

NBER WORKING PAPER SERIES

SEASONAL LIQUIDITY, RURAL LABOR MARKETS AND AGRICULTURAL
PRODUCTION

Günther Fink
B. Kelsey Jack
Felix Masiye

Working Paper 24564
<http://www.nber.org/papers/w24564>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2018

We thank audience members at numerous seminars and conferences for comments and suggestions. We are grateful to 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 for financial support, and to Innovations for Poverty Action for logistical support. Many thanks to Rachel Levenson for careful oversight of the field work and to Daniel Velez Lopez, Chantelle Boudreaux and Carlos Riumallo Herl for assistance with the data. 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. A pilot to this paper was circulated as NBER Working Paper No 20218. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Günther Fink, B. Kelsey Jack, and Felix Masiye. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Seasonal Liquidity, Rural Labor Markets and Agricultural Production
Günther Fink, B. Kelsey Jack, and Felix Masiye
NBER Working Paper No. 24564
April 2018
JEL No. D14,J2,J43,O13

ABSTRACT

Many rural households in low and middle income countries continue to rely on small-scale agriculture as their primary source of income. In the absence of irrigation, income arrives only once or twice per year, and has to cover consumption and input needs until the subsequent harvest. We develop a model to show that seasonal liquidity constraints not only undermine households' ability to smooth consumption over the cropping cycle, but also affect labor markets if liquidity-constrained farmers sell family labor off-farm to meet short-run cash needs. To identify the impact of seasonal constraints on labor allocation and agricultural production, we conducted a two-year randomized controlled trial with small-scale farmers in rural Zambia. Our results indicate that lowering the cost of accessing liquidity at the time of the year when farmers are most constrained (the lean season) reduces aggregate labor supply, drives up wages and leads to a reallocation of labor from less to more liquidity-constrained farms. This reallocation reduces consumption and income inequality among treated farmers and increases average agricultural output.

Günther Fink
Harvard School of Public Health
Department of Global Health and Population
665 Huntington Ave.
Boston, MA 02115
guenther.fink@swisstph.ch

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

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

An appendix is available at
<http://www.nber.org/data-appendix/w24564>

A randomized controlled trials registry entry is available at
<https://www.socialscienceregistry.org/trials/130/>

1 Introduction

A majority of rural households in developing countries continue to rely on small-scale farming as their primary source of income, and on labor as the primary input to agricultural production. In the absence of irrigation and advanced farming technologies, agricultural incomes arrive infrequently, and have to cover both consumption and inputs until the subsequent harvest. Low returns to saving and high costs of borrowing make smoothing from one harvest to the next more costly and raise the relative price of consumption at times of the year most distant from the previous harvest, often referred to as the “lean season” or “hungry season” in rural areas of developing countries.¹ In this paper, we show that seasonal liquidity constraints affect not only hungry season consumption but also local labor markets and agricultural production. In settings where family labor can either be used on-farm or sold locally to other farms, hungry season labor sales can be used to finance short term consumption needs. If these labor sales lower labor inputs on-farm, then hungry season consumption is effectively financed against future harvest income. We formalize this intuition in a simple two-period agricultural household model, and show that a higher cost of accessing liquidity in the first period (the hungry season) lowers on-farm labor inputs, increases labor supply and drives down local wages. To the extent that liquidity during the hungry season varies within the labor market, the result is an inefficient allocation of labor across farms, with lower aggregate agricultural output and greater local income inequality.

We test the key predictions of our model in a two-year randomized controlled trial with 3,139 small-scale farmers in 175 rural villages in Zambia’s Eastern Province. To relax liquidity constraints, a sub-set of farmers in selected villages were offered subsidized cash or food loans during the hungry season, with repayment due after harvest.² Our empirical setting and theoretical framework are 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, and alternatives such as grain storage or livestock holdings are risky. Access to formal credit is reported by less than 5 percent of our sample, while informal credit carries a reported average monthly interest rate of 40 percent. High costs of access to capital markets are also reflected in the savings and consumption patterns observed in our sample: both cash and food reserves and consumption are highly seasonal, peaking after harvest, and reaching their nadir during the hungry season. 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

¹A substantial literature documents seasonality in grain prices and consumption (see, for example, Kaminski et al. (2014); Devereux et al. (2013)). Note that grain prices will exhibit seasonality even where capital markets function well, given storage costs and interest rates, though Gilbert et al. (2017) and Kaminski et al. (2014) argue that the seasonal fluctuations observed in Africa exceed the level explained just by interest rates.

²Offering households a loan rather than a grant reduces the income effect associated with the intervention. Of course, subsidized credit still has a positive income effect, so we include an additional “income effect control” group in our design, in which households receive a small cash grant roughly equal to the size of the income transfer associated with the loans.

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. While these labor flows could maximize aggregate output in principle, they will not if liquidity constraints force farmers to sell their family labor for wages below the marginal product of labor on their own land.³

The intervention we study was designed to isolate the impacts of lowering the cost of accessing liquidity – smoothing seasonal consumption – during the hungry season on local labor market outcomes and agricultural production. In the first year of the study, selected households in two-thirds of villages were eligible for hungry season loans; in the second year, 50 percent of 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).⁴ Despite an implicit interest rate of 4.5 percent per month, more than 98 percent of eligible farmers took up the offer. Close to 95 percent of loans were repaid in the first year, and 98 percent of farmers offered the same loan in the second year signed up again, highlighting both the high demand for hungry season credit and the high cost of alternative financing options.

Consistent with our theoretical framework, we find that the intervention not only increased hungry season consumption but also led to substantial adjustments in household labor allocation among households eligible for treatment. The likelihood that a family sold any *ganyu* (labor) during the hungry season fell by 3 percentage points (10 percent) in response to treatment, with a 23 percent reduction in hours sold, on average. The likelihood of hiring *ganyu* increased by around 2.6 percentage points (25 percent). As a result, average village-level wages, which we measure as daily earnings, increased by about 2 Kwacha, or 15 percent relative to control villages. Agricultural output increased by 8 percent in response to treatment.

We find no change in the use of other inputs, such as pesticides or fertilizer, suggesting that changes in the quantity and quality of labor drive the impact on agricultural production. Back of the envelope calculations indicate that about 50 percent of the loan value was used for additional labor inputs on-farm, i.e. for reallocating family from *ganyu* to on-farm work and for hiring additional labor, while the remainder went to other household expenditures, including increased hungry season consumption. The measured increase in labor inputs fully explain the impact on agricultural output at an marginal product of labor equal to the daily earnings associated with casual labor during the hungry season.⁵ We observe the greatest impact on both labor inputs and agricultural output

³A largely qualitative literature suggests that the local labor markets we study are often associated with deviations from household income maximization (Kerr 2005; Bryceson 2006; Orr et al. 2009; Michaelowa et al. 2010; Cole and Hoon 2013). These papers suggest that off-farm labor supply during the hungry season may instead be driven by immediate consumption needs.

⁴On average, around 50 percent of households, or up to 22 per village, were randomly selected for both data collection and treatment. Treating only selected farmers within the village results in a “partial treatment, general equilibrium” set-up, 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.

⁵A large share of the increase in total labor inputs come from more family labor supplied to the household’s own fields. The magnitude of this increase is larger than the measured reduction in off-farm labor sales, implying an

among the households with the fewest liquid resources at baseline, who report the highest interest rates, and thus benefit most from subsidized loans.⁶

Overall, our results are consistent with the idea that frictions in capital markets contribute to seasonal liquidity constraints, which drive local labor allocation during the hungry season. We examine several alternative interpretations of our findings, and test other margins of adjustment. First, we test whether the anticipation of seasonal constraints affects decisions at planting by informing a subset of farmers about the loans at the beginning of the agricultural season in the second year of the program. While our power to detect differential impacts in this subset is limited, we find that knowing about the loans at planting time led to larger increases in agricultural output, due to more cash crop production and more land cultivated. Second, we find little effect of treatment on other consumption smoothing strategies, including livestock or asset sales, green maize consumption or borrowing from friends and family, though we see some decrease in the likelihood of high interest borrowing from moneylenders. Third, we rule out that the income transfer implicit in the subsidized loans could explain the results. A sub-sample of 172 farmers across 11 villages were given a cash transfer corresponding to the implicit value of the program measured in choice experiments. Finally, we collect a number of checks for bias in our self reported outcome measures, and find no evidence that mis-reporting contributes to our results.

Our paper is closely related to an extensive literature highlighting the links between credit market frictions, agricultural labor markets, and aggregate output. Our theoretical model 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. Studying the general equilibrium effects of a shock to borrowing costs in India, Breza and Kinnan (2016) document reductions in credit supply lead to large reductions in local casual wage rates and consumption. On the extensive margin, Bandiera et al. (2017) find that high borrowing costs prevent poor women in Bangladesh from accessing high income labor market opportunities. 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. Our model differs from this literature by explicitly focusing on households' trade-off between labor inputs on their own land and the period-specific value of revenue obtainable through labor sales on other farms.⁷ The critical role

increase in family labor supply, which contradicts the predictions of the model. Any extension to our model that adds an additional cost to off-farm labor sales can reconcile these results. We speculate that several additional channels may be at work, including search time and psychological costs of casual labor supply (Kerr 2005; Cole and Hoon 2013), and impacts of nutrition and health on energy for on-farm work (e.g., Pitt and Rosenzweig 1986; Strauss 1986; Behrman et al. 1997; Schofield 2013).

⁶Heterogeneity in effective interest rates across farmers can either arise from a wedge between the interest rates on saving and borrowing, where farmers able to consume out of savings face lower effective borrowing costs, or from heterogeneity in the cost of borrowing that depends, for example, on the amount borrowed or on available collateral. These sources of heterogeneity will tend to result in higher interest rates for poorer farmers (Conning and Udry 2007).

⁷Like 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. Our model reflects the sequential nature of agricultural production, which may be subject to period-specific constraints (Behrman et al. 1997; Skoufias 1996).

of family labor sales for smoothing consumption has been documented by others (Kochar 1995, 1999; Rose 2001; Ito and Kurosaki 2009).⁸ 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.⁹ Second, we show that these 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-community inequality.

Our study also relates, more broadly, to recent literature testing the impacts of capital constraints on agricultural productivity. In Ghana, Karlan et al. (2014) find no evidence that capital constraints impede agricultural investments. On the other hand, Beaman et al. (2014) find that relaxing credit constraints through grants increases agricultural investment and yields among rice farmers in Mali, but that the same is not true for loans. These studies focus on capital inputs (seeds, fertilizer or pesticides) as the primary mechanism through which credit impacts yields; we show that labor inputs can also be highly sensitive to capital access.

The results presented here also relate to a recent series of papers describing interventions targeting income, price and consumption seasonality in agricultural markets.¹⁰ Burke et al. (2018) offered farmers in Kenya loans at harvest that allowed them to exploit seasonal variation in maize prices and find significant effects on total maize revenues and household expenditures. Relatedly, Aggarwal et al. (2017) subsidize better storage technologies to provide similar price arbitrage opportunities. Gross et al. (2017) stabilize lean season grain prices through improved (community level) storage and find positive effects on nutrition and health. Bryan, Chowdhury and Mobarak (2014) show that providing credit and grants leads to large increases in seasonal labor migration in Bangladesh, arguing that credit market failures and highly uncertain returns likely keep long-distance labor supply below optimal levels. Basu and Wong (2015) 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 within-community labor allocation or production. Our findings contribute to this growing literature by providing the first direct evidence that hungry season loans aimed largely at consumption smoothing affect agricultural labor market outcomes, aggregate output and inequality within communities. Fink et al. (2014) describe a pilot for the

⁸This 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)).

⁹A small literature shows reductions in labor supply in response to (anticipated) cash transfers in both developing (Edmonds 2006) and developed (Yang 2018) countries, and posits that liquidity constraints offer the best explanation for the observed patterns of behavior.

¹⁰The relationship between income seasonality and consumption smoothing is a subject of some debate. While some studies suggest that precautionary savings are sufficient to smooth consumption even if income is highly seasonal (Paxson 1993; Chaudhuri and Paxson 2002; Jacoby and Skoufias 1998), others have highlighted the pronounced consumption differences over the year as hard to reconcile with optimal smoothing (Dercon and Krishnan 2000; Khandker 2012). We take consumption seasonality as given and are agnostic as to its origins, though both our model and our empirical results show that high borrowing costs contribute to seasonal fluctuations in consumption.

current study, using a different (smaller) sample and slightly different intervention, and show results consistent with those presented here.

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 et al. 2017) about effects outside of our study sample and design. First, we simulate program impacts on untreated households in intervention villages, who are affected by treatment through changes in labor supply and demand among treated farmers.¹¹ Among untreated farmers in intervention villages, consumption rises as a result of higher equilibrium wages, but agricultural output falls due to lower labor inputs.¹² Second, we simulate the effects of scaling up credit access to all households in a community, which leads to larger equilibrium wage adjustments, since additional labor demand must be met by supply from other treated farmers. Greater treatment coverage also implies a more homogenous interest rates with the village, resulting in greater reductions in income inequality. The potential for scaling up seasonal consumption loans is hindered by high transaction costs. Bundling seasonal loans with other technologies, such as digital borrowing platforms, or piggybacking on existing rural networks may help bring down costs. Other strategies for lowering the cost of seasonal consumption smoothing, such as more secure savings, may also decrease reliance on family labor for financing hungry season consumption.

The paper proceeds as follows. In the next section, we present a simple model that highlights the links between seasonal liquidity constraints, labor allocation and agricultural output, and generates testable predictions for our empirical analysis. Section 3 provides information on the local context as well as the experimental design. We present the data and descriptive statistics in Section 4 and the experimental results in Section 5. Section 6 reconciles magnitudes, shows robustness checks and further explores alternative explanations. We discuss simulation results and welfare implications, in Section 7. Section 8 concludes.

2 An agrarian economy with capital market frictions

We study a rural agricultural economy, in which farming households maximize utility over consumption and leisure, and have access to local labor and capital markets. The objectives of the model are to 1) illustrate the links between seasonal liquidity, credit, labor markets and agricultural production in this setting, and 2) to generate predictions regarding household and village-level responses to exogenous shifts in credit market interest rates.

¹¹Changes in grain prices would also affect untreated farmers in treatment villages. We test for and find no impact on grain prices as a result of treatment.

¹²Note that, in the model, consumption is the most appropriate welfare proxy, and captures income from both wages and agricultural production net of hiring costs.

2.1 Setup

Our theoretical framework builds on the agrarian labor market model introduced in Jayachandran (2006).¹³ Each village economy has a finite number N of households that maximize utility over two periods ($t = 1, 2$). Each household i has initial liquid resources S_{i0} .¹⁴ All households have the same endowment of land, k , and time, \bar{h} , which they allocate between labor, h_i , and leisure, l_i , in period 1 (the farming period). In the second period (the harvest period), they consume their harvest production net of outstanding debt. Production y is Cobb-Douglas in labor d_i and land k , and proportional to the farm's productivity A_i :

$$y(d_i, k) = A_i d_i^\beta k^{1-\beta}. \quad (1)$$

Total on-farm labor input (farm-level labor demand) d_i includes both own (family) labor on farm and hired labor; $\beta \in (0, 1)$ defines the relative productivity of labor and land.

Households have Stone-Geary preferences over consumption and leisure. Period-specific utility is given by

$$u(c_{it}, l_i) = \log(c_{it} - \underline{c}) + \frac{1-\alpha}{\alpha} \log(l_i), \quad (2)$$

where $\alpha \in (0, 1)$, and $\underline{c} > 0$ is the minimum (subsistence) level all households must consume. Utility is additive and separable across the two periods; second period utility is discounted by a subjective discount factor $\rho < 1$.

Households can save at a rate r^s and borrow at a rate r_i^b ; due to frictions in local capital markets we assume that $r_i^b > r^s$.¹⁵ All farmers have access to the same saving technology. Interest rates on borrowing may vary across farmers and are assumed to decrease with the farm's initial resources, S_{i0} , i.e. $\frac{\partial r_i^b}{\partial S_{i0}} < 0$.¹⁶ All borrowing needs to be repaid by the end of the second period. The labor

¹³We 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.

¹⁴This 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 predetermined and positively correlated with farm productivity A_i (and verify this assumption empirically, see Appendix 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 of income redistribution would be necessary. We observe a fair amount of churning in within-village hungry season resources: across years, the within-household correlation in hungry season liquid resource ranking is 0.49 and in agricultural output is 0.63.

¹⁵We discuss these frictions in greater detail in Section 4. For example, transaction costs in the formal credit market include high transport costs, the absence of legal infrastructure, and low levels of assets to be used as collateral. Informal saving options include the storage of grain and the purchase of livestock, both of which are subject to substantial risk, including fire, theft and and pests or disease. Alternative borrowing strategies such as the temporary sale of assets or livestock are also costly because of correlated rural shocks and large distances to urban markets.

¹⁶This assumption is consistent with any model where the expected ability to repay increases with collateral (which

market clears at the endogenous wage w such that average farm labor input equals average labor supply:

$$\sum_{i=1}^N d_i(w) = \sum_{i=1}^N (\bar{h} - l_i(w)). \quad (3)$$

2.2 Household utility maximization

Rational households maximize utility from consumption c_{it} and leisure l_i over two periods:

$$\max_{c, l} \log(c_{i1} - \underline{c}) + \frac{1 - \alpha}{\alpha} \log(l_i) + \rho \log(c_{i2} - \underline{c}) \quad (4)$$

subject to

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

$$c_{i2} \leq y_i(d_i) - B_i[(1 + r_i^b)1(B_i > 0) - (1 + r^s)1(B_i < 0)],$$

where B_i is net resources borrowed ($B_i > 0$ implies borrowing and $B_i < 0$ implies saving) during the first period. In period 1, households optimally choose labor inputs d_i on the farm and their own labor supply, $h_i = \bar{h} - l_i$. In period 2, households receive harvest income y_i . Period 2 net income (and consumption) is given by harvest income plus period 1 borrowing or savings times the respective interest rate. Period 1 consumption can be financed through initial savings (S_{i0}), labor income (the labor endowment, \bar{h} , minus time allocated to leisure, l_i , and minus on-farm labor inputs, d_i , times the wage rate, w), as well as through borrowing (B_i).

2.2.1 Optimal labor input

For any given wage and interest rate farmers will always choose labor inputs such that the discounted marginal product of labor earned in period 2 equals the period 1 wage:

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

where

$$r_i^e = \begin{cases} 1 + r^b(S_{i0}) & \text{if } B_i > 0 \\ 1 + r^s & \text{if } B_i < 0 \end{cases}$$

is the effective interest rate. Re-arranging equation 5, optimal labor input can be expressed as function of productivity A_i , wage w , and the effective interest rate:

is proxied by S_{i0}) and decreases with loan size (which is also a function of S_{i0}).

$$d_i^* = k \left(\frac{\beta A_i}{w r_i^e} \right)^{\frac{1}{1-\beta}}. \quad (6)$$

Since farmers choose labor inputs such that the return on labor matches the marginal cost of capital, optimal labor inputs, and thus community level labor demand, decrease with interest rates.

2.2.2 Optimal consumption, labor supply and leisure

To maximize utility, the usual inter-temporal optimality condition for consumption must hold:

$$\frac{c_2 - \underline{c}}{\rho(c_1 - \underline{c})} = r_i^e. \quad (7)$$

Taking the derivation of (4) with respect to leisure and consumption, we can derive an interior solution for the optimal allocation of time toward labor, $h_i = \bar{h} - l_i$:¹⁷

$$h_i^* = \frac{\frac{\alpha}{1-\alpha}(1+\rho)\bar{h} + \frac{1}{w} \left[\underline{c}(1 + \frac{1}{r_i^e}) - S_i^* \right]}{(1 + (1+\rho)\frac{\alpha}{1-\alpha})}, \quad (8)$$

where $S_i^* = S_{i0} + \frac{1}{r_i^e} y_i(d_i^*(r_i^e, w)) - d_i^*(r_i^e, w)w$ are total household resources made up of the initial resources, S_{i0} , and the discounted value of farm production net of labor costs. Optimal consumption levels in period 1 and period 2 are given by

$$c_1^* = \frac{\underline{c}(\frac{1}{r_i^e} + \frac{\alpha}{1-\alpha} + \rho) + S_i^* + \bar{h}w}{\frac{1}{1-\alpha} + \rho} \quad (9)$$

and

$$c_2^* = \frac{\underline{c}(\frac{1}{1-\alpha} - r_i^e \rho) + \rho r_i^e S_i^* + \bar{h} \rho r_i^e w}{\frac{1}{1-\alpha} + \rho} \quad (10)$$

respectively, where seasonality is defined as the ratio of period 2 to period 1 consumption.¹⁸

2.2.3 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, $\hat{r} < r_i^e$.¹⁹ We derived the impacts on labor allocation, wages, agricultural output and consumption among the

¹⁷See Appendix A.1 for derivations of the optimality conditions for household labor supply and period-specific consumption.

¹⁸Higher 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 6.3.1.

¹⁹We use loan take-up rates to validate this assertion.

subset of treated farmers (see Appendix A.1 for further derivations). We discuss the effects on untreated (without access to \hat{r}) farmers in the same labor market, and the effects of offering \hat{r} to all households in the market, in Section 2.2.4.

From equation (6) lower interest rates always increase demand for labor inputs, holding wages constant, i.e., $\partial d_i^* / \partial r_i^e < 0$. Similarly, by equation (8), and again holding wages constant, $\partial h_i^* / \partial r_i^e > 0$ because higher interest rates always reduce net disposable resources of borrowing farmers ($\frac{\partial S_i^*}{\partial r_i^e} < 0$). For the labor market clearing condition to hold, wages must therefore rise in equilibrium.

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

The increase in on-farm labor demand d_i^* will increase agricultural output ($\partial y^* / \partial d^* > 0$), as long as the income effects (reductions in labor supply in response to higher wages) are sufficiently small (see Appendix A.1).

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

Given that both income ($\frac{\partial S_i^*}{\partial r_i^e}$) and substitution (through equation 7) effects are positive for first period consumption, lower interest rates will always increase first period consumption. Deriving c_1^* with respect to r_i^e , yields

$$\frac{\partial c_1^*}{\partial r_i^e} = \frac{-\frac{c}{(r_i^e)^2} + \frac{\partial S_i^*}{\partial r_i^e}}{\frac{1}{1-\alpha} + \rho}.$$

Given that $\frac{\partial S_i^*}{\partial r_i^e} < 0$, first period consumption will always increase as interest rates fall. Consumption seasonality (c_2/c_1) is given by

$$\frac{c_2^*}{c_1^*} = \frac{\underline{c}(\frac{1}{1-\alpha} - r_i^e \rho) + \rho r_i^e S_i^* + \bar{h} \rho r_i^e w}{\underline{c}(\frac{1}{r_i^e} + \frac{\alpha}{1-\alpha} + \rho) + S_i^* + \bar{h} w}. \quad (11)$$

The derivative of this expression with respect to r_i^e is positive because positive income effects for second period consumption are partially offset by temporal substitution toward the first period. Therefore, borrowing at $\hat{r} < r_i^e$ reduces consumption seasonality.

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

Since effective interest rates are 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} ; accordingly, marginal adjustments in consumption, labor inputs and output should be largest for the farmers with the lowest S_{i0} . Since consumption and agricultural output are both increasing in S_{i0} , this also results in a decline in consumption and agricultural income inequality among farmers borrowing at \hat{r} .

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

Among farmers borrowing at $\hat{r} < r_i^e$, both the mean and the variance in effective interest rates decline with the introduction of \hat{r} . Since the marginal product of labor shown in equation (5) depends directly on r_i^e , the variance in the marginal product of labor among farmers with access to borrowing at \hat{r} will also decline, resulting in a more efficient allocation of labor. Of course, if other farmers within the same labor market do not have access to \hat{r} , the overall efficiency impacts are ambiguous. We turn to these next.

2.2.4 Credit access and general equilibrium effects

Our predictions focus on treated farmers in a “partial treatment, general equilibrium” framework, in keeping with our design and data collection. 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 community – requires an assessment of spillovers to untreated farmers. In relative terms, the expected changes on 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 wage response will be more moderate relative to a full treatment scenario.²⁰

In Section 7, we engage in “structured speculation” (Banerjee et al. 2017) to explore the broader welfare implications of our experimental results. We use a numerical simulation of our theoretical model, calibrated to our study setting, to (1) assess program impact on untreated farmers in our partial treatment setup, and (2) to simulate outcomes under a scaled up policy version of the intervention that lowers credit market interest rates for all households in a community.

3 Experimental design and implementation

We turn now to our experimental setting, design and implementation. We offer further detail on our study setting in Section 4.2.

3.1 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 lives in rural areas, with small-scale farming as primary source of

²⁰This 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.

income. According to the 2010 Living Conditions Monitoring Survey (CSO 2010), average monthly expenditure of rural households in Chipata was estimated at USD 122 in 2010 (USD 0.8 per person-day), or about one third of the national average (USD 389).

The study implementation targeted small-scale farmers, i.e., households growing crops on 5 hectares (12 acres) or less. 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, omitting the those that contain Chipata town, using on village lists from the Ministry of Agriculture’s farm registry and omitting villages with less than 20 or more than 100 farms. Study enumerators visited sampled villages to screen for eligibility.²¹ 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.²² 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). 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 be classified as small scale farmer (N=146).

3.2 Experimental design

The study took place over two years and was designed to coincide with the agricultural cycle (see Appendix figure B.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 referred to as the “hungry season” or “lean season”. In April, early crops start to become available and harvest begins in earnest in May. Between August

²¹Villages were ineligible if: (1) other projects had been conducted there in the recent past, (2) the village bordered a village that was in the study pilot, (3) the village bordered a village already listed, (4) the village had fewer than 17 households, or (5) it was impossible to get a 4x4 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.

²²We 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.

and October, few agricultural activities take place. We refer to study year 1 as covering the 2013-14 agricultural cycle and study year 2 covering the 2014-15 agricultural cycle. The study design is summarized in Appendix figure B.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 offer trade-offs. On the one hand, providing the staple food offers a direct way of targeting food shortages. On the other hand, cash offers a more flexible alternative that can better 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 2014. Of the 175 study villages, 58 (1033 farms) were assigned to a control group, which received no intervention, 58 (1092 farms) were assigned to the cash loan treatment, and 59 (1095 farms) were assigned to a maize loan treatment in the first year of the program. 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 year, and to separate the impact of repeated treatment from first time treatment.

To measure the extent to which farmers adjust their production plan with earlier knowledge of hungry season credit, we also varied the timing of the loan announcement in the second year of the program. Half (40) of the treated villages received notification before the start of the planting season, in September, while the other half of treated villages was only informed about the loan program in January. In addition to the loan treatments, a small number of villages (6 villages / 91 farms in year 1 and 5 villages / 81 farms in year 2) were assigned to an *income effect control* group, which provided a cash grant of 60 Kwacha, corresponding to the median value assigned to participation in the loan groups in choice experiments.²³ Cash grant villages were randomly selected from villages initially assigned to the control group.

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 (\sim USD 33), which corresponded approximately to the value of the three maize bags at official government prices (65 Kwacha per bag) at baseline. 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

²³For further details on choice experiments, see Appendix C.2.

260 Kwacha (or a mix at K65 per bag). Villages randomly selected for the “cash only” repayment program in the second year of the study had to repay 260 Kwacha.²⁴ 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 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 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, and on what repayment modality farmers chose.

To make the two loan programs as comparable as possible, we conducted a series of hypothetical choice experiments in villages outside of the study sample but within the study area in November 2013. In these choice experiments, respondents (N=72) were asked a series of dichotomous choice questions on whether they would prefer a loan of three bags of maize over a cash loan of x Kwacha, with x varied between 50 and 600 Kwacha. 84.7 percent of respondents preferred a maize loan over a cash loan of 175 Kwacha, while only 36 percent preferred the maize loan over a cash loan of 250 Kwacha. As part of these choice experiments, we also asked about timing and acceptable interest rates. Specifically, respondents were asked if they would take up a maize (cash) loan that paid 3 bags (200 Kwacha) in January with a repayment of 4 bags (265 Kwacha) due in subsequent months. While only 27.8 (maize) and 20.8 (cash) respondents were interested in a loan with repayments in May, acceptance rate jumped to 81.9 and 83.3 with repayment in June for maize and cash loans, respectively. Responses to these questions determined final loan contract details. Further detail on the implementation of the choice experiments is provided in Appendix C.2.

3.3 Implementation

Both loan treatments were administered under the “Chipata Loan Project (CLP)” to distinguish loan operations from the survey visits conducted by Innovations for Poverty Action (IPA). This distinction between the CLP and IPA brands and staffing was intended to assure participants that survey responses would not affect loan eligibility. We also ensured that staff members working on loan implementation did not do household surveys to minimize the risk of surveyor bias. All study households in villages randomly selected for treatment were eligible for loans in the first year. In year 2, the same rules applied. In villages treated in both years, i.e., in villages where loans were offered again in the second year of the study, eligibility was further restricted to households who fully repaid in year 1. Loan program participation did not affect any of the survey activities. The loan intervention was announced to households during a village meeting to which eligible

²⁴Requiring 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.

households were invited.²⁵ At the meeting, project staff described the terms of the loan and logistics surrounding distribution and repayment. Loan enrollment and consent forms were provided to eligible households.

Loans were distributed between 3 days and one week after the village meeting at a location convenient for transportation. Project staff registered attendees, confirmed their identity using the national registration card,²⁶ and collected their signed enrollment and consent forms. Before finalizing the transaction, project staff confirmed that the participant understood the terms of the loan. The loans (cash or maize) were handed over and a receipt was provided to the household and kept for project records.

Repayment was due in early July. Villages were notified in advance about the date of repayment as well as the central locations at which repayment would be collected. Two attempts at collecting repayment were made. Households were provided with a repayment receipt upon full repayment. 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. Further summary statistics on repayment patterns are described below.

3.3.1 Randomization

In year 1, 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 5.1. The smallest p-value for the pairwise comparisons observed in the final draw was $p = 0.213$. In year 2, treatment assignment 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, with assignment to both the main treatment arms (control, cash loan and maize loan) and the sub-treatments (income effect control, early notification and cash repayment).

3.3.2 Attrition and selection

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

²⁵Ineligible households were not barred from listening in. Eligible 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). In addition, the baseline data for 3 households who were surveyed was lost. They are dropped from the sample.

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.

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. 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 non-random attrition (column 3). In year 1, there was no selection into meeting attendance or eligibility. In year 2, there was some modest selection into meeting attendance (over 90 percent attendance in all treatments and sub-treatments), and no further selection into eligibility. Column 4 of Appendix table B.3 also previews our take-up results, which we turn to next.

3.3.3 Take up and repayment

Table 1 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.

Repayment was substantially lower in year 2, with an average repayment rate of 80 percent in villages receiving the program for the first time. The decline in repayment appears to be driven in part by worse rainfall patterns and lower overall agricultural output in 2015. In addition to differences in harvest incomes, 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 also suggest some learning about enforcement. In the absence of legal consequences – the only punishment was ineligibility in future years – defaulting on the loan may have been a rational choice even if it excluded farmers from future borrowing.

4 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 modeling choices.

4.1 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 collected on an ongoing rolling sample. For these surveys, smaller survey teams visited a random (rotating) 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 some outcomes vary because the length of the short recall survey varied over the course of the project (see Appendix C.1). A total of 15,044 observations from the sample of 3,139 households were collected over the course of the study. Appendix C.1 summarizes sample sizes and key content collected in each survey.

4.1.1 Outcome measures

We focus on three main outcome types, based on the predictions in our conceptual framework:²⁷ (1) labor allocation and daily earnings, (2) agricultural output, and (3) 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 to construct labor allocation measures over the week prior to the survey during the hungry season. Labor allocation outcomes include (a) family labor sold to other farms (ganyu sold), (b) labor labor purchased (ganyu hired) and (c) family labor invested on-farm. We construct measures on both the extensive margin and as a continuous measure, at the household level. Our continuous measure is in hours, summed across all individuals in the household (i.e. a total household hours measure), to account for the fact that ganyu does not always last the full day and a partial day of ganyu sold might still allow for some time invested on-farm.

We also 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 to address outlier observations 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 median earnings within the village. Villages with no ganyu activities reported in a month receive a missing value.

To measure agricultural output, farmers were asked to report, by crop, output in kilograms as well as the the total value of the harvest, including early consumption and crops still on the field at

²⁷Our data collection and analysis follows a pre-analysis plan available at <https://www.socialscisearch.org/trials/130/>.

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 an alternative value measure based on own reported prices to allow treatment to affect effort seeking out better prices, for example.

Our main consumption measure 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 points to severe food shortages, and has the advantage of being relatively easy to measure. We collected this outcome over a two-day recall period during the a short-recall surveys. We supplement the measure of meals consumed with a measure of grain available for consumption, which we code as a binary indicator that equals one if the household had fully depleted their grain reserves at the time of the survey. 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 several additional outcomes in Section 6, which we describe as they arise.

4.1.2 Heterogeneity measures

We categorize households by their liquid resources using baseline savings in 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 use reported interest rates at baseline to examine the robustness of our findings.

Appendix figure B.3 shows the correlation of initial resources with self-reported baseline interest rates. The gradient is monotonically negative (higher resources correspond to lower reported interest rates) for households with above median baseline resources, and flat for households with below median resources (all of whom have very little liquid savings). The shape of this relationship helps determine our functional form assumptions in our analysis of heterogeneity by initial resources.

²⁸We pre-specified analyzing heterogeneous treatment effects by baseline grain and cash resources but not by baseline reported interest rates. We also pre-specified 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 4.2), since the allocation of land depends both on productivity and favoritism. Results are available upon request.

4.2 Descriptive statistics

Seasonality in resources and consumption As illustrated in Figure 1, households draw down their cash and grain reserves early in the year, and begin to replenish them in April and May. The period between January and March is referred to as the “hungry season” throughout rural Zambia. It also coincides with the time when farmers have crops on their fields and on-field activities (particularly weeding) peak, as illustrated in the agricultural calendar shown in Appendix figure B.1. 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. While grain and cash reserves are low on average, we document substantial variation in reserves and in households’ ability to smooth resources and consumption across periods. Figure 1 shows seasonal patterns in cash and grain availability and consumption for the top and bottom quartile of the control group in terms of baseline food and cash reserves (liquid resources). The top panel shows median cash (left) and grain (right) reserves, measured in 100 Kwacha, per month. The bottom panel shows median total liquidity (cash + grain, left) and the mean number of meals consumed per day by adult members of the household (right). These measures show considerable seasonality, which – consistent with the model described in Section 2 – is more severe for households with low resource availability than for high resource households.

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 14 USD. 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 money-lenders, 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 endline). 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. 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 Appendix figure B.3). In the hungry season, sixty percent of control group households report that they would be unable to borrow 50 Kwacha in cash that day.

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.²⁹ 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 to rely on 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 also occur, and are typically completed by families or groups of individuals. Further detail on how ganyu participation varies with age and gender of household members is discussed in Fink et al. (2014), which describes a pilot for the current study.³⁰

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. 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

²⁹In 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”.

³⁰Additional 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).

quartile did.³¹

At baseline, around two-thirds of households anticipate 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. The anticipated nature of ganyu labor is an important departure from most of the existing literature which has focused on family labor sales as an ex post mechanism for smoothing transitory shocks (e.g., Kochar 1999).

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 community 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 and immediately after harvest, 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 still 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. In general, most farmers do not appear to view land access as a constraint to production. At 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. The average acreage reported available for cultivation was 6.3, but on average only 4.5 acres of land were used for growing crops. We therefore choose to abstract from per acre measures of agricultural output

³¹Conditional 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.

(yield) and focus instead on total agricultural output at the household level.

5 Experimental results

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

5.1 Empirical strategy

We estimate intention-to-treat regressions, including all households regardless of whether they selected into the loan. 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 begin with a pooled specification that averages across years and loan treatment arms:

$$y_{ivt} = \alpha + \beta loan_{vt} + \tau_t + X_{i0}\delta + u_{ivt} \quad (12)$$

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 time dummies and X_{i0} are household-level controls, measured at baseline. Time units t are years for our long recall surveys and month-years for our short recall surveys. Treatment assignment varies over time according to the treatment rotation between years, as described in Section 3.2, and our main analysis estimates the average impact of first-time treatment in either year (i.e., we exclude villages treated in year 1 from the year 2 data). Errors are clustered at the level of the randomization unit, the village v , which addresses both unobserved village level shocks and serial correlation across study years. We also show separate results by treatment arm, using estimating equation (12) and allow β to differ for the cash and maize loans.

We provide a number of supplementary analyses to further investigate potential mechanisms and also to address alternative interpretations of our main results. First, we test heterogeneous treatment effects by baseline resource availability. To estimate how treatment effects vary over the distribution of baseline grain and cash reserves, we estimate the following model:

$$y_{ivt} = \alpha + \beta_1 loan_{vt} + \beta_2 S_i + \beta_3 S_i^2 + \beta_4 loan_{vt} \times S_i + \beta_5 loan_{vt} \times S_i^2 + \tau_t + u_{ivt} \quad (13)$$

where S are liquid resources as discussed in Section 4.1.2. We summarize the results by plotting adjusted predictions evaluated at the mean S in each quartile of the baseline distribution.³²

³²We omit baseline controls to allow for other household characteristics correlated with baseline grain and cash reserves to vary by quartile. We 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 presented in this section for year 2 (excluding villages treated in year 1) is shown in Appendix B. To examine robustness of these results, we also use reported interest rates at baseline

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

Second, we examine the degree to which prior knowledge of seasonal shortages affects agricultural decisions by varying the timing of loan announcement (holding fixed the timing of loan delivery) in the second treatment year. Third, we investigate persistence by phasing some farmers into and some farmers out of the program in year 2. In our main results, we focus on villages assigned to receive treatment for the first time and exclude year 2 data from villages that were treated in year 1. We analyze villages with repeated or discontinued programs separately to investigate persistence of the treatment effects by interacting an indicator for a loan treatment in year 2 with treatment status in year 1.³³ Finally, we show treatment effects for the impacts of a small cash grant to rule out income effects.

Our regression estimates identify the causal effect of the loan under the identifying assumption that treatment assignment is orthogonal to u_{iwt} . 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 4-6, respectively). Column 7 shows the largest pairwise t-statistic and column 8 the largest pairwise normalized difference. Normalized differences are small and all fall below the rule of thumb cutoff of 0.25 (Imbens and Wooldridge 2009), despite some marginal differences across treatments. Overall, the randomization successfully balanced households across treatment arms. The variables shown in Appendix table B.5, up through crop diversity, were used in the randomization to test balance; the extended set of control variables are used throughout the analysis. Results are similar if only variables used in the randomization are included as controls.

5.2 Impact of lowering borrowing rates through seasonal loans

The model presented in Section 2 predicts labor, consumption and production impacts from changing borrowing rates. We test the empirical predictions described in Section 2.2.4, based on access to lower interest loans for a sub-set of farmers in the labor market.

Prediction 1: Labor demand increases and labor supply decreases, causing equilibrium wages to increase We examine impacts on labor allocation in Table 2. Panel A shows the average treatment effect for both loan arms, and Panel B estimates separate effects by treatment arm. We present and discuss all of our average treatment effect tables with a similar layout.

Our main prediction is that loans will reduce net labor supply for treated farmers through

instead of S in estimating equation 13.

³³Our pre-analysis plan was not clear about how we would analyze persistence. Our pre-specification of effect of persistence focused on results among villages treated in year 1 but not year 2 rather than the effect of repeated treatment. The analysis presented separately by first-time versus repeated treatment allows for clearer interpretation of the results, particularly given the somewhat surprising counterproductive effect of repeated treatment described below.

decreases in family labor sales (ganyu sold) and/or an increase in hiring (ganyu hired). 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. Labor outcomes are for the entire household. Column 1 shows that the likelihood of selling ganyu falls by around 3 percentage points, on average, off of a mean of 31 percent in the control group. On average, this extensive margin effect is not significantly different from zero, though the 4.2 percentage point decrease in the maize loan treatment arm (Panel B) has a p-value < 0.10 . Turning to the continuous measure of total family hours sold, treated households sell an average of 0.9 fewer hours or 23 percent less ganyu per week (column 2), on average. The effect on hours sold is also (insignificantly) larger in the maize loan treatment arm.

Treatment has positive, and marginally significant, effects on hiring, on both the extensive margin (column 3) and continuous measures (column 4, collected only for a sub-set of households in an extended hungry season survey module). The likelihood of any hiring (column 3) increased by 2.6 percentage points, off of a mean of 10.5 percent, and the number of hours hired increases by around 1.5 hours or 45 percent (column 4). The impact is insignificantly larger in the cash loan arm (Panel B) for both outcomes.

The decreases in net labor supply predicted by the model reflect both lower labor supply and higher on-farm labor demand. Hours of family labor invested on the farm over the past week increase by around 6.1 hours or around 12 percent (column 5), with similar impacts across treatment arms. Column 6 sums the total hours of labor (own and hired) invested on farm over the past week, and shows a 13 percent increase relative to the control group average. The estimated impact on-farm family labor inputs relative to the impact on family labor sales suggest that households actually increase labor supply by over 5 hours per week on average, in contrast with the model's predicted decrease in labor supply due to positive income effects from lowering effective interest rates. Given that ganyu work is strenuous and accompanied by substantial unmeasured search costs (and social stigma), the costs of doing ganyu may be considerably higher than the number of hours spent working; increased consumption may also have led to more energy for on-farm work. We offer further discussion of magnitudes in Section 6.

Table 3 shows our estimated wage impacts; since daily or hourly wages are not defined in the piece rate work arrangements that we observe, we use daily earnings as our proxy for local wage rates. With higher demand and lower supply, we find the expected increases in daily earnings. 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. Earnings increase by 2.6 or 2.2 Kwacha per day on average, corresponding to an 15 to 17 percent increase over the control group mean, on average (Panel A). Increases are larger in the maize arm for both measures (Panel B), where the decrease in off farm labor sales was also higher. Column 3 shows village-level regressions using median reported daily earnings at the village-month level, based on the household measure winsorized at the 95th percentile. The magnitudes line up reasonably well with the household

level data. 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 Appendix Figure B.4. The estimates are noisy, but show a positive gradient, indicating that a higher treatment intensity is associated with larger general equilibrium effects.³⁴

Prediction 2: Agricultural output increases We show average effects on log of agricultural output in Table 4. Panel A shows that agricultural production (as measured by total harvest value using nominal (column 1) and constant (column 2) prices increased by 8.1 to 8.5 percent on average, with larger, but not statistically different, point estimates for the cash loan treatment arm (11 percent vs. 5 to 6 percent increases). Given a control group mean of 3500 Kwacha at constant prices, an 8 percent increase corresponds to a treatment effect of around 280 Kwacha, which is very similar in magnitude to the loan amount, and corresponds to about 15-20 days of daily labor earnings. We discuss the relative magnitudes of the labor and input effects in further detail in Section 6 below.

Prediction 3: Period 1 consumption increases and consumption seasonality decreases Table 5 shows the treatment effect on consumption and food security measures, with a primary focus on the hungry season, when constrained households reduce their consumption most, as shown in Figure 1. To assess the impact of lower effective borrowing rates, we analyze effects on the number of adult meals consumed per day in the past week and food availability in the home, along with an index of food security.

We find substantial average improvements in consumption and food security outcomes in response to loan access. As shown in column 1, the likelihood of having any grain in the home increased by 11 percentage points, an effect that is almost entirely driven by maize loan arm. The food security index improved by around one quarter of a standard deviation, on average, with larger effects in the maize loan arm (column 2). 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 0.09 meals on average, which corresponds to about three quarters of the difference between average meals consumed in the control group between the hungry season (column 3) and the harvest season (column 4).³⁵ In contrast, at harvest, we see no effect of loan access on consumption outcomes (column 4). This means that the treatment increased consumption overall, with a shift in consumption towards the hungry season, and an overall reduction in consumption seasonality.

³⁴The slope on share of village treated should not be interpreted causally: larger villages, which differ on numerous dimensions, will have a smaller share of the village treated, given our sampling strategy. That said, controlling for important village-level labor market determinants, such as distance to road, does not substantially alter the result.

³⁵We 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.

Prediction 4: Impacts are decreasing in liquid resources To the extent that poorer households – as measured by baseline grain and cash resources (see Section 4.1.2) – face higher interest rates, they will also be relatively more affected by a change in average interest rates. This analysis focuses on year 1, for which the measure of baseline resources is most valid, and show results for year 2 in Appendix figures B.5, B.6 and B.7. Robustness checks using baseline reported interest rates and an alternative heterogeneity measure, are shown in Appendix figures B.8, B.9 and B.10.³⁶

We begin with heterogeneity in labor adjustments. Figure 2 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 for both the control and treatment groups in year 1.³⁷ The top panel shows effects on hours of ganyu sold, the middle panel shows the effect on hiring, 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. Overall, households across the distribution increase on-farm labor inputs, which suggests that non-treated farms absorb most of the additional labor demand. 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).³⁸

Second, Figure 3 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, we speculate that the poorest households, who face the highest cost of hungry season borrowing, have the highest marginal product of labor on their farms at baseline. As a result, a reallocation of labor may help equalize the marginal product of labor across farms.

Finally, we compare our main measures of consumption (food availability and adult meals per day) across the distribution of baseline resources for treatment and control villages during year 1 in Figure 4. The top panel of the figure shows that the likelihood that a household had any food reserves during the hungry season is 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 statistically insignificant. The bottom panel of the figure shows similar effects on 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

³⁶Note that baseline interest rates are negatively correlated with liquid resources, so we expect interaction terms of the opposite sign for this alternative heterogeneity measure.

³⁷The underlying regression results for these analyses are shown in Appendix table B.6.

³⁸The difference in baseline resources between the bottom and top quartiles is almost 10 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.

consumption for low resource households, and a reduction in consumption inequality as illustrated by the substantially flatter resource-consumption relationship in Figure 4. The average increase in consumption induced by the intervention in the bottom quartile 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.

6 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. We then turn to a few alternative explanations and robustness checks.

6.1 Magnitudes and mechanisms

Our main results suggest that a relative small loan, timed to coincide with seasonal shortages in food and cash availability, led to substantial increases in on-farm labor inputs, agricultural output and hungry-season consumption. While the consumption response is unsurprising given that loans were subsidized and thus associated with both positive income and substitution effects, the impacts on labor and agricultural output are notable given both the size and timing of the loan. Rather than a decrease in family labor supply due to the positive income effects created by the loan, we observe increases in family inputs on-farm that exceed the magnitude of the reduction in family labor sales off-farm. This additional family labor supply may have been driven by improvements in nutrition and energy (lower disutility from work), but could also reflect additional motivation to repay the loan or supervision needs for hired labor. Our measure of hours of labor sold also neglects search time, which is substantial. During the hungry season, households report spending two hours looking for ganyu, with average work days of only 4 to 5 hours.

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 (Column 6, Table 2). This increase is driven by increased hiring of external labor and well a substantial increase in family labor on-farm. Taking the thirteen 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 our time use questions), an additional 7 hours per week corresponds to approximately 18 additional days of labor on the farm during the hungry season. We measure a treatment effect on agricultural output of around 280 Kwacha (Table 4), which implies an average marginal product of labor of around 15 Kwacha (assuming all other inputs remain fixed), which lines up almost exactly with average daily earnings reported in Table 3.

Other adjustments, such as improvements in health and nutrition or increases in the use of other inputs may have also contributed to additional family labor investments on-farm and higher

agricultural output. While we cannot precisely partition the output effects according to these other adjustments, we can use survey measures of health outcomes and other input use to test for program impacts. Appendix table B.7 shows modest reductions in illness within the home and improvements in mental health, suggesting that family labor inputs may have increased partially due to fewer days lost to illness. Anthropometric measures are noisier and show little consistent pattern. Appendix table B.8 reports impacts on (non-labor) inputs and shows small and mostly negative changes as a result of treatment, suggesting that labor quantity and quality likely drive the impacts on agricultural output.

We can also compare the size of the loan with the cost to the household of the labor market adjustment we observe. On average, households in the control group sold 3.8 hours of ganyu per week (Column 2, Table 2). With average daily earnings in the control group of 15 Kwacha, this implies total earnings of around 148 Kwacha over the 13 weeks of the hungry season. The treatment effect on ganyu labor sales corresponds to an average of around 2.3 few days sold, or 40 Kwacha less in ganyu earnings. Adding the average additional expenditure for ganyu hiring (66 Kwacha), this implies that around 50 percent of the original loan amount (200 Kwacha) went to foregone earnings plus wages paid to hired labor.

In addition to increasing hungry season consumption and on-farm labor, households may also substitute away from other more expensive consumption smoothing strategies. We collect data on several of these and show treatment effects in Appendix table B.9. We analyze program impact on six outcomes: formal loans (column 1), low interest informal loans from friends or family (column 2), informal high interest loans, typically from moneylenders (column 3), sales of assets (column 4) and livestock (column 5) and consumption of green maize (column 6). All but the last of these were reported retrospectively at the end of each season as part of the long recall surveys, and coded into binary extensive margin variables, with villages treated in year 1 omitted from the year 2 data (i.e., we replicate the main specifications for these outcomes). Green maize consumption was recorded during the hungry season with a two week recall period.

We find no effect on input loans (column 1). The baseline mean rate of input loans is close to 40 percent, while all other forms of borrowing are relatively rare. Column 2 shows that informal low interest borrowing declined insignificantly by about 1 percentage points, on average, relative to a mean of 7 percent of control group households. Column 3 shows that the intervention had the largest (relative) impact on high interest borrowing, which declined by around 2.6 percentage points (40 percent) relative to a baseline mean of 6.5 percent, with larger effects in the maize loan treatment arm. Asset sales are unaffected, though column 5 shows small increases in livestock sales; our long recall data collection does not allow us to determine whether these occurred during the hungry season or at harvest (potentially to repay the loan). Green maize consumption (column 6) was also unaffected by treatment.

Given that farmers had to repay 260 Kwacha after harvest, on average – which almost exactly

matches the estimated average increase in agricultural output— changes in farmers’ overall financial situation at the end of the season were likely small, while welfare improvements associated with better consumption during the hungry season may have been much larger.

6.2 Additional results

As outlined in our pre-analysis plan, our study design varied a number of other program features to help understand underlying mechanisms.

6.2.1 Maize versus cash loans

In our main results, we show effects for both the maize and cash treatment arms, and observe generally statistically indistinguishable responses in both arms. 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. On the other hand, the estimated increase in agricultural outputs is slightly higher in the cash treatment arm, due in part to the effects in year 2 associated with early notification in this treatment arm. When we examine the main rationale for maize loans – the lower anticipated likelihood of loan utilization for wasteful consumption – we see very little difference in reported and observed behaviors. In Appendix table B.10 we show estimates on treatment effects with respect to a range of different (potentially wasteful) expenditure categories. We see little evidence of major expenditure adjustments in response to treatment overall, and find no evidence of increased conspicuous consumption.³⁹

6.2.2 Anticipated versus unanticipated loan availability

Table 6 reports the effects of the year 2 sub-treatment 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). If hungry season resource shortages (and interest rates) are anticipated, farmers are likely to adjust production plans to accommodate outside labor activities. 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 (the reference timing is notification at the start of the hungry season) for our main outcomes (columns 1-5) and for two supplementary outcomes relevant to planting season adjustments (columns 6-7): the acres devoted to cash crops and the total value of capital inputs applied to the household’s fields. The top set of results pools the cash and maize loan treatment arms; the bottom results shows effects for each treatment arm interacted with notification timing. To facilitate comparisons with the main

³⁹We see small increases in expenditure on sweets in the cash loan arm, driven entirely by families with young children.

results, and to increase the distinction between early and late announcement timing, villages treated in year 1 are excluded from the analysis.

While most estimates are imprecise due to the small sample size, we observe qualitatively larger impacts on labor allocation and agricultural output value in villages notified at planting about the loans (Panel A). Columns 6 and 7 suggest that some of this may be due to planting stage adjustments: households apply more capital inputs if they know the loan is coming and devote more acres to cash crops. Both the value of inputs used and the acres under cash crops in the early notification treatment (i.e., the sum of the main effect of the loan and the differential effect of early notification) is about 30 percent higher than in the control group. To put these numbers into perspective, the 250 Kwacha of additional inputs is about the value of one 50 kilogram bag of fertilizer (roughly half of the recommended fertilizer for one acre), and the impact on total acres under production is around 5 percent of the average land area under production. Breaking the results out by treatment arm (Panel B) offers suggestive evidence that planting season notification results in higher output value in the cash loan arm because of additional adjustment in inputs and cropping decisions for this subgroup. This may suggest that the cash loan notification at planting enabled some short run borrowing for planting stage investments, though we lack additional data to back up that interpretation.

6.2.3 Persistence of treatment effects

Table 7 shows the persistence of effects from year 1 into year 2. Columns 1-5 show the effects on our main outcomes by year, while columns 6 and 7 show planting season outcomes that may have been affected if year 1 treatment effects led to adjustments in production in the following year. Panel A shows the effects in year 1; Panel B breaks year 2 out by year 1 treatment status. Recall that the main results described above pool year 1 treated villages with those treated for the first time in year 2, i.e., they estimate the effect of receiving the program for the first time.

Panel A of Table 7 shows that many of our main results are driven by outcomes in year 1. Specifically, we see statistically significant decreases in hours sold (column 1), increases in hours hired (column 2), increases in family hours spent on-farm (column 3), increases in output (column 4) and consumption (column 5). Changes associated with planting season decisions show insignificant effects on agricultural inputs and acres devoted to cash crops in year 1, consistent with the timing of loan notification that year. Panel B shows similar signs and magnitudes (given the wide confidence intervals around most estimates) for villages treated for the first time for labor allocation, output and consumption. Inputs and cash crops increase (insignificantly) more in year 2 than in year 1 for villages treated for the first time, consistent with the notification timing results discussed above.

Panel B also shows year 2 results for villages treated in year 1, which are excluded from most of the main analyses. While the estimates are imprecise, we find relatively little evidence of persistence, though most signs are in the direction of the first-time treatment effects. More surprisingly, we find

no effect for villages treated twice, where farmers seem to have lower family labor on-farm and other inputs, and lower output than farmers in villages treated for the first time in year 2. In general the total effect of repeated treatment (sum of all three coefficients in Panel B) is agricultural output that is very similar to the pure control, i.e. villages that were never part of the loan program. It is possible that the lack of repeated impact is simply a result of the small sample across year 2 sub-treatments; it is however also possible that repeated loan availability led to behavioral adjustments that were outside the scope of our data collection. Finally, we note that year 2 was a very different agricultural season than year 1: comparing the control group average output value shows that farmers received about 12 percent lower income, on average, in year 2 than in year 1. Further research is needed to better understand these dynamic effects of lowering borrowing rates.

6.3 Alternative explanations and robustness checks

While our model and main results highlight the impact of lowering effective borrowing rates on households' consumption and labor decisions, the provision of loans may have also have affected other household decisions, which may in turn have influenced both labor allocation and production. We discuss other potential interpretations of our findings in this section, along with results that fall outside of our model, including some robustness checks.

6.3.1 Grain prices and transactions

Households were told that they could do what they liked with the cash and maize provided to them through the loans. Therefore, 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, affecting the value of consumption or of output. Appendix table B.11 shows the treatment effect, in year 2 only, for maize prices and transactions. Overall, we observe no effect on either prices or transaction probabilities, with point estimates very close to zero and standard errors that allow for adjustments of up to around 25 percent of the control group mean price of 1 Kwacha/kg. These results focus on year 2 since we did not start recording transactions carefully until after the year 1 harvest survey. They average across all months to improve power, but if we restrict the regressions to the hungry season months shortly after the loan was delivered, results are similar. A decline in prices does not appear to drive the main consumption findings.

6.3.2 Income effects

To distinguish temporal reallocation effects from positive income effects created by the loan intervention (due to below-market interest rates), we implemented a small cash transfer sub-treatment, as described in Section 3.2. We calibrate the size of the income effect in through a hypothetical choice experiment (see Appendix C.2) in which households were asked to value the loan versus a

cash gift. In Table B.12, we test whether a transfer of 60 Kwacha at the start of the hungry season led to a measurable effect on our main outcomes. While our power is limited due to the small sample size, we do not find any evidence for labor adjustments 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.

6.3.3 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 reporting bias and thus threaten the validity of our analysis. Even though we made an effort to separate data collection from the loan program as described above, it seems likely that at least some participating households associated the survey with the loan program. Given that participants were generally happy with the loans (and wanted more of them), responses could be biased 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, we included a social psychology scale, adapted from Marlow and Crowne (1961), designed to directly measure social desirability bias in a survey conducted during the hungry season and in our endline survey. The module includes a series of questions designed to prime social desirability bias, such as “Are you always courteous, even to people who are disagreeable?” We construct a social desirability index based on responses to the eight questions and test whether it is affected by treatment, and then test whether there are any differences in the response patterns across treatment arms. As a second robustness check, we also collected objective measures of maize output in years 1 and 2 of the study. In year 1, enumerators visited the fields of a sub-sample of respondents and measured the height of a typical maize stalk in the field. In year 2, the approach was more systematic and involved sampling maize cobs and counting the corresponding number of kernels. In both cases, we can directly compare objective measures of output with reported measures, and test whether there are any systematic differences between the two variables.

Results are shown in Table B.13. Panel A shows that the social desirability index is unaffected by treatment in either survey round. The estimated treatment effects are statistically insignificant and very close to zero. We include no baseline controls in these regressions so that we can interpret the R-squared associated with the regression. Treatment explains almost none of the variation in the social desirability index measure.

In order to test whether treatment eligibility affects reported agricultural output, we regress self reported maize production on our objective measure of productivity, the loan treatment and the interaction of the two, along with a measure of the share of the maize output that comes from hybrid maize in Panel B. The interaction term is the regressor of interest, since it shows whether the objective measure is differentially predictive of self reported outcomes among treated farmers.

If treatment leads to biased self reporting, we expect this term to be negative. In year 1, the sign is positive. In year 2, it is negative, but small relative to the raw correlation between the objective measure and the self report. Furthermore, the overall variation in the outcome explained by the regressors is unchanged when the treatment variable is added (move from column 1 to 2 and 3 to 4). Both of these tests suggest that treatment status did not alter the reporting behavior of study participants.

7 Model calibration and welfare implications

While our main results imply substantial benefits for treated farmers, the same is not necessarily true for untreated farmers in intervention communities. Untreated farmers did not benefit from lower borrowing costs, but were affected indirectly through increases in labor demand and reductions in labor supply of treated farmers. Higher equilibrium wages help cover hungry season consumption needs for households that are net sellers of family labor, but also make it more expensive to hire labor inputs. To provide a better sense of the welfare implications for both untreated farmers and to explore effects were the program scaled to cover all farmers in a village,⁴⁰ we calibrate a simulation of our model to the empirical values observed at baseline.

Specifically, we assume that output is determined by the Cobb-Douglas function as outlined in Section 2, 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. Appendix A.2 and Appendix table A.1 provide further details on the parametric assumptions made for the initial calibrations. In Figure 5, we show simulated impacts on net labor supply, agricultural output value, consumption and utility for both treated and untreated farmers in treatment communities, assuming that 50 percent of farmers get access to subsidized loans. Results are summarized in Appendix table A.2. Similar to our empirical findings, our simulations suggest a 10 percent increase in the equilibrium wage from 17.4 to 19.1 Kwacha. Figure 5 (top panel, right) shows the effect on (gross) agricultural output value. On average, agricultural output among treated farmers increases by 8.3 percent in our simulations (almost identical to our treatment effects) while the agricultural output of non-treated farmers falls by 9.1 percent on average, since they supply more labor to the market and less to their own farms in response to higher wages. This is illustrated in the top left panel of Figure 5: net labor supply declines among treated farmers, and increases among untreated farmers in treated communities.

These adjustments by untreated households do not imply reductions in consumption and welfare.⁴¹ As illustrated by the two bottom panels of Figure 5, both hungry season consumption and total utility show increases among untreated households due to positive wage income effects. Hun-

⁴⁰In other words, we engage in “structured speculation” through the model calibration (Banerjee et al. 2017).

⁴¹In 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.

gry season consumption increases across the distribution of baseline resources for both treated and untreated farmers, with substantially larger increases among treated (13 percent) than untreated farmers (2 percent). Among untreated farmers, the utility gains are greatest at the bottom of the initial resource distribution, and slightly negative among better off households. Among treated farmers, utility gains are positive across almost the entire resource distribution, due to increases in consumption and leisure.

While many microcredit programs lower interest rates for only a sub-set 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 6 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 demand and supply effects, and results in a larger aggregate wage impact, with a simulated average daily wage of 20.7 Kwacha (compared to 19.1 Kwacha with partial treatment). Compared to partial treatment, agricultural income impacts are more moderate in this scenario: since all farmers have access to lower interest rates, labor reallocation 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 6 percent, while agricultural incomes in the top 5 percentiles decrease by about 5 percent. Average agricultural output increases by 3.2 percent, due to a more efficient allocation of labor across farms. Output inequality (Gini) falls from 0.37 to 0.35. The increase in wages as well as the positive income and substitution effects of lower interest rates lead to a 15 percent increase in hungry season consumption, on average. Total consumption increases by 3.7 percent, and consumption inequality decreases by 5.7 percent from a Gini of 22.0 to a Gini of 20.7. Total utility also increases across the distribution with the largest relative gains among the poorest.

8 Conclusions

The results presented in this paper highlight the importance of seasonal incomes, credit constraints, and liquidity that varies across the agricultural calendar for labor markets and agricultural production. From a policy perspective, providing rural farmers with subsidized access to 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 and increase average agricultural output, while reducing farmers' reliance on stigmatized piece work 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. As we show both in our theoretical framework and our empirical results, local labor market equilibria respond strongly to decreasing labor market participation by poorer farmers, resulting in higher

wages and reduced farming revenue for farmers with relatively large capital endowments. Overall output effects are positive due to increased labor inputs and an overall more efficient allocation of labor.

In the specific study area analyzed in this paper – 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 (high interest) loans we offer, even in the second year after farmers had repaid their year one 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 running rural lending schemes in settings with limited road infrastructure are large, and enforcing loan repayment can be difficult in the absence of collateral.

Thus, the loan intervention we study was not designed to maximize impacts on agricultural production, but rather as a direct means of relaxing liquidity constraints during the hungry season to test specific predictions around seasonal liquidity and agricultural labor markets. Further work is needed investigate the returns to capital at different points during the agricultural calendar, and to test more cost effective ways of lowering the costs of consumption smoothing in rural settings with poor transportation infrastructure.

References

- Aggarwal, Shilpa, Eilin Francis, and Jonathan Robinson**, “Grain Today, Gain Tomorrow: Evidence from a Storage Experiment with Savings Clubs in Kenya,” 2017.
- Bandiera, Oriana, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman**, “Labor markets and poverty in village economies,” *The Quarterly Journal of Economics*, 2017, 132 (2), 811–870.
- Banerjee, Abhijit V, Sylvain Chassang, and Erik Snowberg**, “Decision Theoretic Approaches to Experiment Design and External Validity,” in “Handbook of Economic Field Experiments,” Vol. 1, Elsevier, 2017, pp. 141–174.
- Basu, Karna and Maisy Wong**, “Evaluating seasonal food storage and credit programs in east Indonesia,” *Journal of Development Economics*, 2015, 115, 200–216.
- Bauer, Michal and Julie Chytilová**, “The impact of education on subjective discount rate in Ugandan villages,” *Economic development and cultural change*, 2010, 58 (4), 643–669.
- Beaman, Lori, Dean Karlan, Bram Thuysbaert, and Christopher Udry**, “Self-Selection into Credit Markets: Evidence from Agriculture in Mali,” *Mimeo*, 2014.
- Behrman, Jere R, Andrew D Foster, and Mark R Rosenzweig**, “The dynamics of agricultural production and the calorie-income relationship: Evidence from Pakistan,” *Journal of Econometrics*, 1997, 77 (1), 187–207.
- Breza, Emily and Cynthia Kinnan**, “Measuring the equilibrium impacts of credit: Evidence from the Indian microfinance crisis,” *Working Paper*, 2016.
- Bruhn, Miriam and David McKenzie**, “In Pursuit of Balance: Randomization in Practice in Development Field Experiments,” *American Economic Journal: Applied Economics*, 2009, pp. 200–232.
- Bryan, G., S. Chowdhury, and A.M. Mobarak**, “Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh,” *Econometrica*, 2014, 82 (5), 1671–1748.
- Bryceson, Deborah Fahy**, “Ganyu casual labour, famine and HIV/AIDS in rural Malawi: causality and casualty,” *Journal of Modern African Studies*, 2006, 44 (2), 173.
- Burke, Marshall, Lauren Falcao Bergquist, and Edward Miguel**, “Sell Low and Buy High: Arbitrage and Local Price Effects in Kenyan Markets,” *Working Paper*, 2018.
- Chaudhuri, Shubham and Christina Paxson**, “Smoothing consumption under income seasonality: Buffer stocks vs. credit markets,” *Mimeo*, 2002.
- Cole, Steven M and Parakh N Hoon**, “Piecework (Ganyu) as an Indicator of Household Vulnerability in Rural Zambia,” *Ecology of food and nutrition*, 2013, 52 (5), 407–426.
- Conning, Jonathan and Christopher Udry**, “Rural financial markets in developing countries,” *Handbook of agricultural economics*, 2007, 3, 2857–2908.
- Dercon, Stefan and Pramila Krishnan**, “Vulnerability, seasonality and poverty in Ethiopia,”

- The Journal of Development Studies*, 2000, 36 (6), 25–53.
- Devereux, Stephen, Rachel Sabates-Wheeler, and Richard Longhurst**, *Seasonality, rural livelihoods and development*, Routledge, 2013.
- Edmonds, Eric V**, “Child labor and schooling responses to anticipated income in South Africa,” *Journal of development Economics*, 2006, 81 (2), 386–414.
- Fink, Günther, B. Kelsey Jack, and Felix Masiye**, “Seasonal Credit Constraints and Agricultural Labor Supply: Evidence from Zambia,” *NBER Working Paper*, 2014, No. 20218.
- Gilbert, Christopher L, Luc Christiaensen, and Jonathan Kaminski**, “Food price seasonality in Africa: measurement and extent,” *Food Policy*, 2017, 67, 119–132.
- Gross, Jeremie, Catherine Guirkingier, and Jean-Philippe Platteau**, “Buy as You Need: Nutrition and Food Storage Imperfections,” *Working Paper*, 2017.
- Hausdorff, Katharine**, “Household Bargaining Power and the Effect of Microloans,” 2016. MSc Thesis, Department of Economics, Tufts University.
- Imbens, G. and J.M. Wooldridge**, “Recent developments in the econometrics of program evaluation,” *Journal of Economic Literature*, 2009, 47, 5–86.
- Ito, Takahiro and Takashi Kurosaki**, “Weather risk, wages in kind, and the off-farm labor supply of agricultural households in a developing country,” *American Journal of Agricultural Economics*, 2009, 91 (3), 697–710.
- Jacoby, Hanan G and Emmanuel Skoufias**, “Testing theories of consumption behavior using information on aggregate shocks: Income seasonality and rainfall in rural India,” *American Journal of Agricultural Economics*, 1998, 80 (1), 1–14.
- Jayachandran, Seema**, “Selling labor low: Wage responses to productivity shocks in developing countries,” *Journal of Political Economy*, 2006, 114 (3), 538–575.
- Kaminski, Jonathan, Luc Christiaensen, and Christopher L Gilbert**, “The End of Seasonality? New Insights from Sub-Saharan Africa,” *World Bank Policy Research Working Paper*, 2014, (6907).
- Karlan, Dean, Robert Darko Osei, Isaac Osei-Akoto, and Christopher Udry**, “Agricultural decisions after relaxing credit and risk constraints,” *Quarterly Journal of Economics*, 2014, 129 (2).
- Kerr, Rachel Bezner**, “Informal Labor and Social Relations in Northern Malawi: The Theoretical Challenges and Implications of Ganyu Labor for Food Security,” *Rural sociology*, 2005, 70 (2), 167–187.
- Khandker, Shahidur R**, “Seasonality of income and poverty in Bangladesh,” *Journal of Development Economics*, 2012, 97 (2), 244–256.
- Kochar, Anjini**, “Explaining household vulnerability to idiosyncratic income shocks,” *The American Economic Review*, 1995, pp. 159–164.
- , “Smoothing consumption by smoothing income: Hours-of-work responses to idiosyncratic agri-

- cultural shocks in rural India,” *Review of Economics and Statistics*, 1999, 81 (1), 50–61.
- Marlow, David and Douglas P Crowne**, “Social desirability and response to perceived situational demands,” *Journal of consulting psychology*, 1961, 25 (2), 109.
- Michaelowa, Katharina, Ralitza Dimova, and Anke Weber**, “Ganyu Labour in Malawi: Understanding Rural Household’s Labour Supply Strategies,” *Mimeo*, 2010.
- Morduch, Jonathan**, “Income smoothing and consumption smoothing,” *Journal of Economic Perspectives*, 1995, 9, 103–103.
- Orr, Alastair, Blessings Mwale, and Donata Saiti-Chitsonga**, “Exploring seasonal poverty traps: the ‘six-week window’ in southern Malawi,” *The Journal of Development Studies*, 2009, 45 (2), 227–255.
- Paxson, Christina H**, “Consumption and income seasonality in Thailand,” *Journal of Political Economy*, 1993, 101 (1), 39.
- Pitt, Mark M and Mark R Rosenzweig**, “Agricultural prices, food consumption and the health and productivity of farmers,” *Economic Development Center*, 1986.
- **and Shahidur R Khandker**, “Credit programmes for the poor and seasonality in rural Bangladesh,” *Journal of Development Studies*, 2002, 39 (2), 1–24.
- Rose, Elaina**, “Ex ante and ex post labor supply response to risk in a low-income area,” *Journal of Development Economics*, 2001, 64 (2), 371–388.
- Rosenzweig, Mark R**, “Neoclassical theory and the optimizing peasant: An econometric analysis of market family labor supply in a developing country,” *The Quarterly Journal of Economics*, 1980, 94 (1), 31–55.
- **and Kenneth I Wolpin**, “Credit market constraints, consumption smoothing, and the accumulation of durable production assets in low-income countries: Investments in bullocks in India,” *Journal of Political Economy*, 1993, pp. 223–244.
- Schofield, Heather**, “The Economic Costs of Low Caloric Intake: Evidence from India,” *Mimeo*, 2013.
- Skoufias, Emmanuel**, “Intertemporal substitution in labor supply: Micro evidence from rural India,” *Journal of Development Economics*, 1996, 51 (2), 217–237.
- Strauss, John**, “Does better nutrition raise farm productivity?,” *The Journal of Political Economy*, 1986, pp. 297–320.
- World Bank**, “World Development Indicators,” Technical Report 2017.
- Yang, Tzu-Ting**, “Family Labor Supply and the Timing of Cash Transfers: Evidence from the Earned Income Tax Credit,” *Journal of Human Resources*, 2018, 53 (2), 445–473.
- Zambian Central Statistics Office (CSO)**, “2010 Census of Population,” *Government of the Republic of Zambia*, 2010.

Figures and Tables

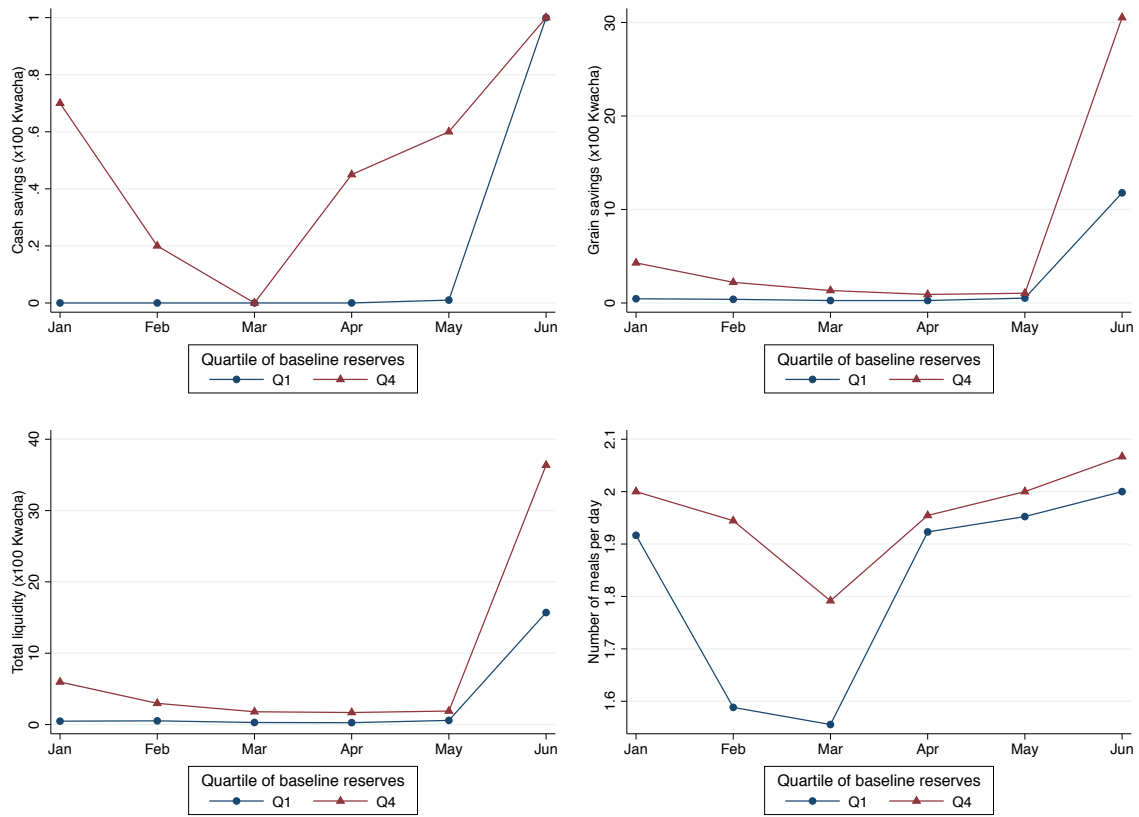


Figure 1: Seasonality in liquidity and consumption, by baseline cash and grain reserves

Notes: Monthly cash and grain holdings (top panel), total liquidity (bottom left) and consumption measured in meals per day (bottom right), for the top and bottom quartile of the distribution of baseline cash and grain reserves in the control group in year 1. Mean consumption in each quartile is shown for the consumption graph; the other three plot quartile median consumption.

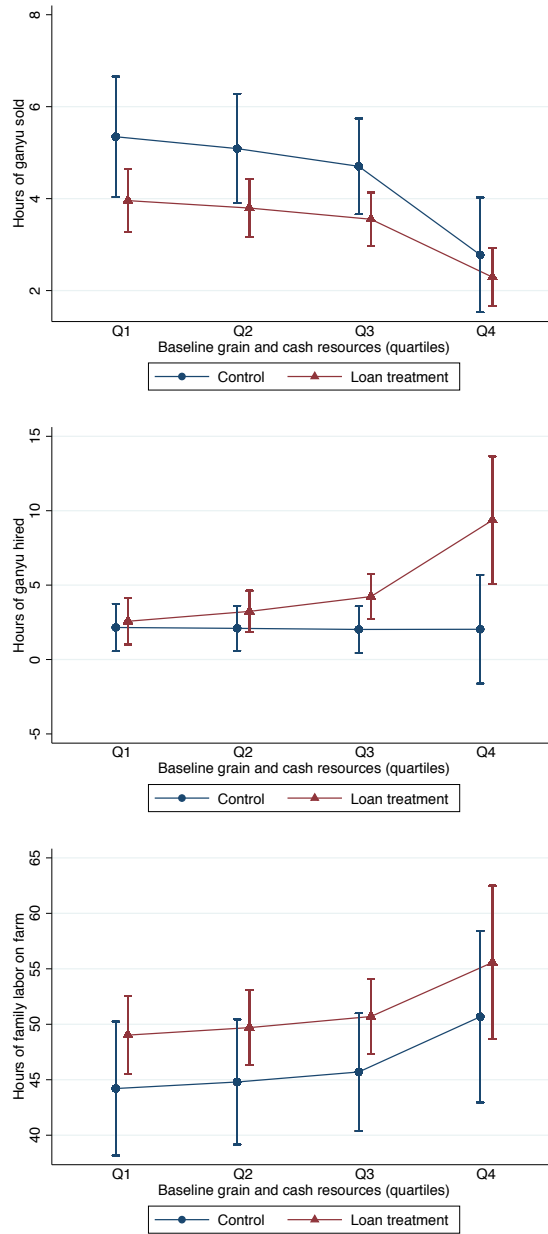


Figure 2: Effect on labor market participation, by baseline reserves

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. 90 percent confidence intervals are plotted based on standard errors clustered at the village level. The red line indicates the treatment group (shifted right). The same figures for year 2 of the project are shown in the appendix.




Figure 3: 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. 90 percent confidence intervals are plotted based on standard errors clustered at the village level. The red line indicates the treatment group (shifted right). The same figure for year 2 of the project is shown in the appendix.

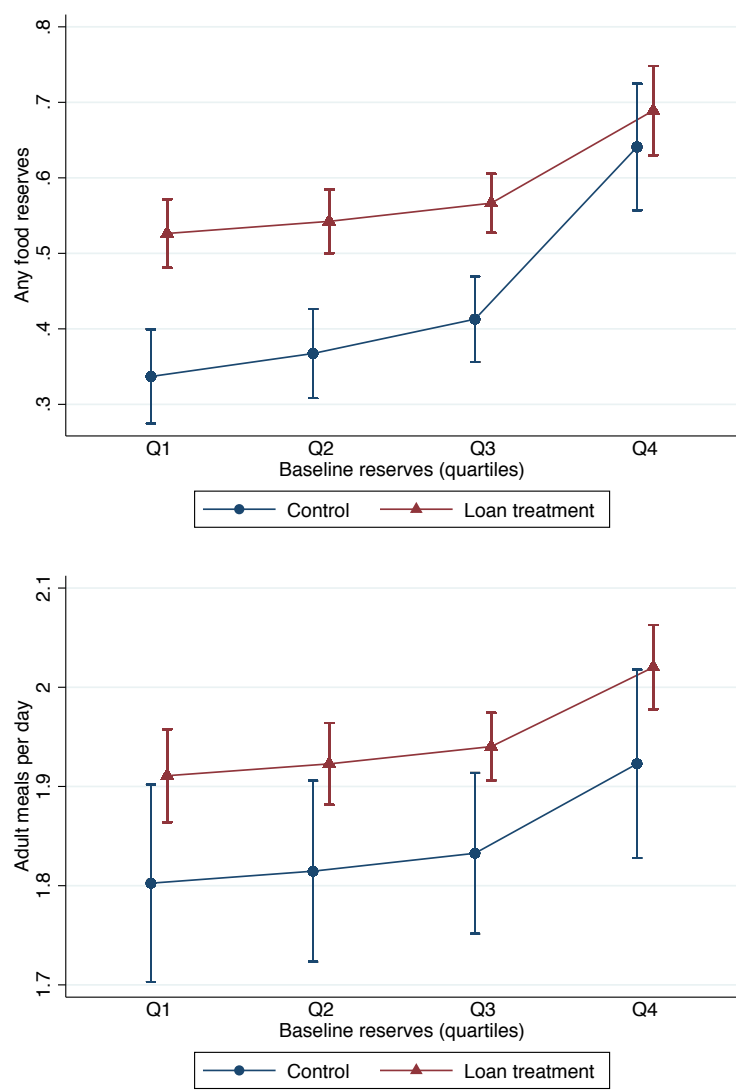


Figure 4: Effect on consumption variables, by baseline reserves, year 1

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. 90 percent confidence intervals are plotted based on standard errors clustered at the village level. The red line indicates the treatment group (shifted right). The same figures for year 2 of the project are shown in the appendix.

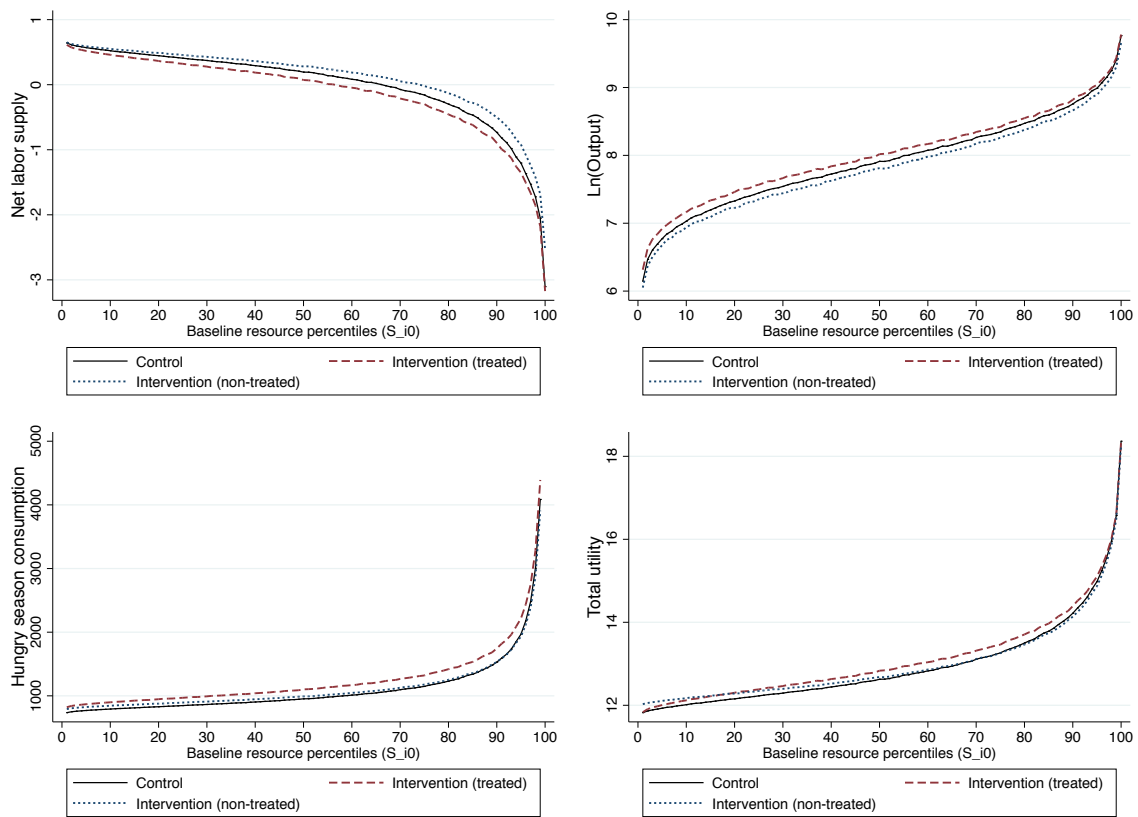


Figure 5: Impacts with partial treatment

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. Simulations calibrate the model to baseline grain and cash reserves, with productivity parameters calibrated to match baseline wage and output levels.

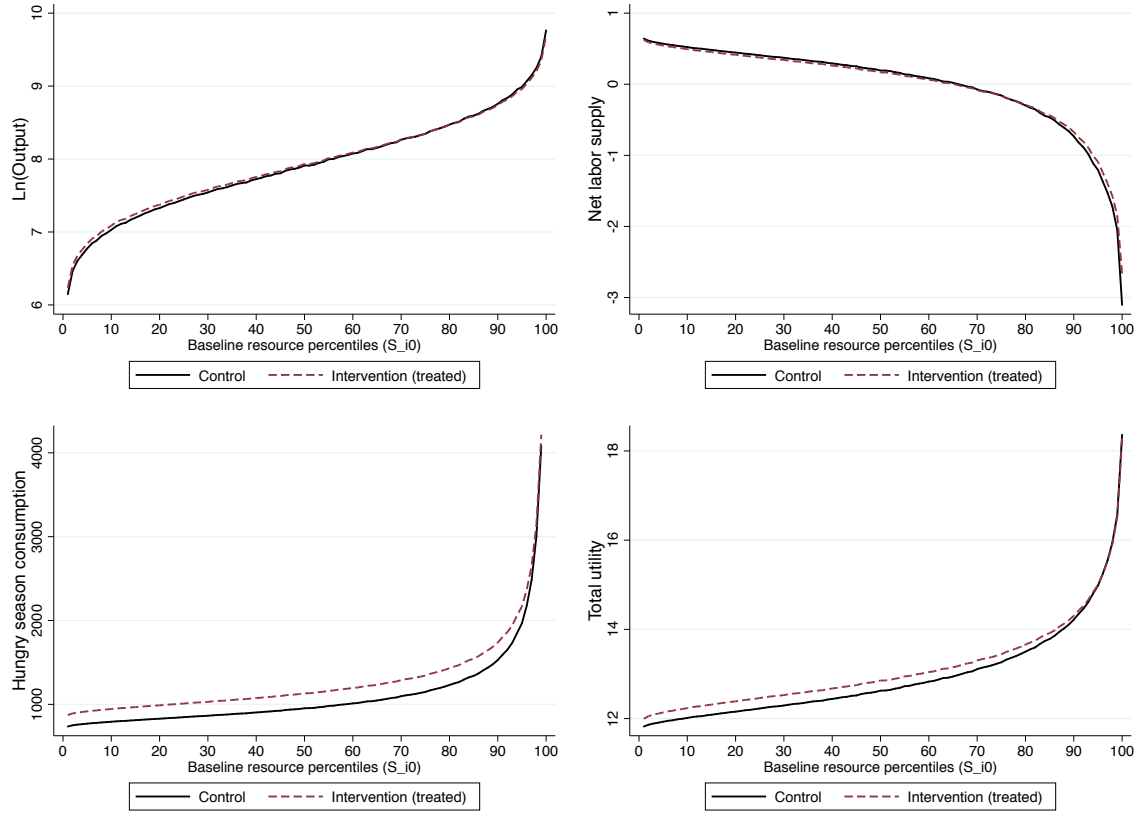


Figure 6: Impacts with full treatment

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 community. Simulations calibrate the model to baseline grain and cash reserves, with productivity parameters calibrated to match baseline wage and output levels.

Table 1: Take up and repayment

	Take up	Full repayment	% repayment	Repaid any cash
A. Year 1				
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)
B. Year 2				
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)
Any default in village in year 1	-0.025 (0.026)	-0.220*** (0.068)	-0.225*** (0.066)	-0.221** (0.099)
Early notification sub-treatment	0.010 (0.016)	-0.019 (0.058)	-0.007 (0.052)	0.012 (0.090)
Cash repayment sub-treatment	-0.008 (0.015)	-0.049 (0.058)	-0.044 (0.052)	0.624*** (0.049)

Notes: Table shows take up and repayment statistics in the cash loan treatment in year 1 (panel A) and year 2 (Panel B). Full repayment is a dummy variable that equals one 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 2, each row corresponds to a separate regression. Standard errors are clustered at the village level.

Table 2: Average treatment effects: Labor

	Any ganyu sold (1)	Hours sold (2)	Any ganyu hired (3)	Hours hired (4)	Family hours on- farm (5)	Total hours on- farm (6)
A. Pooled treatments						
Any loan treatment	-0.030 (0.023)	-0.891** (0.441)	0.026* (0.014)	1.516 (0.994)	6.134** (2.357)	7.075*** (2.542)
B. By treatment						
Cash	-0.020 (0.026)	-0.733 (0.492)	0.029* (0.017)	2.150* (1.180)	6.831** (2.851)	8.149*** (3.029)
Maize	-0.042* (0.025)	-1.061** (0.505)	0.022 (0.017)	0.783 (1.266)	5.382** (2.660)	5.916** (2.880)
Control group mean	0.310	3.855	0.105	3.321	49.753	51.806
Observations	2573	2572	2576	1386	2575	2576

Notes: Treatment effects on labor outcomes, measured during the short recall surveys during the hungry season. 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 99th percentile. All specifications are conditional on month fixed effects and include baseline controls, and cluster standard errors at the village level.

Table 3: Average treatment effects: Daily earnings

	Individual-level daily earnings (top 1% winsorized) (1)	Individual-level daily earnings (top 5% winsorized) (2)	Village median daily earnings (3)
A. Pooled treatments			
Any loan treatment	2.584* (1.483)	2.172* (1.172)	2.324* (1.322)
B. By treatment			
Cash	1.612 (1.858)	1.209 (1.428)	1.842 (1.583)
Maize	3.777** (1.699)	3.354** (1.391)	2.859* (1.628)
Control group mean	15.378	14.846	13.990
Observations	766	766	284

Notes: Treatment effects on reported daily earnings in Kwacha, restricted to the hungry season (Jan-Mar). Column 1 shows the impact on reported earnings using household level data. Column 2 shows the village level mean daily wage. Column 1 winsorizes the top 1 percent, Column 2 winsorizes the top 5 percent. Column 3 uses the individual level winsorized measures in column 2 before averaging at the village-month level. All columns include geographic controls, month-year fixed effects, and cluster standard errors at the village level. Columns 1 and 2 also include household controls and a control for the hours worked per day.

Table 4: Average treatment effects: Agricultural production

	Log harvest value	Log value - constant prices
	(1)	(2)
A. Pooled treatments		
Any loan treatment	0.081** (0.037)	0.085** (0.035)
B. By treatment		
Cash	0.105** (0.041)	0.117*** (0.040)
Maize	0.056 (0.044)	0.053 (0.040)
Control group mean	3462.496	3484.433
Observations	3969	3969

Notes: Treatment effects on agricultural output, measured in log ZMK. All specifications include baseline controls, year effects (panel A), and cluster standard errors at the village level.

Table 5: Average treatment effects: Consumption and food security

	Any grain stocks (1)	Food security (z-score) (2)	Meals per day Hungry season (3)	Meals per day Harvest (4)
A. Pooled treatments				
Any loan treatment	0.114*** (0.028)	0.245*** (0.073)	0.086*** (0.030)	0.025 (0.019)
B. By treatment				
Cash	0.005 (0.030)	0.088 (0.078)	0.072** (0.031)	0.032 (0.024)
Maize	0.232*** (0.034)	0.399*** (0.083)	0.103*** (0.035)	0.017 (0.021)
Control group mean	0.531	0.000	1.874	1.985
Observations	2576	1649	1367	1029

Notes: Treatment effects on consumption outcomes, measured during the short recall surveys during the hungry season (except column 5, which is measured at harvest). Outcome variables are: an indicator for whether the household had any remaining grain reserves (col 1), an index of food security (col 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 (col 3) and at harvest (col 4). All specifications are conditional on month fixed effects and include baseline controls, and cluster standard errors at the village level.

Table 6: Notification timing

	Hours sold (1)	Hours hired (2)	Hours on farm (family) (3)	Log output (4)	Adult meals (5)	Acres cash crops (6)	Input value (7)
A. Pooled treatments							
Any loan treatment	-0.578 (0.659)	-2.470 (1.843)	7.020 (6.926)	0.055 (0.119)	0.092** (0.035)	0.064 (0.165)	127.564 (205.716)
Loan x Early	-0.425 (0.656)	4.609* (2.363)	4.128 (8.435)	0.052 (0.134)	-0.018 (0.042)	0.211 (0.169)	120.079 (216.061)
B. By treatment							
Cash	0.130 (0.916)	-2.858 (1.937)	9.614 (8.575)	0.031 (0.135)	0.038 (0.040)	-0.088 (0.201)	147.412 (340.321)
Maize	-1.329* (0.750)	-1.922 (2.665)	4.621 (8.393)	0.079 (0.184)	0.155*** (0.030)	0.223 (0.241)	83.357 (203.837)
Cash x Early	-0.967 (1.012)	3.684 (2.969)	-0.602 (12.860)	0.105 (0.209)	-0.011 (0.055)	0.497* (0.259)	262.319 (389.260)
Maize x Early	0.131 (0.930)	5.530 (4.603)	8.859 (10.974)	0.002 (0.199)	-0.026 (0.049)	-0.068 (0.240)	23.831 (218.724)
Control group mean	3.855	3.321	49.753	7.700	1.874	1.003	756.987
Observations	606	608	608	958	590	969	961

Notes: Impacts of loan notification timing on main outcomes (columns 1-5) and planting season production decisions (columns 6 and 7), in year 2 for newly treated villages only. The reference group (i.e. not-early loan) is the timing used to announce the loans in year 1. Early notification involved informing farmers of the loan program at planting time in September. The timing of the actual loan delivery was the same across the timing sub-treatments (January). The first set of results pool the loan treatment arms; the second set show results by treatment arm. Note that all coefficients represent treatment effects for the relevant subgroup. All specifications include baseline controls, month or year effects, and cluster standard errors at the village level.

Table 7: Persistence and dynamics

	Hours sold (1)	Hours hired (2)	Hours on farm (family) (3)	Log output (4)	Adult meals (5)	Acres cash crops (6)	Input value (7)
A. Year 1							
Any loan treatment	-1.118** (0.548)	2.563** (1.086)	4.963* (2.616)	0.086** (0.044)	0.058** (0.024)	0.043 (0.054)	72.297 (53.128)
Control group mean	4.531	2.428	46.558	3534.927	1.911	0.918	842.568
Observations	1966	778	1967	3011	1633	3028	3028
B. Year 2							
Any loan treatment	-0.799 (0.489)	0.231 (1.549)	10.458* (5.846)	0.064 (0.092)	0.048** (0.023)	0.186 (0.129)	190.321 (150.369)
Treated in year 1	0.708 (0.520)	0.340 (1.188)	7.872* (4.189)	0.052 (0.070)	0.026 (0.020)	0.096 (0.088)	-43.514 (93.076)
Loan x Treated in year 1	-0.605 (0.646)	-1.201 (1.816)	-14.122** (6.949)	-0.110 (0.107)	-0.017 (0.029)	-0.108 (0.150)	-188.994 (163.545)
Control group mean	2.747	3.885	54.967	3349.127	1.923	0.971	1382.330
Observations	1762	1764	1764	2977	2943	3005	2987

Notes: Impacts of treatment by year on main outcomes (columns 1-5) and planting season production decisions (columns 6 and 7). The main results are restricted to the year 1 and households treated for the first time in year 2 (corresponding to the first line in Panel B). The remainder of Panel B shows the persistent effects of being treated in year 1 on outcomes in year 2. The cash and maize loan treatment arms are pooled throughout. All specifications include baseline controls, month or year effects, and cluster standard errors at the village level.