Seasonal liquidity, rural labor markets and agricultural production: Evidence from Zambia*

Günther Fink†
B. Kelsey Jack‡
Felix Masiye§

PRELIMINARY DRAFT

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 frictions in capital market access not only undermine households' ability to smooth consumption over the cropping cycle, but also distort labor markets by forcing capital-constrained farmers to sell family labor off-farm to meet short-run cash needs. To identify the impact of credit availability 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 borrowing at the time of the year when farmers are most constrained (the lean season) lowers aggregate labor supply, drives up wages and leads to a reallocation of labor from less to more capital-constrained farms. This reallocation reduces differences in the marginal product of labor across farms, increases average agricultural output, and reduces consumption and income inequality.

*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.

†Swiss TPH and Harvard T.H. Chan School of Public Health
‡Tufts University
§University of Zambia

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 tend to be low and infrequent, 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 developing countries.[1] In this paper, we show that capital market frictions also distort local labor markets if interest rates vary within communities, and farming households can sell family labor locally to finance consumption during the hungry season.[2] Since households have discretion over how much labor to use on their own farms, each household will choose labor inputs such that the marginal product of labor on the household's farm equals the household's opportunity cost of capital. Differences in interest rates faced by farmers in the same labor market therefore give rise to differences in the marginal product of labor across farms and result in an inefficient allocation of labor across farms. We formalize this intuition in a simple two-period agricultural household model. We show that, in a setting where farmers optimally choose labor inputs during the hungry season, high interest rates on borrowing result in increased labor sales by the poorest farmers in the community. This credit market induced labor supply drives down local wages, and leads to a suboptimal reallocation of labor from the poorest to the best-off farms.

We test three main predictions from this model in a two-year randomized controlled trial with 3,139 small-scale farmers across 175 rural villages in Zambia's Eastern Province. A sub-set of farmers in selected villages were offered subsidized cash or food loans during the hungry season, with repayment due after harvest. Our findings are largely consistent with the predictions of our stylized model. First, we find that treated households reduce off-farm sales of family labor, which, together with an increase in on-farm labor inputs, drives up local labor market wages. Second, the decrease in off-farm labor and increase in labor inputs is larger for the least well-off farmers, where we also find the highest marginal product of labor at baseline. On average, we find that agricultural output increases as predicted by the model, due in all likelihood to both increases

[1] A 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. Whether predictable intra-annual fluctuations in grain prices translate into seasonal fluctuations in consumption will depend on the cost to households of accessing capital markets and other smoothing strategies.

[2] Note that heterogeneity across farmers in effective interest rates could 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).

in labor inputs overall and improvements in labor allocation across farms, as well as potentially higher labor productivity from better nutrition. Finally, we find substantial increases in hungry season consumption. Both the increases in agricultural output and hungry season consumption are greatest among the worst-off farmers in our sample, who produce and consume the least at baseline, resulting in declines in within-community income and consumption inequality as a result of lower borrowing rates.

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 minimum 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 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 be output-maximizing in principle, they will not be – as highlighted in our model – if capital market frictions distort households' labor allocation and result in a differential marginal product of labor across farms.[3]

The intervention we study was designed to isolate the impacts of lowering the cost of capital access during the hungry season on local labor market outcomes and agricultural production. In the first year of the study, households in two-thirds of communities were eligible for loans; in the second year, 50 percent of communities received the program, with rotation of treatment status between years (i.e., some communities received two years of the program, some one year and some zero years). Despite an implicit interest rate of 30 percent over a six-month period, 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 program in the second year signed up for the loan program 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 loan access affects household labor allocation decisions. The likelihood that a family sold any ganyu during the hungry season fell by 2.9 percentage points (10 percent) in response to treatment, with a 24 percent reduction in

[3] A 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.

hours sold, on average. The likelihood of hiring ganyu increased by around 2 percentage points (18 percent). These labor market adjustments drive up local wages. Average wages, which we measure as daily earnings, increase by about 1.2 Kwacha, or about 9 percent relative to average daily earnings in control villages. Since our intervention targeted half of farmers in each village for the loans, on average, these increases reflect partial adjustments in aggregate labor demand and supply; model simulations show larger wage effects if loan access were provided to all farmers in these communities.

Our model clarifies that production inefficiencies arise from differences in interest rates available to farmers within the same labor market, which leads to differences in the marginal product of labor across farms. Empirically, we show that the variation in effective interest rates is negatively correlated with households baseline liquid resources. We measure baseline resources as the household's cash and grain reserves, and analyze treatment effects across the resource distribution.[4] Reductions in off-farm labor sales come predominantly from the worst-off farmers, while the increase in hiring is driven by the better-off farmers.[5] Total hours of labor input on-farm during the lean season increases at all levels of baseline resources. Consistent with an initial distortion in the marginal product of labor across farms, the estimated relationship between total hungry season labor inputs and agricultural output is declining in baseline resources in the control group, and is flat or slightly increasing in the treatment group. We document an 8 percent increase in the average value of agricultural output among treated farmers, which is concentrated among farmers with relatively small baseline resources. We also find substantial increases in hungry season consumption and reductions in consumption seasonality in response both to the lower borrowing cost (substitution effect) and higher expected harvest income. These effects are largest among the least well-off households in our sample. Together, the heterogeneous effects on agricultural income and consumption are consistent with a reduction in within-community inequality as a result of lower hungry season borrowing rates.

Our main results are consistent with our model of season-specific consumption needs and infrequent harvest income driving labor allocation decisions during the hungry season, with impacts on agricultural production and within-community inequality. We examine alternative interpretations of our findings, and test other margins of adjustment. First, we provide additional support that the smoothing function of the loans pertains to anticipated shortages rather than unanticipated shocks. Our main effects are concentrated among farmers who anticipate running short of food at the beginning of the agricultural cycle. We also test whether the anticipation of cash or food shortages affects decisions at planting (such as plot size and crop mix, which we abstract from in the model), by informing a subset of farmers about the loans at the beginning of the agricultural

[4] Our baseline survey overlaps with the planting season and therefore our measure of the distribution of available resources may be compressed if households that anticipate hungry season shortages also invest less during planting. We conduct robustness checks using the prior year's harvest value as an alternative, more plausibly exogenous, measure of available resources. In our setting land size cannot be taken as exogenous.

[5] Our pre-analysis plan focuses on the average effects of lowering hungry season borrowing costs. The impact of the variance in interest rates within the community was added to the analysis after data collection was complete.

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 leads to larger treatment effects on the value of agricultural output. This difference appears to be driven both by increased capital inputs (fertilizer, seeds, etc.) and the allocation of additional land to cash crops, as well as more pronounced shifts in labor allocation compared to farmers notified of the program at the start of the hungry season. Second, improvements in consumption appear to translate into better physical and mental health in our sample, as measured by fewer illnesses in the family and an index of mental health, respectively. This suggests that some of the improvements in agricultural output may have come not just from an increase in the quantity of labor applied to the family farm, but also an improvement in labor quality. Third, we find little effect of loan access on other consumption smoothing strategies, including livestock or asset sales, temporary migration or borrowing from friends and family, though we see some decrease in the likelihood of high interest borrowing from moneylenders. Fourth, we investigate whether the behavioral changes observed could be the result of the income transfer implicit in the subsidized loans. To test for income effects, 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. While we do find some small increases in consumption in response to the cash gift, the estimated changes are much smaller than the changes observed with the loan treatment and not statistically significant. 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 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. On the extensive margin, Bandiera et al. (2017) show that high borrowing costs prevent poor women in Bangladesh from accessing the labor market opportunities of their richer neighbors, which effectively keeps them poor. 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.[6] The critical role of family labor sales for smoothing consumption has been well documented (Kochar 1995, 1999; Rose 2001; Ito and Kurosaki 2009). 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.[7] Second, we show that these

[6]Like Rosenzweig (1980) and others, we assume that labor markets are well functioning and that land owning households (all of our sample) both buy and sell labor on local markets. Our model reflects the sequential nature of agricultural production, which may be subject to period-specific constraints (Behrman et al. 1997; Skoufias 1996).

[7]This consumption smoothing role of local labor markets is also tied to the substantial literature on informal

credit market induced labor sales result in an inefficient allocation of labor across farms, lowering aggregate output and increasing within-community inequality in income and consumption.

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, with impacts on production associated not only with high average borrowing rates but also by heterogeneity in rates within local labor markets. In addition to affecting the quantity of labor applied to the family farm, we contribute to the literature showing a link between nutrition and labor productivity (Pitt and Rosenzweig 1986; Strauss 1986; Behrman et al. 1997; Schofield 2013).

The results presented here also relate to a recent series of papers describing interventions targeting income, price and consumption seasonality in agricultural markets.[8] Burke et al. (2014) offered farmers in Kenya a loan product 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. 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 yields. Our findings contribute to this growing literature by providing the first direct evidence that capital market interventions timed to coincide with the hungry season not only affect consumption and output, but also affect agricultural labor market outcomes, aggregate output and inequality within communities. Fink et al. (2014) describes a pilot for the current study, with results consistent with those presented here.

From a policy perspective, our findings suggest that the potential welfare gains from reducing capital market frictions are large, particularly among the poorest farmers. We use our model to

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)).

[8]The 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.

simulate the impacts on the roughly 50 percent of farmers in the villages we study who do not receive access to the loans and find mixed results: their consumption increases as a result of higher equilibrium wages, but their own agricultural output falls as a result of lower labor inputs on their own farms. Some of the positive impacts on treated farmers can therefore be attributed to the labor supply pulled into the market from untreated farmers, which dampened the equilibrium wage adjustment, and lowered the opportunity cost of the labor reallocation that we document relative to a scenario with full treatment. Potential welfare gains may further be eroded by high transaction costs, which are at least partially reflected in the high interest rates that we observe in our study setting. Bundling seasonal loans with other technologies, such as mobile-based savings and borrowing platforms, or piggybacking on existing rural networks may offer scaling opportunities at substantially lower cost. Other strategies for decreasing the cost of consumption smoothing, such as more secure savings, may also decrease reliance on family labor for consumption smoothing, and improve agricultural production.

The paper proceeds as follows. In the next section, we present a simple model that highlights the linkages between capital markets, 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 shows robustness checks and further explores alternative mechanisms and explanations. We reconcile and calibrate magnitudes, and discuss 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) derive the relationship between a household's available liquid resources (cash and grain reserves), capital market frictions and local labor market outcomes, and 2) generate predictions regarding household and village-level responses to exogenous shifts in credit market interest rates.

2.1 Setup

Our theoretical framework builds on the agrarian labor market model introduced in Jayachandran (2006).[9] Each village economy has a finite number N of households that maximize utility over two

[9] We 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.

periods ($t = 1, 2$).[10] Each household i has initial liquid resources S_{i0}.[11] 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:

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

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

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

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^b; due to frictions in local capital markets we assume that $r^b > r^s$.[12] All farmers have access to the same saving technology; local borrowing rates are assumed to decrease with the farm's initial resources, S_{i0}, i.e. $\frac{\partial r^b}{\partial S_{i0}} < 0$.[13] All borrowing needs to be repaid by the end of the second period. The labor market clears at the endogenous wage w such that average farm labor input equals average labor supply:

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

[10]In Appendix A.1, we show an extension of the model with three periods, including a "period 0" during which other input investments occur.

[11]This initial distribution of liquid resources can be assumed to be the result of a stochastic process where all farms 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 local pests. We assume that initial resources are predetermined and positively correlated with farm productivity A_i (we also verify the correlation empirically). In the appendix, we show an alternative model where we model S_{i0} as an outcome variable of an initial optimization process.

[12]We 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.

[13]This assumption is consistent with any model where the expected ability to repay increases with collateral (which is proxied by S_{i0}) and decreases with loan size (which is also a function of S_{i0}).

2.2 Household utility maximization

Rational households maximize their utility from consumption and leisure over two periods:

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

subject to

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

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

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 on the farm and their own labor supply. In period 2, households receive harvest income. Period 2 net income (and consumption) is given by harvest production, y_i, plus period 1 borrowing or savings times the respective interest rate. Period 1 consumption is given by initial savings, S_{i0}, net labor income, given by the labor endowment, $\bar{h}$, minus time allocated to leisure, l_i, on-farm labor inputs, d_i, (some of which may be hired) times the wage rate, w, and net 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 in period two equals the wage that can be earned in local markets:

$$\frac{\partial y_i}{\partial d_i} = \beta A_i(\frac{k}{d_i})^{1-\beta} = wr_i^e \tag{5}$$

where

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

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:

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

Optimal labor inputs, and thus labor demand, decrease with the interest rate, lowering equilibrium wages.

2.2.2 Optimal consumption, labor supply and leisure

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

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

Combining optimality conditions (5) and (6), we can derive an interior solution for the optimal allocation of time toward labor, $h_i = \bar{h} - l_i$:[14]

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

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 levels of consumption in period 1 and period 2 are given by

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

and

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

respectively, where seasonality – the ratio of period 2 to period 1 consumption – increases with interest rates, i.e. $\frac{\partial c_2^*/c_1^*}{\partial {r_i}^e} > 0$.

2.2.3 The distribution of resources and local labor markets

The model offers several insights on the relationship between initial resources (S_{i0}), interest rates and labor allocation. First, as shown in equation (6), on farm labor inputs increase with productivity, but decrease with the effective interest rate faced by farmers. Since interest rates are a decreasing function of initial resources, S_{i0}, the marginal product of labor will on average always be higher on farms with low resources than on farms with high resources.[15]

Second, wages increase with average resources (S_{i0}) in the labor market. This can be directly seen from equations (6) and (8): for any given wage, lower interest rates will increase labor inputs (labor demand) d_i and reduce labor supply h_i due to positive income effects from lower borrowing

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

[15] If A_i were the same for all farms, low resources farms would thus employ less labor on their land than farms with high resources. The differences in labor inputs across farms will be even greater under a more realistic assumption that initial resources are positively correlated with productivity; a positive correlation between resources and productivity implies that higher resource farms will on average employ more labor both because of higher A_i and because of their lower opportunity cost of capital (lower r_i^e).

rates. Given that higher initial resources imply lower average interest rates, they also imply higher aggregate demand and reduced aggregate labor supply, so that equilibrium wages must always increase in order for local labor markets to clear.

Third, consumption seasonality– the ratio of period 2 to period 1 consumption – is a decreasing function of initial resources, S_{i0}. As shown in the inter-temporal optimality condition in equation (7), relative consumption levels are driven by the effective interest rate faced by each farmer: the higher the effective interest rate, the more households will cut back on consumption in the hungry season relative to consumption in the harvest season.[16]

2.2.4 The effect of lowering interest rates

Our experimental intervention subsidized access to credit to small-scale farmers. This corresponds to offering a specific effective borrowing rate $\hat{r}$ to farmers. Under the assumption that this rate is below the effective interest rate faced by farmers prior to the intervention,[17] we predict the following adjustments in labor allocation, wages, output, and consumption:

1. **Lowering effective borrowing rates increases aggregate labor demand and reduces aggregate labor supply, and increases wages in equilibrium.** Lower interest rates will increase demand for labor inputs, d_i^*, and decrease labor supply – holding wages fixed – due to positive income effects. In response to these aggregate changes in supply and demand, wages will increase until a new market-clearing equilibrium is reached.

2. **Access to lower borrowing rates reduces within-community differences in the marginal product of labor, increases output and lowers income inequality.** Access to uniformly lower borrowing rates also reduces differences in the effective interest rates faced by farmers. If all such differences were eliminated, the marginal product of labor would be the same on all farms due to optimality condition (6). To equalize the marginal product of labor, labor is reallocated from farms with high S_{i0} (who also have low r_i^e) to farms with low S_{i0}. This implies that high initial resource farms will decrease hiring, and low initial resource farms will decrease selling. As differentials in the marginal product of labor are reduced, the overall allocation of labor becomes more efficient, increasing aggregate output. Both higher wages and the reallocation of labor towards farms with low initial resources will contribute to reduced income inequality.

3. **Lowering effective borrowing rates increases hungry season consumption (c_1), reduces consumption seasonality (c_2/c_1) and reduces consumption inequality.** Given

[16]Higher food prices in the hungry season may further increase 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 6.1.5.

[17]We test this assumption empirically by analyzing loan uptake rates.

that both income and substitution effects are positive for first period consumption for farms benefitting from lower borrowing costs, lower interest rates will always increase first period consumption for low resource farmers. The same is not necessarily true for the second period because positive income effects for second period consumption are partially offset by a temporal substitution effect towards the first period. Given this, seasonality in consumption (c_2/c_1) will decrease among low initial resource farmers (who otherwise exhibit the most seasonal consumption patterns), resulting in lower consumption inequality.

2.2.5 The effect of lowering interest rates for some farmers

In our intervention, discussed in greater detail below, only a fraction of households in each community were eligible for lower interest loans. As a result, effective interest rates remained unchanged for roughly half of the population, while effective interest rates fell for the other half. Our empirical predictions therefore differ from the theoretical predictions outlined above in the following ways:

Prediction 1 is unchanged: increases in aggregate labor demand and reductions in aggregate labor supply will simply be scaled by the proportion of farmers treated, and thus result in smaller wage increases compared to programs targeting all farmers. Among treated farmers, Prediction 3 is also unchanged. Untreated farmers will also increase hungry season consumption as a result of higher hungry season wages.

Prediction 2, on the other hand, requires some more substantial revision. In a setting where all farmers have access to lower borrowing rates, reductions in net labor supply by low initial resource farmers is always matched by increases in net labor supply from high resource farms in order for labor markets to clear. This is not the case with partial treatment because increased labor demand by treated farms will at least partially be absorbed by untreated farms whose interest rates do not change, but who face higher wage rates.

Figure 1 illustrates simulated changes in labor supply with partial and full treatment. In both, we assume a log-normal initial resource distribution, calibrated to what we observe empirically at baseline.[18] In the top panel, we show labor adjustments in response to a reduction in interest rates: as in Prediction 2, net labor supply decreases at the bottom of the S_{i0} distribution (lower net sales on average), while net labor supply increases in the top two quintiles (reduced net hiring). In the new equilibrium, wages are 18 percent higher. With partial treatment –when only 50 percent of farmers receive lower interest rates (bottom panel) –the impact on wages is more moderate (+9.7 percent). In equilibrium, all but the top 5 percent of farmers reduce net labor supply. These labor adjustments reflect a reduction in net sales at the bottom of the S_{i0} distribution, and an increase in net hiring in the 60th-90th percentiles. For untreated farmers, net labor supply increases in response to the higher wage rate. Given our study design, these changes among non-treated farmers in treatment

[18] Specifically, we set initial resources to a mean of 400 Kwacha and a median of 50 Kwacha then calibrate A_i until we match the observed baseline daily wage of 15 Kwacha.

communities are not observed in our sample, but are important for the welfare implications of the program, as discussed in Section 7.

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, where we examine the descriptive implications of our model and how they match our data.

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). Approximately 100,000 people live in Chipata town, the district and provincial capital; the remaining population lives in rural areas, with small-scale farming as primary source of income. According to the 2010 Living Conditions Monitoring Survey (CSO 2010), 63 percent of households in rural Chipata were classified as very poor compared to 32 percent in Zambia overall. Average monthly expenditure of rural households was estimated at US$ 122 in 2010 (US$ 0.8 per person-day), corresponding to about one third of the national average (US$ 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 overwhelming majority of households in rural villages in Zambia. In our study villages, we document that over 95 percent of households meet this definition.

Study sample

The study sample was drawn from the population of small-scale farmers living in rural areas of Chipata District. The district is divided into 8 administrative blocks, each of which contains a number of camps. We randomly sampled 5 villages from 50 of the 53 camps in the district, omitting the camps that contain Chipata town. The village list was assembled from the Ministry of Agriculture's farm registry, which included 99,000 registered farms in the district in 2013. To facilitate sampling, villages with less than 20 or more than 100 farms listed in the registry were excluded from the initial village selection.

Study enumerators visited sampled villages to record the number of households, and screen for eligibility.[19] Enumerator screening visits stopped once 201 villages met all eligibility criteria. During

[19]Villages were ineligible if: (1) other projects had been conducted there in the recent past, (2) the village bordered

the baseline survey, 25 additional villages 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. This left us with a sample of 175 villages for the study.

Within each eligible village, households were sampled from the village rosters collected during the initial screening visits. Only small farms – less than 5 hectares according to the Zambian Ministry of Agriculture – were eligible for the program.[20] Eligible households were randomly sorted and the first 22 selected for the baseline survey. This resulted in 53 percent of households on average being selected for the project; across all villages, the share of households enrolled in the study ranged from 15 and 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 A.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 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 Figure 2.

The study included two main loan treatment arms: a *cash loan treatment* and a *maize loan treatment*, both offered at the start of the hungry season (January). Repayment was due at harvest (July), and loans could be repaid in either cash or maize (or both). The two treatment arms offer tradeoffs. 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

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.

[20] We 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.

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. In our main results, we focus on the treatment effects of being treated for the first time in either year, and estimate the effects of repeated treatment and persistence in separate analyses.

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, which was the median value assigned to participation in the loan groups in our choice experiments.[21] Cash grant villages were randomly selected within geographic blocks 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 (~ 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 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.[22] While both treatment arms were designed to reflect an interest rate of about 30 percent (over 5 months), 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 Table 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 choose.

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,

[21] For further details on choice experiments, see Appendix C.1.

[22] Requiring 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.

with x varied between 50 and 600 Kwacha.[23] 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 to 81.9 and 83.3 with repayment in June for maize and cash loans, respectively. Responses to these questions on value and timing determined final design decisions for the treatments. Further detail on the implementation of the choice experiments is provided in Appendix C.1.

3.3 Implementation

Both loan programs 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 households were invited.[24] At the meeting, project staff began by describing eligibility for the program to clarify why only some households were invited to the meeting. The terms of the loan were then described, followed by details on how the loan distribution would be organized. Loan enrollment and consent forms were provided to eligible households. If a household wished to join the program, they were required to present both forms with a signature from the household head when picking up the loan.

Loans were distributed between 3 days and one week after the village meeting at a location convenient for transportation, which was selected in cooperation with the village headman. Project staff registered attendees, confirmed their identity using the national registration card,[25] and collected their signed enrollment and consent forms. Before finalizing the transaction, project staff

[23]Hypothetical loan dates were consistent with program offered (pay out in January and repayment in June), but the hypothetical loans involved no interest.

[24]Ineligible 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.

[25]In 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.

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 14 household level 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 A.2 reports the number of households sampled in each survey round, and the probability of being in the survey round as a function of treatment. Panel A shows year 1 treatments and panel B year 2 treatments. The coefficients and standard errors are from OLS regressions for each survey round, with errors clustered at the village level. Overall, attrition rates are low: 3,030 out of the 3,139 households (96.5 percent) enrolled at baseline completed the endline survey. We do not find any differences in attrition overall or 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 A.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 the loan programs, 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 A.3 also previews our take-up results, which we turn to next.

3.3.3 Take up and repayment

As described above, loan conditions were established through focus groups and choice experiments to be both feasible from an implementation perspective and attractive for farmers from an interest rates perspective. Table 2 shows 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 values, 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 participating in (highly uncertain) future programs.

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 compare our setting to the descriptive implications of the model.

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 2013), harvest of year 1 (August 2014) and harvest of year 2 (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 between long recall survey rounds. We refer to these as short recall surveys since they primarily ask about activities in the past two days to two weeks. A total of 15,044 observations from the sample of 3,139 households were collected over the course of the study. Appendix A.3 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: (1) measures of labor allocation and daily earnings, (2) measures of agricultural output, and (3) consumption indicators. In many cases, we focus on data collected during the hungry season (January - March) of each year, since this is the period of interest in our conceptual model and the time most likely to be directly affected by the loan intervention. We describe the measurement of each of the main outcome types, then discuss the construction of our primary heterogeneity measure.

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 (when most ganyu is reported). 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 daily earnings based on days worked as well as total earnings. We winsorize the top 1 or 5 percent of household-level responses to address outlier observations and analyze daily ganyu earnings – as our proxy for local wage levels – 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 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 consistently across survey rounds, and rely, in our main analyses, on the short-recall measures from the labor surveys, with a recall period of two days. 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 initial 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. As we show in Appendix A.1, endogenous consumption, savings and planting investments should suppress both the level and the variation in S_{i0}, but should not affect household rankings or resource quartiles.

As an alternative and more plausibly predetermined measure of available resources, we use the total value of harvest at baseline. Unfortunately, we did not collect measures of outstanding debts at harvest during the baseline survey, so this measure may overstate net harvest income for households that are productive because of borrowing. We use this alternative measure to examine the robustness of our findings.

Appendix figure A.2 shows the correlation of each measure of initial resources with self-reported interest rates. For our preferred measure, the gradient is monotonically negative (higher reserves correspond to lower reported interest rates) for households with above median baseline reserves, and flat for households with below median reserves (all of whom have very little liquid savings). For our alternative measure, confidence intervals are wider, and we again observe a negative gradient for above median baseline harvest values (i.e. more harvest income is associated with lower interest rates), though we also see a modest positive slope below the median, which may be due to some variation in harvest values arising from borrowing behavior. The shape of these relationship helps determine our functional form assumptions in our analysis of heterogeneity by initial resources.

For our analysis, we construct quartiles of the distribution of these baseline measures of available resources. Our predictions about heterogeneity in the effects of the loan treatments depend both on how the loan affects the household's own interest rate, and also how it affects the variance in the interest rate within-village. The quartiles we construct are based on the distribution in the entire sample; within village distributions are suppressed because of our relatively small sample sizes within village. The quartile based on within village variation is very similar to the quartile coming off of the full baseline dataset (a correlation of 0.84 for the grain and cash reserves measure).

The correlation between our preferred measure of quartiles of baseline cash and grain reserves and quartiles of baseline harvest value is 0.5.

4.2 Descriptive results

Seasonality in resources and consumption As illustrated in Figure 3, households accumulate grain and cash reserves after harvest, which are mostly depleted by early January. 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 A.1. This shortage of resources during the hungry season is anticipated by farmers: at baseline, 76 percent of households stated that they did not think their maize reserves would last to 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 within-community variation in reserves and in households' ability to smooth consumption across periods. Figure 4 shows seasonal patterns in food availability and consumption for the top and bottom quartile of our sample in terms of baseline food and cash reserves (available resources). The left panel shows the average number of meals consumed per day by adult members of the household in the control group. The right panel shows the share of households in that quartile with any grain stock at home. Both measures are coarse proxies for consumption, but still show considerable seasonality in consumption, which – consistent with the model described in Section 2 – is more severe for 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 formal 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, which typically occurs 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 the maize needs of a typical family of five about 2-3 months. Sixty percent of households report storage losses in the past season.

Credit markets also appear highly imperfect: only 5 percent of household respondents 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, locally referred to as *kaloba*, with monthly interest rates of 40-50 percent. 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). To measure interest rates, we ask households how much they would have to repay in one month if they borrowed 50 Kwacha today from a source other than a friend or family member. The implied monthly interest rate is around 40 percent for households with below median baseline grain and cash reserves; implied interest rates decline slightly with baseline reserves for households with above median reserves, though rates are still high (around 34 percent) for even the "richest" decile of our study sample (see Appendix figure A.2). In the hungry season, sixty percent of households report that they would be unable to borrow 50 Kwacha in cash that day. These descriptive patterns justify the modeling choice that effective borrowing rates (r_i^e) are a function of initial resources (S_{i0}).

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: 39.7 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.[26] 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 include borrow from friends or family (28 percent), using savings (22 percent) and sell assets or livestock (17 percent), all of which may be difficult to rely on during the hungry season. On average, ganyu contracts specify relatively small tasks (such as weeding an acre of land), which can be completed by an individual within a working day; larger contracts also occur, and it is common for families or groups of individuals to complete specific tasks together. 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.

Almost two thirds of farmers in our sample reported having engaged in ganyu activities in the previous season at baseline, and about the same fraction of farmers reported that they would engage in ganyu labor in the coming season. Households appear reasonably accurate in their forecast of whether they will have to engage in ganyu in a given year. Among control group households that

[26] In 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".

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

Ganyu wage rates are typically negotiated on a case-by-case basis, and anecdotally are highly responsive both to demand and supply shocks. 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). Consequently, as described in the model, wages are determined by local market clearing, which either means within a village or within a small group of villages.

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 Though not a focus of our data collection or analysis, it is worth noting a few details of the land market that inform our decisions for how to model household decisions and analyze the data. First, land constraints are typically not binding. At baseline, over 60 percent of respondents said that they could have farmed more land than than did in the previous season. The average acreage cultivated was 4.5; the average acreage reported available for cultivation was 6.3. Second, land markets for sale or lease are largely absent and most tenure is customary with annual land allocations by the village headperson. Anecdotally, these allocations are determined by a complex set of factors, including past production and favoritism. Some land is reallocated every year. Partly as a result of these allocation procedures, respondents tend not to know their land area with very much precision. As a result, we leave land area out of most of our analyses, except as a baseline control (see below).

The distribution of resources and local labor markets One of the central implications of our theoretical model is that households' relative levels of initial resources (S_{i0} in our model) shape households' net labor supply in local markets. To examine this empirically, we rank households with respect to their baseline cash and food reserves, and then compare household labor activities and transactions during the hungry season for control group households in year 1. The left bar of Figure 5 shows the number of hours household members worked for other farms; the middle bar shows hiring of external labor; the middle bar shows labor effort on their own farm land.[27] Hours sold is substantially higher in the bottom quartile than in the top quartile, while households in the top quartile supply slightly more labor to their own farms overall, but less per acre of land. Hiring is most common among households with top-quartile initial resources. These patterns imply that, consistent with our model, low resource households sell labor to households with relative high levels of baseline resources.

Our model also implies that the marginal product of labor will vary with baseline resources. To identify the empirical relationship between labor inputs and outputs, we use a production function approach and regress log agricultural output value on a quadratic function of the log number of hours of labor inputs on-farm during the hungry season by households in the control group in year one. The marginal effects at the mean are plotted for the bottom and top initial resource quartiles in Figure 6. While the estimates are noisy – they come from a sample of 223 bottom quartile and 258 top quartile households in the control group in year 1, and include geographic controls but no other contemporaneous inputs – they are consistent with a higher marginal product of labor in the bottom quartile households than in the top quartile households.

Next, we consider the relationship between baseline resources and equilibrium earnings to examine whether villages where labor supply driven by high hungry season borrowing costs have more depressed wages. We calculate the average baseline resources at the village level, using the same baseline measure of household cash and grain reserves. Regressing the median daily earnings on the median resources (in 100 Kwacha) in the 69 control village-months in year 1 that have non-missing earnings data, controlling for month, survey round, and geographic controls indicates that a 100 Kwacha increase in the average level of resources in the village is associated with a 1.2 (s.e. 0.62) Kwacha increase in equilibrium earnings. Figure 7 shows both the raw data and best fit line (excluding the geographic controls included in the regression).[28] Of course, this relationship may reflect omitted village-level characteristics associated both with equilibrium wages and average baseline resources. We control for two of these – distance to a paved road and village size – and the coefficient remains largely unchanged and marginally significant (1.1, s.e. 0.62).

[27]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, while they work around 6 hours per week, conditional on working. The average hungry season ganyu work day in the control group lasts for just over four hours, not including search time.

[28]If the outlier observations shown in the figure (median daily earnings above 40 Kwacha and median reserves above 1500 Kwacha) are dropped from the regression, the coefficient falls to 0.67 and remains marginally significant.

5 Experimental results

In this section, we present our main experimental results. We start by outlining our empirical strategy and discussing some important differences between our model and our empirical analysis, before presenting key results on labor allocation, agricultural output and consumption.

5.1 Empirical strategy and identification

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. We order our analysis as follows. For our main outcomes, we begin with a pooled specification that averages across years and loan treatment arms:

$$y_{ivt} = \alpha + \beta loan_{vt} + \tau_t + X_{i0}\delta + u_{ivt} \tag{11}$$

where y_{ivt} is an outcome of interest for household or individual i located in village v at time t, where t is a year for our long recall surveys and a month-year for our short recall surveys. $loan_{vt}$ 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. Treatment assignment varies over time according to the treatment rotation between years, as described in Section 3.2. 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. We define each year's treated period as lasting from loan distribution (January) until harvest in June. We focus both data collection and analysis on the hungry season and harvest period, with an emphasis on the former. 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 separate the treatment effect by treatment arm, using estimating equation (11) and allow β to differ for the cash and maize loans.

We also estimate how treatment effects vary over the distribution of baseline grain and cash reserves, following the descriptive results presented above.

$$y_{ivt} = \alpha + \beta_1 loan_{vt} + \beta_2 S_i + \beta_3 S_i^2 + \beta_4 loan_{vt} \times S_i + \beta_4 loan_{vt} \times S_i^2 + X_{i0}\delta + \tau_t + u_{ivt} \tag{12}$$

allows the treatment effects to vary with the baseline distribution of initial resources, S (which corresponds to S_{i0} in our model). To illustrate heterogeneity in the treatment effects, we plot the estimated quartile-specific coefficients, evaluated at the mean S in each quartile of the baseline distribution: $\beta_2 + \beta_3$ for the control group and $\beta_1 + \beta_2 + \beta_3 + \beta_4$ for the treatment group. Standard errors are calculated using the delta method.

We provide a number of supplementary analyses to further investigate potential mechanisms and also to address alternative interpretations of our main effect. In a first step, we examine the degree to which hungry season shortages and treatment effects are anticipated. We first examine heterogeneous treatment effects based on anticipation of resource shortages at baseline, restricting attention to year 1, for which baseline measures are most relevant. We then exploit the variation in the timing of loan announcement (holding fixed the timing of loan delivery) in the second treatment year. We also investigate persistence by estimating cross-sectional OLS models for each year of the project. For year 2, we interact an indicator for a loan treatment in year 2 with treatment status in year 1. 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_{ivt}. Appendix table A.1 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 A.1 are our household-level controls, used throughout the analysis.

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 2.2.5, based on access to lower interest loans for a sub-set of farmers in the labor market.

Prediction 1: Lowering effective borrowing rates decreases labor demand, increases labor supply and increases wages We examine impacts both on the extensive and intensive margins of labor in Table 3. 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 the majority of treated farmers. As illustrated in Figure 1, these shifts in labor supply will happen both through 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 30 percent among sampled farmers. 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) is significant at $p < 0.10$. Turning to the continuous measure of total family hours sold, treated households sell an average of 0.9 fewer hours or 24 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 also has positive, but statistically insignificant, effects on hiring, on both the extensive margin (column 3) and continuous measures (column 4). The likelihood of any hiring (column 3) increased by 2 percentage points, off of a mean of 11 percent, and the number of hours hired increases by around 1 hour or 27 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 4.5 hours or around 8.5 percent (column 5), with similar impacts across treatment arms. This results in a similar measure of total hours invested on the family farm (column 6). The estimated magnitudes – a 4.5 hour increase on-farm relative to a 0.9 hour decline in labor sales – suggest that households actually work more on average (increase their own labor supply), in contrast with the model's predicted decrease in labor supply due to positive income effects. Given that ganyu work is strenuous and accompanied by substantial search costs (and social stigma), it may be that households no longer obliged to do ganyu are both better rested and more motivated for field work; the consumption impacts we present below highlight the further possibility that higher caloric intake led to more energy for on-farm work. At the same time, the loan program itself may motivate farmers to work harder, or the (insignificant) increase in hiring results in additional supervision needs on the farm.

Table 4 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. As expected, we find modest increases in daily earnings, which are sensitive to the right hand tail of the distribution. 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 2.4 or 1.2 Kwacha per day on average, corresponding to an 9 to 16 percent increase over the control group mean, on average (Panel A). Increases are larger in the maize arm for both measures (Panel B). Column 3 shows village-level regressions using median reported daily earnings at the village-month level. The magnitudes line up reasonably well with the household level data reported in column 2. 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 A.3. The estimates are noisy, but show the expected positive gradient.

Prediction 2: Lowering effective borrowing rates reduces differences in the marginal product of labor across farms and increases output on average Figure 8 plots treatment effects at the mean of each quartile of the baseline distribution of grain and cash reserves on our continuous measures (total hours over the past week) of labor market adjustments for both the control and treatment groups in year 1.[29] 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 created. While the increases among high baseline resource farms are inconsistent with predictions associated with full treatment, they are consistent with our partial treatment setting where wage adjustments are dampened by the labor supply of untreated households, as shown in Section 2.2.5.

Higher labor inputs among households at the bottom of the distribution of initial resources, who were seen to have the highest estimated relationship between hungry season labor inputs and agricultural output in the control group (see Figure 6), may help equalize the marginal product of labor across farms. While we cannot directly observe the marginal product of labor, we use the same approach as in Section 4.2 to estimate the effect of additional hungry season labor (in log of hours) on log agricultural output value. Figure 9 shows that this proxy for the marginal product of labor changes slope as a result of treatment. Consistent with the model's predictions, the estimated relationship is decreasing in baseline reserves in the control group and flat (or in this case slightly increasing) in the treatment group. These regressions do not hold constant other inputs, and the (unobserved) quality of labor inputs may be increasing at the same time as the quantity. Thus, this analysis should be interpreted as suggestive support for the prediction that a reallocation of labor helps equalize the marginal product of labor across farms.

Increases in labor inputs on-farm may increase agricultural output, on average, and particularly among low baseline resource households if the marginal product of labor is, indeed, larger for low resource farms. We show average effects on log of agricultural output in Table 5. Panel A shows that agricultural production (as measured by total harvest value using nominal (column 1) and constant (column 2) prices increased by around 8 percent on average, with slightly larger point estimates in the cash loan treatment arm. Figure 10 shows that the average treatment effect on agricultural output in year 1 is largest in the bottom quartile, and declines with initial resources.

[29]The underlying regression omits 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 figures A.4, A.5 and A.6. Robustness of our main heterogeneity measure of available hungry season resources uses baseline harvest value as an alternative to baseline cash and grain reserves. Appendix figures A.7, A.8 and A.9 show the results. See Section 4.1 for further discussion of the heterogeneity measures.

Prediction 3: Lowering effective borrowing rates increases consumption during the hungry season, reduces consumption seasonality, and reduces consumption inequality Table 7 shows the treatment effect on consumption and food security measures, with a primary focus on the hungry season, when constrained households have to compress their consumption most as shown in Figure 4. To assess the impact of lower effective borrowing rates, we rely on the consumption outcomes used in Figure 4: 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-third 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 0.08 meals or 0.19 standard deviations, on average, which is comparable to the difference between average meals consumed in the control group between the hungry season (column 3) and the harvest season (column 4). In contrast, at harvest, we see no effect of loan access on consumption outcomes (column 4). This means that the treatment resulted in a reallocation of consumption towards the hungry season, and generating the reduction in consumption seasonality predicted by our theoretical model. Though it is outside of the scope of the model's predictions, Appendix table A.5 shows modest reductions in illness within the home and in mental health, with no changes in treatment seeking probabilities. Thus, the changes in family labor inputs may be driven in part by fewer days lost to illness, and the quality of labor may also have improved.

A final implication of the model is a reduction in consumption inequality between low and high baseline resource households. We compare our main measures of consumption (food availability and adult meals per day) across baseline endowment quartiles for treatment and control villages during year 1 in Figure 11. 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 consumption for low resource households, and a reduction in consumption inequality as illustrated by the substantially flatter resource-consumption relationship in Figure 11.

6 Interpretation 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.

6.1 Interpretation

6.1.1 Anticipated versus unanticipated shortages

Much of the existing literature linking labor markets to capital availability in rural agricultural settings has focused on smoothing unanticipated shortages arising from aggregate or idiosyncratic shocks (e.g., Jayachandran (2006); Kochar (1999)). Our model and results suggest that capital market frictions shape labor market outcomes even in the absence of unanticipated shocks; farms with low reserves sell labor locally and reduce own farm labor inputs in response to high borrowing costs. Households may, of course, also use loans to smooth unanticipated shortages. To further investigate whether labor adjustments are "structural" and anticipated, we use baseline information on households' anticipation of food shortages in the coming season, and estimate treatment effects by subgroup. If loan availability primarily modified household responses to unanticipated shocks, we should see that treatment impact is independent of households' anticipated shortages. If, on the other hand, loan availability primarily matters for those who anticipate running short of resources, we should see largest treatment impacts among those farms most dependent on labor sales to finance consumption.

We show results in Table 6, pooling across years and treatment arms, with villages treated in year 1 omitted from the year 2 results (i.e., the same sample as we use in the main results, above). The correlations between the anticipation measure and the outcome suggest that many outcomes are predictable: households that anticipate food shortages at baseline consume fewer meals (column 1), sell more ganyu (column 2), and have lower agricultural output (column 5). Treatment effects also appear to depend on anticipated shortages for these outcomes: the effect on consumption is driven entirely by households who anticipated food shortages at baseline. Effects on ganyu sales are also driven by this sub-group, though the interaction term is imprecisely estimated. Finally, treatment effects on agricultural output are driven entirely by households that anticipated running out of food. Interestingly, the effect on family labor inputs or hiring does not appear to be concentrated in the sub-group that anticipated shortages, though the inclusion of household controls imposes a strong *ceteris paribus* assumption that may impact the results. Overall, it appears that anticipated shortages play an important role in shaping treatment effects, making it less likely that our main results are driven by unanticipated shocks. We turn next to timing of when farmers learned about the program in year 2 as a further check on whether adjustments are anticipated.

6.1.2 Anticipated versus unanticipated loans

Table 8 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 total value of capital inputs applied to the household's fields and the acres devoted to cash crops. 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 results, 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 in villages notified at planting about the loan program, and larger effects on agricultural output value (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.

6.1.3 Persistence of treatment effects

Table 9 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 9 shows that many of our main results are driven by outcomes in year 1. Specifically, we see statistically significant increases in meals consumed (column 1), decreases in hours sold (column 2), increases in hours hired (column 3), increases in family hours spent on-farm (column 4) and increases in output (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 patterns and magnitudes for villages treated for the first time for consumption, ganyu sales, family labor investments and output. Unlike in year 1, hours hired (column 3) is statistically insignificant with a negative coefficient. Inputs and cash crops increase 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 results for villages treated in year 1, which are excluded from most of the main analyses. While the estimates are noisy, we find relatively little evidence of persistence. On average, villages treated in year 1 but not in year 2 look similar to the year 2 control group, aside from large but imprecise increases in hours of ganyu sold and in output value. 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 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 overlooked by 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.

6.1.4 Other margins of adjustment

In addition to increasing hungry season consumption and on-farm labor, it also seemed likely that at least some farmers would use the loans as substitutes for other more expensive consumption smoothing strategies. We collect data on several of these and show treatment effects in Appendix table A.6. We begin by assessing the degree to which the program substitutes for other types of borrowing. We analyze program impact on three outcomes: formal loans (column 1), low interest informal loans from friends or family (column 2) or informal high interest loans, typically from moneylenders (column 3). All types of borrowing were reported retrospectively at the end of each season as part of the long recall surveys, and coded into basic 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).

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. While we observe little detail about the conditions around these informal loans, these results suggest that at least some households reduced their exposure to high interest borrowing. We also conduct additional analysis for livestock or asset sales as alternative smoothing strategies (columns 4 and 5), and observe small increases in livestock sales of around 3.5 percentage points or around 10 percent of the control group mean. Both the effects on high interest borrowing (column 3) and livestock sales (column 5) are largest among households in the maize loan treatment arm.

6.1.5 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 A.7 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.

6.1.6 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. 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 A.8 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.[30]

6.2 Robustness checks

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. In Table A.9, 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.

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.

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

Results are shown in Table 10. 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 Calibration and welfare implications

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 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. On average, we estimate that the loan program resulted in approximately 5 additional hours of labor inputs per week on the family farm (Column 6, Table 3). This increase is driven both by more work by family members and by increased hiring of external labor for specific tasks. Taking the thirteen weeks of the hungry season from January to March as our time reference, and assuming an average five hour work day on the field (as reported in our time use questions), 5 hours per week means approximately 13 additional days of labor on the farm during the hungry season. Our main results show that the loans increased the value of agricultural output of by about 8 percent, which corresponds to approximately 260 Kwacha (Column 1, Table 5).[31] Combining these two estimates, we get an implicit average marginal product of labor of around 20 Kwacha, which is of

[31]We use nominal prices for this calculation, to allow for comparisons with the nominal daily earnings, and to reflect actual income for farmers, which may have included adjustments in cropping in response to price variation between years.

similar magnitude but about 30 percent higher than the treatment group average daily earnings reported in Table 4. Given that we see largest labor increases among farms with the lowest capital endowment, this can be interpreted as evidence for the marginal product of labor being particularly large in this subgroup as predicted by the model. However, other adjustments including changes in the crop mix in year 2, improvements in nutrition, and the addition of some other types of inputs may have further enhanced agricultural production in response to the loans. We lack the data to quantify the relative contributions of more labor inputs, improvements in labor quality and other adjustments made by treated households.

We can also calibrate the earnings implications for households in our study. On average, households in the control group sold 3.7 hours of ganyu per week (Column 2, Table 3) and received daily earnings of 13.2 Kwacha, which implies total hungry season earnings of around 127 Kwacha (assuming 5 hour work days and 13 weeks during the hungry season). Households in treated villages sold an average of 2.8 hours of ganyu for daily earnings of 15.4 Kwacha, resulting in 112 Kwacha of extra labor market income over the course of the hungry season. Thus, while treated households earned less on the labor market during the hungry season, some of their reduced hours were offset by higher daily earnings, and the estimated loss of labor market earnings (~15 Kwacha) is relatively small. However, treated households also spent around 61 Kwacha more on labor, which implies a reduction in net labor revenue of 76 Kwacha.

Our results suggests that a substantial proportion of the original 200 Kwacha loan was used for financing consumption otherwise covered by earnings from ganyu labor, and for hiring additional labor. We also find some evidence that households borrowed less from informal channels potentially reducing their outstanding debts. On the other hand, our limited data on assets, livestock and net savings suggest that the overall impact on measures of wealth is small. Given that treated farms had to repay 260 Kwacha at the end of the season, the net impact on household's financial position is ambiguous, and depends on numerous adjustments, many of which we do not measure. Our strongest revealed preference measure of overall welfare impacts is the take up rate of 98 percent among repeat borrowers in the second year (after repayment), which suggests that the overwhelming majority of treated households viewed these loans as welfare improving.

Our data collection focused on treated farmers. Untreated farmers in the village were also affected, primarily through the effect of the program on wages. Higher wages improve consumption outcomes for these farmers. In Figure 12, we show simulated impacts on first period consumption and agricultural income by calibrating the model to our baseline data (see Section 2.2.5 for further description). The model simulations show increases in hungry season consumption both for treated and for untreated farmers (top panel). Given that these improvements are particularly large for untreated farmers who benefit from higher wages, overall consumption inequality decreases substantially. At the same time, the effect on untreated farmers' agricultural output is negative (Figure 12, bottom panel), because untreated farmers lower on-farm labor investment in response to higher wages with unchanged interest rates. As a result, income inequality declines only marginally due to increased inequality between treated and untreated farms, which partially offsets the inequality reductions among treated farmers.

Studying the effects of partial access to subsidized hungry season loans is relevant for policy, given that many rural lending programs only enroll a subset of households in a community. However, we could imagine two alternative scenarios, one in which all households were given access to our loans and another in which unrestricted borrowing at frictionless market rates were available to all farmers. As described in Section 2.2.5, effects would likely differ substantially with full loan access for all households. Unrestricted access to loans will further eliminate differences in the marginal product of labor, lead to larger increases in wages and restrict differentials in agricultural output to generic differences in farm productivity.

8 Conclusions

The results presented in this paper highlight the importance of access to credit markets for agricultural labor markets in general, and for the within-community allocation of labor, income and consumption in particular. With limited initial resources, capital market frictions increase income and consumption inequality within communities, and lower agricultural output overall through an inefficient allocation of labor across farms. 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 are 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. However, since increased access to capital reduces differences in the marginal product of labor across farms, the net effect of increased capital access on output should always be positive.

In the specific study area analyzed in this paper – and presumably in many others – the room for improvement in local credit markets is large. 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.

The loan intervention we study was not designed to maximize impacts on agricultural production, but rather to relax borrowing costs during the hungry season and allow us to test specific predictions around seasonal incomes and credit market frictions. Further work is needed investigate the returns to capital at different points during the agricultural calendar, and to test more cost effective ways of extending credit 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.

Basu, Karna and Maisy Wong, "Evaluating seasonal food storage and credit programs in east Indonesia," *Journal of Development Economics*, 2015, *115*, 200–216.

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.

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

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

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

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.

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.

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 agricultural 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.

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

Figures and Tables

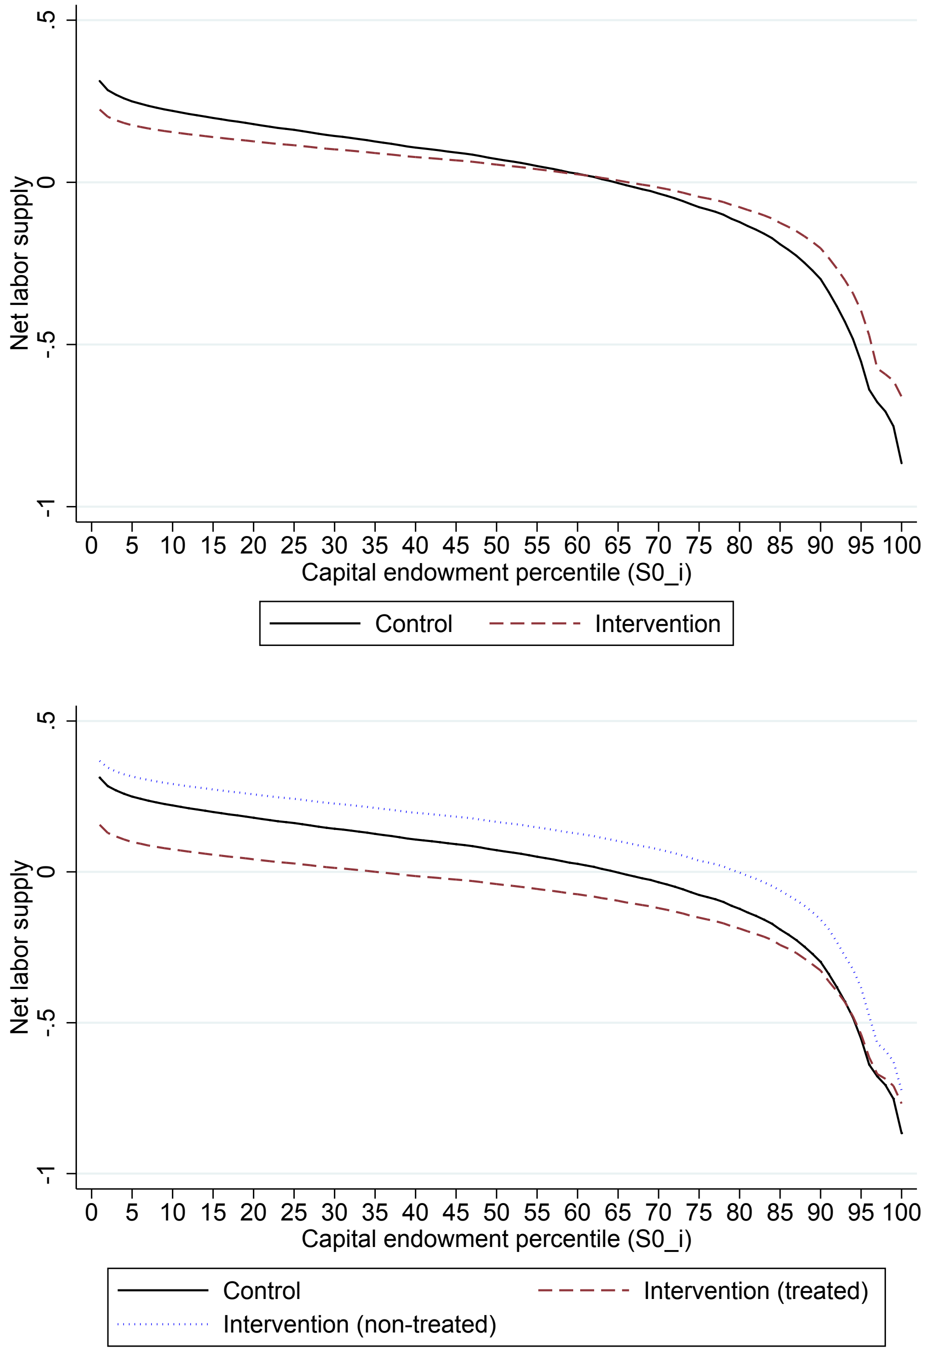

Figure 1: Labor adjustments with partial and full treatment

Notes: Model simulations based on either 100 percent access to lower interest rates (top) or 50 percent access to lower interest rates (bottom). Simulations calibrate the model to baseline grain and cash reserves, with productivity parameters calibrated to match local wage and output levels.

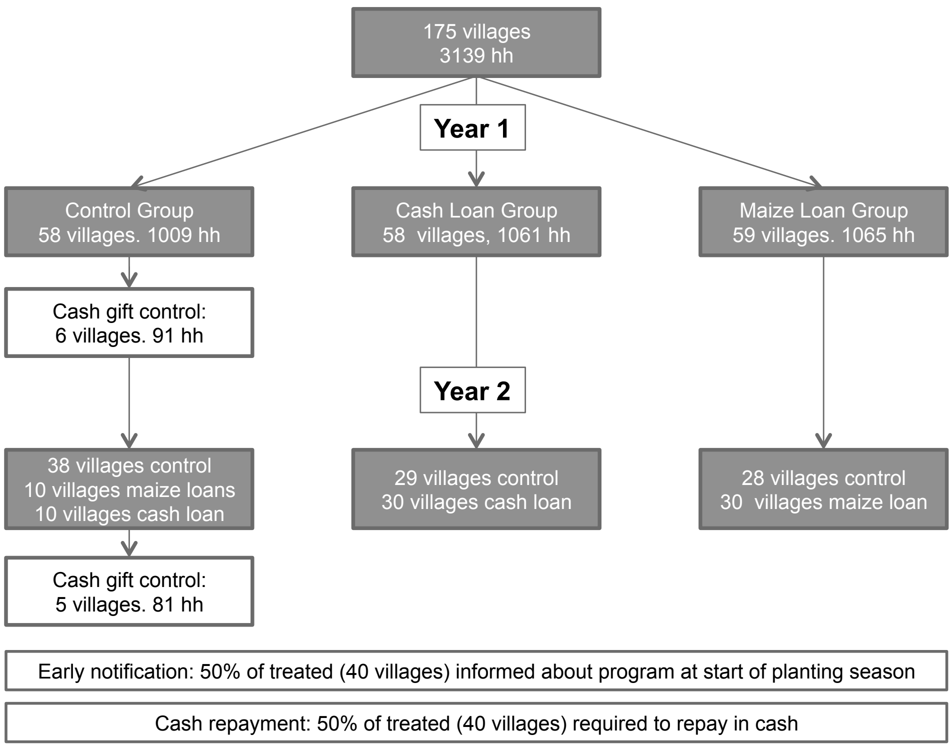

Figure 2: Study design

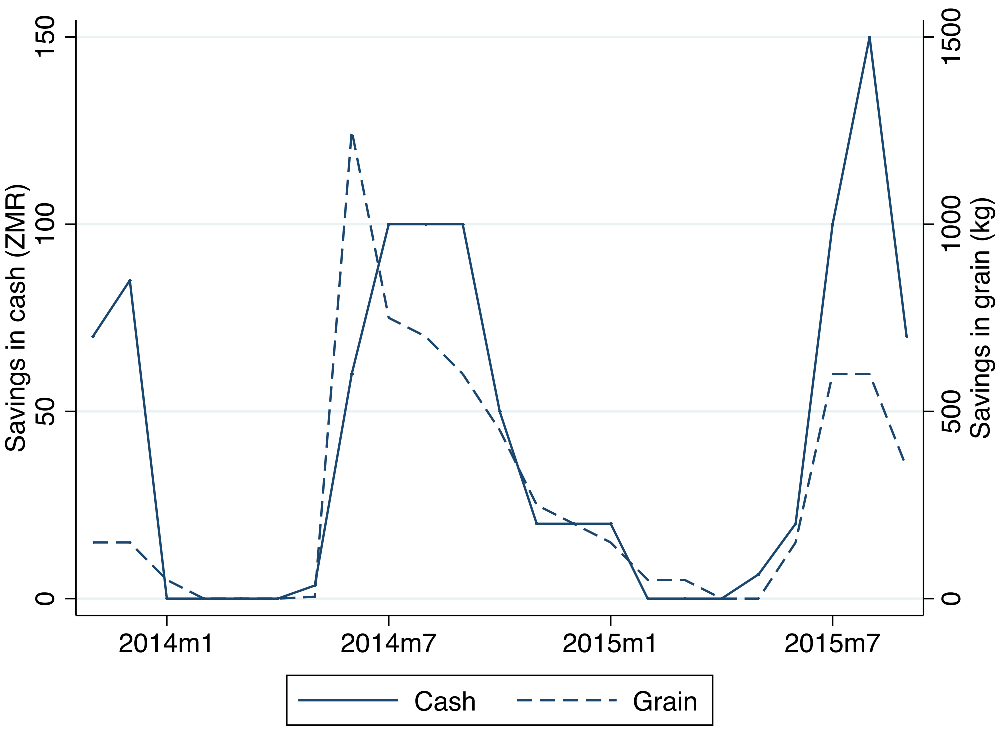

Figure 3: Cash and grain savings (median) in the control group

Notes: Cash savings and the value of stored maize (in Kwacha) held by the household at the time of survey, by month of the year in the control group. The months shown on the x-axis cover our full study, beginning with baseline data collection in November 2013 and ending with endline data collection in September 2015. For additional details on the survey rounds and data collection, see Appendix A.3.

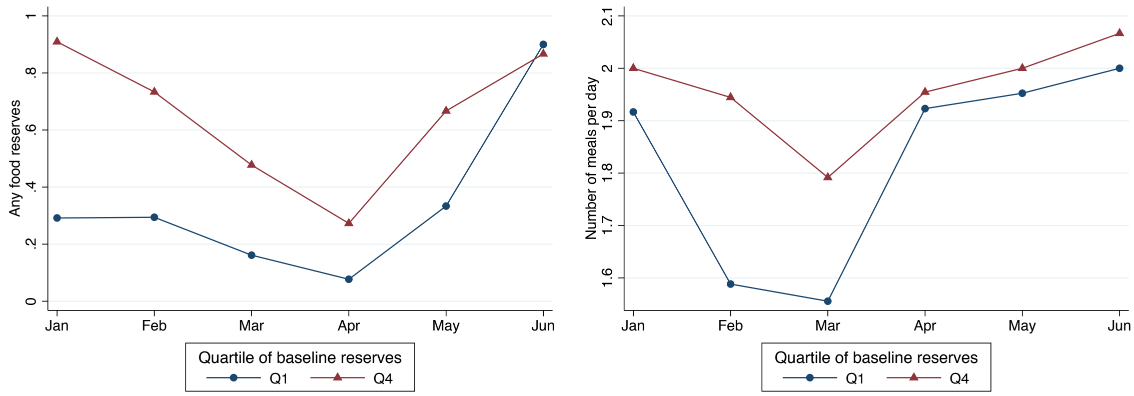

Figure 4: Seasonality in consumption, by baseline cash and grain reserves

Notes: Monthly consumption outcomes, for the top and bottom quartile of the distribution of baseline cash and grain reserves in the control group in year 1. The left graph shows an indicator for any food reserves available in the home and the right graph shows a measure of meals per day consumed by adults within the household.

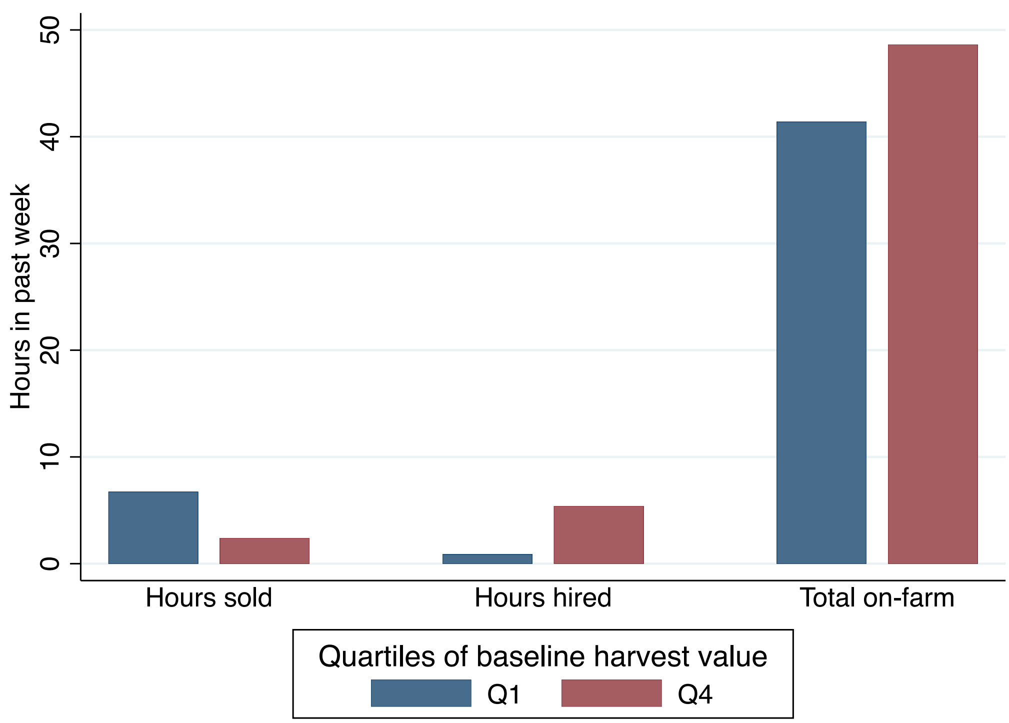

Figure 5: Labor market participation, by baseline cash and grain reserves

Notes: Hours of labor market activity over the past week, during the hungry season, for the control group in year 1. The bars represent mean winsorized (top 1 percent) hours in the top and bottom quartile of the distribution of baseline cash and grain reserves.

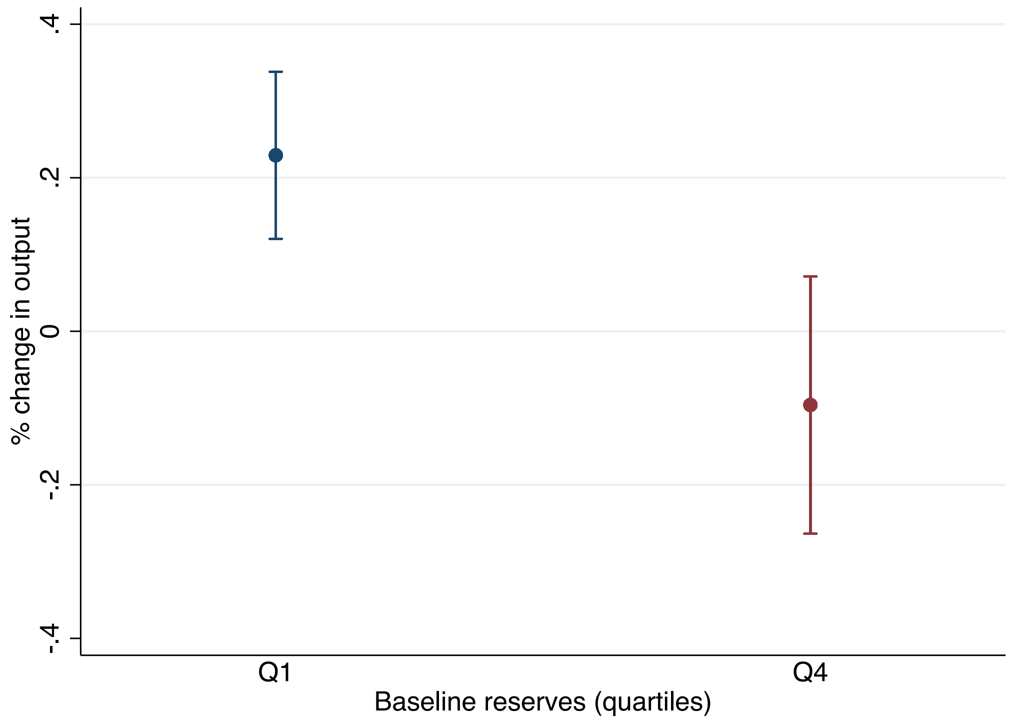

Figure 6: Estimated relationship between labor inputs and agricultural output, by baseline cash and grain reserves

Notes: Point estimates and 90 percent confidence intervals are from a regression of log agricultural output value on a quadratic in log total hours of labor inputs on-farm during the hungry season, for the control group in year 1. Results are the percent change in output value for a one percent increase in hours of labor input for the top and bottom quartile of the distribution of baseline cash and grain reserves.

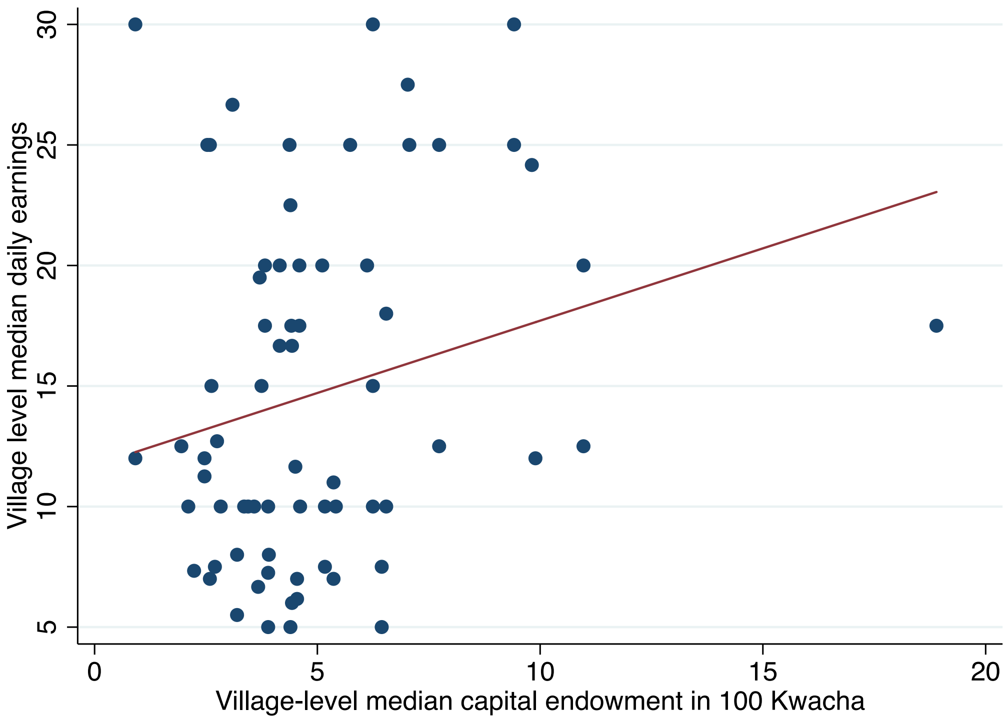

Figure 7: Village average daily earnings, by baseline cash and grain reserves

Notes: Village level median daily earnings in the hungry season, for the control group in year 1, versus the village median baseline value of grain and cash reserves.

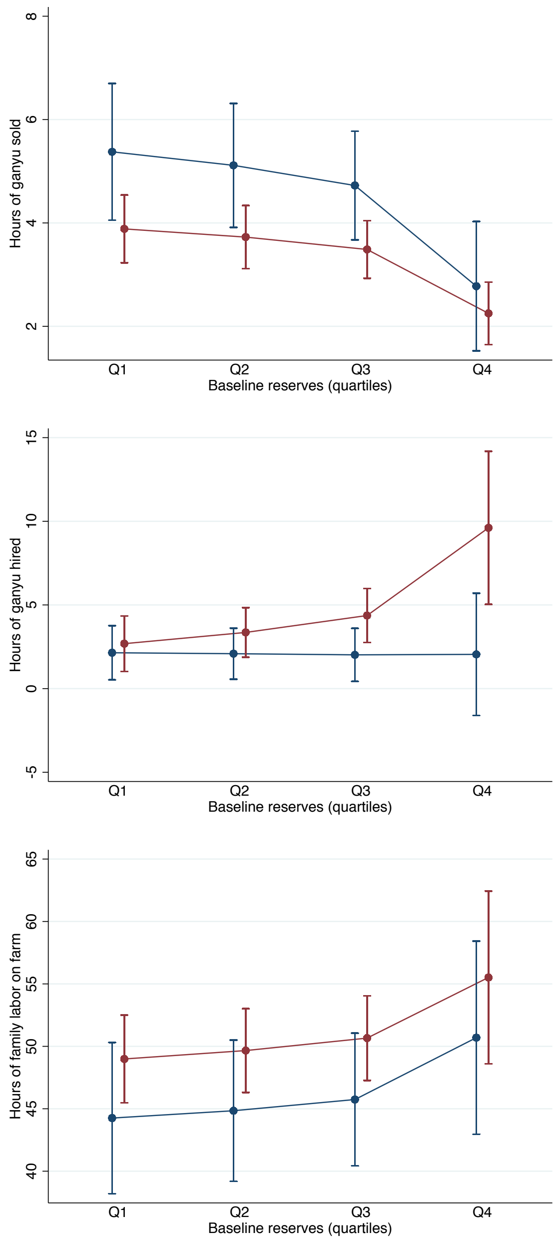

Figure 8: Effect on labor market participation, by baseline reserves

Notes: Heterogeneous impacts on labor market outcomes in the hungry season of year 1, estimated using a quadratic in baseline reserves. Plots show estimated means in each quartile of the distribution of baseline reserves, 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.

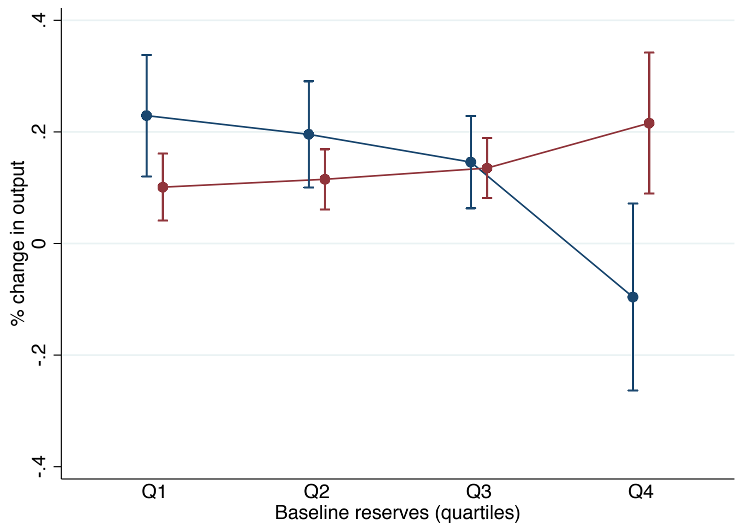

Figure 9: Estimated relationship between hungry season labor and agricultural output, by baseline reserves, year 1

Notes: Heterogeneous impacts on the relationship between labor inputs (measures as log hours of hungry season on-farm labor) and (log) agricultural output value in year 1, estimated using a quadratic in baseline reserves. Plots show estimated means, based on separate regressions for the treatment and control groups, 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.

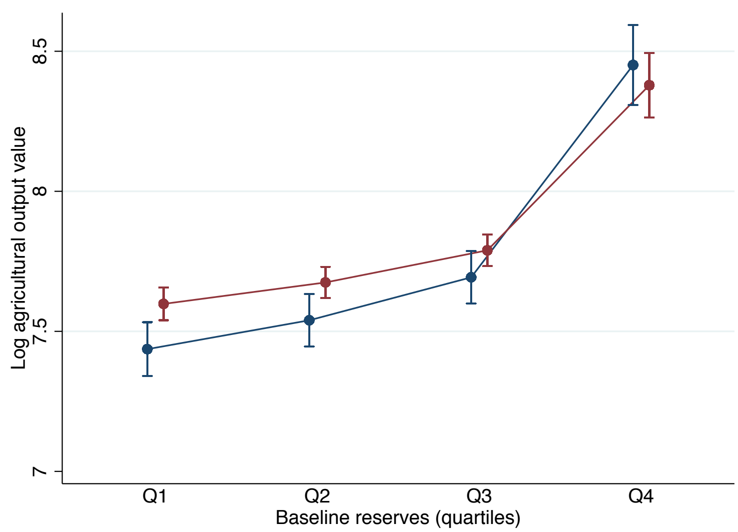

Figure 10: Effect on log agricultural output, by baseline reserves, year 1

Notes: Heterogeneous impacts on agricultural output value in year 1, estimated using a quadratic in baseline reserves. Plots show estimated means, 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.

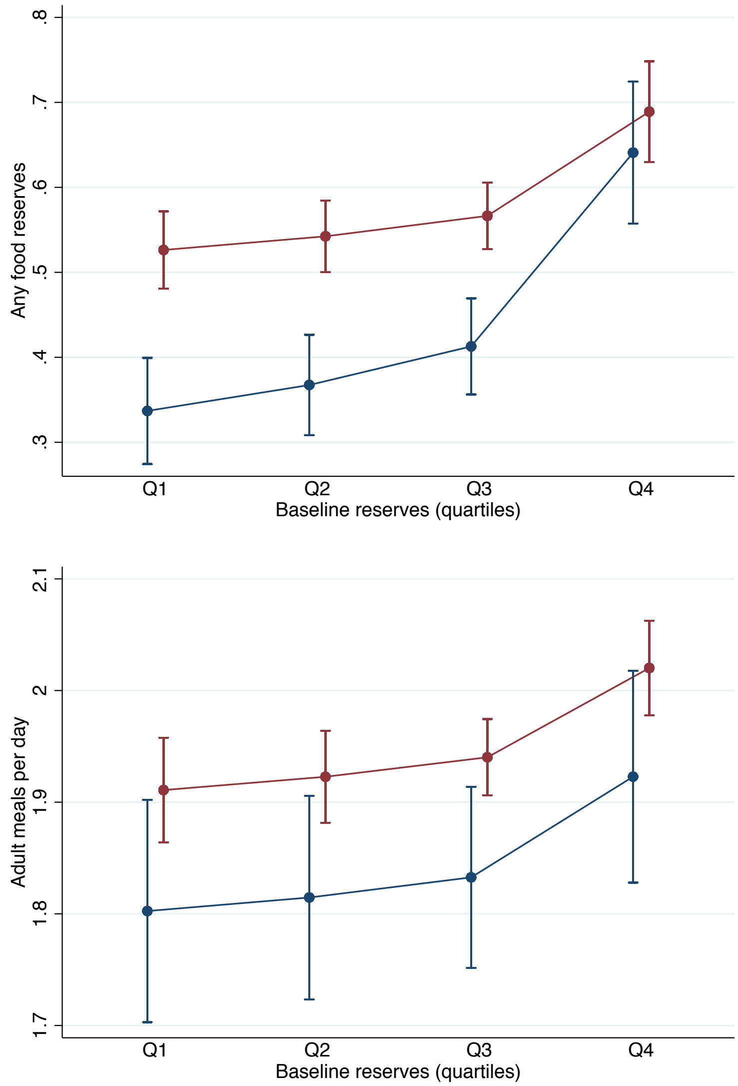

Figure 11: Effect on consumption variables, by baseline reserves, year 1

Notes: Heterogeneous impacts on consumption outcomes in the hungry season of year 1, estimated using a quadratic in baseline reserves. Plots show estimated means, 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.

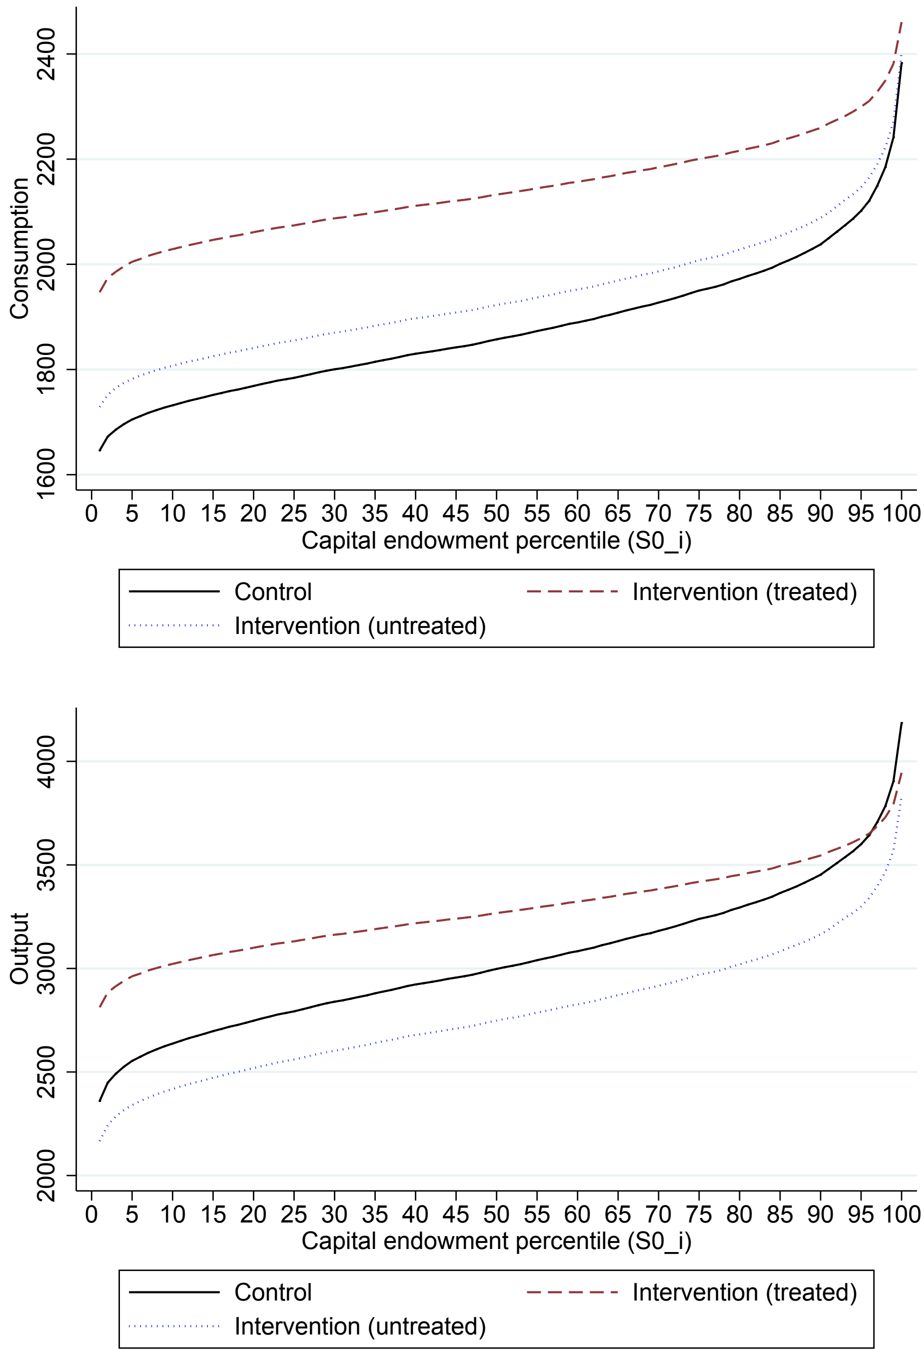

Figure 12: Welfare effects

Notes: Figure shows estimated welfare (consumption and income) effects on treated and untreated farmers. Simulations calibrate the model to baseline grain and cash reserves, with productivity parameters calibrated to match local wage and output levels.

Table 1: Loan treatments

	Loan (January)	Repayment (July)	Implied interest
		A. Maize Loan	
Offer	3 bags (50 kg ea)	4 bags (50 kg ea)	30%
Value (official)	K 195	K 260	33%
Value (reported)	K 261	K 234	-10%
		B. Cash Loan	
Offer	K 200	K 260	30%

Notes: Columns describe the loan and repayment terms, and the implied interest rate for the maize and cash loan treatment arms. The official value is the government-set maize price. The reported value is the average reported in the harvest survey for buying and selling maize.

Table 2: 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	0.013	0.015	-0.225***
	(0.007)	(0.020)	(0.018)	(0.035)
		B. Year 2		
Cash loan mean	0.98	0.76	0.78	0.55
Maize loan	-0.012	0.008	0.026	-0.084
	(0.015)	(0.058)	(0.052)	(0.089)
Treated in year 1	0.018	-0.046	-0.062	-0.142
	(0.016)	(0.069)	(0.055)	(0.098)
Any default in village in year 1	-0.025	-0.220***	-0.225***	-0.221**
	(0.026)	(0.068)	(0.066)	(0.099)
Early notification sub-treatment	0.010	-0.019	-0.007	0.012
	(0.016)	(0.058)	(0.052)	(0.090)
Cash repayment sub-treatment	-0.008	-0.049	-0.044	0.624***
	(0.015)	(0.058)	(0.052)	(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 dummary 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 3: Average treatment effects: Labor

	Any ganyu sold	Hours sold	Any ganyu hired	Hours hired	Family hours on-farm	Total hours on-farm
	(1)	(2)	(3)	(4)	(5)	(6)
A. Pooled treatments						
Any loan treatment	-0.029	-0.898**	0.020	0.996	4.484*	4.407**
	(0.021)	(0.408)	(0.013)	(0.894)	(2.285)	(1.939)
B. By treatment						
Cash	-0.016	-0.736	0.022	1.525	4.514	4.574*
	(0.025)	(0.458)	(0.016)	(1.113)	(2.806)	(2.401)
Maize	-0.042*	-1.066**	0.018	0.429	4.452*	4.234*
	(0.023)	(0.470)	(0.015)	(1.108)	(2.552)	(2.202)
Control group mean	0.303	3.731	0.114	3.714	53.258	51.463
Observations	2824	2823	2827	1637	2826	2827

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 4: Average treatment effects: Daily earnings

	Individual-level daily earnings (top 1% winsorized)	Individual-level daily earnings (top 5% winsorized)	Village median daily earnings
	(1)	(2)	(3)
A. Pooled treatments			
Any loan treatment	2.437*	1.197	1.277
	(1.249)	(0.740)	(0.969)
B. By treatment			
Cash	1.455	0.491	0.948
	(1.546)	(0.896)	(1.216)
Maize	3.399**	1.888**	1.617
	(1.551)	(0.913)	(1.127)
Control group mean	14.772	13.224	13.855
Observations	1079	1079	296

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 5: Average treatment effects: Agricultural production

	Log harvest value (1)	Log value - constant prices (2)
A. Pooled treatments		
Any loan treatment	0.081**	0.086**
	(0.037)	(0.034)
B. By treatment		
Cash	0.103**	0.114***
	(0.041)	(0.040)
Maize	0.060	0.058
	(0.043)	(0.039)
Control group mean	3187.586	3492.094
Observations	4015	4015

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 6: Heterogeneity by anticipated shortage

	Adult meals	Hours sold	Hours hired	Family hours on-farm	Log ag output
	(1)	(2)	(3)	(4)	(5)
Any loan treatment	-0.010	0.084	2.933	5.380	-0.007
	(0.031)	(0.491)	(2.088)	(3.865)	(0.037)
Anticipated food shortage	-0.110***	0.966**	2.138	-0.604	-0.191***
	(0.031)	(0.467)	(1.450)	(2.957)	(0.023)
Loan x Shortage	0.110***	-0.773	-3.136	-3.495	0.124***
	(0.040)	(0.632)	(2.200)	(4.176)	(0.037)
Control group mean	1.840	4.549	2.428	46.595	3448.630

Notes: Heterogeneous treatment effects on main outcomes, based on whether the household reported at baseline that it would run out of food before the coming harvest. Years 1 and 2 are pooled (restricted to villages newly treated in year 2). All specifications are conditional on month or year fixed effects and include baseline controls, and cluster standard errors at the village level.

Table 7: Average treatment effects: Consumption and food security

	Any grain stocks	Food security (z-score)	Adult meals - Hungry season	Adult meals - Harvest
	(1)	(2)	(3)	(4)
A. Pooled treatments				
Any loan treatment	0.114***	0.241***	0.086***	0.004
	(0.028)	(0.073)	(0.029)	(0.013)
B. By treatment				
Cash	0.005	0.084	0.072**	0.017
	(0.030)	(0.078)	(0.030)	(0.014)
Maize	0.232***	0.397***	0.103***	0.015
	(0.034)	(0.082)	(0.035)	(0.013)
Control group mean	0.530	0.000	1.874	1.993
Observations	2578	1651	1369	603

Notes: Treatment effects on consumption outcomes, measured during the short recall surveys during the hungry season (except column 4, 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 8: Notification timing

	Adult meals	Hours sold	Hours hired	Hours on farm (family)	Log output	Input value	Acres cash crops
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. Pooled treatments							
Any loan treatment	0.061*	-0.143	-1.696	3.121	0.054	114.584	0.060
	(0.034)	(0.571)	(1.490)	(5.663)	(0.111)	(187.540)	(0.150)
Loan x Early	0.006	-0.863	2.428	1.602	0.062	137.155	0.219
	(0.037)	(0.597)	(1.944)	(7.451)	(0.124)	(198.867)	(0.160)
B. By treatment							
Cash	0.004	0.559	-1.682	0.906	0.012	72.313	-0.074
	(0.034)	(0.700)	(2.118)	(7.633)	(0.129)	(313.500)	(0.186)
Maize	0.115***	-0.733	-1.618	5.330	0.101	147.407	0.204
	(0.035)	(0.712)	(1.822)	(7.125)	(0.167)	(187.781)	(0.218)
Cash x Early	0.029	-1.467*	1.095	-1.925	0.119	332.362	0.450*
	(0.046)	(0.838)	(2.545)	(12.087)	(0.198)	(359.561)	(0.247)
Maize x Early	-0.009	-0.410	3.674	5.283	0.001	-26.472	-0.008
	(0.050)	(0.778)	(3.093)	(8.525)	(0.181)	(203.939)	(0.221)
Control group mean	1.874	3.731	3.714	53.258	7.707	883.541	0.947
Observations	592	857	859	859	1004	1008	1018

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 9: Persistence and dynamics

	Adult meals	Hours sold	Hours hired	Hours on farm (family)	Log output	Input value	Acres cash crops
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. Year 1							
Any loan treatment	0.072**	-1.200**	2.742**	4.876*	0.088**	72.297	0.043
	(0.031)	(0.546)	(1.136)	(2.619)	(0.040)	(53.128)	(0.054)
Control group mean	1.957	4.549	2.428	46.595	3640.440	842.568	0.918
Observations	2171	1966	778	1967	3011	3028	3028
B. Year 2							
Any loan treatment	0.056**	-0.702*	-0.329	4.662	0.088	177.529	0.183
	(0.025)	(0.406)	(1.233)	(4.835)	(0.087)	(141.648)	(0.121)
Treated in year 1	0.036	0.885	-0.322	0.790	0.081	-31.121	0.104
	(0.027)	(0.548)	(1.132)	(3.973)	(0.064)	(92.038)	(0.088)
Loan x Treated in year 1	-0.029	-0.953	-0.421	-6.393	-0.144	-179.041	-0.108
	(0.034)	(0.608)	(1.584)	(5.810)	(0.100)	(154.930)	(0.142)
Control group mean	1.905	2.733	4.319	61.358	3202.031	1370.579	0.963
Observations	2405	1762	1764	1764	2976	2987	3005

Notes: Impacts of loan notification timing on main outcomes (columns 1-5) and planting season production decisions (columns 6 and 7), in year 2. Hungry season notification is the timing used in year 1. Planting season 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 and notification timing. 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 10: Reporting bias

	A. Social desirability bias			
	Labor survey		Endline	
Pooled treatments				
Any loan treatment	-0.041		0.041	
	(0.143)		(0.099)	
By treatment				
Cash		-0.124		-0.024
		(0.156)		(0.123)
Maize		0.043		0.104
		(0.195)		(0.118)
Observations	1387	1387	2992	2992
Control group mean		21.639		20.578

	B. Self-reported maize yields			
	Year 1		Year 2	
Objective measure	0.870***	0.775**	0.051***	0.053***
	(0.316)	(0.384)	(0.011)	(0.009)
Any loan treatment		-31.009		19.513
		(123.080)		(60.638)
Loan treatment x Objective measure		0.150		-0.002
		(0.623)		(0.019)
Share of maize yield from hybrid	292.001***	293.166***	171.017***	171.573***
	(55.952)	(57.899)	(47.567)	(47.920)
Observations	362	362	438	438
Control group mean		563.367		600.645

Notes: Tests for self-reporting bias by treatment. Panel A regresses an index of social desirability bias on treatment, with cross sectional data from two survey rounds: labor survey round 3 (hungry season, immediately after receiving year 2 loans) and endline survey (harvest survey, immediately after repaying year 2 loans). Panel B regresses self reported maize yields on an objective measure of maize productivity, the loan treatment and an interaction of the two, along with a control for the share of the self reported yield that comes from hybrid maize. In year 1 (columns 1 and 2), the objective measure is a measure of maize height during the hungry season. In year 2 (columns 3 and 4), the objective measure is based on the number of maize kernels counted during a systemic on-field sampling. See text for detail. No baseline controls are included in these regressions. Standard errors are clustered at the village level.

A.1 Model Appendix

We provide an extension to the model, which adds a "period zero" corresponding to the planting season. In future drafts, we will also use the model appendix to show We derivations of the optimality conditions for labor demand, labor supply and consumption, and show the comparative statics for these conditions with respect to the effective interest rate.

A.1.1 Model extension: Planting period investments

Our theoretical framework abstracts from initial saving and investment decisions prior to the start of the hungry season. We extend this framework to include an initial stage ($t = 0$), where household make decisions about agricultural non-labor inputs I along with consumption and saving. In such an extended framework, production y is given by:

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

where $k_i = (kI_i)$ is now heterogeneous across farms and the product of the fixed land endowment k and farm specific capital investment I_i.

Period-specific utility would be the same as before with

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

where $\alpha \in (0, 1)$, and $\underline{c} > 0$ is the minimum (subsistence) level all households must consume. Utility is now additive and separable across three periods; second and third period utility is discounted at a subjective discount factor $\rho < 1$. Savings and borrowing rates are defined as in the main text, as is the market clearing condition.

A.1.1.1 Household utility maximization

Rational households maximize their utility from consumption and leisure over three periods:

$$\max_{c,l} log(c_{i0} - \underline{c}) + \rho log(c_{i1} - \underline{c}) + \frac{\rho 1-\alpha}{\alpha} log(l_{i1}) + \rho^2 log(c_{i2} - \underline{c}) \tag{A.3}$$

subject to

$$c_{i0} \leq S_{i-1} + I_i + B_{i0}$$

$$c_{i1} \leq B_{io}[(1 + r^b)1(B_i > 0) - (1 + r^s)1(B_{i0} < 0)] + (\bar{h} - l_i - d_i)w + B_{i1}$$

$$c_{i2} \leq y_i(d_i, k_i) - B_{i1}[(1 + r^b)1(B_i > 0) - (1 + r^s)1(B_i < 0)],$$

where S_{i-1} are net resources available at the end of the previous seasons. In period 0, households consume and optimally choose on-farm investment, I_i. For simplicity we set the price of the consumption and investment goods equal to 1. In period 1, households optimally choose consumption, labor input on the farm as well as their own labor supply. In period 2, households receive harvest income. Period 2 net income (and consumption) is given by harvest production, y_i, plus period 1 borrowing or savings times the respective interest rate. Period 1 consumption is given by initial savings, S_{i0}, net labor income, given by the labor endowment, $\bar{h}$, minus time allocated to leisure, l_i, and on-farm labor inputs, d_i, (some of which may be hired) times the wage rate, w, and net borrowing, B_i. Period 0 consumption is given by savings from the previous harvest minus spending on I_i plus net borrowing.

A.1.1.2 Optimal labor and capital input

Since farmers still face the same tradeoffs in the labor markets, optimal labor inputs do not change:

$$\frac{\partial y_i}{\partial d_i} = \beta A_i (\frac{k_i}{d_i})^{1-\beta} = w r_i^e \tag{A.4}$$

while optimal investment must always be such that

$$\frac{\partial y_i}{\partial k_i} = (1 - \beta) A_i (\frac{d_i}{k_i})^{\beta} = (r_i^e)^2. \tag{A.5}$$

It is easy to see that the main insights of the model will now also apply to capital investment: high interest rates will result in low capital investment; differential interest rates resulting from frictions in the capital market will result in different marginal product of capital across farm. The results of this setup are almost identical to the main model; the main difference is that income differentials across farms with different initial reserves (and interest rates) are further increased by the addition of the capital investment channel.

Investment in I is increasing in S_{i-1} which is just the value of agricultural income from the previous harvest. After investments and any endogenous savings decisions, S_{i0} will be lower variance than S_{i-1} since lower S_{i-1} households have a lower demand for inputs and a stronger incentive to save, though the rank order will generally be preserved.

A.2 Appendix tables and figures

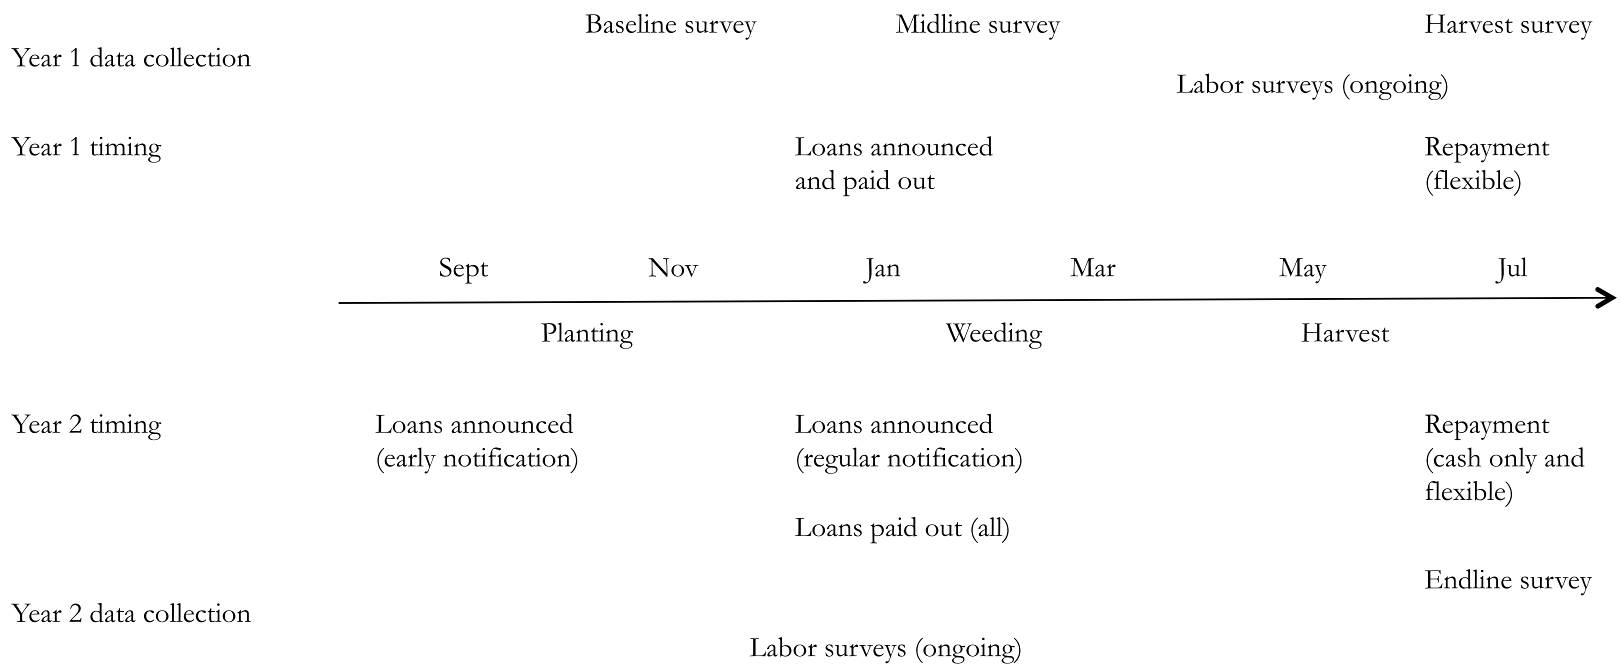

Figure A.1: Study implementation

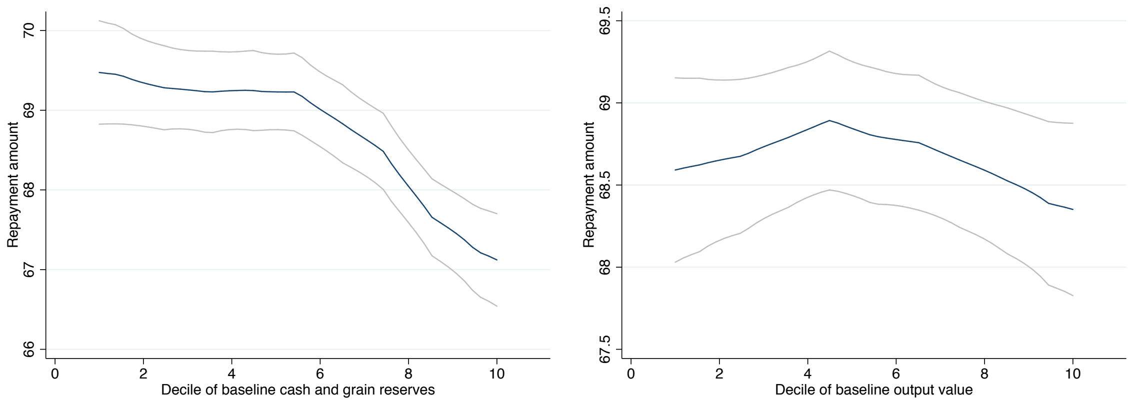

Figure A.2: Interest rates by baseline measures of liquid resources

Notes: Responses to survey question asking how much respondent would have to repay in a month for 50 Kwacha borrowed today from a source other than friends or family. The sample is restricted to the control group and the confidence intervals are from a local polynomial smoothing. The x-axis shows deciles of a measure of baseline cash and grain reserves (left figure) and baseline agricultural output value (right figure).

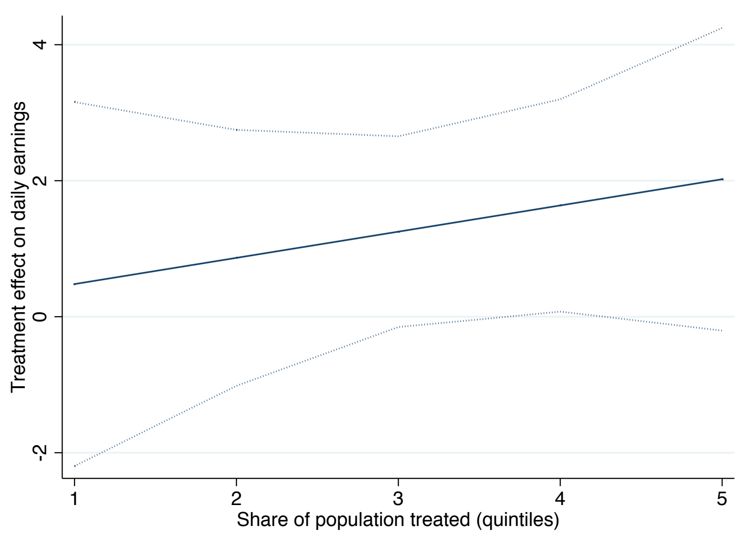

Figure A.3: Treatment effect on daily earnings, by share of village treated

Notes: Village level median daily reported earnings during the hungry season, in villages treated for the first time (pooled across years) relative to the control group. Regressions control for geographic variables, including distance to the nearest paved road. Figure shows 90 percent confidence intervals based on standard errors clustered at the village level.

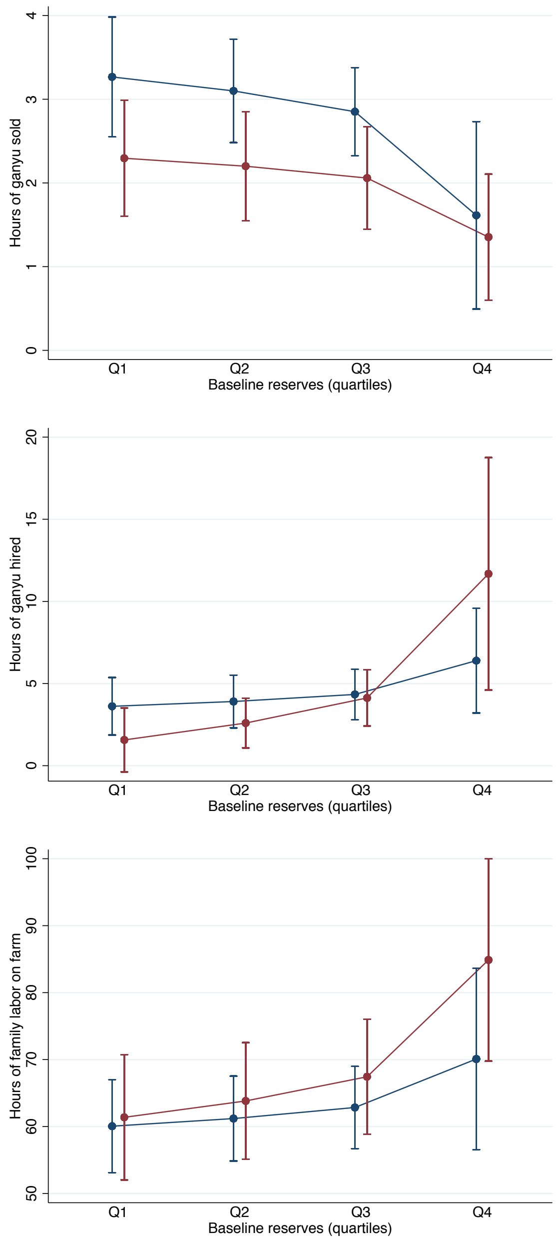

Figure A.4: Effect on labor market participation, by baseline reserves, year 2

Notes: Plots are the same is in Figure 8, for year 2.

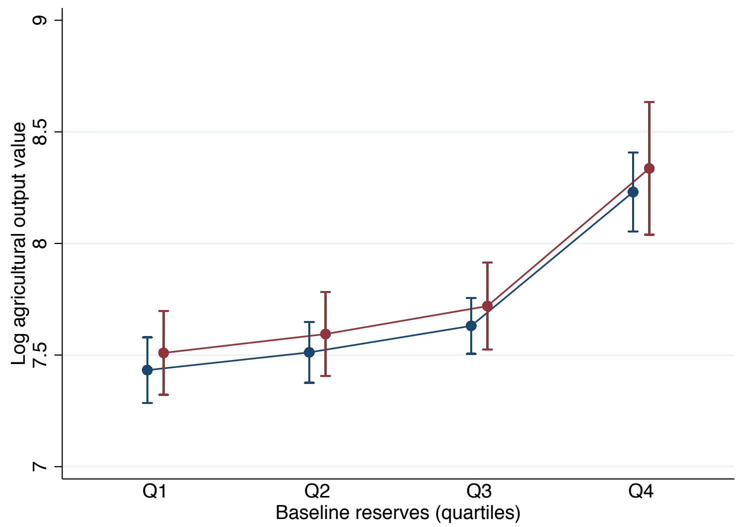

Figure A.5: Effect on log agricultural output, by baseline reserves, year 2

Notes: Plots are the same is in Figure 10, for year 2.

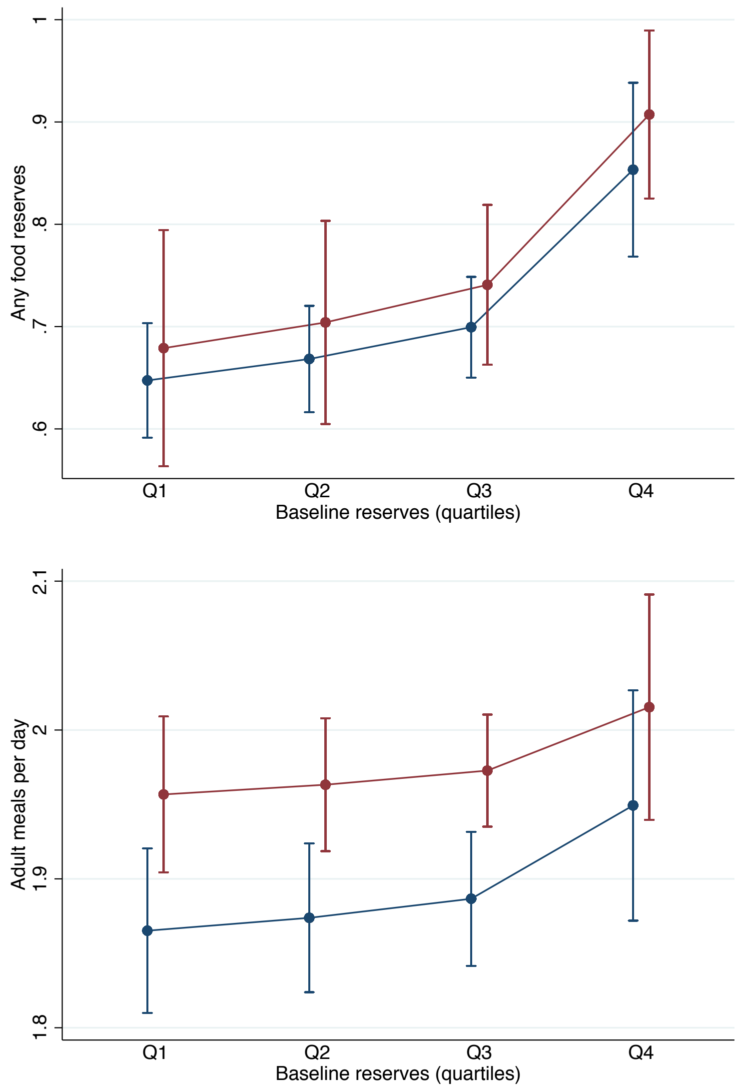

Figure A.6: Treatment effect on consumption variables, by baseline reserves, year 2

Notes: Plots are the same is in Figure 11, for year 2.

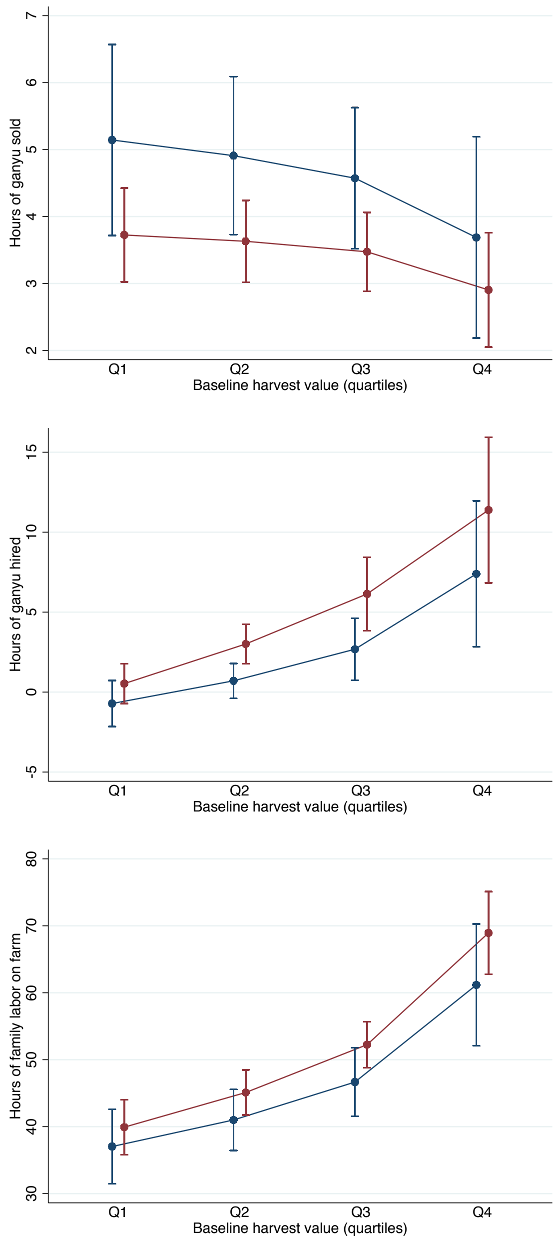

Figure A.7: Effect on labor market participation, by baseline harvest

Notes: Plots are the same is in Figure 8, using an alternative proxy for initial resources (baseline harvest value).

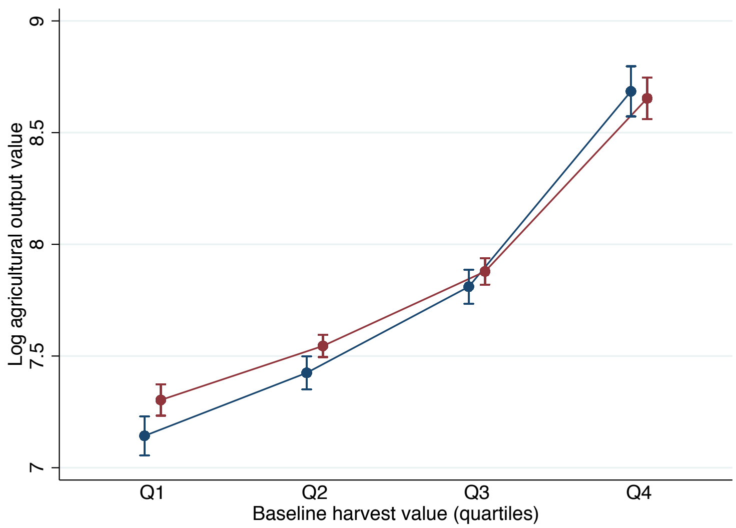

Figure A.8: Effect on log agricultural output, by baseline harvest

Notes: Plots are the same is in Figure 10, using an alternative proxy for initial resources (baseline harvest value).

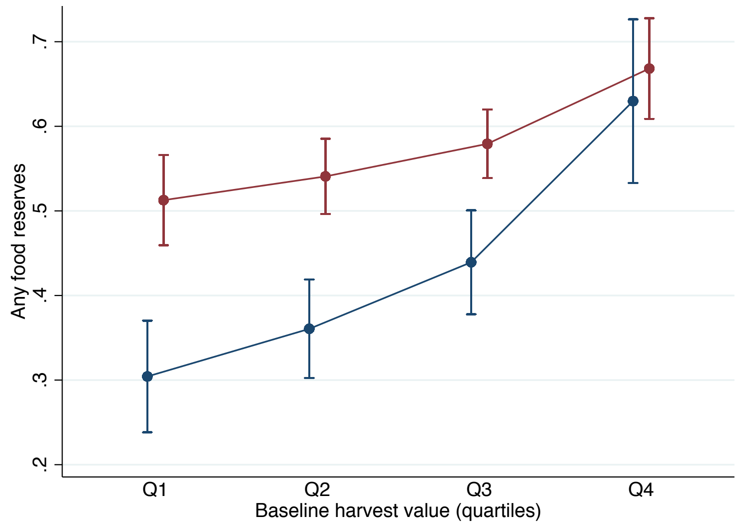

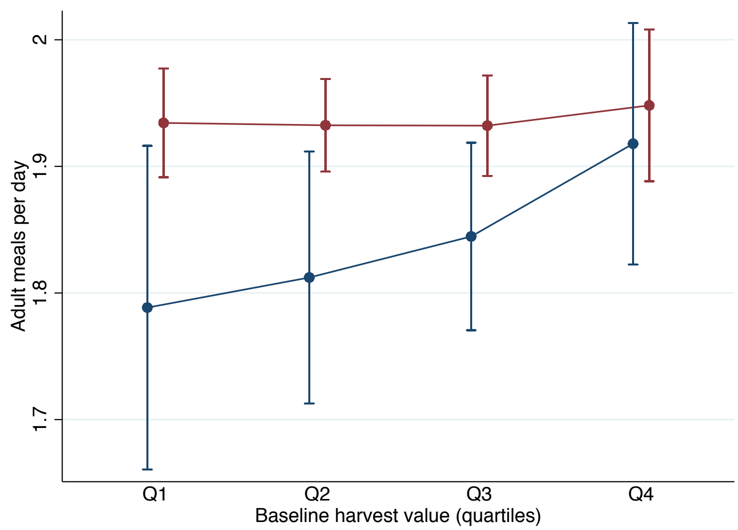

Figure A.9: Effect on consumption variables, by baseline harvest

Notes: Plots are the same is in Figure 11, using an alternative proxy for initial resources (baseline harvest value).

Table A.1: Randomization balance

	Year 1			Year 2			Largest pairwise:	
	Control	Cash loan	Maize loan	Control	Cash loan	Maize loan	t-statistic	normalized difference
Age of HH head	42.785	42.875	42.439	42.623	42.994	42.574	0.783	0.026
	[14.990]	[15.280]	[14.960]	[15.335]	[14.669]	[14.901]		
Imputed the age of the household he	0.004	0.002	0.004	0.003	0.001	0.005	1.331	0.049
	[0.063]	[0.043]	[0.061]	[0.055]	[0.037]	[0.073]		
Female HH head	0.243	0.271	0.254	0.255	0.261	0.255	1.459	0.045
	[0.429]	[0.445]	[0.436]	[0.436]	[0.440]	[0.436]		
HH members < 5	0.925	0.975	0.928	0.922	0.941	0.989	1.609	0.051
	[0.907]	[0.940]	[0.940]	[0.910]	[0.941]	[0.959]		
HH member 5-14	1.775	1.72	1.758	1.721	1.747	1.82	1.474	0.046
	[1.530]	[1.545]	[1.492]	[1.506]	[1.549]	[1.531]		
HH members 15-64	2.473	2.424	2.41	2.429	2.478	2.408	1.137	0.038
	[1.259]	[1.325]	[1.264]	[1.279]	[1.311]	[1.267]		
HH members 65+	0.172	0.187	0.173	0.183	0.167	0.174	0.897	0.03
	[0.441]	[0.473]	[0.457]	[0.457]	[0.459]	[0.457]		
Household asset quintile	3.057	3.013	2.934	2.984	3.038	3.001	1.991	0.062
	[1.392]	[1.426]	[1.422]	[1.395]	[1.443]	[1.428]		
Baseline livestock value	3586.024	3304.445	3423.568	3250.009	3479.293	3805.57	1.921	0.06
	[6220.081]	[6023.122]	[6590.769]	[6326.978]	[5759.958]	[6657.904]		
HH did ganyu last year	0.614	0.636	0.611	0.617	0.634	0.616	1.19	0.036
	[0.487]	[0.481]	[0.488]	[0.486]	[0.482]	[0.487]		
Hired ganyu last season	0.315	0.337	0.313	0.312	0.333	0.33	1.356	0.038
	[0.465]	[0.473]	[0.464]	[0.464]	[0.472]	[0.471]		
HH plans to do ganyu this year	0.622	0.64	0.634	0.631	0.642	0.624	0.856	0.029
	[0.485]	[0.480]	[0.482]	[0.483]	[0.480]	[0.485]		
# of adults working on farm	2.687	2.688	2.646	2.671	2.717	2.635	1.187	0.044
	[1.356]	[1.361]	[1.347]	[1.377]	[1.343]	[1.315]		
# of adults working in other IGA	1.126	1.132	1.181	1.137	1.148	1.165	1.264	0.039
	[1.000]	[0.965]	[0.982]	[0.983]	[0.981]	[0.983]		
Acres of maize	2.368	2.3	2.299	2.342	2.246	2.351	1.816	0.062
	[1.411]	[1.294]	[1.403]	[1.332]	[1.368]	[1.449]		
Acres of cash crops	0.967	1.04	1.099	1.06	0.949	1.069	2.809	0.094
	[1.080]	[1.168]	[1.181]	[1.162]	[1.064]	[1.185]		
Baseline harvest value	3096.843	3037.518	3032.011	3092.466	3031.347	2993.88	0.849	0.027
	[2811.876]	[2607.473]	[2765.994]	[2712.900]	[2888.188]	[2597.123]		
Crop diversity index	2.978	3.042	3.011	3.05	2.955	2.98	2.071	0.064
	[1.120]	[1.018]	[1.084]	[1.109]	[1.001]	[1.063]		
Input value	533.167	494.109	472.324	482.011	556.276	481.579	1.603	0.05
	[915.210]	[1365.201]	[807.085]	[808.145]	[1652.910]	[759.373]		

Notes: Baseline means for each major treatment group by year with standard deviations in brackets. The final two columns show the largest t-stat for a pairwise test of equal means and the largest pairwise normalized difference.

Table A.2: Attrition, by survey round

Year 1 treatments	Baseline	Y1 Harvest	Y2 Endline	Midline	Labor 1	Labor 2	Labor 3	Labor 4
Sampled	3139	3139	3139	1223	1531	1679	1673	827
Cash loan treatment	0.34	0.009	-0.009	-0.008	-0.014	-0.003	0.029	0.010
		(0.007)	(0.008)	(0.010)	(0.021)	(0.025)	(0.022)	(0.021)
Maize loan treatment	0.34	-0.005	0.008	0.000	-0.008	-0.022	-0.011	-0.055***
		(0.008)	(0.008)	(0.010)	(0.024)	(0.028)	(0.028)	(0.020)
Year 2 treatments	**Baseline**	**Y1 Harvest**	**Y2 Endline**	**Midline**	**Labor 1**	**Labor 2**	**Labor 3**	**Labor 4**
Cash loan treatment	0.23	-0.000	-0.012	-0.010	-0.003	0.049*	0.026	0.003
		(0.008)	(0.009)	(0.013)	(0.025)	(0.026)	(0.024)	(0.021)
Maize loan treatment	0.24	-0.008	0.010	-0.001	0.015	-0.059	-0.013	-0.015
		(0.008)	(0.008)	(0.010)	(0.029)	(0.036)	(0.036)	(0.023)
Conditional on treatment								
Early notification sub-treatment	0.50	0.012	-0.025**	0.007	0.003	-0.048	0.008	-0.013
		(0.010)	(0.010)	(0.015)	(0.035)	(0.043)	(0.040)	(0.028)
Cash repayment sub-treatment	0.51	-0.003	0.002	-0.000	-0.017	-0.091**	-0.037	-0.033
		(0.010)	(0.011)	(0.015)	(0.035)	(0.041)	(0.040)	(0.027)

Notes: Column 1 shows the share of sampled baseline households by treatment arm in Year 1 (Panel A) and Year 2 (Panel B). Each subsequent column reports the number of sampled households and the regression of whether the household was in the survey round on treatment indicators, conditional on being sampled for the survey. Standard errors are clustered at the village level.

Table A.3: Attrition, by participation stage

Year 1		Invited	At meeting	Eligible	Take up
Cash loan treatment	N	1023	1023	1023	1009
	Share		1	1	0.99
Maize loan treatment	N	1019	1019	1016	999
	Share		1	1	0.98

Year 2		Invited	At meeting	Eligible	Take up
Cash loan treatment					
Pooled	N	701	660	658	643
	Share		0.94	1	0.98
Notification timing sub-treatment					
Standard notification	N	356	328	328	319
	Share		0.92	1	0.97
Early notification	N	345	332	330	324
	Share		0.96	0.99	0.98
Cash repayment sub-treatment					
Standard repayment	N	336	320	319	311
	Share		0.95	1	0.97
Cash only repayment	N	365	340	339	332
	Share		0.93	1	0.98
Maize loan treatment					
Pooled	N	718	663	662	639
	Share		0.92	1	0.97
Notification timing sub-treatment					
Standard notification	N	351	327	327	314
	Share		0.93	1	0.96
Early notification	N	367	336	335	325
	Share		0.92	1	0.97
Cash repayment sub-treatment					
Standard repayment	N	365	333	332	324
	Share		0.91	1	0.98
Cash only repayment	N	353	330	330	315
	Share		0.93	1	0.95

Notes: Table reports stages of household self-selection into eligibility. To be eligible, households had to attend the meeting (before learning treatment status) and hand in a consent form (after learning treatment status).

Table A.4: Capital inputs and cash crops

	Seed	Fertilizer	Chemicals	Total input value	Acres cash crops
	(1)	(2)	(3)	(4)	(5)
			A. Both years		
Pooled treatments					
Any loan treatment	-39.443***	-10.754	22.251	-33.617	0.052
	(13.929)	(45.412)	(25.428)	(64.577)	(0.049)
By treatment					
Cash	-41.197**	4.171	-1.083	-41.366	0.065
	(15.835)	(54.740)	(10.056)	(72.975)	(0.056)
Maize	-37.695**	-25.587	45.519	-25.914	0.038
	(15.180)	(50.598)	(49.763)	(82.513)	(0.059)
Control group mean	137.976	812.084	58.479	1051.607	1.005
Observations	4030	4034	4026	4036	4046

Notes: Treatment effects on input value (col 1-4) and on acres devoted to cash crops (col 5). All regressions cluster standard errors at the village level. Outcomes in columns 1-4 are self-reported expenditures (or value when received on credit) in ZMK over the past year. Column 5 includes acres under cotton, tobacco and soya.

Table A.5: Average treatment effects: Health outcomes (year 2)

	Any illness (1)	Sought treatment (2)	Mental health problem (3)
A. Pooled treatments			
Any loan treatment	-0.071	-0.002	-0.027*
	(0.045)	(0.040)	(0.016)
B. By treatment			
Cash	-0.102	-0.065	-0.030*
	(0.061)	(0.049)	(0.017)
Maize	-0.037	0.063	-0.024
	(0.054)	(0.049)	(0.025)
Control group mean	0.782	0.916	0.040
Observations	462	348	462

Notes: Treatment effects on self-reported health outcomes in year 2, measured during the short recall surveys during the hungry season, with a 2-week recall window in columns 1 and 2. Outcome variables are: an indicator for whether the household had any illnesses in the past two weeks (col 1), whether the household sought treatment conditional on an illness (col 2), and in indicator for whether the respondent scored above a threshold for mental health problems on a 19 point scale (col 3). All specifications are conditional on month fixed effects and include baseline controls, and cluster standard errors at the village level.

Table A.6: Other consumption smoothing

	Input loan	Low interest informal loan	High interest informal loan	Sold asset	Sold livestock
	(1)	(2)	(3)	(4)	(5)
A. Pooled treatments					
Any loan treatment	-0.011	-0.009	-0.026***	0.000	0.035*
	(0.021)	(0.007)	(0.007)	(0.011)	(0.018)
B. By treatment					
Cash	-0.004	-0.004	-0.017**	-0.005	0.028
	(0.024)	(0.008)	(0.008)	(0.011)	(0.021)
Maize	-0.018	-0.013	-0.034***	0.005	0.042*
	(0.025)	(0.008)	(0.007)	(0.013)	(0.022)
Control group mean	0.438	0.072	0.065	0.072	0.316
Observations	4045	4046	4046	4045	4046

Notes: Treatment effects on other consumption smoothing strategies, measured during the long recall surveys at year 1 and 2 harvest (omitting villages treated in year 1 in the year 2 data). Outcome variables are: an indicator for whether the household accessed a formal loan including in-kind inputs (col 1), or either low interest or high interest informal loan (col 2 and 3), or sold any assets or livestock. All specifications include baseline controls, geographic controls and year fixed effects, and cluster standard errors at the village level.

Table A.7: Grain prices

	Purchase price (1)	Sales price (2)	Any purchase (3)	Any sale (4)
Pooled treatments				
Any loan treatment	0.011	-0.017	0.003	-0.002
	(0.064)	(0.032)	(0.011)	(0.011)
By treatment				
Cash	0.078	-0.028	0.008	-0.006
	(0.080)	(0.051)	(0.014)	(0.012)
Maize	-0.081	-0.007	-0.002	0.002
	(0.078)	(0.028)	(0.014)	(0.015)
Control group mean	1.035	0.977	0.048	0.061
Observations	131	110	2444	2444
By year 1 treatment status				
Any loan treatment	0.033	-0.012	-0.003	0.003
	(0.076)	(0.034)	(0.015)	(0.015)
Treated in year 1	-0.062	-0.025	0.019	0.011
	(0.078)	(0.040)	(0.013)	(0.013)
Loan x Treated in year 1	-0.045	-0.024	0.008	-0.012
	(0.111)	(0.096)	(0.020)	(0.016)
Control group mean	1.175	0.926	0.068	0.054
Observations	131	110	2444	2444

Notes: Treatment effects on maize market outcomes, measured during the short recall surveys in year 2 only (January to June). Panel A shows effects pooled across years, first for the pooled treatment effect, then by treatment arm, and Panels B and C show results by year. Outcome variables are: the price paid per kilogram for maize purchased (col 1) and sold (col 2), and indicators for whether the household made any purchase (col 3) or sale (col 4). All specifications are conditional on month fixed effects and include baseline controls, and cluster standard errors at the village level.

Table A.8: Spending on temptation and luxury goods

	Clothing	Beer	Tobacco	Sweets	Tea	Meat	Airtime
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Pooled treatments							
Any loan treatment	0.038*	0.017	-0.012	0.040**	0.013	0.022	0.002
	(0.022)	(0.018)	(0.018)	(0.020)	(0.022)	(0.019)	(0.023)
By treatment							
Cash	0.040	0.030	-0.010	0.039*	0.014	0.022	0.004
	(0.026)	(0.021)	(0.020)	(0.023)	(0.025)	(0.022)	(0.025)
Maize	0.037	0.003	-0.013	0.041*	0.012	0.021	0.000
	(0.024)	(0.021)	(0.020)	(0.023)	(0.026)	(0.022)	(0.027)
Control group mean	0.509	0.196	0.190	0.493	0.432	0.740	0.510
Observations	3937	3937	3937	3937	3937	3937	3937

Notes: Treatment effects on indicators of temptation spending over the past two weeks, measured during during the hungry season and at harvest in year 1 only. All columns are conditional on month fixed effects and include baseline controls, and cluster standard errors at the village level.

Table A.9: Income effect control

	Adult meals	Hours sold	Hours hired	Family hours on-farm	Log ag output
	(1)	(2)	(3)	(4)	(5)
Cash grant	0.029	-0.035	-1.035	-3.726	-0.039
	(0.117)	(1.731)	(1.573)	(6.529)	(0.085)
Observations	250	648	250	648	1597

Notes: Impact of a small (60 Kwacha) cash grant given to 5-6 villages, relative to the pure control group, pooling across both years of the program. Year 2 results exclude villages treated in year 1. All specifications are conditional on month or year fixed effects and include baseline controls, and cluster standard errors at the village level.

A.3 Survey descriptions

1. Baseline survey (November-December 2013, N=3139): Survey of up to 22 households per village, conducted with household heads. The baseline survey includes sections on household demographics (including individual roster, employment roster of working household members, general household information about assets owned and food insecurity faced, farming information for 2012-2013 season, expected farming activity for 2013-2014 season, risk and time preferences).

2. Labor surveys (January 2014-August 2015, N=4679): Rolling survey of ~70 households per week (7 of the baseline households in 2 villages per day). The list of baseline households for each village were randomized and the first ~7 households interviewed, in cases where a household can't be interviewed (temporarily busy, moved, etc.), the household is skipped and the next household on list visited. Survey asks one week and one to two day recall questions on household labor allocation, ganyu earnings, and consumption. Four rounds of labor surveys were conducted (a new round started once all villages were visited). The third round coincided with the hungry season in year 2 and serves as a midline survey.

3. Employer survey (January 2014-August 2015): Rolling survey of ~10 ganyu employers per week. Sampling is based on Labor survey records of where households in a village report doing ganyu. Additional sampling is done in a snowball method where employers interviewed then provide names of other employers of ganyu that they know. The employer survey tracks the labor survey by geographic block and rotates through villages rather than targeting an explicit sample.

4. Midline maize assessment (February-March 2014, N=380): On-field assessments of maize height (measurement) and visual records (photographs) for a sample of 380 households in 64 villages. Only households with their nearest field within a 30 minute walk were eligible.

5. Midline survey (February-March 2014, N=1193): Hungry season survey of 1200 randomly selected households, stratified on treatment. One week and one month recall questions on labor supply, ganyu earnings, consumption, basic strength and anthropometric measurement.

6. Harvest survey, year 1 (July-September 2014, N=3028): Survey of all baseline households. Includes sections on changes to household composition, shocks experienced by the household, agricultural productivity. Includes anthropometric measures for adults and children.

7. Endline survey, year 2 (July-September 2015, N=3005): Survey of all baseline households. Similar structure to harvest survey.

C.1 Choice experiments

Hypothetical choice experiments were conducted on a convenience sample of participants in November and December 2013. In the initial wave of questions, 72 respondents were interviewed, one-third of which were female. The surveys took place in villages in and around the study area, but not eligible for the study either because they were too large (>100 households) or they had participated in the pilot program. Respondents were approached by an enumerator who explained the exercise, emphasizing that the offers were hypothetical and that responses would not affect any future programs they might be offered. In spite of these disclaimers, which were intended to minimize strategic responses and avoid building expectations, respondents took the decision tasks seriously.

Six scenarios were presented to respondents, involving different dichotomous choices that varied a relevant parameter of the loan offer. The ordering of the parameter set were varied across respondents.

Scenario 1: Maize loan versus cash loan

Script: *Suppose that we had two loans available that would start in January. The first would offer three (3) bags of [50 kg maize] in January that you have to repay in June. The second would offer cash that you would have to repay in June. Please take your time to make your choice, as I will be going through different categories. Would you prefer a cash loan that paid ____ Kwacha that you would pay back in June or would you prefer the [maize] loan that you would pay back in June?*

Parameters: 50, 110, 150, 175, 250, 275, 350, 375, 425, 450, 600 Kwacha

Scenario 2: Cash repayment

Script: *Now, supposed the loan changed so that you could still receive three (3) bags of [mealie meal / maize] in January. But instead of repaying in maize in June, you had to repay in cash. I'm going to go through some different repayment amounts. You should tell me whether you would choose to take up a loan that gave you [maize] in January and had to repay that amount of cash in June. Would you be willing to take up a loan that gave you 3 bags of [maize] in January and required that you repay ___ Kwacha in June?*

Parameters: 600, 450, 400, 325, 275, 250, 200, 175, 125, 100, 75, 50 Kwacha

Scenario 3: Cash gift vs. maize loan

Script: *Again, suppose, we were to offer a loan that offered three (3) bags of [maize] in January that you had to repay in June. Would you prefer to take that loan or would you prefer to receive ____ Kwacha in January, which you would not require to pay back?*

Parameters: 10, 30, 60, 80, 100, 110, 130, 150, 175, 200, 250 Kwacha

Scenario 4: Cash gift vs. cash loan

Script: *Suppose now that the loan was cash instead and we were to offer a loan that provided 200 Kwacha in January that you had to repay in June without any interest (repay 200 Kwacha in June). Would you prefer to take that loan or would you prefer to receive _ _ _ _ Kwacha in January which you would not require to pay back.*

Parameters: 10, 30, 60, 80, 100, 110, 130, 150, 175, 200, 250 Kwacha

Scenario 5: Maize loan repayment month

Script: *Suppose, we were to offer a loan that offered three (3) bags of [maize] in January that required you to repay four (4) bags. I'd like you to think about whether you would choose to take that loan. I will list different months when the repayment would be due. Would you be willing to take a loan of three bags of [mealie meal / maize] in June that required you repay 4 bags if the repayment were due in _ _ _ _ _ _?*

Parameters: February, March, April, May, June, July, August, September, October, November, December

Scenario 6: Cash loan repayment month

Script: *Again, let's look at this activity but considering a loan in cash instead of maize: Suppose, we were to offer a loan that offered 200 Kwacha in cash in January that required you to repay 330 Kwacha in cash. Would you be willing to take that loan for 200 Kwacha in cash that repaid 265 Kwacha if the repayment were due in _ _ _ _ _ _?*

Parameters: February, March, April, May, June, July, August, September, October, November, December