary outcomes in Section V, which we describe as they arise.

Heterogeneity Measures.—We categorize households by their liquid resources using baseline reserves of grain (valued in Kwacha) and cash. Conceptually, this is

intended to represent the liquid resources available to the household at the beginning of period 1 in the model, S_{i0} . Baseline survey data collection coincided with planting, so this measure is net of early season investment and consumption decisions since the previous harvest. Since the primary role of S_{i0} in the model is to generate variation in hungry season interest rates, we also show results based on heterogeneity in reported interest rates on borrowing at baseline. Of course, numerous other factors may influence households' cost of borrowing during the hungry season, and interest rates are likely to be reported with considerable error.²⁸ With this in mind, we rely on baseline resources as our main heterogeneity measure, and use reported interest rates at baseline to examine the robustness of our findings.

B. Descriptive Statistics

Seasonal Variation in Resources and Consumption.—As illustrated in Figure 3, which uses data from the control group, households draw down their cash and grain reserves after harvest, and begin to replenish them in April and May. The period with the lowest reserves (top-left panel) and lowest food consumption (top-right panel), between January and March, is referred to as the “hungry season” throughout rural Zambia.²⁹ This shortage of resources during the hungry season is anticipated by farmers: at baseline, 76 percent of households did not expect their maize reserves to last until the next harvest, and most expected to run out of maize in January or February. This period of restricted consumption is also the time when farmers have crops on their fields and on-field activities (particularly weeding) peak, as illustrated in the agricultural calendar shown in Figure 1 and the bottom-left panel of Figure 3, which shows total labor inputs on-farm. Total labor inputs on-farm are lowest after harvest in July and August, and then slowly increase to peak with weeding in January; after a slower period in March and April, a second intense labor period occurs during harvest.³⁰ In spite of the high on-farm demand in January, the same households show a peak in family labor sold in local markets in January and February (lower-right panel of Figure 3).³¹

²⁸We prespecified analyzing heterogeneous treatment effects by baseline grain and cash resources but not by baseline reported interest rates. We also prespecified analyzing heterogeneous responses to labor availability measured as the baseline ratio of workers per acre. However, additional information about local land markets suggests that this is not a good measure of labor constraints (see Section IIIB), since the allocation of land depends both on productivity and favoritism.

²⁹In online Appendix Figure B.1 we use data from the Demographic and Health Surveys (CSO 2002, 2007, 2014) to provide suggestive evidence of the health and human capital costs of these fluctuations across Zambia. An estimated 19 percent of children under the age of five were estimated to be underweight (weight-for-age z -score < -2) in the post-harvest period. During the hungry season, this number increases to 25 percent, a roughly 25 percent increase relative to harvest months. While sampling is not representative by month, we find no evidence that adult height displays a similar pattern.

³⁰Increases in family labor sales are not driven by seasonality in wages: in the control group, daily earnings are low during the hungry season relative to planting and harvest time.

³¹Grain markets are also highly seasonal in rural Zambia (see online Appendix Figure B.2) with prices falling after harvest when most households sell grain and increasing during the hungry season. Consequently, we avoid assigning monetary values to grain reserves or to consumption, except in constructing our main heterogeneity variable (measured at baseline) and in constructing Figure 3 (which uses baseline prices). Unlike labor markets, grain markets encompass multiple villages, and our study was not designed to impact equilibrium maize prices (see Burke, Bergquist, and Miguel 2019 for an example of market-level interventions to smooth seasonal grain price fluctuations).

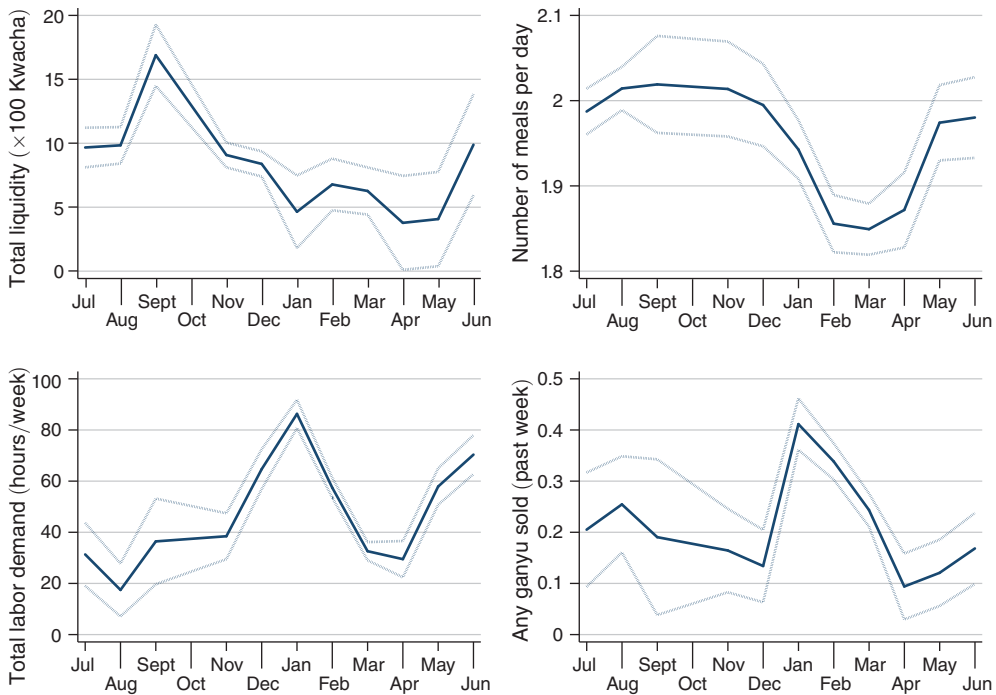


FIGURE 3. SEASONAL VARIATION IN LIQUIDITY, LABOR, AND CONSUMPTION

Notes: Figures plot means and 95 percent confidence intervals by month for control group households, using data from short recall surveys. The top-left figure shows total liquidity for control group households in the study sample, where liquidity is the total cash value of grain and cash holdings at the time of survey. The top-right figure shows number of meals per day consumed by adults in control group households in the two days preceding the survey. The bottom-left figure shows total hours of labor (family and hired) on-farm in the week preceding the survey. The bottom-right figure shows the share of households that report selling any family labor to the casual labor market in the week preceding the survey.

The loans we offer were thus designed to inject resources during the most constrained period, from January to March, with repayment at a time of relative abundance, in July.

Savings, Credit, and Effective Interest Rates.—As in many rural developing country settings, access to formal savings and credit markets is limited in rural Zambia. At baseline, only 5.6 percent of households report saving in a bank; slightly more (9.1 percent) report saving with friends, family, or employers. By far the most common savings strategy, reported by 76.7 percent of households, is saving money at home. The median self reported cash savings (a measure likely to be reported with substantial error) at the start of the planting season of the first year was 80 Kwacha or around US\$14. Non-cash savings also occurs through grain storage, typically in a bamboo (62 percent of respondents) or thatch (28 percent of respondents) granary. The median grain storage amount at baseline was four bags, which would meet the maize needs of a typical family of five for about three months. Sixty percent of households report storage losses in the past season.

Credit market access is also limited: only 5 percent of household respondents at baseline report accessing formal cash loans from banks, credit unions, NGOs, or government sources. Informal borrowing channels are slightly more common: around 7 percent of baseline respondents report taking high interest loans from moneylenders, locally referred to as *kaloba*. Informal loans from friends and family are reported by around 8.5 percent of baseline respondents, though reported interest rates on these are also high (around 30 percent per month, measured at end-line). Transfers between households are common at planting and harvest time, when around 40 percent of households report recent transfers. During the hungry season, this number drops to 15 percent.

To measure interest rates, we ask households at baseline how much they would have to repay in one month if they borrowed 50 Kwacha that day from a source other than a friend or family member. In the hungry season, 60 percent of control group households report that they would be unable to borrow 50 Kwacha in cash that day. The implied monthly interest rate is around 40 percent for households with low baseline grain and cash reserves; measured interest rates decline slightly with baseline reserves for households with above-median reserves, though rates are still high (around 34 percent) for even the best-off of our study sample (see online Appendix Figure B.3). We allow for this nonlinear relationship in our analysis of heterogeneous treatment effects.

Participation in microcredit institutions, rotating savings and credit associations (ROSCAs), and village savings and loan associations (VSLAs) are each reported by around 1 percent of baseline respondents. In-kind input loans are more common than opportunities for borrowing in cash or food: 40 percent of baseline respondents accessed an in-kind input loan, typically seeds and chemicals provided by outgrower companies or agro-dealers.

Local Labor Markets.—Local wage earning opportunities for study households are defined largely by piece-wise labor contracts locally referred to as *ganyu*. In focus groups, respondents described *ganyu* both as the most common strategy to cope with temporary cash needs and as an activity most farmers would rather avoid if possible.³² In the baseline survey, the most common response to why an individual in the household worked *ganyu* during the previous agricultural season was to obtain food. The second most common reason was to access cash for a personal purchase, and the third was to deal with an emergency. When asked what the household would do in the coming year if they ran out of food, 56 percent reported that they would do *ganyu*. The next most common answers included borrowing from friends or family (28 percent), using savings (22 percent), and selling assets or livestock (17 percent), all of which may be difficult during the hungry season. Typical *ganyu* contracts specify relatively small tasks (such as weeding an acre of land) that can be completed by an individual within a working day; larger contracts are typically completed by families or groups of individuals. Further detail on how *ganyu*

³²In our baseline survey, around 90 percent of households disagreed with the statements “Doing *ganyu* increases people’s respect for you in the community” and “Successful farmers do lots of *ganyu*.” Around 60 percent of households agreed with the statements “Lazy people do lots of *ganyu*” and “People who can’t budget do lots of *ganyu*.”

participation varies with age and gender of household members is discussed in Fink, Jack, and Masiye (2014), which describes a pilot for the current study.³³

Almost two-thirds of farmers in our sample reported having engaged in ganyu activities in the previous season at baseline. This number is decreasing in baseline grain and cash resources. Online Appendix Table B.4 shows how hours of ganyu sold over the previous week (and other baseline variables) vary across quartiles of the distribution of baseline resources: 75 percent of households in the bottom quartile did ganyu the previous season while only 46 percent of households in the top quartile did.³⁴

At baseline, around two-thirds of households anticipated doing ganyu in the coming agricultural season. These expectations appear reasonably accurate. Among control group households that predicted at baseline that they would have to do ganyu in the coming year, around 76 percent did; among those that predicted not doing ganyu, around 41 percent ended up working off-farm. At the same time, the likelihood that a household sells ganyu is not constant across years. Among control group households that did not engage in ganyu the year before the study, 40 percent sold ganyu the following year.

Ganyu wage rates are typically negotiated on a case-by-case basis. The majority of casual labor transactions take place in or near the worker's own village, which may be explained by low population density and a general absence of motorized transport. The majority (>80 percent) of farms hiring ganyu are small (i.e., fewer than 5 hectares of land), with some farms acting as both buyers and sellers during a single season (though typically at different points in the season). Given that most labor transactions happen within a village, each rural village can be thought of as its own labor market, where wages are determined endogenously.

Seasonal migration is uncommon in rural Zambia. In our sample, in any given month, only around 3 percent of households report that someone who is typically a member of the household moved away temporarily. This number peaks around harvest time, and is lowest during the hungry season when around 2 percent of households report temporary migration. Permanent migration is more common, as suggested by data on remittances: around 20 percent of households report that someone who does not live in the village contributes regularly to household income.

Local Land Markets.—Land is relatively abundant in Zambia, with a population density of 22 inhabitants per square kilometer (World Bank 2017). Land markets for sale or lease are largely absent and most tenure is customary with annual land allocations by the village head person. Anecdotally, these allocations are determined by a complex set of factors, including past production, ability to use the land, and favoritism. At baseline, the average acreage reported available for cultivation was 6.3, but only 4.5 acres of land were used on average for growing crops. In general, most farmers do not appear to view land access as a constraint to production. At

³³ Additional analysis of patterns of labor allocation by gender and the relationship between intrahousehold decision-making measures and outcomes for this intervention can be found in Hausdorff (2016).

³⁴ Conditional on working, households sold an average of 12 person-hours of ganyu labor per week during the hungry season. These numbers ignore the substantial additional time burden associated with searching for ganyu. Surveys during the hungry season of year 1 of the project indicate that control group households spend around 3 hours searching for ganyu each week, conditional on searching.

baseline, over 60 percent of respondents said that they could have farmed more land than they did in the previous season if they had more inputs. Given that land appears to be a relatively disposable production factor, we do not model it explicitly. In our analysis, we focus on total output to measure agricultural productivity, and condition on baseline land size. Normalizing output by field size would add noise to the estimation given substantial measurement error in reported land size.

IV. Experimental Results

In this section, we present our main experimental results. We start by outlining our empirical strategy and then present results on take-up and repayment, labor allocation, agricultural output, and consumption.

A. Empirical Strategy

We estimate intention-to-treat regressions. High take-up in both years means that these estimates are very close to the treatment on the treated effect. For our main outcomes, we first estimate separate regressions for each study year, pooling across loan treatment arms. For year 1, we estimate

$$(7) \quad y_{ivt} = \alpha + \beta \text{loan}_{vt} + \tau_t + X_{i0} \delta + u_{ivt},$$

where y_{ivt} is an outcome of interest for household i located in village v at time t , loan_{vt} indicates that the village was assigned to one of the two loan treatments, τ_t are month fixed effects, and X_{i0} are household-level controls, measured at baseline. Errors are clustered at the level of the randomization unit, the village v , which addresses both unobserved village level shocks and correlation across study years.

Treatment assignment is rotated between years, as described in Section IIB, and we allow treatment effects in year 2 to differ by year 1 treatment status by estimating

$$(8) \quad y_{ivt} = \alpha + \beta_1 \text{loan}_{vt} + \beta_2 \text{loan}_{vt-1} + \beta_3 \text{loan}_{vt} \times \text{loan}_{vt-1} + \tau_t + X_{i0} \delta + u_{ivt},$$

where loan_{vt} indicates a loan treatment in year 2, and loan_{vt-1} captures year 1 loan treatment status. To facilitate interpretation of the interacted model, we also estimate total year 2 treatment effects on farmers selected for the program in both years ($\beta_1 + \beta_2 + \beta_3$). We also show separate results by treatment arm, pooling across years, using estimating equation (8) and allowing effects to differ for the cash and maize loans. We report the average marginal effects, pooled across the two years, relative to the pure control group.

Our model predicts that treatment effects will vary with interest rates and the household's available liquid resources. We test this by estimating heterogeneous treatment effects by baseline grain and cash reserves using the following model:

$$(9) \quad y_{ivt} = \alpha + \beta_1 \text{loan}_{vt} + \beta_2 S_i + \beta_3 S_i^2 + \beta_4 \text{loan}_{vt} \times S_i + \beta_5 \text{loan}_{vt} \times S_i^2 + \tau_t + u_{ivt},$$

where S are liquid resources as discussed in Section IIIA. We summarize the results by plotting adjusted predictions evaluated at the mean S in each quartile of the

baseline distribution.³⁵ Standard errors for these adjusted predictions are calculated using the delta method. We examine heterogeneous treatment effects for each of our main outcome measures.

Our regression estimates identify the causal effect of the loan under the identifying assumption that treatment assignment is orthogonal to u_{ivt} . Online Appendix Table B.5 presents the means and standard deviations of baseline survey characteristics among study households, by treatment arm for years 1 and 2 (columns 1–3 and 6–8, respectively); p -values are shown for a test of equal means for each treatment group relative to the control, for each year of the study. Overall, the randomization successfully balanced households across treatment arms, with 5 out of 72 tests with $p > 0.05$. Online Appendix Table B.6 shows that the treatment rotation in year 2 similarly resulted in balanced samples across arms. The variables shown in online Appendix Table B.5, up through crop diversity, were used in the randomization to test balance and were specified as controls in the pre-analysis plan; the extended set of control variables are used throughout the analysis. Results are similar if only prespecified controls or if no controls are included in the analysis.

B. Take-Up and Repayment

Table 2 shows loan take-up, which was over 98 percent in both years, suggesting that the borrowing rates available through the intervention were well below those associated with comparable borrowing opportunities in local markets. High repayment rates (94 percent) in year 1, followed by high take-up rates in villages treated in both years, indicate that high take-up was not driven by expectations of default, and provide a strong revealed preference measure of the benefits of the loans.

Repayment was substantially lower in year 2, with an average repayment rate of 80 percent in villages receiving the program for the first time. This decline in repayment appears to be driven in part by worse rainfall (see online Appendix Figure B.4) and lower average agricultural output in 2015. In addition to differences in harvest, we also observed behavioral differences in villages treated for the second time in year 2, with a 6 percentage point decline in repayment rates in villages where nobody had previously defaulted, and a 29 percentage point decline in repayment in villages where at least one farmer had defaulted in year 1. The particularly large drop in repayment in villages with prior default may suggest some learning about enforcement, though it may also arise through other channels such as serially correlated shocks.

Overall, repayment rates in both years are very high compared to many micro-credit programs. Three features of the setting and implementation may have aided repayment. First, from the outset of the study, farmers expressed a strong desire to

³⁵ We omit baseline controls to allow for other household characteristics correlated with baseline grain and cash reserves to vary by quartile. We primarily focus on year 1 in these figures because of the less precisely estimated main effects in year 2 and because baseline measures of available resources continue to predict year 2 outcomes in many cases, but less strongly than for year 1. The set of figures for year 2 matching the estimates presented in this section (excluding villages treated in year 1) is shown in online Appendix Section B. To examine robustness of these results, we also show alternative estimates using reported interest rates at baseline instead of S_{i0} as the source of heterogeneity.

TABLE 2—TAKE-UP AND REPAYMENT

	Take-up (1)	Full repayment (2)	Percent repaid (3)	Repaid any cash (4)
<i>Panel A. Year 1</i>				
Cash loan mean	0.99	0.93	0.94	0.35
Maize loan	-0.003 (0.007)	0.013 (0.020)	0.015 (0.018)	-0.225 (0.035)
<i>Panel B. Year 2</i>				
Cash loan mean	0.98	0.76	0.78	0.55
Maize loan	-0.012 (0.015)	0.008 (0.058)	0.026 (0.052)	-0.084 (0.089)
Treated in year 1	0.018 (0.016)	-0.046 (0.069)	-0.062 (0.055)	-0.142 (0.098)
Early notification subtreatment	0.010 (0.016)	-0.019 (0.058)	-0.007 (0.052)	0.012 (0.090)
Cash repayment subtreatment	-0.008 (0.015)	-0.049 (0.058)	-0.044 (0.052)	0.624 (0.049)
<i>Panel C. Year 2, repeat treatment</i>				
Any default in village in year 1	-0.025 (0.026)	-0.220 (0.068)	-0.225 (0.066)	-0.221 (0.099)

Notes: Table shows take-up and repayment statistics in the cash loan treatment in year 1 (panel A), year 2 (panel B), and year 2 conditional on treatment in both years (panel C). Full repayment is a summary variable that equals 1 if the loan was fully repaid. The means for the cash loan treatment are shown in the first row of each panel and coefficients on a dummy for the maize loan treatment are in each subsequent row. In year two, each row corresponds to a separate regression. Treatment variables in panels A and B are randomly assigned; any default in the village in year 1 (panel C) is not. Standard errors are clustered at the village level.

reciprocate the assistance they received during the hungry season.³⁶ Second, loan collection was coordinated and supported by local head persons, who also received a small reward (one bag of maize) if repayment was over 90 percent. Finally, because the program was implemented by a research organization rather than a large-scale lending operation, greater diligence may have been paid to making the repayment process easy for farmers.

C. Impact of Lowering Borrowing Rates through Seasonal Loans

We test the empirical predictions described in Section IA.

Prediction 1: Labor Demand and Equilibrium Wages Increase.—We examine impacts on labor allocation in Table 3. Panel A shows results for year 1, panel B for year 2, and panel C estimates separate effects by treatment arm, pooling across the two years.

Our main prediction is that loans will increase labor demand among treated farmers. Additional on-farm labor can come from decreases in family labor sales (less ganyu sold), through an increase in hiring (more ganyu hired), or an increase

³⁶ A qualitative audit of the program conducted by IPA three years after the end of the program confirmed high levels of satisfaction: average scores on three questions of satisfaction and perceived program impact were all over 9 on a 10-point Likert scale.

TABLE 3—AVERAGE TREATMENT EFFECTS: LABOR

	Any ganyu sold (1)	Hours sold (2)	Any ganyu hired (3)	Hours hired (4)	Family hours on-farm (5)
<i>Panel A. Year 1: pooled treatment arms</i>					
Any loan treatment	-0.048 (0.026)	-1.137 (0.551)	0.039 (0.015)	2.003 (1.231)	4.953 (2.618)
<i>Panel B. Year 2: pooled treatment arms</i>					
Any loan treatment	-0.021 (0.042)	-0.799 (0.489)	-0.006 (0.030)	0.455 (1.507)	11.467 (5.658)
Treated in Y1	0.045 (0.036)	0.708 (0.520)	0.001 (0.026)	0.325 (1.098)	7.908 (3.827)
Loan × treated in Y1	-0.058 (0.051)	-0.605 (0.646)	0.020 (0.040)	-1.210 (1.765)	-14.367 (6.765)
Loan + Y1 + loan × Y1	-0.034 (0.033)	-0.696 (0.419)	0.015 (0.029)	-0.430 (1.061)	5.008 (4.194)
<i>Panel C. By treatment arm: pooled years</i>					
Cash	-0.040 (0.021)	-0.966 (0.360)	0.027 (0.015)	0.515 (0.861)	2.712 (2.474)
Maize	-0.052 (0.021)	-1.169 (0.407)	0.018 (0.015)	-0.131 (0.925)	4.597 (2.373)
Year 1 control mean	0.34	4.55	0.08	3.03	46.57
Year 2 control mean	0.27	2.75	0.14	3.73	54.38
Year 1 = year 2 new	0.21	0.20	0.04	0.44	0.88
Year 1 = year 2 repeat	0.35	0.20	0.15	0.22	0.24
Cash = maize	0.57	0.59	0.61	0.53	0.51
Observations	3,729	3,728	3,732	2,542	3,731

Notes: Treatment effects on labor outcomes, measured during the short recall surveys during the hungry season (3,732 observations). Hours of labor hired (column 4) was not collected during the midline survey. Extensive margin outcomes (columns 1 and 3) indicate whether the household engaged in any of the labor activity over the past week. Other columns show the number of hours allocated to each activity, include zeros, and are winsorized at the ninety-ninth percentile. All specifications are conditional on month fixed effects, include baseline controls, and cluster standard errors at the village level. Panel C also conditions on year. Panel B also shows the total effect for repeat treatment in year 2, relative to the year 2 control group: loan + Y1 + loan × Y1.

in family labor supply if improvements in first period consumption relax constraints on family labor supply. We test for these responses using data from the short-recall surveys during the hungry season, where households report on labor activities in the previous week. Family labor is aggregated over the entire household. Column 1 shows that the likelihood of selling ganyu falls by around 4.8 percentage points in year 1 ($p = 0.066$, panel A), on average, off of a mean of 34 percent in the control group. The effects are negative but imprecisely estimated in year 2 ($p = 0.626$ for first time treatment and $p = 0.316$ for repeat treatment, panel B) and similar in the maize arm and the cash arm (panel C). Turning to the continuous measure of total family hours sold (column 2), treated households sell an average of 1.14 fewer hours (25 percent) of ganyu per week in year 1 ($p = 0.041$, panel A), again with smaller and less precise effects in year 2 ($p = 0.104$ for first time treatment and 0.099 for repeat treatment, panel B). The effect on hours sold is also similar in the maize loan treatment arm and the cash arm (panel C). The differences between years and between treatment arms are not statistically significant at conventional levels, as shown by the p -values reported at the bottom of the table.

Treatment has positive effects on hiring, on both the extensive margin (column 3) and the continuous measure (column 4, collected only for a subset of households in an extended hungry season survey module) in year 1 (panel A). The likelihood of any hiring increases by 3.9 percentage points, a 50 percent increase relative to the control ($p = 0.009$), and the number of hours hired increases by 2.0 per week, a 67 percent increase in average hours hired relative to the control group ($p = 0.107$). In year 2, ganyu hiring is more common in the control group than in year 1, and effects on the loan treatment on hiring are inconsistently signed, with large standard errors (panel B). Hiring effects are similar in the cash loan arm and the maize loan arm (panel C).

As shown in column 5, hours of family labor invested on the farm over the past week increase by 4.95 hours or around 11 percent in year 1 ($p = 0.060$, panel A). In year 2, the increase in family hours on-farm is comparable for repeat treatment farmers (5.0 hours, $p = 0.234$), and even larger among newly treated farms (11.5 hours, $p = 0.040$). Increases in family labor on-farm are precisely estimated in the maize treatment arm only. These increases in family labor are consistent with consumption-related constraints on initial labor supply. We show increases in first period consumption associated with prediction 3, and provide a further discussion of these results in Section V.

Table 4 shows effects on wages; since daily or hourly wages are not defined in the piece rate work arrangements that we observe, we use reported daily earnings as our proxy for local wages. Columns 1 and 2 show regression results using household level data, winsorizing the top 1 and 5 percent of observations within treatment and year, respectively, and controlling for the hours worked in the day. Earnings increase by 2.91 or 2.52 Kwacha per day on average, corresponding to a 17 to 19 percent increase over the control group mean in year 1 ($p = 0.116$ and $p = 0.084$, panel A). In year 2, increases are of similar magnitude but with larger standard errors (panel B). Increases are slightly larger in the maize arm than the cash arm in columns 1 and 2 (panel C), where the decrease in off farm labor sales was also higher. Column 3 shows village-level regressions using mean reported daily earnings at the village-month level, based on the household measure winsorized at the ninety-fifth percentile. The magnitudes line up reasonably well with the household level data, particularly in panels A and C.

We measure general equilibrium wage effects by comparing reported earnings between treated and control households. However, as shown in Table 3, the likelihood of supplying labor to the labor market also decreases with the loan treatment, changing the composition of the observed labor force. To control for these selection effects, we estimate hypothetical impacts under different selection assumptions. Following Lee (2009), we estimate bounds by trimming the proportion of the treatment group not observed in the daily earnings data in the control group, first from the bottom and then from the top of the distribution, and re-estimate treatment effects in the trimmed sample. We compute bootstrapped standard errors, and report the resulting treatment bounds in columns 4 and 5 of Table 4. Column 4 shows the lower bound on the effect, which is small, positive, and imprecise in year 1 ($p = 0.514$, panel A), as well as in villages treated for the first time in year 2 ($p = 0.716$, panel B). In villages treated for the second time in year 2 (panel C), the lower bound on the treatment effect is negative ($p = 0.206$). Column 5 reports the upper bound on the

TABLE 4—AVERAGE TREATMENT EFFECTS: DAILY EARNINGS

	Individual-level daily earnings		Village mean daily earnings (3)	Treatment bounds	
	(winsorize 1%) (1)	(winsorize 5%) (2)		Lower (4)	Upper (5)
<i>Panel A. Year 1: pooled treatment arms</i>					
Any loan treatment	2.913 (1.844)	2.522 (1.448)	2.480 (1.621)	1.127 (1.541)	5.908 (1.859)
<i>Panel B. Year 2: pooled treatment arms</i>					
Any loan treatment	3.060 (2.495)	2.080 (2.243)	4.253 (2.577)	0.994 (2.528)	4.186 (2.592)
Treated in Y1	2.682 (1.848)	1.158 (1.588)	2.302 (1.633)		
Loan × treated in Y1	−1.885 (3.480)	−1.545 (2.916)	−3.525 (3.065)		
Loan + Y1 + loan × Y1	3.858 (2.084)	1.693 (1.709)	3.030 (1.810)	−2.645 (2.015)	6.683 (2.262)
<i>Panel C. By treatment arm: pooled years</i>					
Cash	1.866 (1.609)	1.366 (1.253)	2.288 (1.467)		
Maize	3.219 (1.571)	2.510 (1.280)	2.232 (1.266)		
Year 1 control mean	15.58	14.81	15.18		
Year 2 control mean	14.92	14.92	13.84		
Year 1 = year 2 new	0.30	0.28	0.81		
Year 1 = year 2 repeat	0.37	0.16	0.52		
Cash = maize	0.48	0.45	0.97		
Observations	1,083	1,083	402		

Notes: Treatment effects on reported daily earnings in Kwacha, measured during the short recall surveys. The number of observations is the number of households reporting any ganyu during the hungry season (columns 1 and 2) or the number of village-months with non-missing ganyu observations (column 3). Column 1 winsorizes the top 1 percent of daily earnings; column 2 winsorizes the top 5 percent. Column 3 averages the individual level measures in column 2 at the village-month level. Columns 1–3 include a village size control, geographic controls, month fixed effects, and cluster standard errors at the village level. Columns 1 and 2 include household controls and the hours worked per day. Panel C also conditions on year. Panel B also shows the total effect for repeat treatment in year 2, relative to the year 2 control group: loan + Y1 + loan × Y1. Bounds on the estimated treatment effects are presented in columns 4 and 5, following Lee (2009). The bounds follow column 2, but omit household and geographic controls other than the selection variable: any reported ganyu over the past week.

treatment effect, which is positive in panels A ($p = 0.002$) and C ($p = 0.004$).³⁷ This bounding exercise suggests that at least part of the adjustment in daily earnings that we observe is due to actual equilibrium shifts in wages, particularly in year 1. As an additional manipulation check, and given that equilibrium effects should be proportional to treatment intensity, we show estimated earnings impacts by the share of the village eligible for treatment in online Appendix Figure B.5. The estimates show a positive gradient, indicating that a higher treatment intensity is associated with larger general equilibrium effects. The slope on share of village treated should not be interpreted causally: larger villages, which differ from smaller villages on numerous dimensions, will have a smaller share of the village treated, given our

³⁷The upper bound in panel B lies below the estimated effect in column 2 because the regression analysis includes controls while the bounding analysis does not. In this case, inclusion of the hours worked per day affects the magnitude of the effect reported in column 2.

TABLE 5—AVERAGE TREATMENT EFFECTS: AGRICULTURAL PRODUCTION

	Log harvest value (1)	Log harvest value constant prices (2)	Total input value (3)	Acres cash crops (4)
<i>Panel A. Year 1: pooled treatment arms</i>				
Any loan treatment	0.086 (0.044)	0.087 (0.040)	67.929 (52.688)	0.043 (0.054)
<i>Panel B. Year 2: pooled treatment arms</i>				
Any loan treatment	0.064 (0.092)	0.081 (0.091)	184.591 (149.005)	0.186 (0.129)
Treated in Y1	0.052 (0.070)	0.068 (0.067)	-30.387 (95.214)	0.096 (0.088)
Loan × treated in Y1	-0.110 (0.107)	-0.132 (0.104)	-204.837 (161.456)	-0.108 (0.150)
Loan + Y1 + loan × Y1	0.006 (0.069)	0.017 (0.068)	-50.634 (88.849)	0.173 (0.091)
<i>Panel C. By treatment arm: pooled years</i>				
Cash	0.057 (0.038)	0.053 (0.035)	61.280 (50.858)	0.059 (0.053)
Maize	0.039 (0.040)	0.039 (0.038)	67.960 (58.729)	0.095 (0.060)
Year 1 control mean	3,534.93	3,640.44	816.71	0.92
Year 2 control mean	3,349.13	3,240.25	1,311.82	0.97
Year 1 = year 2 new	0.71	0.57	0.13	0.14
Year 1 = year 2 repeat	0.64	0.55	0.07	0.35
Cash = maize	0.69	0.74	0.92	0.61
Observations	5,988	5,987	6,015	6,033

Notes: Treatment effects on agricultural output and inputs. Columns 1–3 are measured in Zambian Kwacha during the long recall surveys covering the past season's agricultural production (6,033 observations). All specifications include baseline controls, and cluster standard errors at the village level. Panel C also conditions on year. Panel B also shows the total effect for repeat treatment in year 2, relative to the year 2 control group: loan + Y1 + loan × Y1.

sampling strategy. All analysis of impacts on daily earnings control for village size. In addition, controlling for important village-level labor market determinants, such as distance to road, does not substantially alter the result.

Prediction 2: Agricultural Output Increases.—We report effects on agricultural production in Table 5, including both output and inputs. In year 1 (panel A), agricultural output (as measured by total harvest value using reported (column 1) and constant (column 2) prices) increased by around 0.086 to 0.087 log points on average ($p = 0.050$ and $p = 0.030$). Villages treated for the first time in year 2 show effects of a similar magnitude, though with larger standard errors ($p = 0.487$ and $p = 0.375$), and the persistent effect of having been treated in year 1 on output in year 2 is also positive and imprecisely estimated ($p = 0.457$ and $p = 0.307$, panel B). The coefficient on repeat treatment is large, negative, and very imprecisely estimated; together, the effect of repeat treatment is indistinguishable from the pure control group, which is never treated ($p = 0.931$ and $p = 0.801$). This pattern closely mirrors the year 2 results on labor inputs (columns 4 and 5 of Table 3). The average treatment effect, pooled across years, is imprecisely estimated in both the

cash and maize loan arms (panel C). Given a control group mean of around 3,600 Kwacha in year 1, an 0.087 log point increase corresponds to a treatment effect of around 325 Kwacha, which is slightly above the amount owed on the loan, and corresponds to about 20 days of hungry season daily labor earnings. We reconcile the magnitudes of the labor and output effects in Section V.

Columns 3 and 4 of Table 5 report impacts on two summary measures of inputs. Column 3 is the total value of non-labor inputs, including seeds, fertilizer, and chemicals (pesticides and herbicides). Column 4 is the acres devoted to cash crops.³⁸ In year 1 (panel A), we observe a small and imprecise increase in input use ($p = 0.199$, column 3) and acres devoted to cash crops ($p = 0.428$, column 4). In year 2 (panel B), effects are less consistent but mirror the impacts on agricultural output: the effect is large and positive in newly treated villages for both inputs and acres of cash crops, and the interaction term is large and negative. All are imprecisely estimated. The average input use in the control group was around 65 percent higher in year 2 than in year 1.

Prediction 3: Period 1 Consumption Increases and Consumption Seasonality Decreases.—Table 6 shows the treatment effect on consumption and food security measures. We find substantial average improvements in consumption and food security outcomes in response to lower interest rates. In year 1 (panel A), the number of months with adequate food increased by 0.33 or around 5 percent relative to the control group ($p = 0.014$, column 1), and an index of food security improved by 0.3 standard deviations ($p < 0.001$, column 2). In year 2, treatment had smaller effects in newly treated villages ($p = 0.520$ and $p = 0.158$ in columns 3 and 4, respectively) and similar magnitude effects in villages treated for the second time ($p = 0.186$ and $p < 0.001$, panel B). Effects are larger for both food security measures in the maize treatment arm than the cash arm (panel C).

Adults eat more meals of the staple food during the hungry season as a result of treatment, as shown in column 3. Daily meals increase by around 0.10 meals in year 1 ($p = 0.028$, panel A) and 0.079 in both newly treated and repeat villages in year 2 ($p = 0.008$ and $p = 0.005$, panel B). At harvest, consumption in the treatment groups is similar to that in the control group in both years (all p -values > 0.40 , column 4). The magnitude of the treatment effect in year 1 is very similar to the difference between average meals consumed in the control group between the hungry season (column 3) and the harvest season (column 4).³⁹ This difference between hungry season and harvest season consumption in the control group has a p -value < 0.01 in both years. At harvest, we see no effect of loan access on consumption outcomes (column 4), though the estimated coefficient is not different from the treatment effect in the hungry season. This means that the treatment increased consumption overall, with a shift in consumption toward the hungry season, and an overall reduction in consumption seasonality.

³⁸ As described in Section IIIB, land is not typically the limiting factor in production in this setting. We therefore focus on acres devoted to cash crops as opposed to total acres in production, some of which may be very low yield.

³⁹ We also see positive but statistically insignificant effects on child consumption of the staple food, and on both adult and child consumption of protein. Note that child grain consumption demonstrates considerably less seasonality than does adult consumption in the control group. Results are available upon request.

TABLE 6—AVERAGE TREATMENT EFFECTS: CONSUMPTION AND FOOD SECURITY

	Months with enough food (1)	Food security (z-score) (2)	Meals per day hungry season (3)	Meals per day harvest season (4)
<i>Panel A. Year 1: pooled treatment arms</i>				
Any loan treatment	0.331 (0.133)	0.305 (0.079)	0.100 (0.045)	0.012 (0.014)
<i>Panel B. Year 2: pooled treatment arms</i>				
Any loan treatment	0.073 (0.132)	0.174 (0.123)	0.079 (0.029)	0.011 (0.025)
Treated in Y1	-0.055 (0.132)	-0.016 (0.100)	0.046 (0.033)	0.006 (0.019)
Loan × treated in Y1	0.141 (0.173)	0.285 (0.149)	-0.037 (0.041)	-0.002 (0.031)
Loan + Y1 + loan × Y1	0.159 (0.120)	0.442 (0.099)	0.087 (0.031)	0.015 (0.020)
<i>Panel C. By treatment arm: pooled years</i>				
Cash	0.129 (0.098)	0.175 (0.067)	0.058 (0.025)	0.010 (0.014)
Maize	0.384 (0.101)	0.471 (0.065)	0.074 (0.025)	0.010 (0.016)
Year 1 control mean	6.83	0.00	1.85	1.99
Year 2 control mean	8.40	-0.10	1.90	1.98
Year 1 = year 2 new	0.03	0.01	0.26	0.89
Year 1 = year 2 repeat	0.05	0.31	0.30	0.98
Cash = maize	0.03	0.00	0.55	0.98
Observations	6,032	2,578	2,541	2,299

Notes: Treatment effects on consumption outcomes. Column 1 is measured during the harvest and endline survey rounds (6,033 observations). Column 2 is measured during the midline and labor survey round 3 (2,578 observations). Column 3 is measured during labor survey rounds 1–3 (2,541 observations). Column 4 is measured during labor survey rounds 1, 2, and 4 and the endline survey (2,299 observations). Outcome variables are: an indicator for whether the household had any remaining grain reserves (column 1), an index of food security (column 2), and the number of adult meals per day, where a meal is defined by consumption of the staple food, nshima, in the hungry season (column 3) and at harvest (column 4). All specifications are conditional on month fixed effects and include baseline controls, and cluster standard errors at the village level. Panel C also conditions on year. Panel B also shows the total effect for repeat treatment in year 2, relative to the year 2 control group: loan + Y1 + loan × Y1.

Grain Prices and Transactions.—To interpret the results on consumption and food security, it is important to consider whether the cost of consumption was also affected by treatment status. Households were told that they could do what they liked with the cash and maize provided to them through the loans. While maize markets typically encompass multiple villages, and are therefore unlikely to show the same equilibrium price adjustments as village-level labor markets, an increase in maize sales in maize loan villages and/or an increase in maize purchases in cash loan villages may have changed the price of the staple crop within the village, affecting the value of consumption or of output.⁴⁰ Table 7 shows the treatment effect on maize prices and transactions, measured with a recall period of two weeks. During the first year, these outcomes were measured during the post-harvest season (July–November 2014). During the second year, they were measured during the hungry

⁴⁰ Online Appendix Figure B.2 shows seasonal fluctuations in grain prices in the district trading center, Chipata town.

TABLE 7—GRAIN PRICES

	Year 1 post-harvest season			Year 2 hungry season		
	Any purchase (1)	Any sale (2)	Price (3)	Any purchase (4)	Any sale (5)	Price (6)
<i>Panel A. Pooled treatment arms</i>						
Any loan treatment	0.000 (0.028)	0.008 (0.063)	0.231 (0.211)	-0.036 (0.021)	0.021 (0.012)	0.072 (0.133)
Treated in Y1				-0.014 (0.017)	0.015 (0.009)	
Loan × treated in Y1				0.038 (0.025)	-0.030 (0.016)	
Loan + Y1 + loan × Y1				-0.013 (0.016)	0.006 (0.009)	
<i>Panel B. By treatment arm</i>						
Cash	-0.009 (0.035)	0.029 (0.071)	0.138 (0.238)	-0.010 (0.011)	0.000 (0.007)	0.241 (0.147)
Maize	0.013 (0.041)	-0.020 (0.081)	0.319 (0.180)	-0.012 (0.014)	0.002 (0.011)	-0.155 (0.190)
Control mean	0.09	0.10	0.39	0.05	0.01	0.50
Cash = maize	0.67	0.58	0.20	0.88	0.86	0.10
Observations	552	552	121	1,764	1,764	98

Notes: Treatment effects on maize prices, measured during the short recall surveys from July–November 2014 (year 1 post-harvest season, 552 observations) and January–March 2015 (year 2 hungry season, 1,764 observations). Outcome variables are: whether the household reports any purchase (columns 1 and 4) or any sale (columns 2 and 5) of maize, and the price paid or received conditional on any transaction (columns 3 and 6). All specifications are conditional on month fixed effects, include baseline controls, and cluster standard errors at the village level. Columns 3 and 6 are conditional on reporting any transaction; column 6, panel A pools across year 1 treatment status, to accommodate the small number of observations. Panel B also conditions on year. Panel A also shows the total effect for repeat treatment in year 2, relative to the year 2 control group: loan + Y1 + loan × Y1.

season (January–March 2015). Overall, we observe no measurable effect on either prices or transaction probabilities during the post-harvest period in year 1, though estimates are imprecise and we cannot rule out large (proportional) changes in transaction patterns. During the year 2 hungry season, newly treated households are slightly less likely to buy ($p = 0.092$) and slightly more likely to sell ($p = 0.076$) grain, though the likelihood of any transaction remains small (5 percent of control households purchase and 1 percent sell maize in the hungry season in the control group). These low transaction probabilities result in a very small number of observations in columns 3 and 6, which are restricted to observations in which a transaction is observed.⁴¹ Taken together, this evidence suggests that a decline in prices does not drive the main consumption findings.⁴²

⁴¹ Given the low likelihood of reporting a transaction, the results on grain prices may also suffer from selection, yet we lack the data to implement a bounding exercise similar to what we report for wages in Table 4. With this caveat in mind, we observe larger increases in prices during the year 1 post harvest season, with the biggest effect in the maize loan arm ($p = 0.082$), which works in the opposite direction of the effect in maize arm in the year 2 hungry season ($p = 0.418$), as shown in column 6. This is consistent with repaying at harvest exerting upward pressure on grain prices and borrowing maize during the hungry season leading to a downward pressure on grain prices. Given the small number of transactions in the sample, these estimates should be interpreted with caution.

⁴² We might have also expected an increase in hungry season purchases in the cash loan treatment arm and/or a decrease in hungry season purchases in the maize loan arm. Given our short recall window, designed to improve

Prediction 4: Impacts Are Decreasing in Liquid Resources.—In order to assess the relationship between initial resources and loan impacts, we focus our heterogeneity analysis on year 1, for which we have an exogenous measure of baseline resources (see Section IIIA). We show additional results for year 2 in online Appendix Figures B.6, B.7, and B.8. Robustness checks using baseline reported interest rates as an alternative heterogeneity measure, are shown in online Appendix Figures B.9, B.10, and B.11.⁴³

We begin with heterogeneity in labor adjustments. Figure 4 plots adjusted predictions at the mean of each quartile of the baseline distribution of grain and cash reserves for our continuous measures (total hours over the past week) of labor market adjustments during the hungry season for both the control and treatment groups in year 1.⁴⁴ The top panel shows effects on hours of ganyu sold, the middle panel shows the effect on hours hired, and the bottom panel shows total family hours spent on-farm. In the bottom two quartiles, the primary adjustments are a shift from selling ganyu to working on-farm; in the top two quartiles, increases in labor inputs come both from more ganyu hiring and more family labor invested on-farm.⁴⁵ Overall, households across the distribution increase on-farm labor inputs. The estimated average program impact in the bottom quartile is 5 additional hours of on-farm inputs (hired plus family labor) per week, which is slightly below the cross sectional difference in labor inputs between the bottom and the top quartile in the control group (6 hours).

Figure 5 shows that the average treatment effect on agricultural output in year 1 is largest in the bottom quartile, and declines with initial resources. Higher labor inputs among households at the bottom of the distribution of initial resources may help equalize the marginal product of labor across farms. While we cannot directly observe the marginal product of labor, poorer households report higher interest rates at baseline and therefore experience the largest decline in interest rates from treatment, potentially explaining the larger impact of additional labor on output in lower quartiles of the reserves distribution.

Finally, in Figure 6, we compare our main measures of consumption (months with enough food and adult meals per day during the hungry season) across the distribution of baseline resources for treatment and control villages during year 1. The top panel shows that months of adequate food is strongly correlated with baseline resources, and that the treatment effect is largest at the bottom end of the resource distribution. In the top quartile, the treatment effect on food availability is close to zero and statistically insignificant. The bottom panel shows similar effects on the absolute number of meals consumed during the hungry season in all quartiles. Given that base consumption is substantially lower for the bottom quartile, this indicates a larger proportional improvement in consumption for low resource households. The average increase in consumption induced by the intervention in the bottom quartile

precision in the measurement of prices, we are likely to have missed most transactions that occurred immediately following delivery of the loan; in year 2, the minimum lag between loan delivery and surveying was 13 days.

⁴³Note that baseline interest rates are negatively correlated with liquid resources, so we expect interaction terms of the opposite sign for this alternative heterogeneity measure.

⁴⁴The underlying regression results for these analyses are shown in online Appendix Table B.7.

⁴⁵The difference in baseline resources between the bottom and top quartiles is almost ten times the loan amount. Baseline resources must cover a diverse set of expenditures including inputs, schooling, and other household needs; higher expenditure elasticities for unanticipated seasonal loans therefore seem plausible.

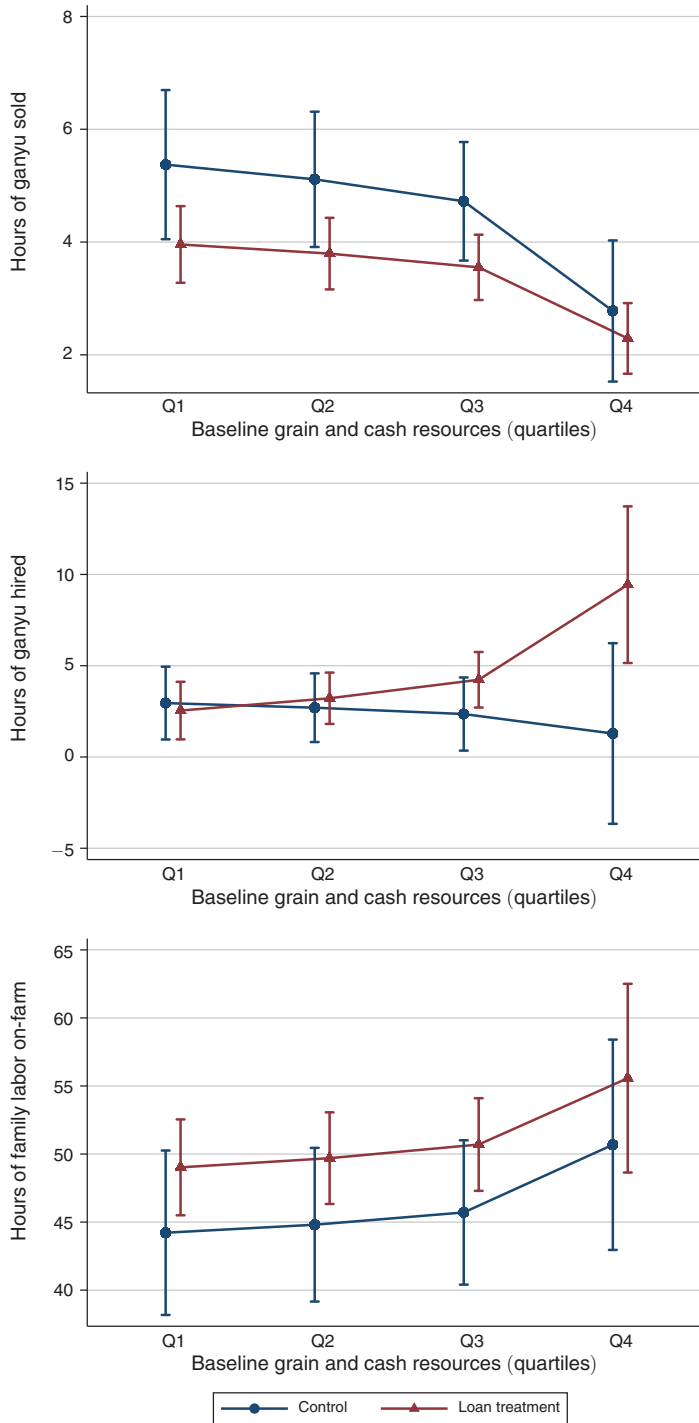


FIGURE 4. EFFECT ON LABOR MARKET PARTICIPATION, BY BASELINE RESERVES

Notes: Heterogeneous impacts on labor market outcomes in the hungry season of year 1, estimated using a quadratic in baseline reserves. Plots show adjusted predictions at the mean in each quartile of the baseline distribution, based on regressions that control for geographic variables only. Ninety percent confidence intervals are plotted based on standard errors clustered at the village level. The red triangles indicate the treatment group (shifted right). The same figures for year 2 of the project are shown in the online Appendix.

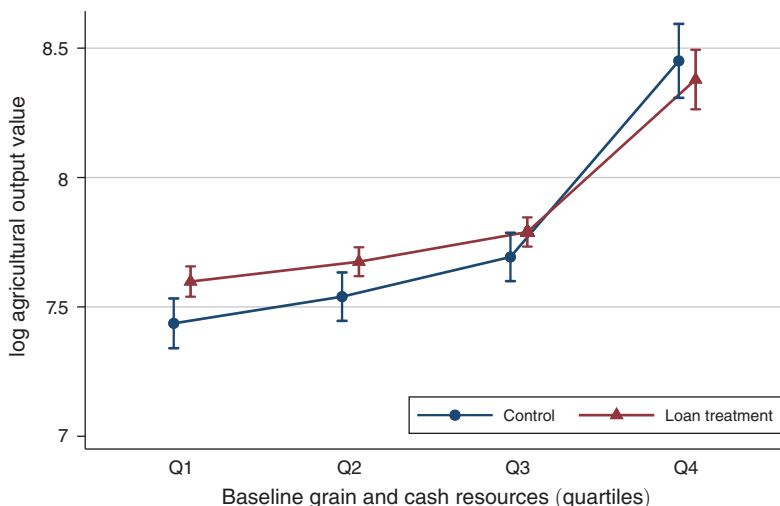


FIGURE 5. EFFECT ON LOG AGRICULTURAL OUTPUT, BY BASELINE RESERVES, YEAR 1

Notes: Heterogeneous impacts on agricultural output value in year 1, estimated using a quadratic in baseline reserves. Plots show adjusted predictions at the mean in each quartile of the baseline distribution, based on regressions that control for geographic variables only. Ninety percent confidence intervals are plotted based on standard errors clustered at the village level. The red triangles indicate the treatment group (shifted right). The same figure for year 2 of the project is shown in the online Appendix.

of the resource distribution is similar in magnitude to the cross-sectional difference in hungry season consumption between the top and the bottom quartile in the control group.

V. Interpretation and Robustness Checks

To aid interpretation of the main results, we evaluate the plausibility of the magnitudes of our findings and provide further evidence in this section.

A. Magnitudes

Our main results suggest that a relatively small loan, timed to coincide with the period of high labor demand and limited food and cash availability, led to substantial increases in on-farm labor inputs, agricultural output, and hungry-season consumption. We discuss the plausibility of the effect sizes with a focus on the results in year 1, which are more precise and easier to interpret. On average, we estimate that the loan program resulted in approximately 7 additional hours of labor inputs per week on the family farm during the hungry season in year 1 (columns 4 and 5, panel A of Table 3). This increase is driven by increased hiring of external labor (2 hours) as well as a substantial increase in family labor on-farm (5 hours). Taking the 13 weeks of the hungry season from January to March as our temporal reference frame, and assuming an average five-hour work day on the field (as reported in time use questions), an additional 7 hours per week corresponds to approximately 18 additional days of labor on the farm during the hungry season in year 1. We estimate

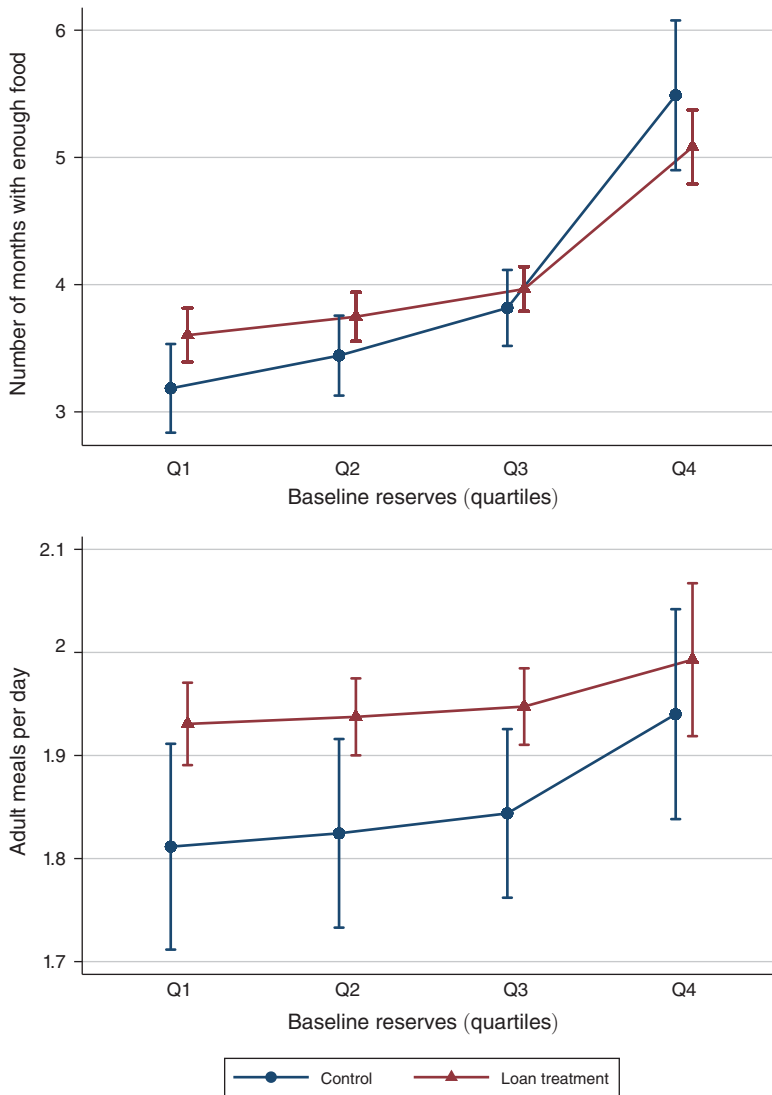


FIGURE 6. EFFECT ON CONSUMPTION VARIABLES, BY BASELINE RESERVES, YEAR 1

Notes: Heterogeneous impacts on consumption outcomes in the hungry season of year 1, estimated using a quadratic in baseline reserves. Plots show adjusted predictions at the mean in each quartile of the baseline distribution, based on regressions that control for geographic variables only. Ninety percent confidence intervals are plotted based on standard errors clustered at the village level. The red triangles indicate indicates the treatment group (shifted right). The same figures for year 2 of the project are shown in the online Appendix.

a treatment effect of around 325 Kwacha on agricultural output (column 2, panel A of Table 5). With 18 additional days, this implies an average marginal product of labor of around 18 Kwacha (assuming all other inputs remain fixed), which is very close to the average daily earnings in the treatment group reported in Table 4 (18.5 Kwacha). This suggests that the increase in hours of labor on the family farm at the equilibrium wage is sufficient to explain most of the estimated impact on agricultural output, even without discounting. Consistent with this, the small and imprecise

increase in inputs and acres devoted to cash crops in year 1, reported in Table 5 (columns 3 and 4, panel A), are unlikely to have majorly affected the results.⁴⁶

We can also compare the size of the loan with the cost to the household of the labor reallocation we observe. On average, in year 1 of the study, households in the control group sold 4.6 hours of ganyu per week (column 2, panel A of Table 3). With average daily earnings in the control group of 15.6 Kwacha, this implies total earnings of around 220 Kwacha over the 13 weeks of the hungry season, assuming a 5 hour work day. The treatment effect on ganyu labor sales (-1.1 hours per week) corresponds to an average of around 2.9 fewer days sold, or 46 Kwacha less in ganyu earnings, valued at the control group wages. Adding the average additional expenditure for ganyu hiring (85 Kwacha, valued at the wage in the control group), this implies that about 65 percent of the loan amount (130 Kwacha) went to foregone earnings plus wages paid to hired labor.

In online Appendix Table B.8, we show effects on other consumption smoothing strategies. We find no effect on input loans, which is the most common type of borrowing reported at baseline. Both informal low interest and high interest borrowing declined in year 1 (panel A: $p = 0.226$ and $p < 0.001$, respectively). Asset and livestock sales and green maize consumption were all unaffected by treatment (columns 5, 6, and 7), though villages treated in both years increased livestock sales and green maize consumption slightly (panel B). Together, these results suggest that treatment primarily affected labor allocation, with some reduction in high-cost borrowing.

B. Mechanisms: Family Labor Supply

Our main results suggest that treated households substantially increased their labor supply in response to treatment. We consider three potential mechanisms that could result in a positive labor supply response in this setting: (i) nutritional status and physical health, (ii) cognition and decision-making, and (iii) stress, affect, and motivation.

In online Appendix Table B.9 we analyze a range of physical health proxies, including self-rated health, waist and biceps circumference, and grip strength, all measured during the hungry season. We only find improvements on grip strength test endurance, which could be driven both by increased strength and by increased motivation. In online Appendix Table B.10 we analyze various measures of decision-making. We see no treatment effect on expenditures on temptation goods such as clothes, beer, tobacco, sweets, and tea (all $p > 0.20$, columns 1–5). We also collected data on cognitive ability in the second year of the study. Performance on a Ravens Progressive Matrices test was unaffected ($p = 0.473$ for villages treated for the first time and $p = 0.704$ for villages treated for the second time, column 6), while performance on two measures from a numerical Stroop test decreased in

⁴⁶The (insignificantly) larger effects on inputs and cash crops in newly treated villages in year 2 appear to be associated, at least in part, with the early notification of the loan in half of the villages in year 2. We describe these results in Section VC.

newly treated villages ($p = 0.056$ and $p = 0.051$) and remained the same in villages treated for the second time in year 2 ($p = 0.391$ and $p = 0.121$).⁴⁷

Food shortages during the hungry season may cause stigma and stress, and negatively affect motivation to work. Online Appendix Table B.11 examines loan impacts on a number of measures of stress, affect, and motivation. We administered a 19-item mental health module to survey respondents during the year 2 hungry season. In column 1, we analyze effects on the total score, where a higher score implies more mental health problems. This measure decreases significantly, by 75 percent of the control group mean, in newly treated villages in year 2 ($p = 0.090$). Columns 2–5 show that relevant index components (a sample of the 19 questions on the instrument) all decline, with a particularly large and significant decrease in the likelihood that respondents report feeling easily tired ($p = 0.006$, column 5). In both years, as part of the food security module, we asked whether respondents worried about food in the past two weeks (column 6). The likelihood of worrying about food decreased in year 1 ($p < 0.001$) and in year 2 ($p = 0.170$ for villages treated for the first time and $p < 0.001$ for villages treated for the second time).

Overall, our evidence appears most consistent with mental health and motivation acting as constraints on family labor supply that can be relaxed by access to the subsidized loans. However, other mechanisms may also have contributed to a positive labor supply responses to lower interest rates. For example, complementarities between other agricultural inputs and labor might result in additional labor inputs if lower borrowing costs increase the use of other inputs. While we do find some small (and imprecisely estimated) increases in spending on inputs and acres under cash crops in both years (Table 5), households could easily satisfy additional labor demand with hired labor. Of course, the complementarity could be specific to family labor, and we lack the necessary data to test such a prediction. We view the relationship between liquidity, consumption, and labor (constraints) as promising area for future research.

C. Additional Results

Our study design varied a number of other program features to help understand underlying mechanisms.

Year 1 versus Year 2 Treatment Effects.—Our main results are split by treatment year, with year 2 further divided between villages treated for the first time in year 2 and those treated in both years. Here, we summarize differences in the results by year and discuss potential explanations for these differences.

Many outcomes are similar in villages treated for the first time in both years (i.e., all treated villages in year 1 and newly treated villages in year 2): estimated impacts on ganyu sales, labor supply, labor demand, daily earnings, agricultural output, and consumption are all of the same sign and of similar magnitude, though year 2 is less precisely estimated in most cases, due in part to the smaller sample size. The one exception is ganyu hiring, which shows no extensive margin adjustment in newly

⁴⁷Fehr, Fink, and Jack (2019) report on these and other measures of cognition and decision-making in the study setting, and find broad evidence for improvements in decision-making when resources are more constrained.

treated villages in year 2. Villages treated for a second time in year 2 saw similar impacts on ganyu labor sales and consumption. However, neither labor demand nor agricultural output show measurable increases in this subgroup.

One important difference across study years is rainfall, which shapes both the subsequent year's hungry season (potentially affecting both take-up and treatment effects) and resources available for repayment. In online Appendix Figure B.4, we show the long-run rainfall distribution in Chipata district.⁴⁸ Rainfall patterns match the average agricultural output across years, which is highest in year 1, and lowest in year 2. This means that farmers came into year 2 with unusually high levels of reserves, making the year 2 hungry season less severe, consistent with the control group means shown in Table 6. This implies that high take-up rates in both years are not due to worse-than-average hungry seasons; take-up was close to 100 percent even after an unusually good harvest. On the other hand, higher default rates in year 2 are consistent with worse farming conditions in that year, though we cannot rule out alternative explanations. The different patterns of investment and output in villages treated for a second time in year 2 may be specific to a bad rainfall realization: with greater savings and access to additional resources during the hungry season, households may have chosen to divert these additional resources away from agriculture once the poor rainfall patterns became clear. Of course this is a purely speculative interpretation and we lack the high resolution data on rainfall or the detailed time use data necessary to further investigate this interpretation.

Maize versus Cash Loans.—In our main results, we show effects for both the maize and cash treatment arms, and observe generally similar responses. That said, effects on consumption and some labor outcomes are slightly larger in the maize treatment arm, perhaps reflecting the higher value of the loans during the hungry season, or transaction costs associated with converting cash into grain and vice versa. On the other hand, the estimated increase in agricultural output is slightly higher in the cash treatment arm. In online Appendix Table B.10 we show estimated treatment effects on a range of different (potentially wasteful) expenditure categories to test whether cash loans increase the likelihood of wasteful consumption. We find no evidence of increased wasteful spending in the cash treatment arm.

Anticipated versus Unanticipated Loan Availability.—Online Appendix Table B.12 reports the effects of the year 2 subtreatment that varied whether eligible households learned about the loan program at the start of the hungry season (the same timing that was used in year 1) or at planting time (early notification). Notifying farmers at the time they are making their production plans could increase the impact of lower effective borrowing rates, since farms can adjust crop mix and crop timing in anticipation of loan availability. We interact treatment status with an indicator for early notification for our main outcomes (columns 1–5) and for two supplementary outcomes relevant to planting season adjustments (columns 6–7): the total value of capital inputs applied to the household's fields and the acres

⁴⁸ Unfortunately, the scale of our experiment and local paucity of ground stations to measure rainfall make modeled precipitation products too coarse to be useful for heterogeneity analysis. Instead we use rainfall gauge measures obtained from the Msekera Research Station in Chipata District (Msekera 2015).

devoted to cash crops.⁴⁹ Results are imprecisely estimated, but suggestive of larger effects on labor allocation and agricultural output value in villages notified at planting time, potentially driven by adjustments to capital inputs and planting decisions (columns 6 and 7).⁵⁰

D. Alternative Explanations and Robustness Checks

We discuss other potential interpretations of our findings in this section, including some robustness checks.

Income or Insurance Effects.—To rule out that our results are driven by income effects (due to below-market interest rates), we implemented a small cash transfer sub-treatment, as described in Section IIB. In online Appendix Table B.13, we test whether a transfer of 60 Kwacha at the start of the hungry season led to a measurable effect on our main outcomes, relative to the pure control group. While our power is limited due to the small sample size, we do not find any evidence for labor adjustments, or any effects on output or consumption, in response to the cash gift, which suggests that the treatment effects we measure are not driven by the relatively small net transfers embedded in the loans.

It is also possible that the loan programs offered some implicit insurance value to farmers. In the absence of an immediate penalty for default, the loans might allow farmers to invest upfront but default in the event of bad harvest outcomes (consistent with one interpretation for the higher default rates in year 2 of the program). This might explain some of the impacts on agricultural output. However, given that the total loan amount only corresponded to about 6 percent of the average harvest value at baseline, the implicit insurance value is small and unlikely to explain an important share of the results.

Reporting Bias.—Given that we mostly rely on self-reported outcomes in our analysis, one obvious concern is that household responses may have been affected by treated households responding more positively or by making more effort to provide the socially desired answer. We test for bias in our self-reported survey measures in two ways. First, during selected surveys, we administered a social psychology scale (Marlow and Crowne 1961), designed to measure social desirability bias. We test for differences in the response patterns across treatment arms. We find no treatment effects on responses to this social desirability module (panel A of online Appendix Table B.14). Second, we collected objective measures of maize output in years 1 and 2 of the study.⁵¹ We directly compare objective measures of output with reported measures, and test whether there are any systematic differences between the two

⁴⁹Treatment and timing are also interacted with year 1 treatment status. These interaction terms are suppressed. As a result, the reported coefficients can be interpreted as the effects of treatment and early notification for villages treated for the first time in year 1.

⁵⁰To put the magnitudes into perspective, the 250 Kwacha of additional inputs is about the value of one 50 kg bag of fertilizer (roughly one-half of the recommended fertilizer for one acre), and the impact on acres under cash crops is around 7 percent of the average land area under production.

⁵¹In year 1, enumerators visited the fields of a subsample of respondents and measured the height of a typical maize stalk in the field. In year 2, enumerators sampled maize cobs and counted the corresponding number of kernels.

variables by treatment status (shown by the interaction term in panel B of online Appendix Table B.14). Both of these tests suggest that treatment status did not alter the reporting behavior of study participants in a systematic way.

VI. Model Calibration, Welfare, and Policy Implications

A. Model Calibration and Simulation: Spillover Effects and Scale-Up

Our main results imply substantial benefits for eligible households; the same is not necessarily true for ineligible households in treated villages. Untreated farmers did not benefit from lower interest rates, but were affected indirectly through increases in labor demand by treated farmers and resulting equilibrium wage increases. Higher equilibrium wages help cover hungry season consumption needs for households that are net sellers of family labor, but also increased the cost of labor inputs on their own farms. We calibrate a simulation of our model from Section I to the empirical values observed at baseline;⁵² this allows us to engage in “structured speculation” (Banerjee, Chassang, and Snowberg 2017) about the welfare implications for untreated farmers and the effects were the program scaled to cover all farmers in a village. Results are summarized in online Appendix Table A.2.

Beginning with the simulated spillovers effects, assuming that 50 percent of households get access to subsidized loans, we see a roughly 10 percent (0.096 log points) increase in the equilibrium wage from 13.9 to 15.3 Kwacha, slightly smaller than our estimated treatment effect of 16 percent. Figure 7 shows agricultural output value, net labor supply, hungry season consumption, and utility for both treated and untreated households in treatment villages. On average, agricultural output among treated farmers increases by 9.6 percent in our simulations (only marginally larger than the estimated treatment effect of 8.7 percent) while agricultural output of untreated farmers falls by 9.5 percent on average in intervention villages, since they invest less labor on their own farms in response to higher wages and unchanged interest rates. This is illustrated in the top-right panel of Figure 7: net labor supply (labor supply minus labor demand) declines among treated farmers, and increases among untreated farmers. Untreated farmers in intervention villages experience modest declines in consumption during the hungry season (bottom-left panel); their additional wage income is less than their loss in agricultural output value resulting from less on-farm labor investment, while the share of resources they allocate to first period consumption (as a function of interest rates) is unchanged. Total utility falls by 14 percent in this group.⁵³

Utility losses among untreated households are greatest in the middle of the initial resource distribution, mimicking the losses in agricultural output. Among treated households, utility gains are positive across almost the entire resource distribution,

⁵²Specifically, we assume that output is determined by the Cobb-Douglas function outlined in Section I, and that the distribution of farm productivity A_i is log-normal with mean μ_a and standard deviation σ_a . We then calibrate μ_a and σ_a such that the distance between the empirically observed and simulated baseline outputs is minimized. We also assume A_i to be correlated with baseline resources. We estimate a correlation coefficient of 0.40 using our panel of agricultural output and household fixed effects. Online Appendix Section A.2 provides further details on the parametric assumptions made for the initial calibrations.

⁵³In our model framework, consumption over the two periods is a close proxy for welfare, since it accounts for income from both wage earnings and agricultural output.

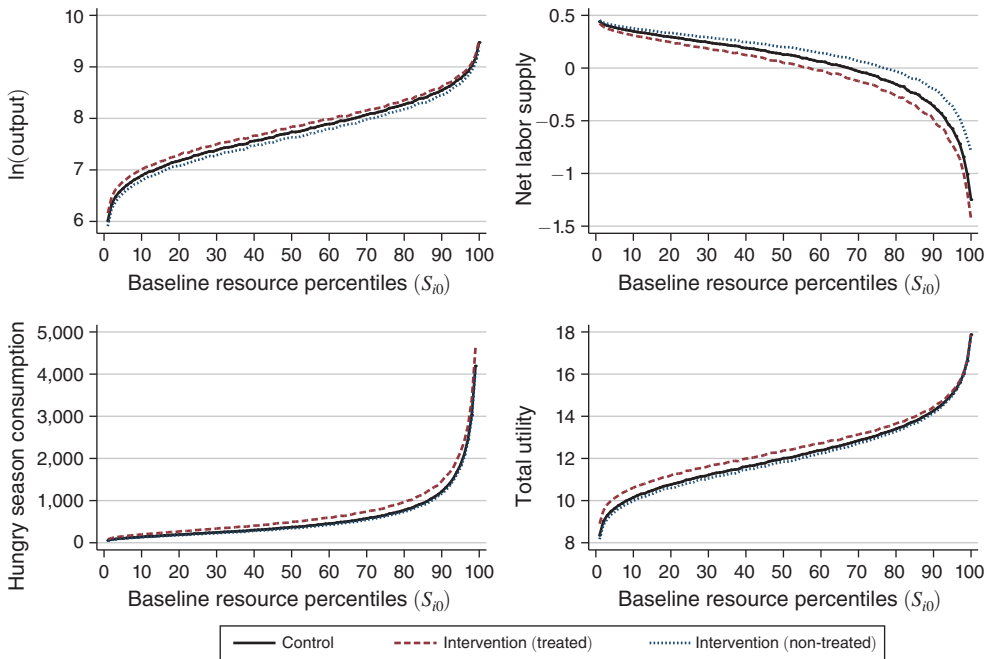


FIGURE 7. IMPACTS WITH PARTIAL TREATMENT

Notes: Figure shows estimated impacts on agricultural output (top-left), net labor supply (top-right), period 1 consumption (bottom-left), and utility (bottom-right), for treated and untreated farmers. The productivity distribution was calibrated so that the resulting wage and output distributions match the average values observed in the control group.

and largest among those with the fewest baseline resources who experience the largest utility gains from increased consumption. Since we assume that interest rates are negatively correlated with baseline reserves, poorer households will always experience the largest reductions in interest rates. This direct interest rate effect is, however, partially offset by (assumed) lower productivity at the bottom of the distribution.

While many microcredit programs lower interest rates for only a subset of the village (women only or members of an agricultural cooperative, for example), an ideal policy intervention would lower interest rates for all households through, for example, improvements in monitoring and enforcement technologies or changes in regulation. Figure 8 shows simulated impacts of lowering interest rates for all households. Relative to the partial treatment case covered by our predictions, full treatment further increases aggregate labor demand effects, and results in a larger aggregate wage adjustment, with a simulated average daily wage of 20.2 Kwacha (compared to 15.3 Kwacha with partial treatment). Compared to partial treatment, effects on agricultural income are more moderate in this scenario: since all households have access to lower interest rates, labor reallocation primarily consists of a transfer from households with more baseline resources to those with fewer baseline resources. As a result, agricultural incomes in the bottom quintile of baseline resources increase by about 13 percent, while agricultural incomes in the top 5 percentiles decrease

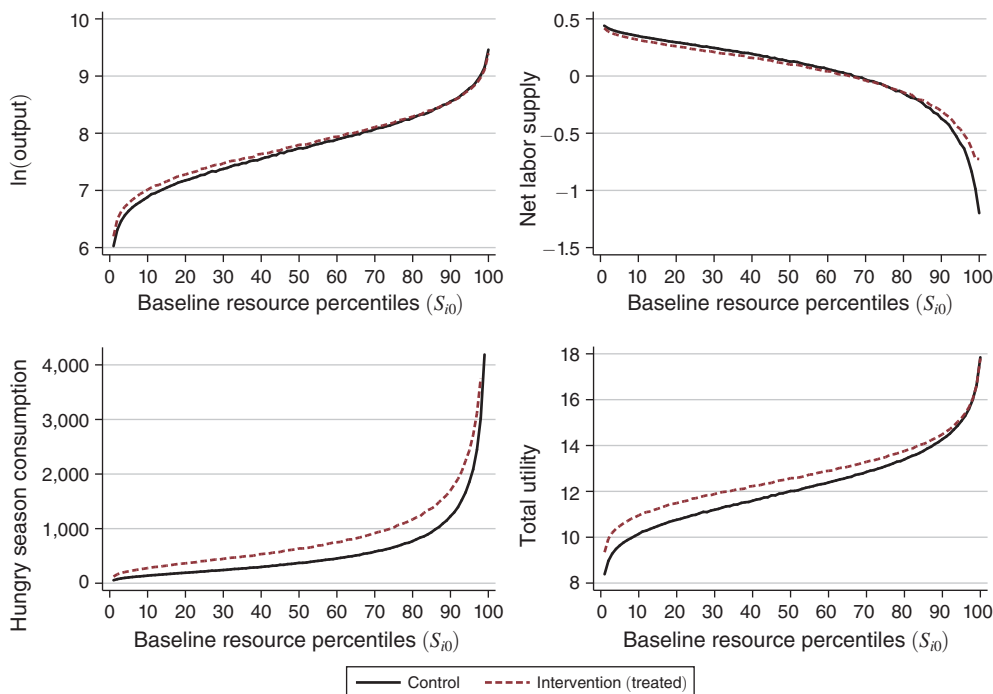


FIGURE 8. IMPACTS WITH FULL TREATMENT

Notes: Figure shows estimated impacts on agricultural output (top-left), net labor supply (top-right), period 1 consumption (bottom-left), and utility (bottom-right) under a scenario of full treatment, i.e., access to subsidized loans for all households in the village. The productivity distribution was calibrated so that the resulting wage and output distributions match the average values observed in the control group.

by 3.9 percent. Average agricultural output increases by 6.1 percent, due both to increased (less constrained) labor supply and a more efficient allocation of labor across farms. Output inequality (Gini) falls from 0.35 to 0.33. Positive income and substitution effects lead to a 38.6 percent increase in hungry season consumption, on average. Total utility increases by 0.54 log points, and hungry season consumption inequality decreases by 14.6 percent from a Gini of 52 to a Gini of 45.

VII. Conclusion

The results presented in this paper highlight the importance of seasonal incomes, credit access, and liquidity for labor markets and agricultural production. From a policy perspective, providing rural farmers with subsidized credit has several attractive features: in addition to increasing financial inclusion, the results presented in this paper suggest that loans targeted to poor farmers during the hungry season can improve food security, increase wages, and raise average agricultural output, while reducing farmers' reliance on stigmatized piecework labor. These improvements in welfare resulting from lower borrowing costs during the hungry season appear particularly large among farmers with fewer resources going in to the hungry season, and result in an increase in aggregate agricultural output due to a more efficient allocation of labor.

In the specific location we study, and presumably in many other similar settings, the room for improvement in local credit markets is large, but not without challenges. Most farmers in our study report monthly interest rates of around 40 percent. These self-reports are backed up by the 98 percent take-up of the loans we offer, even in the second year after farmers had repaid their first loans. High interest rates should however not be equated with inefficient or failing credit markets. As our own project experience shows, transaction costs involved in rural lending in settings with limited road infrastructure are high, and enforcing loan repayment can be difficult in the absence of collateral. Ignoring transaction costs, the 30 percent interest rate on the approximately six-month loan was sufficient to cover default rates of up to 30 percent, well in excess of even the default rates in year 2. However, interest rates must cover both default and implementation costs and we calculate the latter to be high in our setting. Even in the cheapest treatment arm (cash loan with cash repayment), the average cost per loan was over 100 Kwacha, while the value of the interest was only 60 Kwacha.⁵⁴ Even with zero default and expanded eligibility to include all farmers in the village, these implementation costs are likely too high to make small consumption loans at reasonable interest rates financially viable.

Higher interest rate or larger loan sizes could help offset the high transaction costs of administering consumption loans in rural areas, but may undermine some of the benefits that we document. Instead, new and cheaper platforms, such as mobile money or piggybacking on existing contractual arrangements, are needed in order to be able to run similar interventions in a financially sustainable manner. Alternatively, other interventions that require less infrastructure, such as approaches that target savings, present promising alternative ways of increasing hungry season liquidity and decreasing dependence on labor markets as a smoothing device.

REFERENCES

- Aggarwal, Shilpa, Eilin Francis, and Jonathan Robinson.** 2018. "Grain Today, Gain Tomorrow: Evidence from a Storage Experiment with Savings Clubs in Kenya." *Journal of Development Economics* 134: 1–15.
- Banerjee, Abhijit V., Sylvain Chassang, and Erik Snowberg.** 2017. "Decision Theoretic Approaches to Experiment Design and External Validity." In *Handbook of Economic Field Experiments*. Vol. 1, edited by Abhijit V. Banerjee and Esther Duflo, 141–74. Amsterdam: Elsevier.
- Banerjee, Abhijit, Dean Karlan, Hannah Trachtman, and Christopher Udry.** 2020. "Does Poverty Change Labor Supply? Evidence from Multiple Income Effects and 115,579 Bags." CEPR Discussion Paper 14812.
- Banerjee, Abhijit, Dean Karlan, and Jonathan Zinman.** 2015. "Six Randomized Evaluations of Microcredit: Introduction and Further Steps." *American Economic Journal: Applied Economics* 7 (1): 1–21.
- Basu, Karna, and Maisy Wong.** 2015. "Evaluating Seasonal Food Storage and Credit Programs in East Indonesia." *Journal of Development Economics* 115 (July): 200–216.
- Behrman, Jere R., Andrew D. Foster, and Mark R. Rosenzweig.** 1997. "The Dynamics of Agricultural Production and the Calorie-Income Relationship: Evidence from Pakistan." *Journal of Econometrics* 77 (1): 187–207.

⁵⁴We focus on the cash loan treatment arm in calculating implementation costs, given that it was considerably cheaper to implement, and most treatment effects (and repayment rates) are similar across treatment arms. The average cost of loan implementation, including sign-up, delivery and repayment, was 1,800 Kwacha (~US\$300) per village, with an average of 17 farmers per village taking out loans.

- Bruhn, Miriam, and David McKenzie.** 2009. "In Pursuit of Balance: Randomization in Practice in Development Field Experiments." *American Economic Journal: Applied Economics* 1 (4): 200–232.
- Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak.** 2014. "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh." *Econometrica* 82 (5): 1671–1748.
- Burke, Marshall, Lauren Falcao Bergquist, and Edward Miguel.** 2019. "Sell Low and Buy High: Arbitrage and Local Price Effects in Kenyan Markets." *Quarterly Journal of Economics* 134 (2): 785–842.
- Casaburi, Lorenzo, and Jack Willis.** 2018. "Time versus State in Insurance: Experimental Evidence from Contract Farming in Kenya." *American Economic Review* 108 (12): 3778–813.
- Central Statistical Office (CSO), Central Board of Health, and ORC Macro.** 2002. "Zambia Demographic and Health Survey 2001–02: Dataset." Data Extract from ZMKR42.SAV. IPUMS Demographic and Health Surveys (IPUMS DHS), version 7, IPUMS and ICF [Distributors]. <http://idhsdata.org>.
- Central Statistical Office (CSO), Ministry of Health, and ICF International.** 2014. "Zambia Demographic and Health Survey 2013–14: Dataset." Data Extract from ZMKR61.SAV IPUMS Demographic and Health Surveys (IPUMS DHS), version 7, IPUMS and ICF [Distributors]. <http://idhsdata.org>.
- Central Statistical Office (CSO), Ministry of Health, Tropical Diseases Research Centre, University of Zambia, and ORC Macro.** 2007. "Zambia Demographic and Health Survey 2007: Dataset." Data Extract from ZMKR51.SAV. IPUMS Demographic and Health Surveys (IPUMS DHS), version 7, IPUMS and ICF [Distributors]. <http://idhsdata.org>.
- Cull, Robert and Jonathan Morduch.** 2018. "Microfinance and Economic Development." In *Handbook of Finance and Development*, edited by Thorsten Beck and Ross Levine, 550–572. Cheltenham, UK: Edward Elgar Publishing.
- Dasgupta, Partha, and Debraj Ray.** 1986. "Inequality as a Determinant of Malnutrition and Unemployment: Theory." *Economic Journal* 96 (384): 1011–34.
- Dasgupta, Partha, and Debraj Ray.** 1987. "Inequality as a Determinant of Malnutrition and Unemployment: Policy." *Economic Journal* 97 (385): 177–88.
- Devereux, Stephen, Rachel Sabates-Wheeler, and Richard Longhurst.** 2013. *Seasonality, Rural Livelihoods and Development*. New York: Routledge.
- Duflo, Esther, Michael Kremer, and Jonathan Robinson.** 2011. "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya." *American Economic Review* 101 (6): 2350–90.
- Fehr, Dietmar, Gunther Fink, and Kelsey Jack.** 2019. "Poverty, Seasonal Scarcity and Exchange Asymmetries." NBER Working Paper 26357.
- Field, Erica, Rohini Pande, John Papp, and Natalia Rigol.** 2013. "Does the Classic Microfinance Model Discourage Entrepreneurship among the Poor? Experimental Evidence from India." *American Economic Review* 103 (6): 2196–2226.
- Fink, Günther, B. Kelsey Jack, and Felix Masiye.** 2014. "Seasonal Credit Constraints and Agricultural Labor Supply: Evidence from Zambia." NBER Working Paper 20218.
- Fink, Günther, B. Kelsey Jack, and Felix Masiye.** 2018. "Food Constraints, Yield Uncertainty and Ganyu Labour." AEA RCT Registry. February 12. <https://doi.org/10.1257/rct.130-3.0>.
- Fink, Günther, B. Kelsey Jack, and Felix Masiye.** 2020. "Replication Data for: Seasonal Liquidity, Rural Labor Markets and Agricultural Production: Dataset." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E119649V1>.
- Hausdorff, Katharine.** 2016. "Household Bargaining Power and the Effect of Microloans." MSc thesis. Tufts University.
- Ito, Takahiro, and Takashi Kurosaki.** 2009. "Weather Risk, Wages in Kind, and the Off-Farm Labor Supply of Agricultural Households in a Developing Country." *American Journal of Agricultural Economics* 91 (3): 697–710.
- Jayachandran, Seema.** 2006. "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries." *Journal of Political Economy* 114 (3): 538–75.
- Kaminski, Jonathan, Luc Christiaensen, and Christopher L. Gilbert.** 2014. "The End of Seasonality? New Insights from Sub-Saharan Africa." World Bank Policy Research Working Paper 6907.
- Kaur, Supreet, Sendhil Mullainathan, Suanna Oh, and Frank Schilbach.** 2019. "Does Financial Strain Lower Worker Productivity?" Unpublished.
- Kochar, Anjini.** 1995. "Explaining Household Vulnerability to Idiosyncratic Income Shocks." *American Economic Review* 85 (2): 159–64.

- Kochar, Anjini.** 1999. "Smoothing Consumption by Smoothing Income: Hours-of-Work Responses to Idiosyncratic Agricultural Shocks in Rural India." *Review of Economics and Statistics* 81 (1): 50–61.
- Kremer, Michael, Gautam Rao, and Frank Schilbach.** 2019. "Behavioral Development Economics." In *Handbook of Behavioral Economics: Foundations and Applications 2*, edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, 345–458. Amsterdam: Elsevier.
- Lee, David S.** 2009. "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects." *Review of Economic Studies* 76 (3): 1071–102.
- Marlow, David, and Douglas P. Crowne.** 1961. "Social Desirability and Response to Perceived Situational Demands." *Journal of Consulting Psychology* 25 (2): 109–115.
- Morduch, Jonathan.** 1995. "Income Smoothing and Consumption Smoothing." *Journal of Economic Perspectives* 9 (3): 103–14.
- Msekera Research Station.** 2015. "Monthly Rainfall Data 1970-2015: Dataset." Data obtained in-person from the Msekera Research Station, Chipata District, Ministry of Agriculture, Zambia.
- Pitt, Mark M., and Shahidur R. Khandker.** 2002. "Credit Programmes for the Poor and Seasonality in Rural Bangladesh." *Journal of Development Studies* 39 (2): 1–24.
- Rose, Elaina.** 2001. "Ex Ante and Ex Post Labor Supply Response to Risk in a Low-Income Area." *Journal of Development Economics* 64 (2): 371–88.
- Rosenzweig, Mark R.** 1980. "Neoclassical Theory and the Optimizing Peasant: An Econometric Analysis of Market Family Labor Supply in a Developing Country." *Quarterly Journal of Economics* 94 (1): 31–55.
- Rosenzweig, Mark R., and Kenneth I. Wolpin.** 1993. "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investment in Bullocks in India." *Journal of Political Economy* 101 (2): 223–44.
- Skoufias, Emmanuel.** 1996. "Intertemporal Substitution in Labor Supply: Micro Evidence from Rural India." *Journal of Development Economics* 51 (2): 217–37.
- World Bank.** 2017. "World Development Indicators."
- Zambian Central Statistics Office (CSO).** 2010. "2010 Census of Population." Government of the Republic of Zambia.