CREDIT AND WELFARE ACROSS THE LEAN SEASON

ETHAN LIGON & JED SILVER

Abstract. Consumption expenditures in rural areas of low-income countries are highly variable across seasons, yet the literature still lacks a standard framework for asking whether seasonal poverty reflects local credit market failure or poor integration with the broader economy. We develop an intertemporal model of farmers' portfolio choices under seasonal price risk and borrowing constraints, and derive a sign diagnostic: The amount a farmer is willing to pay for a small risk-free bond (call this price q) must rise in response to a positive income shock if credit constraints bind, but fall if precautionary motives dominate. We apply this framework to a randomized post-harvest loan program in Gombe, Nigeria, and supplement the experiment by also collecting high-frequency data on prices, stocks, and expenditures. The loan sharply reduces the marginal utility of expenditure around delivery, but q never rises over the full follow-up $\Delta \log q < 0$. Precautionary savings, not credit constraints, govern the intertemporal allocation. Receipt of the loan leads to a portfolio rebalancing, as farmers adjust their grain stores, and increase investment. But maize prices increase little after harvest and season-average consumption expenditure effects are small, though the MUE—a more sensitive welfare measure—detects a large and significant improvement around delivery that expenditures miss. We fail to reject the null of well-functioning local financial markets. The binding constraint is poor spatial integration rather than inefficient local allocation—promoting market integration may improve lean-season welfare more than would the local provision of credit.

Date: March 26, 2026.

We are grateful to Justin Graham, Parth Ahya, Abubakar Umar, and others at the Taimaka Project for their collaboration, dedication and helpful discussions. We also wish to express our appreciation to Muhammad Meki, Kathleen Beegle, and Albert Rodriguez-Sala for their helpful comments, to seminar audiences at Ben-Gurion, Berkeley, Haifa, Hebrew U, Notre Dame, and Madison, and to conference attendees in Philadelphia, Nairobi, Bergen, Bordeaux, Marseille and Blantyre. We gratefully acknowledge funding from STEG. This study was registered as "Offering Postharvest Credit to Improve Farmer Welfare" in the AEA RCT Registry (ID number AEARCTR-0008022).

1. Introduction

In isolated rural areas, the nature of the agricultural calendar often creates seasons of relative plenty, and its opposite—the "lean season." Food prices rise between harvests and consumption falls, with consequences for both short-term welfare and long-run outcomes like health and human capital (Christian and Dillon 2018). But designing an effective policy response requires two things: credible measures of the welfare cost—hard to come by when high-frequency expenditure data are scarce (Merfeld and Morduch 2023) and mapping them into welfare isn't straightforward—and an understanding of the root cause. Do lean seasons persist because local markets fail to allocate resources efficiently across households, or because local markets are poorly connected to the broader economy? In particular post-harvest loans (PHLs) and other forms of subsidized credit have been proposed as a remedy for lean seasons, but their justification hinges on local credit markets actually being the problem. Despite a growing empirical literature on seasonality, the field still lacks a disciplined framework for distinguishing these explanations—one that embeds farmers' storage and consumption decisions in a fully articulated intertemporal optimization problem and derives testable predictions about market functioning. We apply such a framework to randomized post-harvest loan (PHL) program in Gombe, Nigeria, combined with high-frequency data on expenditures, stocks, prices and experimental elicitations of households' intertemporal prices.

It seems difficult to justify PHLs on theoretical grounds,[1] especially without a diagnosis of credit market frictions, and yet the empirical literature reports encouraging results (Burke, Bergquist, and Miguel 2019; Channa et al. 2022). How can one make sense of this? RCTs are informative only about treatment effects in realized states of the world (Rosenzweig and Udry 2020); they can't tell us whether the treatment is good policy in expectation. To shed light on this, we partnered with the Taimaka Project, an NGO in Gombe State in northeastern Nigeria, to implement a randomized PHL program[2] modeled after the One Acre Fund scheme studied by Burke, Bergquist, and Miguel (2019). We investigate two treatments: a cash loan and a similarly-valued in-kind loan of maize. But we go further: We collect high-frequency survey data on stocks and disaggregate expenditures, supplemented

[1]One might be surprised by the idea of offering farmers credit right after harvest—isn't that when resources are plentiful and prices low? It would seem to be a time to save, not borrow. Advocates reply that if farmers could borrow at harvest to cover consumption needs, they could avoid selling their crop at its seasonal nadir. But framed this way, PHL programs are really encouraging farmers to make leveraged bets on the price of the crop they grow. Leveraged bets are inherently risky. In this case, local year-to-year variation in prices is large relative to any predictable seasonal component; Cardell and Michelson (2022) show that across a wide range of sub-Saharan African settings, the seasonal price spike frequently fails to materialise, and that even mild risk aversion can rationalize the apparent failure to arbitrage. Our contribution is complementary but distinct: rather than asking whether risk aversion explains low storage, we test directly whether local financial markets allocate resources efficiently.

[2]See pre-specified descriptions of our tests at `https://www.socialscienceregistry.org/trials/8022`.

with experimental elicitations of the intertemporal price q, and use conditional moment restrictions from the asset-pricing literature to test whether (in expectation) PHLs are suited to address the relevant market failures. These tests are informative about the general state of credit markets during the lean season, independent of any particular loan program.

What do we find? Two broad patterns stand out. First, seasonal variation dwarfs any treatment effect. Scarcity rises sharply through the lean season and falls at harvest, and these swings are several times larger than anything the loans induce. This is what one would expect if the problem is aggregate, not allocative. Second, the cash and maize loans are almost never distinguishable from each other; the *form* of the loan matters much less than whether a household gets one at all, undermining the specific "encourage maize storage" theory of change behind PHLs.

Because the loan's interest rate (15% over nine months) is well below elicited shadow rates, it functions as a positive-NPV transfer for virtually every recipient. We develop some theoretical predictions to help us understand how such a positive wealth shock should affect outcomes given different frictions. The causal chain we seem to observe runs as follows. The transfer sharply lowers the marginal utility of expenditure around delivery ($\Delta \log \lambda \approx -0.22$): households are genuinely better off on impact. They don't immediately consume the windfall—instead they save it, building up grain stores and increasing business investment (cash arm). But grain prices stay relatively flat, and most households that sell grain do so at prices that don't cover the 15% interest rate. The leveraged bet on storage doesn't pay off in this year: By endline, estimated household MUEs are significantly higher for both treatments—consistent with households cutting consumption to repay a loan that financed an investment that didn't pan out. Season-average effects on consumption expenditure are small, though the MUE picks up welfare changes that expenditures miss. We find no significant effect on total grain sales.

So, at least in Gombe in 2021–22, investments in storing maize were a bad idea *ex post*. But the question that matters for policy is whether PHLs are viable *in expectation*, with expected benefits that exceed the costs of the program? It may be that giving a loan or a transfer to a credit-constrained farmer would allow farmers to improve their expected utility by changing their portfolio of stocks and financial assets.[3] Adapting tools from finance, we test for systematically differential responsiveness of intertemporal prices to grain price changes. If local financial markets function well and households optimize, these prices may respond to shocks, but they'll do so uniformly across households. Credit constraints, by contrast, would cause different households to hold different grain portfolios, and their intertemporal prices would deviate from the economy average depending on returns. If these are not the major

[3]This necessary condition is far from sufficient. For example, there are also supply-side issues: prices *within* a market are highly correlated, so a lender needs sufficient diversification to withstand correlated default when returns are low.

constraints, a local PHL program is unlikely to address the root cause of seasonal hunger—it is simply injecting a fixed quantity of capital at below the market-clearing interest rate.[4] The losses from such programs would then be better spent on lowering *aggregate* risk, perhaps by improving market integration.

We provide two complementary pieces of evidence against credit constraints as the dominant friction. The first is a within-household test: we show that under standard preference restrictions, a credit-constrained household's intertemporal price q (the idiosyncratic price they'd be willing to pay for a risk-free bond, if such an asset was traded) must *rise* in response to a positive income shock, while an unconstrained household's must fall. An accident of timing (maize loans arrived after the first round of bond-price elicitations) lets us test this directly. For recipients of the PHL transfers the intertemporal price never rises; over the full follow-up $\Delta \log q < 0$ ($t = -2.4$), ruling out binding credit constraints and pointing instead to consumption smoothing and precautionary motives for saving. The second is a cross-sectional test: we look for systematically differential responsiveness of intertemporal prices to grain price changes using lagged bond prices (both estimated and experimentally elicited), lagged stocks, and random assignment to treatment. In each case, and jointly, we fail to reject the hypothesis that local financial markets function efficiently. This may be surprising, but it's consistent with evidence that village economies smooth idiosyncratic risk fairly well while struggling with aggregate risk (Samphantharak and Townsend 2018).[5] The main culprit for low and highly variable consumption expenditures appears to be seasonal variability in *aggregate* resources—consistent with Gombe being poorly integrated with the rest of Nigeria, which we illustrate using high-frequency price data. This implies that facilitating *spatial* rather than *intertemporal* arbitrage would more effectively address the underlying causes of seasonal poverty, without the risks and logistical challenges of PHLs.

Beyond the specific policy evaluation, we view the framework itself as a contribution. Most of the empirical literature on seasonality, including the influential work of Burke, Bergquist, and Miguel (2019), provides compelling evidence on seasonal price patterns and treatment effects but does not offer a formal model of intertemporal price behavior, expectations, and welfare. Our model embeds farmers' storage and consumption decisions in a fully articulated

[4]Such programs can be thought of as cash transfers with an aggregate value equal to the interest rate subsidy times the quantity of capital, but a highly inefficient transfer given their administrative costs and default risk.

[5]One might argue that the overwhelming demand for Taimaka's loan product is evidence of poorly functioning financial markets. But excess demand for artificially cheap credit isn't a sign of market failure—and since nearly every farmer borrows the maximum amount, the PHL is functioning as a take-it-or-leave-it transfer, not a marginal credit decision. More fundamentally, with pronounced seasonality and poor market integration the local risk-free rate must vary over time to clear debt markets—and our elicited rates confirm substantial seasonal variation. A fixed subsidized rate will naturally generate excess demand regardless of how well local markets function.

intertemporal optimization problem: households choose portfolios of storable assets with uncertain returns, subject to borrowing constraints and season-specific distributions of weather and prices. The resulting moment conditions let us test whether local asset markets function efficiently, compare classes of policy interventions (crop-specific credit, general liquidity, improved storage technology), and ask which margins a given policy relaxes. Because the large cross-crop and cross-year variation in seasonal prices we document is likely not unique to Gombe, this framework should be portable to other sub-Saharan African settings where high-frequency price and household data are available.

1.1. **Related Literature.** These results contribute primarily to the literature on seasonality of income and consumption in agricultural settings, including work on PHL and storage interventions. Basu and Wong (2015) find that storage drums and in-kind staple food loans increase expenditure and income, but that only credit smooths lean-season consumption. Aggarwal, Francis, and Robinson (2018) and Omotilewa et al. (2018) find that pure storage interventions using hermetic PICS bags increase maize storage volumes, duration, and revenues in Kenya and Uganda, respectively. Channa et al. (2022) also include a treatment arm with only PICS bags and find slightly smaller, insignificant effects on storage compared to households receiving both credit and bags (but can't reject equality). Channa et al. (2022) and Burke, Bergquist, and Miguel (2019)—who study the program on which Taimaka's PHLs were modeled—find that postharvest credit improves farmers' incomes but doesn't affect consumption or investment.[6] We provide not just a cautionary tale about PHLs in years with minimal price increases, but show how credit provision more generally may fail to address the root causes of seasonal hunger. Nevertheless, other schemes such as warrantage have shown promise at smoothing intertemporal consumption in similar environments (Delavallade and Godlonton 2023).

We also contribute to a broader literature on lean-season consumption variability. Barrett and Dorosh (1996) argue that poor households bear the brunt of price variability in Madagascar. Fink, Jack, and Masiye (2020) find that lean-season credit enables a productive reallocation of labor from off-farm to farm in Zambia. Also in Zambia, Augenblick et al. (2023) find that a simple budgeting exercise induces households to save maize significantly longer into the lean season. A series of papers (Bryan, Chowdhury, and Mobarak 2014; Meghir et al. 2022) studies seasonal migration as a risky but profitable means of smoothing consumption. While these experimental studies point to a more positive role for credit, our results differ—likely because of how local financial markets function in Gombe and the particular state of the world realized during our study.

[6]The idea that credit can alleviate "sell low, buy high" is also supported by Stephens and Barrett (2011) and Dillon (2021), while the price-risk rationalization is supported by Cardell and Michelson (2022).

Finally, we relate to the trade and development literature on market integration and income volatility (Fafchamps 1992; Burgess and Donaldson 2010; Caselli et al. 2020; Allen and Atkin 2022). While we don't provide direct evidence on the effects of market integration, we're among the first to consider the relative benefits of spatial and intertemporal arbitrage. Gombe's apparent disconnection from the broader economy is consistent with the local general equilibrium price effects found by Burke, Bergquist, and Miguel (2019) in Kenya, and the spatial price variation across Africa documented in Porteous (2019). Bergquist, McIntosh, and Startz (2021) provide an example of a low-cost intervention that improved market integration in Uganda—an approach that might facilitate spatial arbitrage in settings like Nigeria as well.

The rest of the paper proceeds as follows. Section 2 develops the intertemporal model and derives the moment conditions we take to the data. Section 3 describes the experimental design, while Section 4 describes the economic environment in Gombe, including the cross-crop and cross-year price variation that motivates our framework. Section 5 presents the RCT results. We then turn to what these results imply about local financial markets. Section 6 tests the strong null of full insurance and rejects it. Section 7 tests the weaker null that households' portfolios are efficient given available local assets; we fail to reject this, suggesting that the binding constraint is poor market integration rather than inefficient local allocation. Section 8 concludes with implications for policy design and for applying this framework to other settings.

2. Model

We want a model that gives us sharp, testable predictions of what farmers' decisions would look like under a null hypothesis of no-arbitrage, and that can be used to evaluate alternative policies within a common structure. We develop three results in this section. First, an asset pricing equation that links each household's portfolio choices to the covariance of asset returns with its intertemporal marginal rate of substitution. This is utterly standard. Second, an aggregation that tells us what no-arbitrage looks like in the local economy; we believe this may be novel. Third, a moment condition that lets us test whether individual households' IMRSs respond differently to returns than does the local average—which is what we'd expect to see if local financial markets weren't functioning well.

Consider a collection of farm-households in Gombe. Each household has a per-period utility function $U(c)$ over a J-vector of goods c, assumed to be increasing, concave, and continuously differentiable. Goods can be bought and sold on local spot markets at prices p; some can also be stored across periods, and some can be produced by the farm-household.

Our principal concern is seasonality, so we focus on a single agricultural year consisting of a sequence of periods indexed by $s = 1, \ldots, S$. Agricultural output depends on weather,

and weather—while random—varies across seasons, with season s having a realization w_s drawn from a cumulative distribution function $G_s(w)$. Over the course of the year the weather history $w^s = (w_1, w_2, \ldots, w_s)$ affects agricultural output, drawn from a conditional distribution $F_s(y \mid w^s)$. For seasons when there's no possible harvest this distribution may be degenerate (i.e., $F_s(0 \mid w^s) = 1$). From the farm-households' point of view,[7] prices may also be random but drawn from season-specific distributions, with a cumulative distribution function $H_s(p \mid w^s)$ that may depend on the weather history.

In any season, the farm-household enters the period holding a vector of stocks k, to which any new production y is added. These stocks $k + y$ can be adjusted—subject to possible bounds, *infra*—via purchases or sales at prices p. From this the household consumes a vector c and carries forward a vector of stocks k'. The flow budget constraint is

$$p(k' + c) = p(k + y). \tag{1}$$

Some stocks must be non-negative—think of physical maize, which you can't hold in negative quantities—though this needn't be true for financial assets (think of bonds, which can be negative for a net debtor). So we also have a set of J constraints

$$k' \geq b, \tag{2}$$

where b is a vector of lower bounds on stocks of each good (some elements could be minus infinity), assumed fixed.

So: in any season s, the household takes as given the stocks k brought into the period, prices p, and the weather history w^s. Its problem can be written as a dynamic program:

$$V_s(k, p, w^s) = \max_{k', c} U(c) + \beta \int V_t(k', p', w^t)\, dF(y \mid w^t)\, dH_t(p' \mid w^t)\, dG_t(w_t),$$

with $t = (s + 1) \mod S$, subject to the budget constraint (1) and bounds (2).

This yields the first order conditions

$$u_j(c)/p_j = \lambda \qquad j = 1, \ldots, J,$$

where λ is the Lagrange multiplier on the flow budget constraint, $u_j(c)$ is short-hand for $\frac{\partial U}{\partial c_j}(c)$, and

$$p_j \lambda = \beta E \left[p'_j \frac{\partial V}{\partial k_j} \right] + \eta^j$$

where η_j is the multiplier on the lower bound for stocks of good j.

Adding time subscripts and applying the envelope theorem delivers the standard asset pricing equation for each asset j, allowing for the possibility that the household is at a

[7] From the point of view of the economy more broadly, prices may *not* be random, but instead may adjust to clear markets. We could close the model by describing distributions of prices H_s as equilibrium objects.

corner (i.e., (2) is binding), in which case the multiplier η_t^j will be positive:

$$\beta \mathbb{E}_t \frac{\lambda_{t+1}}{\lambda_t} \frac{p_{t+1}^j}{p_t^j} + \frac{\eta_t^j}{p_t^j \lambda_t} = 1 \qquad \text{for all } j = 1, 2 \ldots, J. \tag{3}$$

The realized intertemporal marginal rate of substitution (IMRS) between periods t and $t+1$ is $m_{it+1} = \beta\lambda_{t+1}/\lambda_t$. Since λ_t is the marginal utility of expenditures at time t, and expenditures depend on income, prices, and consumption via (1), Equation (3) tells us that the farm-household's portfolio choices are governed by the covariance of returns $R_{t+1}^j = p_{t+1}^j/p_t^j$ with potential realizations of the IMRS across the available assets j. That's our first result. Note that an upper bound on holdings of asset j—a loan ceiling, say—would introduce a multiplier analogous to η_t^j but of opposite sign. When this ceiling binds, the household is again at a corner, and (3) doesn't hold at equality for that asset.

Now assume there are $i = 1, 2, \ldots, N$ households in a given location (e.g., Gombe). With data on common prices p_t^j and estimates of λ_{it}, the orthogonality conditions implied by Equation 3 can be used to understand the functioning of local financial markets.

To develop this, let $\delta_{it}^j \in \{0, 1\}$ indicate whether household i holds positive stocks of asset j at time t. Let $R_{t+1}^j = p_{t+1}^j/p_t^j$ denote the one-period return to holding asset j from t to $t+1$, and let $m_{it+1} = \beta\lambda_{it+1}/\lambda_{it}$ denote i's *ex post* IMRS. By complementary slackness:

$$\delta_{it}^j \mathbb{E}_t \left(m_{it+1} R_{t+1}^j \right) = \delta_{it}^j \qquad \text{for all } j = 1, 2 \ldots, J. \tag{4}$$

Now let $\bar{m}_{t+1}^j = \sum_{i=1}^N m_{it+1}\delta_{it}^j / \sum_{i=1}^N \delta_{it}^j$ denote the average IMRS among households holding positive stocks of the asset. In the aggregate, an absence of arbitrage opportunities implies our second result,

$$\mathbb{E}_t \bar{m}_{t+1}^j R_{t+1}^j = 1 \qquad \text{for all } t. \tag{5}$$

How is this consistent with seasonality? If seasonality affects local prices for asset j in a way that's at all predictable, then we can't have $\mathbb{E}_t R_{t+1}^j = 1$, so for (5) to hold $\bar{m}_{t+1}^j$ must covary negatively with R_{t+1}^j—marginal utilities must decrease as prices of j increase. The logic is simple: if a storable asset will likely command higher prices next period, money can be made by investing in more of it. On the other hand, if there's increased aggregate storage of asset j—"buy low, sell high"—this may retard the rate of price increase, particularly if the local economy is poorly integrated with the broader economy. Note too that even if no risk-free debt is exchanged locally, each farmer has a shadow price individual rate $R_{it}^f = [\mathbb{E}_t m_{it+1}]^{-1}$; under full insurance these would coincide, but since we reject full insurance (Section 6), each farmer in general faces a different cost of credit. With pronounced seasonality and poor integration, these rates must vary over time and in response to aggregate shocks to clear local debt markets. We elicit R_{it}^f directly using a multiple price listing experiment described below.

We can therefore look at evidence that local prices for storable goods like maize change predictably over the season. If the change is only in *local* prices—if they don't covary with world prices or with prices elsewhere in Nigeria—then this is evidence not only that maize markets are poorly integrated, but that Gombe's *other* markets are too. Why? Refer back to (5). Suppose prices for good j are seasonal and predictably increasing, and the market for j is poorly integrated. But suppose the market for some other good *is* well-integrated. Then Gombe farm-households could take a short position in the latter market—since they're small relative to the broader economy, this would only modestly affect prices—and use those funds to store Gombe maize. This bet on maize prices would be profitable up to the point at which the aggregate supply response (less maize for consumption now, more stored for next period) drives down returns enough to satisfy (5).

Now back to individual households. Substituting (5) into (4), we obtain

$$\mathbb{E}_t(m_{it+1} - \bar{m}^j_{t+1})R^j_{t+1}\delta^j_{it} = 0. \tag{6}$$

What is this object? Focus first on $(m_{it+1} - \bar{m}^j_{t+1})$: the deviation of household i's IMRS from the average IMRS of households holding asset j. If all households were fully insured, this term would always be zero—aggregate shocks to Gombe might change IMRSs, but they'd change in precisely the same way for everyone. However, (6) can hold even without full insurance. If we think of $(m_{it+1} - \bar{m}^j_{t+1})$ as measuring idiosyncratic shocks, then (6) tells us these shocks are orthogonal to returns on assets the household holds. Put differently, this idiosyncratic risk is *uninsurable* using the assets held in positive quantity.

Where i's IMRS differs from the aggregate, we might suspect this is because it held a different portfolio than other households. For example, if household i held proportionally more maize than others, then $m_{it+1} - \bar{m}^j_{t+1}$ would be negatively correlated with returns to maize. So we're looking for evidence that some households have IMRSs which predictably respond differently to returns than the average household does. If these predictable differences exist, given local prices, that could be interpreted (Burke, Bergquist, and Miguel 2019) as evidence of unexploited arbitrage opportunities.

This brings us to our third result: If we can't reject (6), there's no evidence of local arbitrage opportunities—though *spatial* arbitrage may still exist if prices vary differently across localities. The central prediction implied by the null of no-arbitrage is that realizations of $y^j_{it+1} = (m_{it+1} - \bar{m}^j_{t+1})R^j_{t+1}\delta^j_{it}$ should be uncorrelated with any variable z^j_{it} in the time t information set. Or expressed conversely, if local financial markets do not function well, then lags of IMRS and stocks, as well as treatment assignment, would plausibly be correlated with y^j_{it+1}.

2.1. **Credit Constraints, Precautionary Premia, and the Intertemporal Price.** The intuition is straightforward. A household that receives a positive income shock can do one of

two things: consume it today or save it for the future. A credit-constrained household—one whose borrowing constraint binds—will predominantly consume the transfer now, raising current consumption and lowering current marginal utility while leaving future marginal utility unchanged or higher (especially if, as in our setting, a repayment obligation falls due later). This pushes the intertemporal price q up. An unconstrained household, by contrast, saves part of the shock, raising expected future resources; under standard preference restrictions, this shrinks the precautionary premium by enough to push q down. The sign of Δq thus identifies which friction dominates.

We establish this in four steps. First, we decompose the Euler equation to show how Δq relates to changes in the risk premium and the relaxation of borrowing constraints (Proposition 1). Second and third, we show that under progressively stronger preference restrictions—Decreasing Absolute Prudence, then Decreasing Relative Prudence—a positive income shock shrinks the precautionary premium and the uninsurable variation in the IMRS (Propositions 2 and 3). Fourth, we combine these results to show that an unconstrained household's intertemporal price must fall in response to a positive shock (Proposition 4), so that an observed increase identifies binding credit constraints (Corollary 1).

Definition 1. We define the intertemporal price q as the expected intertemporal marginal rate of substitution (IMRS):

$$q \equiv \beta \frac{\mathbb{E}[\lambda(x', p')]}{\lambda(x, p)}$$

Definition 2. We define the precautionary premium ψ as the wedge generated by uninsurable resource variation, evaluated by conditioning on the realization of future prices p':

$$\psi \equiv \mathbb{E}\left[\mathbb{E}[\lambda(x', p')|p'] - \lambda(\mathbb{E}[x'|p'], p')\right]$$

2.1.1. *Credit Constraints.* Let $q_{it} = \mathbb{E}_t[m_{it+1}]$ be the intertemporal price for household i at time t. Consider a discrete, exogenous positive endowment shock $\Delta y > 0$. For any asset $j \in \{1, \ldots, J\}$, we define the normalized shadow wedge $\mu^j_{it} = \frac{\eta^j_{it}}{p^j_t \lambda_{it}}$, where $\eta^j_{it} \geq 0$ is the multiplier on the lower bound constraint for stocks of good j.

Proposition 1. *Consider a positive perturbation of the income vector $\Delta y \geq 0$ (where $\Delta y_j > 0$ for at least one good j), and assume that returns are positive and independent of Δy. If this perturbation strictly decreases the household's intertemporal price ($\Delta q < 0$), then for any asset j in the household's choice set, the perturbation strictly increases the covariance between the IMRS and the asset return. Furthermore, this increase strictly bounds the magnitude of any liquidity constraint relaxation:*

$$\Delta Cov(m, R^j) > -\Delta \mu^j \geq 0$$

Proof. For any asset j, the household's Euler equation accommodating a potential lower-bound constraint is given by:

$$\mathbb{E}[mR^j] = 1 - \mu^j$$

Applying the standard covariance decomposition to the expected product and substituting the intertemporal price $q = \mathbb{E}[m]$, we obtain:

$$q\mathbb{E}[R^j] + \text{Cov}(m, R^j) = 1 - \mu^j$$

Taking the discrete difference of this equation with respect to the strictly positive exogenous shock Δy, and noting that expected returns are invariant to the household's endowment, yields:

$$\Delta q\mathbb{E}[R^j] + \Delta\text{Cov}(m, R^j) = -\Delta\mu^j$$

Because a positive wealth perturbation weakly expands the household's choice set and weakly increases current consumption, it can only reduce or leave unchanged the shadow penalty of a borrowing constraint. Therefore, the change in the shadow wedge is weakly negative ($\Delta\mu^j \leq 0$), which implies the liquidity relief term is weakly positive:

$$-\Delta\mu^j \geq 0$$

Rearranging the differenced Euler equation to isolate the change in the covariance yields:

$$\Delta\text{Cov}(m, R^j) = -\Delta\mu^j - \Delta q\mathbb{E}[R^j]$$

Since the expected asset return is strictly positive ($\mathbb{E}[R^j] > 0$), the empirical condition $\Delta q < 0$ implies that the risk-free price effect, $-\Delta q\mathbb{E}[R^j]$, is strictly positive.

Adding this strictly positive term to the weakly positive liquidity relief term guarantees that:

$$\Delta\text{Cov}(m, R^j) > -\Delta\mu^j \geq 0$$

□

2.1.2. *Precautionary Premia.*

Assumption 1. Decreasing Absolute Prudence (DAP): Absolute prudence $P(x) \equiv -\lambda_{xx}(x, p)/\lambda_x(x, p)$ is strictly decreasing in x.

Proposition 2. *Suppose the household holds at least one asset j in a positive quantity, and consider a positive perturbation of the income vector $\Delta y \geq 0$ (where $\Delta y_j > 0$ for at least one good j). If this perturbation strictly decreases the intertemporal price ($\Delta q < 0$), then the precautionary premium strictly decreases ($\Delta\psi < 0$).*

Proof. The intertemporal price can be written as:

$$q = \beta\frac{\mathbb{E}[\lambda(x', p')]}{\lambda(x, p)} = \beta\frac{\mathbb{E}[\lambda(\mathbb{E}[x'|p'], p')] + \psi}{\lambda(x, p)}$$

A strictly positive perturbation Δy unambiguously increases the current consumption budget ($\Delta x > 0$). Because $\lambda_x < 0$, current marginal utility falls, so the denominator of q strictly decreases. In isolation this would force q to rise. For $\Delta q < 0$ to hold, the numerator $\mathbb{E}[\lambda(x', p')]$ must fall by strictly more—which requires the household to allocate part of the perturbation to future resources. Since this saved portion is deterministic, the expected conditional future budget increases for every realization of future prices:

$$\Delta x > 0 \quad \text{and} \quad \Delta \mathbb{E}[x'|p'] > 0 \ \ \forall p'.$$

It remains to show that ψ strictly falls. Taking the partial derivative of the inner wedge of the precautionary premium with respect to the conditional expected resources yields:

$$\frac{\partial}{\partial \mathbb{E}[x'|p']} \left(\mathbb{E}[\lambda(x', p')|p'] - \lambda(\mathbb{E}[x'|p'], p')\right) = \mathbb{E}[\lambda_x(x', p')|p'] - \lambda_x(\mathbb{E}[x'|p'], p')$$

By Jensen's inequality, this expression is strictly negative if $\lambda_x(x, p)$ is strictly concave in x. The assumption of Decreasing Absolute Prudence (DAP) implies this strict concavity. Hence the precautionary premium is strictly decreasing in expected conditional resources for any given p'. Because $\Delta \mathbb{E}[x'|p'] > 0$ for all p', integrating over the outer expectation guarantees $\Delta\psi < 0$. □

Remark 1. The condition $\Delta q < 0$ does real work: it ensures that the contraction of the precautionary premium is large enough to dominate the standard consumption-smoothing effect on current marginal utility. Without it, we can't rule out a household that consumes the entire perturbation today.

2.1.3. *Income Perturbations and the Variation of the IMRS.* Above we assumed that preferences exhibited decreasing absolute prudence. Here we propose a strengthening of this assumption:

Assumption 2. Decreasing Relative Prudence (DRP): Relative prudence $xP(x) = -x\lambda_{xx}(x, p)/\lambda_x(x, p)$ is strictly decreasing in x. Note that DRP implies DAP.

Now, let

$$\Omega \equiv q - \beta \frac{\mathbb{E}[\lambda(\mathbb{E}[x'|p'], p')|p']}{\lambda(x, p)}.$$

This is the uninsurable variation in the IMRS, the wedge between the intertemporal price $q = \mathbb{E}m$ and the certainty-equivalent IMRS. Note that this variation is directly proportional to the precautionary premium ψ, as $\Omega = \frac{\beta\psi}{\lambda(x,p)}$.

Proposition 3. *Suppose the household holds at least one asset j in a positive quantity, and consider a positive perturbation of the income vector $\Delta y \geq 0$ (where $\Delta y_j > 0$ for at least one good j). If preferences exhibit decreasing relative prudence and either (a) the perturbation strictly decreases the intertemporal price ($\Delta q < 0$), or (b) the household is at an interior*

solution for some asset, then the uninsurable variation in the intertemporal marginal rate of substitution strictly decreases ($\Delta\Omega < 0$).

Proof. Since $\Omega = \beta\psi/\lambda(x,p)$, we need the ratio ψ/λ to fall. Under either condition, the perturbation strictly increases both current resources ($\Delta x > 0$) and the expected conditional future budget ($\Delta\mathbb{E}[x'|p'] > 0$) for all p': under (a) this follows from Proposition 2; under (b) it follows directly from standard consumption smoothing under strict concavity, since an unconstrained household must allocate part of any positive perturbation to future resources.

Because $\lambda_x < 0$, the increase in x strictly decreases $\lambda(x,p)$, which in isolation would push Ω up. Simultaneously, the increase in expected future resources strictly reduces ψ via Decreasing Absolute Prudence (which is implied by DRP). For the ratio to fall, ψ must decline proportionally faster than λ.

Decreasing Relative Prudence provides exactly this. In the local approximation developed below, $\Omega \propto P(x)A(x)$, the product of absolute prudence and absolute risk aversion. DRP implies that both $P(x)$ and $A(x)$ are decreasing in x, so their product declines strictly as resources grow. Globally, DRP ensures that the elasticity of the precautionary premium with respect to resources exceeds the elasticity of marginal utility, so the ratio ψ/λ falls and $\Delta\Omega < 0$. □

2.1.4. *The Intertemporal Price for Unconstrained Households.* The preceding results establish that DRP shrinks the precautionary wedge Ω for unconstrained households. Whether the intertemporal price q itself falls depends on a second channel: the certainty-equivalent component q_{CE}.

Proposition 4. *Suppose the household is at an interior solution for some asset and preferences exhibit DRP. Consider a positive perturbation $\Delta y > 0$. The change in the intertemporal price decomposes as:*

$$\Delta q = \Delta q_{CE} + \Delta\Omega$$

where $q_{CE} \equiv \beta\mathbb{E}_{p'}[\lambda(\mathbb{E}[x'|p'],p')]/\lambda(x,p)$ is the certainty-equivalent intertemporal price and $\Delta\Omega < 0$ by Proposition 3(b). If additionally $\Delta q_{CE} \leq 0$, then $\Delta q < 0$.

Proof. Immediate from the decomposition and $\Delta\Omega < 0$. □

The condition $\Delta q_{CE} \leq 0$ requires that expected future marginal utility falls proportionally at least as fast as current marginal utility. This can't be established from preference conditions alone—it depends on how the household allocates the perturbation between current consumption and saving. Two settings deliver it, both involving a further assumption on preferences:

Assumption 3. Decreasing Relative Risk Aversion (DRRA): Relative risk aversion $\gamma(x) \equiv xA(x) = -x\lambda_x(x,p)/\lambda(x,p)$ is strictly decreasing in x.

First, if preferences exhibit DRRA and the perturbation scales all resources proportionally, the savings decision drops out. Under CRRA (the boundary case), q_{CE} is scale-invariant: $\Delta q_{CE} = 0$ exactly. Under DRRA, $\Delta q_{CE} \leq 0$ whenever the expenditure-weighted average future relative risk aversion $\bar{\gamma}$ weakly exceeds the current level $\gamma(x)$—that is, whenever the household expects future resources to be no higher than current resources.

Second, and more relevant to our setting: if the perturbation is a positive-NPV loan disbursed at harvest with repayment due during the lean season, the loan structure itself disciplines the allocation. An unconstrained household invests the proceeds (e.g., in maize storage), generating returns that exceed the repayment obligation. Net future resources rise. Under DRRA, lean-season households—who expect future consumption to fall—have $\bar{\gamma} \geq \gamma(x)$, ensuring $\Delta q_{CE} \leq 0$.

Remark 2. The repayment obligation sharpens the diagnostic. A constrained household that consumes the loan proceeds faces reduced future resources ($\Delta \mathbb{E}[x'|p'] < 0$), since the repayment arrives with no offsetting investment return. This pushes $\Delta q_{CE} > 0$ unambiguously: the numerator of q_{CE} rises (higher future marginal utility from tighter future budgets) while the denominator falls (lower current marginal utility from the cash infusion). Combined with the weakened precautionary motive, this produces $\Delta q > 0$—a strong signature of binding liquidity constraints. The loan's repayment structure thus amplifies the separation between regimes: unconstrained households exhibit $\Delta q < 0$, constrained households exhibit $\Delta q > 0$.

2.1.5. *Corollary: Liquidity Constraint.*

Corollary 1. *Consider a positive perturbation satisfying the conditions of Proposition 4, so that $\Delta q < 0$ for any unconstrained household. If the perturbation strictly increases the observed intertemporal price ($\Delta q > 0$), the household cannot be at an interior solution for all assets. The perturbation must be relaxing a strictly binding liquidity constraint ($\Delta \mu^j < 0$ for some j).*

Proof. By Proposition 4, an unconstrained household with DRP and $\Delta q_{CE} \leq 0$ has $\Delta q < 0$. The observation $\Delta q > 0$ contradicts this, so the household cannot be at an interior solution for all assets. A positive income perturbation can only relax a borrowing constraint ($\Delta \mu^j \leq 0$), and the differenced Euler equation (Proposition 1) requires $\Delta \mu^j < 0$ for some j. □

2.1.6. *Discussion: Local Approximation and the $A \times P$ Decomposition.* While Proposition 3 establishes the reduction of uninsurable variation Ω globally without relying on local approximations, a second-order expansion provides economic intuition by linking the uninsurable variation directly to standard measures of risk sensitivity.

Evaluating the precautionary premium ψ locally around the expected conditional resources $\bar{x}' = \mathbb{E}[x'|p']$ yields:

$$\psi \approx \frac{1}{2}\text{Var}(x'|p')\lambda_{xx}(\bar{x}', p')$$

Substituting this into the definition of Ω and evaluating in a neighborhood where $x \approx \bar{x}'$:

$$\Omega \approx \frac{\beta}{2}\text{Var}(x'|p')\frac{\lambda_{xx}(x,p)}{\lambda(x,p)}$$

We can decompose this ratio by multiplying and dividing by the first derivative $-\lambda_x(x,p)$:

$$\frac{\lambda_{xx}}{\lambda} = \left(\frac{-\lambda_{xx}}{-\lambda_x}\right)\left(\frac{-\lambda_x}{\lambda}\right) \equiv P(x)A(x)$$

where $P(x)$ is Absolute Prudence and $A(x)$ is Absolute Risk Aversion. Thus the uninsurable variation in the IMRS is locally proportional to the product of prudence and risk aversion:

$$\Omega \propto P(x)A(x)$$

This decomposition isolates the economic mechanisms driving the shrinking of the wedge. A positive income perturbation $\Delta y > 0$ strictly increases current and expected resources. If preferences exhibit Decreasing Absolute Risk Aversion (DARA) alongside Decreasing Absolute Prudence (DAP)—both implied by the stronger DRP assumption—then both $A(x)$ and $P(x)$ strictly decline as resources grow. The simultaneous decline of both risk aversion and prudence forces their product to fall, providing a (local) logic for the decline in uninsurable variation Ω.

2.1.7. *Discussion: Identifying the Dominant Margin of Intertemporal Choice.* In principle q is observable—and we elicit it directly using the bond-price exercise described in Section 3. If we're willing to suppose that preferences satisfy DRP and DRRA, then the sign of Δq in response to an exogenous income shock becomes a diagnostic for the dominant friction governing households' intertemporal choices; we test this in Section 5.2. The sign cleanly separates two economic regimes.

In the first regime, a positive income shock strictly decreases the intertemporal price ($\Delta q < 0$). The household saves part of the transfer, raising expected future resources and shrinking the precautionary wedge Ω. Under the conditions of Proposition 4, this contraction dominates any change in the certainty-equivalent price. While liquidity constraints may exist, $\Delta q < 0$ guarantees they don't dictate the household's decisions at the margin; any relaxation of a shadow borrowing penalty is dominated by a reduction in the precautionary premium.

In the second regime, the intertemporal price strictly increases ($\Delta q > 0$). Risk resolution can't explain this. By Corollary 1, the household must be liquidity constrained, consuming the perturbation today rather than investing it. In our setting, the loan's repayment structure amplifies this signal: a constrained household that consumes the loan proceeds faces higher

future marginal utility (from the repayment obligation) and lower current marginal utility (from the cash infusion), pushing q up from both sides.

Note that this diagnostic identifies the dominant friction *at the time of the transfer*. A household that saves the shock as a buffer stock against credit constraints that might bind later still exhibits $\Delta q < 0$—because it is saving, not consuming today. Indeed, the logic extends: if the household is smoothly passing resources forward into the lean season, it isn't constrained then either, since a constrained household would have consumed the savings rather than carrying them further. What the test can't address is whether a differently timed transfer—say, a short-term loan disbursed midway through the lean season—would encounter binding constraints. But for the harvest-time PHLs we study, $\Delta q < 0$ says unambiguously that precautionary motives, not credit constraints, govern the intertemporal allocation.

3. Experimental Design and Data

This section describes three things: how Taimaka selected and randomized farmers, what the loans looked like, and how we measured welfare at high frequency throughout the season. The design has four key features. First, randomization was finely stratified within sites, following Athey and Imbens (2017). Second, there were two treatment arms—a cash loan and a similarly-valued in-kind maize loan—plus a control group. Third, we collected eight rounds of survey data at roughly two-month intervals, giving us unusually high-frequency measures of consumption expenditures, stocks, and welfare. Fourth, we supplemented these surveys with an experimental elicitation of farmers' intertemporal marginal rates of substitution, using a simple bond-purchase exercise.

3.1. **Sample Frame and Household Selection.**

Taimaka began by identifying the poorest 10 percent of locations in Gombe State using satellite-based poverty predictions following (Aiken et al. 2020)—about 50 sites. From these, 20 were randomly drawn for a rapid rural appraisal. Half were then chosen for household listing based on perceived need and accessibility, and six of those were ultimately selected for program implementation.

Within each of the six sites, Taimaka advertised the loan program and met with traditional leaders to obtain their assent. After two days, they returned for an informational session emphasizing the loan terms, the program's theory of change, and the group indemnity structure. Interested farmers were told to form groups of five.

Each group completed an enrollment form that included the 10-question Poverty Probability Index (PPI) for Nigeria, basic farming and demographic information, and photographs. Taimaka then ranked groups by ability to repay and need for the loan.

Traditional leaders verified that group members were local residents, genuine farmers, and creditworthy. If a single member was deemed unqualified, the group could replace that individual; if more than two were disqualified, the group was struck. Loan officers then

went down the updated list, visiting three households in each group to verify application information. Any material misrepresentation led to the group being dropped. Groups that survived this process were enrolled and eligible for the sample.

The ordered list was then partitioned into strata of six adjacently ranked groups. Within each stratum, one pair of groups was randomly assigned to the cash loan, another to the in-kind maize loan, and the remaining pair to the control group—a finely stratified design following Athey and Imbens (2017). This yielded a sample of 935 individuals from 187 groups.

3.2. **Treatment.** Both treatments offered a joint-liability loan of up to 50,000 Naira (roughly USD 100) to each of the five members in a group. The terms differed somewhat between the two arms.

For the cash loan, each farmer chose how many bags of maize to "commit" to Taimaka—that is, to have on hand (whether from own harvest, existing stocks, or purchase) and promise to store until the loan's due date of July 15, 2022.[8] The loan amount scaled with this commitment: 4 bags yielded 50,000 Naira (47,600 in cash plus 4 PICS bags at 600 Naira each), 3 bags yielded 40,000 Naira, and 2 bags yielded 30,000 Naira (see Appendix A for full details). All but three farmers chose to commit the maximum number of bags, receiving the largest available loan. In practice, then, the loan ceiling was binding for nearly everyone, making the PHL a take-it-or-leave-it transfer rather than a marginal borrowing decision. Note that this design narrows the distinction between the cash and in-kind arms considerably: both effectively channel liquidity toward maize storage, limiting the scope for cross-crop arbitrage. Indeed, the cash arm committed farmers to holding *more* maize (4 bags) than the maize arm provided (3 bags), so the total maize-linked position was similar across arms—and may partly explain the larger stock effects we observe for the cash arm. This is worth bearing in mind when interpreting the similarity of treatment effects across arms below. Loans were to be repaid with a 15% user fee in July 2022. Delivery took place between September 16 and October 10, 2021, and farmers were also required to make monthly repayments of at least 3,000 Naira starting in December 2021. Table 1 summarizes the present value of each arm's cash flows as a function of the household's geometric-mean monthly risk-free rate $\bar{R}_i^f$. The present value of the cash loan is zero at an annualized rate of about 25%, and the maize loan at about 27%—close to the median "prime" rates for commercial agricultural projects in Nigeria at this time.[9] Any borrower whose discount rate exceeds ~25% receives a loan with positive net present value, and our IMRS elicitations confirm that these farmers' shadow rates far exceed that threshold, implying a substantial subsidy even before accounting for administrative costs and default risk.

[8]Though a promise was made, Taimaka made no particular effort to monitor it or to ensure it was kept.
[9]See `https://cbn.gov.ng/Documents/depositandlending.asp`

For the in-kind maize loan, each farmer received a flat allocation of three bags of maize plus 3 PICS bags. Repayment was also due in July 2022 with a 15% user fee, calculated on the price Taimaka paid for the maize plus transport and the cost of the PICS bags—in total, 62,000–63,000 Naira depending on the group's distance from market (see Appendix A). These loans weren't disbursed until November 15–29, 2021, with minimum monthly repayments beginning in January 2022. Otherwise, terms were identical to the cash arm.

The control group received no loan from Taimaka.

3.3. **Surveys.** We collected 8 rounds of household surveys at roughly two-month intervals between August 2021 and November 2022. Each round included modules on household composition, grain stocks, food acquisitions, other expenditures, and measures of seasonal hunger. Baseline and endline surveys also covered agricultural inputs and output from the previous season, plus assets including livestock.

3.3.1. *Food Acquisition and Stocks.* In each wave, we elicited information on recent acquisitions of 22 goods that households in this setting consume. Some of these—milk, sugar—are fairly clearly for near-term consumption. Others may be held as a combination of present consumption and investment. In particular, households often hold positive stocks of maize, millet, beans, guinea corn (sorghum), and less frequently rice, Bambara nut, and groundnut; a few also reported holding cassava.

We asked about these stocks in each wave, but the instrument evolved over time. At baseline, we asked for the amount of each crop stored. In waves 1–3, we also asked how much of each crop households had harvested, purchased in bulk, and sold or given away since the previous visit, in addition to current stocks. After wave 3, enumerators told us that households found questions about their stocks sensitive—so for waves 4–6 we instead asked how much of their stocks they'd consumed since the previous visit, giving us an account of grain flows that lets us impute stocks. At endline, we returned to asking directly about stocks.[10]

3.3.2. *Bond Price Elicitation.* We also elicited each household's intertemporal price q—its willingness to pay for future purchasing power—using a simple bond-purchase exercise. After each survey wave t, enumerators asked whether the individual would invest 500 Naira (roughly one dollar) at the time of the following wave $(t+1)$ to receive $500R$ Naira in wave $t+2$. We elicited willingness to invest for values of R between 0.9 and 2, in increments of 0.1—so the lowest return corresponded to an interest rate of -10%, and the highest to

[10]Unsurprisingly, households' reported stocks and flows don't always balance in an accounting sense. We describe the resulting imputations in the Appendix.

TABLE 1. Present value of loan cash flows by treatment arm

	Formula	Cash	Maize
Transfer received[a]	$v/\bar{R}_i^{\tau}$	$50{,}000/\bar{R}_i$	$50{,}000/\bar{R}_i^2$
PICS bags	$n_b \cdot 600/\bar{R}_i^{\tau}$	$2{,}400/\bar{R}_i$	$1{,}800/\bar{R}_i^2$
Transport	$\delta/\bar{R}_i^{\tau}$	—	$\approx 2{,}000/\bar{R}_i^2$
Monthly payments[b]	$-p\sum_{t=\tau+k}^{\tau+k+n-1}\bar{R}_i^{-t}$	$-3{,}000\sum_{t=3}^{9}\bar{R}_i^{-t}$	$-3{,}000\sum_{t=4}^{9}\bar{R}_i^{-t}$
Final settlement[c]	$-(1.15V-np)/\bar{R}_i^{T}$	$-(36{,}500)/\bar{R}_i^{10}$	$-(\approx 43{,}870)/\bar{R}_i^{10}$
Net present value	**Sum**		
Breakeven: $\bar{R}_i^{12} \approx 1.25$ (monthly $\bar{R}_i \approx 1.019$)			
Net PV		≈ 0	≈ 0
Elicited p10: $\bar{R}_i^{12} \approx 2.53$ (monthly $\bar{R}_i \approx 1.080$)			
Net PV		$\approx 17{,}400$	$\approx 17{,}400$
Elicited median: $\bar{R}_i^{12} \approx 5.99$ (monthly $\bar{R}_i \approx 1.161$)			
Net PV		$\approx 30{,}100$	$\approx 31{,}000$
Elicited p90: $\bar{R}_i^{12} \approx 15.39$ (monthly $\bar{R}_i \approx 1.256$)			
Net PV		$\approx 37{,}900$	$\approx 39{,}800$

All values in Naira. $v = 50{,}000$ is the base value of the transfer (cash or maize); n_b is the number of PICS bags (4 cash, 3 maize); δ is transport cost (maize arm only); $p = 3{,}000$ is the monthly payment; n is the number of monthly payments before final settlement (7 cash, 6 maize); $V = v + n_b \cdot 600 + \delta$ is the total value on which the 15% fee is assessed; $T = 10$ is the month of final settlement; τ is months from application to disbursement (1 cash, 2 maize); $k = 2$ is months from disbursement to first payment.

[a] Cash arm: 47,600 Naira in cash plus PICS bags. Maize arm: three bags of maize at market price (≈50,000 Naira).

[b] Minimum 3,000 Naira/month; December–June for cash (7 payments), January–June for maize (6 payments). The July payment is subsumed in the final settlement row.

[c] Total owed = $1.15V$; final settlement = $1.15V - np$. For the maize arm, V includes transport costs that varied by site (62,000–63,000 Naira total obligation).

Notes: Present values computed using $\bar{R}_i^f$, the geometric mean of household i's monthly risk-free rates over the loan period. Under no-arbitrage, $(\bar{R}_i^f)^T$ equals the return on a T-month zero-coupon bond for household i. In practice these rates are not known in advance and vary with seasonal conditions; the table should be read as an ex post accounting identity, not an ex ante prediction. "Elicited" rows use quantiles of the per-farmer geometric mean of experimentally elicited two-month gross returns (Section 3.3.2), converted to monthly rates via $R^{30/d}$ where d is the inter-wave holding period in days. The first elicitation wave is excluded (learning period), as are responses at the menu cap (right-censored) and those implying $R \leq 1$. Even at the 10th percentile of elicited rates, the loan has substantial positive net present value, confirming that it functions as a transfer for virtually all participants.

100%.[11] An app then randomly selected one of the offered interest rates. If the individual

[11]Starting in Wave 3, we added a return of 3 (200% over about six weeks), reasoning that refusing this option would likely indicate misunderstanding or refusal to participate. Exactly one participant took this option and no other, suggesting that null responses don't contain meaningful information about the shadow rate.

had agreed to the selected rate, they were required to bring 500–2,500 Naira at the next survey wave; the enumerator collected the money and repaid it with interest one wave later.[12] If the individual hadn't agreed to the drawn rate, no deal was implemented. Households that reneged on an accepted deal were barred from future payouts. The minimum accepted return $\underline{R}$ identifies the bond price $q = 500/(500\underline{R}) = 1/\underline{R}$ in each survey round.

Implementing this exercise in the field wasn't straightforward. Many individuals initially refused to participate, citing religious concerns that the exercise was too similar to gambling or usury. We asked local imams in each subsite to explain the exercise to participants and confirm it was permissible. After a refusal rate of 33% in the first round, refusal rates fell to 14–24% in subsequent rounds. We also took care to distinguish between refusal to participate and choosing not to accept any offered interest rate. Comprehension and trust were additional concerns, though both appeared to improve over time. In Wave 1, 38% of responses exhibited inconsistencies—through misunderstanding or error, respondents indicated they would prefer a smaller return on their savings to a larger one—while in each subsequent round fewer than 3% of responses were inconsistent, save for April 2022 when inconsistencies occurred in 7% of cases.[13] Payouts from deals implemented in Wave 1 were made in Wave 3, which may have helped establish credibility.

Interpreting observations in which the individual didn't agree to any interest rate—or refused to participate at all—is also tricky. In our preferred specifications, we include only observations in which an interest rate was selected, because we believe null choices are more likely to reflect reluctance to participate for religious reasons, lack of trust, or misunderstanding rather than true intertemporal preferences.

3.3.3. *Other Variables.* We also asked whether households had any large non-food purchases of at least 10,000 Naira (\(20); whether they engaged in another business and, if so, their expenses and revenues; and how much money they borrowed and repaid. We also prespecified the Reduced Coping Strategies Index (RCSI) (Maxwell, Vaitla, and Coates 2014), which asks households about the number of days in a month they restricted food consumption because of lack of resources, as a measure of seasonal hunger.

3.4. **Other Data Sources.** We use two sources of price data. Taimaka collected prices of staple commodities from local markets in Gombe State on a roughly monthly basis. We also use weekly price data from the Famine Early Warning System Network (FEWS-Net) for Nigeria, which covers major regional markets including Gombe Town. In addition, we draw on Taimaka's administrative records on loan applications (including group scoring),

[12]Choices at endline were made with the knowledge that enumerators would return after one month to implement the choices as usual. However, households were unexpectedly told during the subsequent visit that payouts would be made instantaneously for operational reasons. The choices made at endline remain valid elicitations of q, assuming households expected the usual procedure.

[13]In our preferred specifications, we drop inconsistent observations.

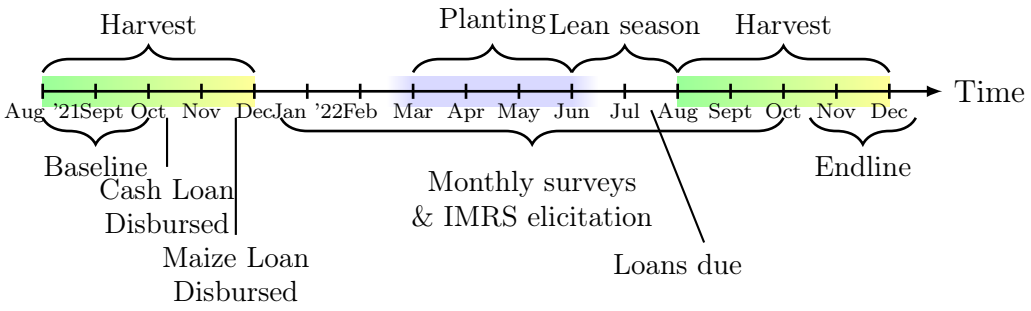


FIGURE 1. Intervention and Survey Timeline

Notes: Timeline of loan disbursement, repayment schedule, and survey waves.

disbursement, and repayment. In supplementary analysis , we also compare our sample to representative samples of rural households from Northeastern Nigeria in the Nigeria General Household Survey 2018-19 wave and the National Longitudinal Phone Survey, which covers our 2021-22 study period.

3.5. **Attrition.** We faced attrition in two forms: individual households missing from a given wave, and entire groups permanently dropping out. 7 groups (35 individuals) assigned to the cash loan dropped out after treatment assignment, either after being disqualified for misrepresenting their applications or after voluntarily withdrawing. In addition, 5 maize groups and 2 control groups withdrew consent to be surveyed after Round 1. The finely stratified design helps here, since we can drop the entire stratum of any group that attritted (Bruhn and McKenzie 2009). We present results using this approach in Appendix C.1.1; our main estimates don't change substantially.

Conditional on the group remaining in the program, we achieve a response rate of about 90% for the following rounds, which is evenly balanced across the three arms (Table A8). Moreover, as shown in Table A9, baseline covariates are not predictive of subsequent individual attrition — in a regression of probability of being surveyed on 21 covariates, only one is statistically significant. We can also adapt the canonical Lee (2009) procedure to our application of the finely stratified Athey and Imbens (2017) estimator to produce bounds on the treatment effects under extreme cases of selective attrition. We describe this procedure further in Appendix C.3, which produces relatively tight bounds for most specifications, on the order of the 95% confidence intervals from the main specifications. One caveat: due to time constraints on the loan program's rollout, we were only able to survey slightly over 60% of the sample at baseline (Round 0), though this doesn't affect our main estimates.

4. Descriptive Statistics

Before turning to the experimental results, we document the economic environment in which the experiment took place. Three features of this environment matter for interpreting what follows: the volatility and poor spatial integration of maize prices, the seasonal pattern in elicited shadow rates, and the seasonal pattern in consumption-based welfare.We also discuss the composition of our sample of Taimaka applicants relative to representative households in the rural Northeast.

4.1. **Prices.** Using over 10 years of monthly FEWS Network price data, together with prices collected from markets in Gombe State during 2021–22, Figure 2 shows that the magnitude of seasonal price increases is highly variable. On average prices increase by 67% from floor to peak, in the median year, the increase is 25%. The gap between mean and median is driven by extreme price increases of 260% and 188% in 2016 and 2020, respectively. On the other hand, price increases were 15% or below in three of 11 years. So intertemporal arbitrage typically yields moderate positive *expected* returns—but the distribution has fat tails in both directions. As Cardell and Michelson (2022) argue, even very modest risk aversion could rationalize the supposed failure to arbitrage.

The FEWS data gives some useful history, but what happened during our study period? Maize prices in the 2021–2022 season increased by about 20% (from 16,500 to 22,000 Naira per 100kg bag) between November and January, then stayed between 18,000 and 20,000 Naira throughout the loan period. By weight, 73% of the maize in the sample was sold after the January peak, meaning households would have obtained a return on stored maize of 11–23% depending on the week they sold, not factoring in depreciation. In practice, by weight, 31% of maize sold in the sample was sold below 18,000 Naira per bag and 83% was sold for below 20,000 Naira per bag. Few households made significant profits from arbitrage; many made negative profits net of Taimaka's interest.

The price data also exhibit limited spatial correlation, whether we consider within Gombe, across markets in Nigeria, or between Nigeria and world markets. This indicates poor market integration and suggests some scope for spatial arbitrage. Using FEWS data at weekly frequencies, we can reject the null of no spatial correlation in maize prices (Skillings-Mack test; $p = 0.039$): when prices go up in one location they're very likely to go up in another (see Figures A12 and A13 for national and local price co-movement, and Figures A1 and A2 for the underlying price series). But the magnitude of this correlation isn't particularly large. The R^2 from regressing the time series of local prices on that of international prices is .26, with higher correlation in more recent years (see Figure A3). We provide additional results showing the limited spatial correlations in price movements in Appendix˜D.1.

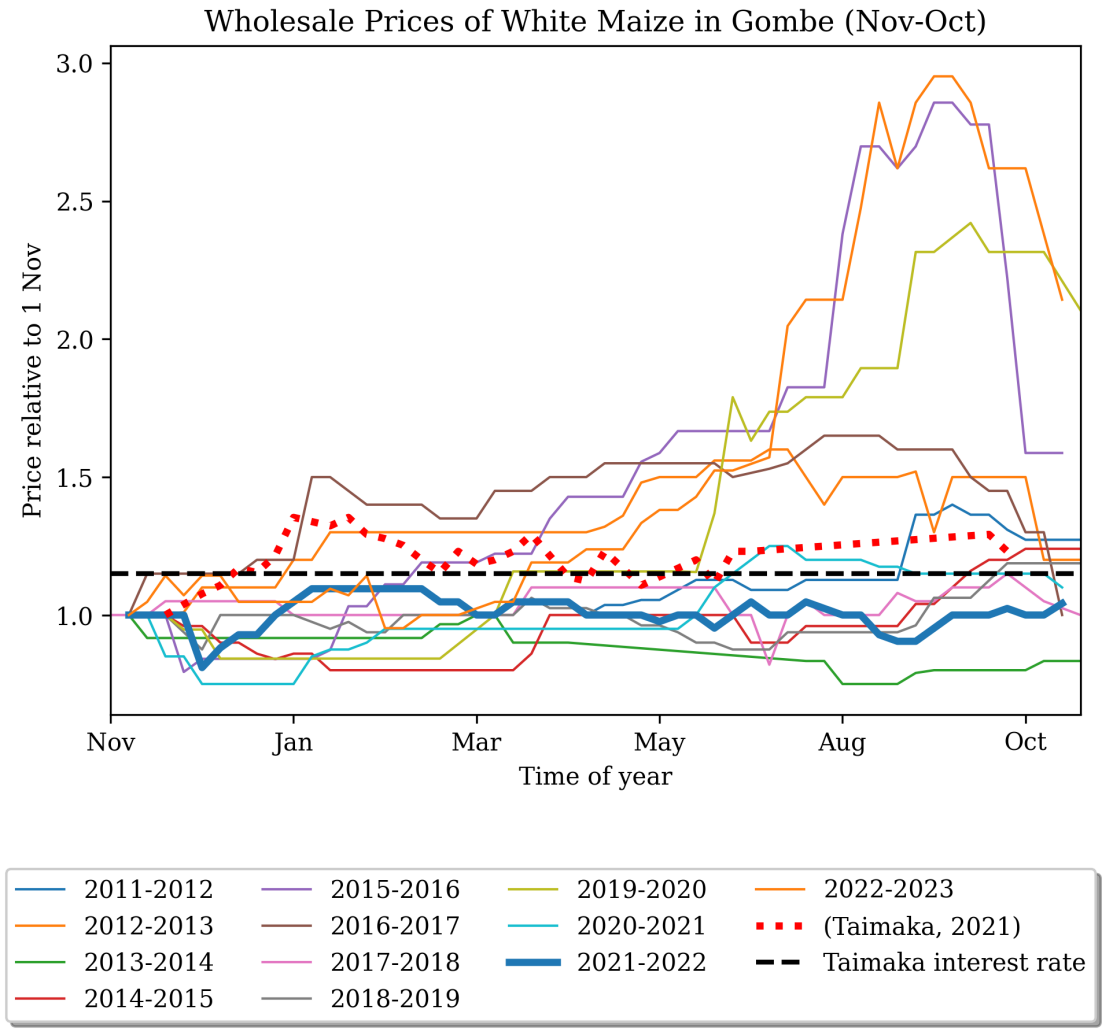


FIGURE 2. Yearly maize price increases in Gombe

Altogether, this paints a picture of volatile maize prices and poorly integrated markets—suggesting scope for spatial arbitrage in addition to, or perhaps instead of, intertemporal arbitrage.

To quantify the predictable seasonal component, we estimate a semi-linear model using the full FEWS panel (October 2004–October 2023, 16 markets, two maize varieties):

$$\log p_y^j(d) = a^j(d) + b_y^j + c_m + e_y^j(d),$$

where d is the day of the year, j indexes product (white and yellow maize), y indexes year, and m indexes market. The year-product fixed effects b_y^j absorb level shifts (e.g., the 2016 spike), and the market effects c_m absorb cross-sectional price differences. The seasonal component $a^j(d)$ is estimated using a Nadaraya-Watson kernel regression with a Gaussian kernel and a bandwidth of one week.

Figure 3 plots the estimated seasonal component. Maize prices rise steadily from January through August, peaking at roughly 30 log points above their post-harvest floor—well outside the 99% acceptance region constructed by permuting prices across dates. The predictable seasonal increase is thus statistically unambiguous. But Figure 4 overlays the actual year-by-year Gombe price paths: while the *average* pattern is clear, individual years deviate enormously. Some years see price increases exceeding 100 log points; others see prices fall. This is the fundamental tension: the expected return to storage is positive, but the variance is very large.

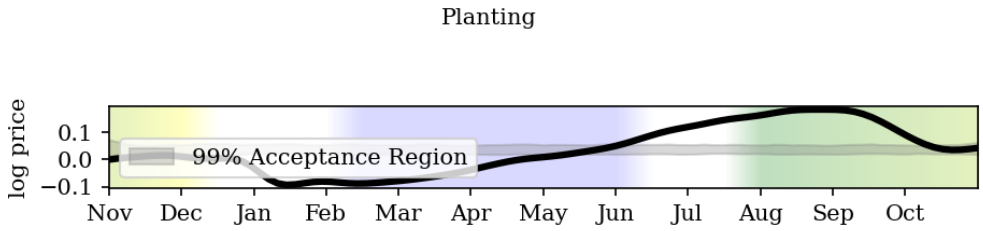


FIGURE 3. Mean seasonal component of maize prices (kernel regression), with 99% acceptance region under the null of no seasonality.

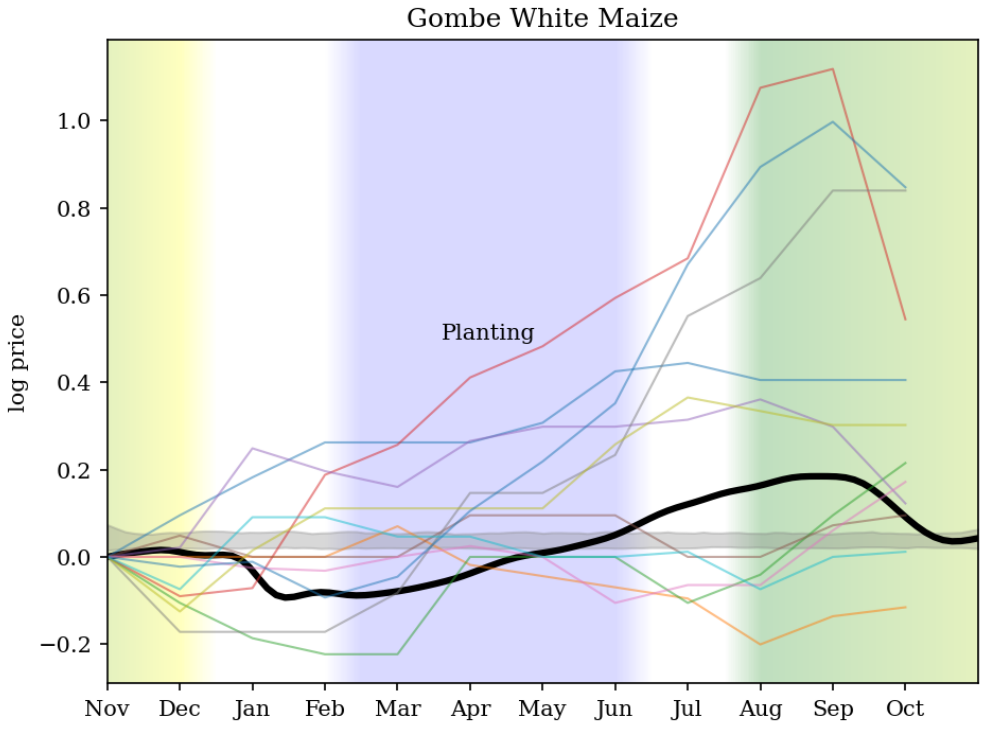


FIGURE 4. Year-by-year realizations of white maize prices in Gombe (thin lines) overlaid on the estimated seasonal component (thick line) and acceptance region.

4.2. **Elicited Shadow Rate.** Recall our elicitation of the smallest interest rate at which individual i was willing to save. Let $\underline{R}_{it}$ denote the smallest rate elicited at time t, and let M_{it} denote individual i's IMRS between dates t and $t-1$. The realization of M_{it} won't be known before time t, but if i regards the savings opportunity as risk-free, then our elicitation implies $\mathbb{E}_t M_{it+2}\underline{R}_{it} = 1$. (At time t we're asking about the minimum return at which i would commit to purchase a bond at $t+1$ that matures at $t+2$.)

Of course, the farmer may not regard this as risk-free. The obvious risk is that Taimaka might fail to redeem the bond. Suppose the probability of redemption is π, and that otherwise the return is zero. Then $\pi\mathbb{E}_t M_{it+2} = 1/\underline{R}_{it}$, and we can use the right-hand side to identify the left.

Interpreting observations where the individual didn't agree to any interest rate—or refused to participate entirely—is tricky. In our preferred specifications we include only observations where an interest rate was selected, because null choices more likely reflect reluctance to participate for religious reasons, lack of trust, or misunderstanding than true intertemporal preferences. We also show results coding null choices, and refusals, as the maximum available rate.

Conversely, we may not believe that individuals would choose a -10% or 0% interest rate, assuming they have a discount factor less than 1.[14] We therefore also present results dropping observations that select a weakly negative rate.

The elicited shadow rates exhibit an increasing pattern over the season, as shown in Figure 5. During the lean season—which becomes increasingly acute during the March–September growing season—households demand increasingly high returns on bonds (i.e., q falls). The median household goes from requiring a 10% return in January to a 30% return from June onward. This is consistent with resources becoming increasingly scarce. Results from the first few waves are somewhat sensitive to topcoding and to dropping selections of the 0% rate.

Figure 5 maps these elicitations onto calendar dates using a Nadaraya-Watson kernel regression (solid line), with a 99% permutation-based acceptance region for the null of no seasonal pattern. Because survey waves were unevenly spaced (43–69 days apart), the gross return R_t elicited in each wave corresponds to holding periods of different lengths. To facilitate comparison, the dotted line shows the equivalent monthly (30-day) yield, $R_t^{30/d}$, where d is the number of days between the relevant waves.[15]

[14] In theory, individuals might accept negative rates due to intrahousehold taxation or present focus, e.g., (Dupas and Robinson 2013), though misunderstanding the exercise seems more likely in our case.

[15] Farmers did not know the exact date of the next survey visit when making their elicitation choices, so the horizon d was uncertain ex ante. We use realized inter-wave gaps to compute monthly yields, but the raw R_t is the quantity directly elicited.

The pattern mirrors the agricultural calendar. R_t is lowest around harvest (October–November) and rises sharply through the planting season, peaking around June–July at the height of the lean season.[16] The monthly yield ranges from roughly 3% near harvest to 13% during the lean season peak—implying annualized shadow rates well in excess of the roughly 25% at which the Taimaka loan breaks even (Table 1), consistent with universal take-up. We return to these shadow rates in Section 5, where we show that the loan shifts treated households' intertemporal price.

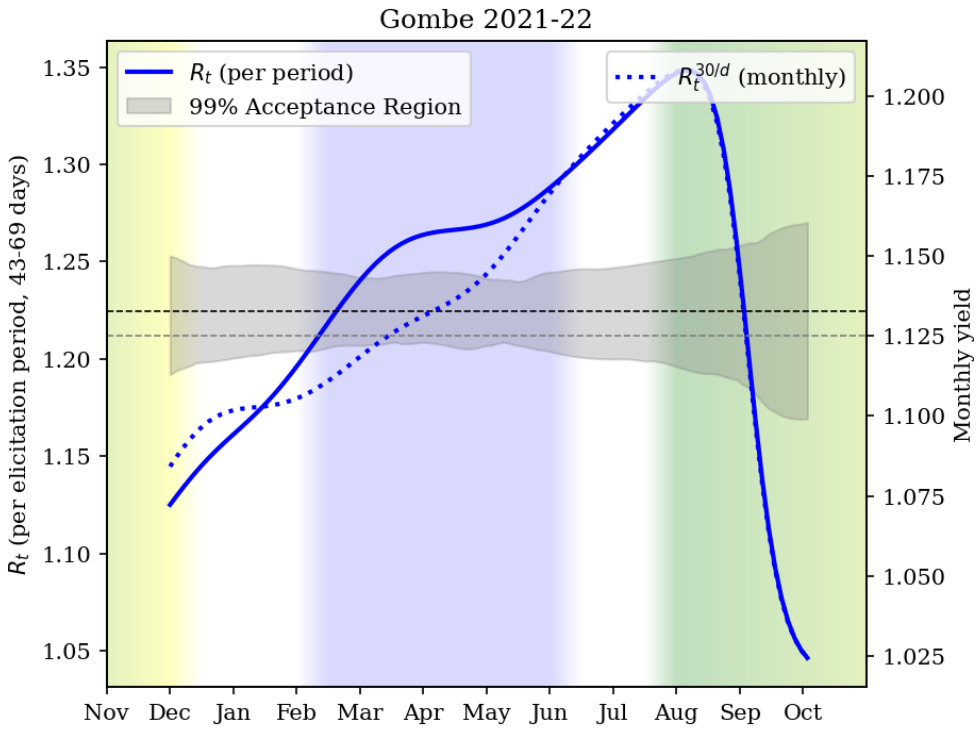


FIGURE 5. Kernel regression of elicited R_t on calendar date. Solid line: raw gross return per elicitation period (43–69 days); dotted line: equivalent monthly (30-day) yield. The gray band shows the 99% acceptance region under the null of no seasonal variation; the dashed line marks the sample mean. Shading indicates planting (blue) and harvest/drying (green–yellow) seasons.

4.3. **Consumption-based λ.** Our second measure of seasonal welfare is the marginal utility of expenditure estimated from survey data. We estimate a Constant Frisch Elasticity (CFE) demand system following Ligon (2020). CFE nests demand systems commonly used in the literature and flexibly accommodates non-homotheticity and a wide range of substitution elasticities. With panel data it permits estimation of an Index Marginal Utility

[16]April 2022 coincided with Ramadan (April 2–May 1) in this predominantly Muslim sample. Several figures show unusual patterns in April that we attribute to the agricultural calendar, but Ramadan-related changes in food expenditure, meal timing, and social transfers could contribute. We have not attempted to disentangle the two.

of Expenditure, which under certain cardinalizations of utility can be interpreted as $\log \lambda$. Since marginal utility is decreasing in wealth, higher $\log \lambda$ indicates greater scarcity; we use this as our preferred measure of seasonal hardship.

As shown in Figure 6, our consumption-based scarcity measure rises over the course of the season. Values are particularly high during the lean season (March–August) before falling at endline, after harvests began to come in.

Figure 6 presents a kernel-smoothed version of $\log \lambda$ plotted against calendar time, using food-acquisition dates for each household rather than discrete wave indicators. The seasonal component is estimated via Nadaraya-Watson kernel regression with a Gaussian kernel and a bandwidth of one month. The continuous curve reveals a pattern strikingly aligned with the agricultural calendar: $\log \lambda$ rises from November through May—corresponding to the depletion of grain stocks and the planting season—before declining as the fresh harvest becomes available around August–September. The amplitude of the seasonal swing (roughly 0.25 log points) implies that marginal utility at the lean-season peak is about 28% higher than at the post-harvest nadir.

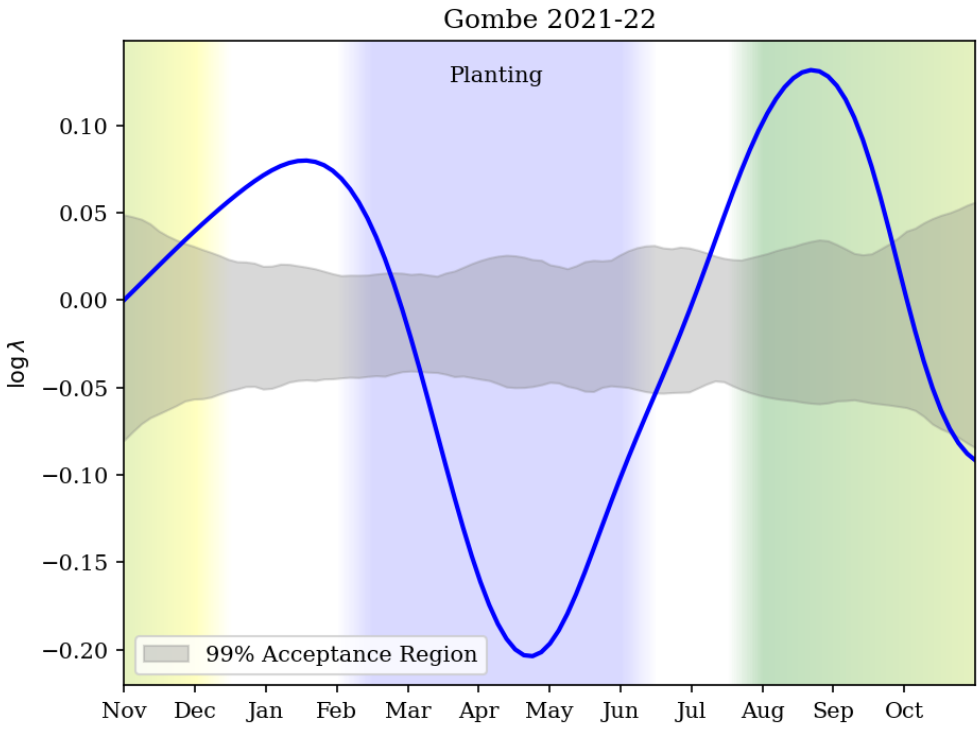


FIGURE 6. Kernel regression of $\log \lambda$ on day of year, Gombe 2021–22. Shading indicates planting (blue), fresh harvest (green), and grain availability (yellow) seasons.

4.4. **Relation Between Marginal Utilities and Shadow Rates.** The elicited shadow rate $1/q$ and the consumption-based $\log \lambda$ aren't directly comparable,[17] but we'd still expect them to be systematically correlated. Figure 7 presents box plots of $\log \lambda$ at each of the possible interest rates from the bond-price elicitation. These are raw correlations pooled across all survey waves, but households willing to accept lower returns (higher q) do tend to have higher estimated $\log \lambda$ (greater scarcity). Given their correlation and similar seasonal patterns, it's reasonable to believe both measures are capturing something similar.

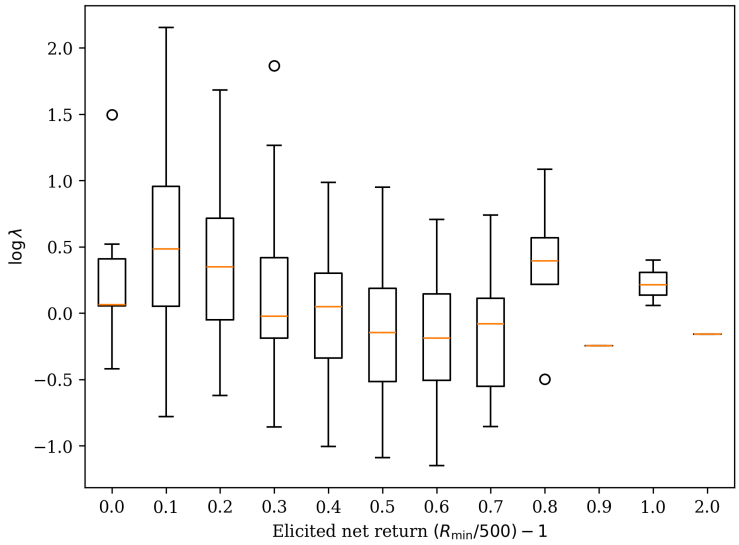


FIGURE 7. Estimated $\log \lambda$ at elicited shadow rates $1/q$. Box plots show median, interquartile range, and whiskers at $1.5 \times$ IQR.

4.5. **Socioeconomic Narrative and Sample Representativeness.** By design, our sample consists of households living in areas targeted by Taimaka's program who were deemed eligible and chose to apply for its PHLs. While randomizing among PHL applicants helps contribute to the internal validity of our study, how comparable are these households to others in the rural Northeast? In Appendix D.2 we present a detailed narrative of households' socioeconomic conditions using our survey data supplemented by comparisons to representative samples of rural households from Northeastern Nigeria in the Nigeria General Household Survey (GHS) 2018-19 and National Longitudinal Phone Survey (NLPS) 2021-24. These descriptive statistics further illustrate the seasonal volatility households in this region, and especially in our sample, face amid limited access to formal credit. Households in our sample tend to be more reliant on agriculture, earn less income from off-farm sources and are less educated than the representative rural household.

[17]These measures are hard to validate quantitatively against each other, since the bond price involves expectations over future marginal utility, which precludes linear comparisons.

5. Experimental Results

The overall takeaway from the experiment is simple: precautionary savings, not credit constraints, govern the intertemporal allocation of these households. The theory delivers a clean diagnostic (the sign of Δq in response to the loan; Section 2), and the data speak clearly. The intertemporal price never rises in response to the loan; over the full follow-up $\Delta \log q < 0$ ($t = -2.4$), ruling out the constrained-household prediction ($\Delta q > 0$) and confirming the precautionary one. The loan sharply lowers the marginal utility of expenditure around delivery ($\Delta \log \lambda \approx -0.22$, $t > 4$): households are genuinely better off, but they smooth the shock intertemporally rather than consuming it today. Season-average effects on consumption expenditure are small, though the MUE—a more sensitive welfare measure—reveals significant effects that expenditures miss. Meanwhile, the cash loan leads households to store more grain and invest more, consistent with portfolio rebalancing driven by precautionary motives rather than the relaxation of a binding borrowing limit.

Two features of the results deserve emphasis at the outset. First, seasonal variation in our outcome variables is large, much larger than the treatment effects we estimate. The wave-by-wave figures that follow make this visually obvious: the time path of each outcome moves far more than the treatment–control difference at any point. This foreshadows our finding that the lean season is fundamentally an aggregate phenomenon. Second, where either arm produces a significant effect relative to control, the two treatments almost always move together—we can't reject equality of their effects. Access to credit, not the specific encouragement to store maize, drives the results.

As described in Section 3, farmer groups in each of the six study sites were partitioned into strata of six adjacently ranked groups, with two offered the cash loan, two the maize loan, and two serving as controls. Where the number of groups in a site wasn't divisible by 6, treatment was randomly assigned to the remaining groups; for the analysis we merge these with the preceding stratum in their respective sites. In our main specifications we use the Neyman estimator suggested by Athey and Imbens (2017). While simple, this estimator can offer precision gains over OLS with fixed effects when the number of clusters is small. In Appendix E.3.1, we show results from the regression-based approach, which are qualitatively similar but noisier.

Denote the treatment assignment of household i in stratum s as $D_{is} \in \{C, M, 0\}$, corresponding to Cash, Maize, and Control. Within each stratum we have a fully randomized experiment and obtain the cash treatment effect on outcome Y_{ist} (which may be observed at different periods t) as $\hat{\tau}_s^C = \bar{Y}_s^C - \bar{Y}_s^0$ and the maize treatment effect as $\hat{\tau}_s^M = \bar{Y}_s^M - \bar{Y}_s^0$. The variances are

$$\hat{\mathbb{V}}(\hat{\tau}_s^C) = \frac{s_{C,s}^2}{N_{C,s}} + \frac{s_{0,s}^2}{N_{0,s}} \ , \hat{\mathbb{V}}(\hat{\tau}_s^M) = \frac{s_{M,s}^2}{N_{M,s}} + \frac{s_{0,s}^2}{N_{0,s}},$$

where $s^2_{D,s}$ is the sample variance of Y_{ist} for households in arm D. These stratum-level estimates are aggregated as

$$\hat{\tau}^C = \sum_{s=1}^{S} \hat{\tau}^C_s \frac{N_s}{N} \ , \hat{\mathbb{V}}(\hat{\tau}^C) = \sum_{s=1}^{S} \hat{\mathbb{V}}(\hat{\tau}^C_s) \left(\frac{N_s}{N}\right)^2$$

and

$$\hat{\tau}^M = \sum_{s=1}^{S} \hat{\tau}^M_s \frac{N_s}{N} \ , \hat{\mathbb{V}}(\hat{\tau}^M) = \sum_{s=1}^{S} \hat{\mathbb{V}}(\hat{\tau}^M_s) \left(\frac{N_s}{N}\right)^2$$

Note that $N_{D,s} = 2$ and $N_s = 6$ except in the last stratum of each subsite when the number of groups isn't divisible by 6. For outcome variables measured in each survey round, we estimate separate τs for each round as well as summing or averaging over the season, as pre-specified.

Table 5 summarizes treatment effects across the main pre-specified outcomes. The first two columns report the cash and maize treatment effects relative to control; the third tests whether the two treatments differ; the fourth pools both arms. In the subsections that follow we present wave-by-wave figures that illuminate *when* effects arise.

5.1. **The Loan as an Investment.** How attractive is the Taimaka loan as an investment? Recall from Figure 5 that the seasonal shadow rate—elicited by offering households a menu of returns on a two-month saving instrument—ranges from about 3% to 13% per month over the agricultural calendar. The loan's breakeven rate is roughly 1.9% per month (about 25% annualized; see Table 1), well below these shadow rates. Universal take-up is therefore unsurprising: the loan represents a positive-NPV transfer for virtually every recipient. The question is what this wealth shock does to the household's intertemporal allocation—and in particular to the intertemporal price q.

5.2. **Effects on the Intertemporal Price.** What does the loan do to the household's willingness to trade resources across time? The bond-price exercise described in Section 3 elicits a forward-looking object: the intertemporal price $q_t = \beta\, \mathbb{E}_t[\lambda_{t+1}/\lambda_t]$, the household's expected marginal rate of substitution between current and future expenditure. This isn't the *realized* IMRS, which depends on states of the world the household hasn't yet observed, but the price it would pay today for a unit of future purchasing power. The theory in Section 2 delivers a sharp sign diagnostic: under DRP and DRRA, $\Delta q < 0$ for an unconstrained household receiving a positive income shock (Proposition 4), while $\Delta q > 0$ identifies binding credit constraints (Corollary 1).

An accident of timing lets us test this. Cash loans were delivered September 16–October 10, 2021, before Wave 1 elicitations began on October 5, so we have no pre-receipt bond prices for cash-arm households. Maize loans didn't arrive until November 15–28, *after* the

TABLE 2. Summary of treatment effects

	Cash − Control	Maize − Control	Cash − Maize	Any Loan	Control mean
Panel A: Grain Stocks and Flows (000s Naira)					
Grain stocks (cum.)	82.779***	4.260	78.518***	24.929	723.737
	(23.966)	(25.018)	(27.588)	(20.044)	
Grain sales	20.847	21.910	−1.064	21.598	306.740
	(30.207)	(28.139)	(34.273)	(23.460)	
Harvested grain	167.214**	44.135	123.079	92.746	841.500
	(76.078)	(75.491)	(86.367)	(61.977)	
Own-stock consumption	131.759**	61.458	70.301	88.372**	674.250
	(55.981)	(56.657)	(66.100)	(44.954)	
Purchases for storage	17.649	−28.391*	46.040***	−11.974	79.940
	(15.772)	(14.689)	(15.136)	(13.181)	
Panel B: Consumption and Welfare					
Log food expenditure	−0.037	0.011	−0.048	−0.019	7.865
	(0.043)	(0.041)	(0.042)	(0.036)	
Log non-storable exp.	0.010	0.066**	−0.056*	0.032	6.845
	(0.033)	(0.031)	(0.033)	(0.026)	
MUE ($\log \lambda$)	−0.005	0.004	−0.009	−0.009	0.033
	(0.021)	(0.020)	(0.022)	(0.017)	
Shadow rate $1/q$	−0.028***	−0.028***	0.000	−0.028***	1.283
	(0.011)	(0.010)	(0.011)	(0.009)	
Hunger index (−RCSI)	−0.091*	−0.124**	0.033	−0.115***	0.370
	(0.052)	(0.049)	(0.055)	(0.041)	
Panel C: Investment (000s Naira)					
Business expenditure	44.021***	2.440	41.581***	20.640**	75.557
	(14.563)	(8.942)	(15.074)	(9.319)	
Business revenue	0.011**	0.002	0.009	0.005*	0.018
	(0.005)	(0.003)	(0.005)	(0.003)	
Durable expenditure	1.892	−0.056	1.948	0.325	38.791
	(6.260)	(5.149)	(5.964)	(4.878)	

Results obtained using the estimator proposed by Athey and Imbens (2017) for tuple-wise randomized experiments. Panel A reports total grain flows summed across all crops and survey waves (except grain stocks, which are averaged across waves), in 000s of Naira. Panel B reports averages across post-baseline waves. Panel C reports totals summed across waves, in 000s of Naira. $^{***}p < .01$,$^{**}p < .05$,$^{*}p < 0.10$.

first round of elicitations. This gives us a within-arm pre/post comparison for the maize arm and, after removing the seasonal trend common to all arms, a difference-in-differences.

Table 3 reports the results using two post-delivery windows: "adjacent" (only the first post-delivery elicitation per household, the tightest comparison) and "all" (every post-delivery elicitation, capturing the cumulative effect). We report both Δq and $\Delta \log q$; the latter is

immune to the Jensen's-inequality ambiguity that arises when comparing means of q with means of $1/q$, and in practice the two agree in sign throughout. The adjacent-wave DiD for $\log q$ is null ($+0.008$, $t = 0.5$): at receipt, the intertemporal price doesn't move relative to controls. Over the full follow-up the DiD turns negative and significant (-0.030, $t = -2.4$ on the log scale). Meanwhile, the DiD for $\log \lambda$ is large and negative at both horizons: -0.220 ($t = -4.1$) adjacent and -0.312 ($t = -6.3$) pooled. The loan durably reduced the marginal utility of expenditure relative to controls. Among 68 maize-arm households with both a pre- and post-delivery elicitation, a Wilcoxon signed-rank test confirms the within-household decline in q ($p = 0.0001$).[18]

TABLE 3. Intertemporal Price and Marginal Utility Around Maize Loan Delivery

	Δq (all)	$\Delta \log q$ (all)	$\Delta \log \lambda$ (all)	Δq (adj)	$\Delta \log q$ (adj)	$\Delta \log \lambda$ (ad
Maize	-0.0842^{***}	-0.1015^{***}	-0.0988^{**}	-0.0228^{**}	-0.0270^{**}	-0.2499^{**}
	(0.0086)	(0.0102)	(0.0422)	(0.0102)	(0.0124)	(0.0472)
Control	-0.0603^{***}	-0.0717^{***}	0.2129^{***}	-0.0276^{***}	-0.0346^{***}	-0.0302
	(0.0062)	(0.0070)	(0.0252)	(0.0065)	(0.0074)	(0.0255)
Cash	-0.0789^{***}	-0.0946^{***}	-0.1244^{***}	-0.0458^{***}	-0.0547^{***}	-0.0820^{**}
	(0.0077)	(0.0100)	(0.0319)	(0.0092)	(0.0119)	(0.0369)
DiD (maize−ctrl)	-0.0239^{**}	-0.0298^{**}	-0.3118^{***}	0.0048	0.0076	-0.2198^{**}
	(0.0106)	(0.0124)	(0.0492)	(0.0120)	(0.0144)	(0.0536)

Notes: Each cell reports the Athey–Imbens stratified pre→post change, stratified by cluster. "All" pools every post-delivery elicitation (waves 2–6); "adj" restricts to the first post-delivery elicitation per household. $\Delta \log q$ avoids the Jensen inequality ambiguity that arises when comparing means of q and $1/q$. $\Delta \log \lambda$ uses inverse-variance weights from the CFE demand system. The DiD row subtracts the control pre→post change from the maize pre→post change. Among $N = 68$ maize-arm households with both a pre- and post-delivery elicitation, a Wilcoxon signed-rank test rejects the null of no change ($W = 551.5$, $p = 0.0001$, mean $\Delta q = 0.0689$). $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

How should we read these two results together? The decomposition $q = q_{CE} + \Omega$ from the theory is helpful. At receipt, both current and expected future λ fall together; the household is smoothing the shock intertemporally, allocating part of the loan to future resources. This drives $\Delta q_{CE} \approx 0$, which is why q doesn't jump. The large drop in $\log \lambda$ confirms the level shift: the household is genuinely better off. Over subsequent months, the precautionary channel cumulates. Under DRP, each period of higher resources shrinks the precautionary wedge Ω; eventually Δq turns negative, exactly as Proposition 4 predicts.

[18]The three-way comparison—pre-notification, notified but pre-delivery, post-delivery—is suggestive but underpowered: the pre-notification cell has only two observations, making it uninformative. For what it's worth, notified-pre-delivery households show $q = 0.871$, *above* both the post-delivery mean (0.823) and the control (0.828)—the opposite of what a borrowing-constraint story would predict.

The absence of the constrained signature ($\Delta q > 0$) at *any* horizon is telling. A credit-constrained household would consume the loan proceeds today, lowering current λ_t but leaving future λ_{t+1} unchanged or higher, pushing q up from both sides. We don't see this. Instead, the pattern is consistent with precautionary motives dominating: the loan gives households a buffer, they save part of it, and the resulting reduction in risk exposure gradually lowers q.

5.3. **Effects on Stocks.** We now turn to the downstream effects on grain flows, which illuminate the portfolio rebalancing through which the IMRS reduction operates. Figure 8 shows the effects of the cash and maize treatments—cumulatively and in each survey wave—on the value of grain flows (sales, consumption from own stock, harvests, and purchases) summed across crops. Effects on the sum of each flow variable across the season appear to the left of the black bar in each figure; wave-by-wave coefficients appear to the right.[19] Total effects by crop are disaggregated in Appendix Tables B–B.

The headline result for stocks is striking. Total stocks increase by large amounts for the cash loan group relative to controls throughout the season (Panel A). Our preferred measurement is simply the cumulative sum of flows, though we also compute results with additional imputations to ensure accounting consistency (see Appendix B.2 for the full algorithm). The qualitative finding—that the cash arm stored more grain—is robust across specifications, but the quantitative magnitude is sensitive to how aggressively one reconciles reported flows with reported stocks. Under stricter accounting rules that enforce non-negative stocks at every period, the estimated treatment effects shrink considerably. This is likely connected to the field-drying issue discussed below: grain still in the field generates "harvest" flows that may be reported with timing errors, inflating the cumulative sum. The estimated treatment effect increases throughout the season, surpassing 100,000 Naira (twice the value of the loan) and becoming significant at the 10% level starting in June 2022—corresponding to the onset of the lean season. The maize treatment, on the other hand, tells a different story: point estimates are small and insignificant in all rounds, though they're also quite large during the lean season onset. Despite Taimaka priming households to use the loans for maize storage, we don't see increases in maize-specific storage for either arm, though we do see increases in maize sales for both arms early in the season. Instead, most of the results are driven by beans and millet. We do see an increase in maize sales for the in-kind group immediately

[19]As shown in Figure 1, the cash loans were disbursed in October 2021 and the maize loans in November 2021, after most surveys f or that wave were completed. In principle, loans were due by end of July 2022, but Taimaka allowed some flexibility in repayment. The endline survey took place in November 2022, during the bean and millet harvests and toward the beginning of the maize harvest. Since flow variables weren't elicited at baseline[20], baseline coefficients aren't shown.

[20]We only elicited initial stocks at baseline. Since this was before the harvest and few households had grain stocks, there would have been little meaningful variation in flows.

after treatment (Figure A25a), but the absence of increased maize *storage* for a group that received bags of maize demands explanation. The key is timing and the practice of field drying. In Gombe's hot, dry post-harvest climate, farmers commonly leave harvested maize to dry in the field rather than bringing it home immediately. Households that received dried, bagged maize from Taimaka could use it for immediate consumption needs—freeing them from the urgency of bringing their own field maize into home storage. Cash-arm households, by contrast, had no ready-to-eat maize on hand, so they went out and purchased dried grain to store at home. The net effect is that in-kind recipients likely held similar *total* maize (home stores plus field stocks), but our survey instrument—which asked about grain brought home—did not capture maize still drying in the field. Together, these results suggest that offering in-kind loans of a single crop is less effective at encouraging arbitrage than unrestricted cash, and that measurement of grain stocks is complicated by the distinction between home and field storage.

Panel B shows no statistically significant effects of either loan on the total value of grain sold during the season. The point estimates for both loans are positive but insignificant—about 21,000 Naira each (about $40 or one bag of maize), and confidence intervals are wide. Estimates during the course of the season are more precise and show a marginally significant increase in sales in January offset by an equivalent decrease in March and another increase in August for the cash loan. The maize loan shows a similar January spike but no significant total effect over the season. Breaking the total effects out by crop in Table B, we don't see significant effects across the season for any crop.

So if households in the cash treatment stored more grain and sold the same amount as controls, where did the extra grain come from? We see no effect of the cash treatment on additional grain purchased for storage, and a marginally significant decrease for the maize arm—the baseline level of which was already quite low (Panel E of Figure 8). Instead, Panel D shows that they reported "harvesting" more grain. Treatment was assigned toward the end of the growing season and acres planted are balanced, so higher physical yields aren't plausible. This is a measurement artifact. There are effectively two forms of grain storage—at home and in the field—and our survey only captured the first. Cash-arm households, having received no dried grain, had less in home storage than maize-arm households and so brought grain from the field to the homestead sooner—showing up as more "harvesting." When we enforce accounting consistency between reported flows and stocks, the effect disappears (see Appendix B.2). At endline—which occurred during the bean and millet harvests but early in the maize harvest—we also find that the cash arm reported bringing significantly (at the 10% level) more harvested grain into storage, consistent with higher output during the 2022 season discussed below.

These results imply that cash-arm households increased consumption of their own grain. The challenge is that up to Wave 3, the recall period from the general consumption module doesn't align with the stocks module, so we infer it as a residual. From Wave 4 onward, consumption from own stocks was directly elicited. The results appear in Panel C of 8. Households receiving the maize loan increase the number of own-produced crops consumed (and are more likely to consume maize) shortly after treatment, but this effect fades by the lean season. Meanwhile, we see households in the cash arm consuming more of their own crops during the lean season (April–June), but no differences by endline (Figure A4).

Despite the difficulty of eliciting stocks that add up in an accounting sense, the picture is fairly clear: the loan helped farmers maintain positive stocks further into the lean season but didn't turn them into arbitrageurs—not in a season when prices didn't appreciate.

One natural reading of these null results is that they reflect the intervention's design or just bad luck rather than a general failure of credit: both arms tied liquidity to maize, and maize prices happened to be flat. Had households been free to reallocate across crops—rice and groundnuts, for instance, exhibited stronger lean-season price increases—the story might have been different. We take this possibility seriously. But the Euler-equation tests in Section 7 speak directly to it. Those tests ask whether *any* locally available asset offered systematically unexploited arbitrage opportunities, not just maize. Our failure to reject efficient local pricing suggests that even untied credit would not have generated large risk-adjusted gains from cross-crop arbitrage within Gombe. The binding constraint appears to be aggregate—poor integration with the broader economy—rather than a failure to exploit the local asset menu.

5.4. **Effects on Consumption Expenditures.** We now turn to consumption expenditures and welfare. Having addressed the intertemporal price above, we focus here on three remaining pre-specified measures: log monthly expenditure; the MUE estimated following (Ligon 2020); and the RCSI measure of seasonal hunger (Maxwell, Vaitla, and Coates 2014). We also use log non-storable expenditure as an outcome, since expenditure on storable goods may be ambiguously correlated with welfare for crops households can store.

The short answer is: few overall effects on direct measures of consumption expenditure, but some notable seasonal patterns. Figure 9a shows a null effect on total expenditure, but we see a significant negative effect of nearly 20 log points for both treatment arms in April 2022—roughly the time at which prices typically begin to climb in years with large seasonal increases, which may have prompted treated households to cut consumption in anticipation of returns that didn't materialize. Figure 9c shows a positive but insignificant overall effect on non-storable food expenditure for the maize loan arm, with a significant increase in the first post-treatment wave. Figure 9e shows no significant effect of either treatment on durable expenditure. Figure 9f shows modest but significant reductions in the RCSI measure

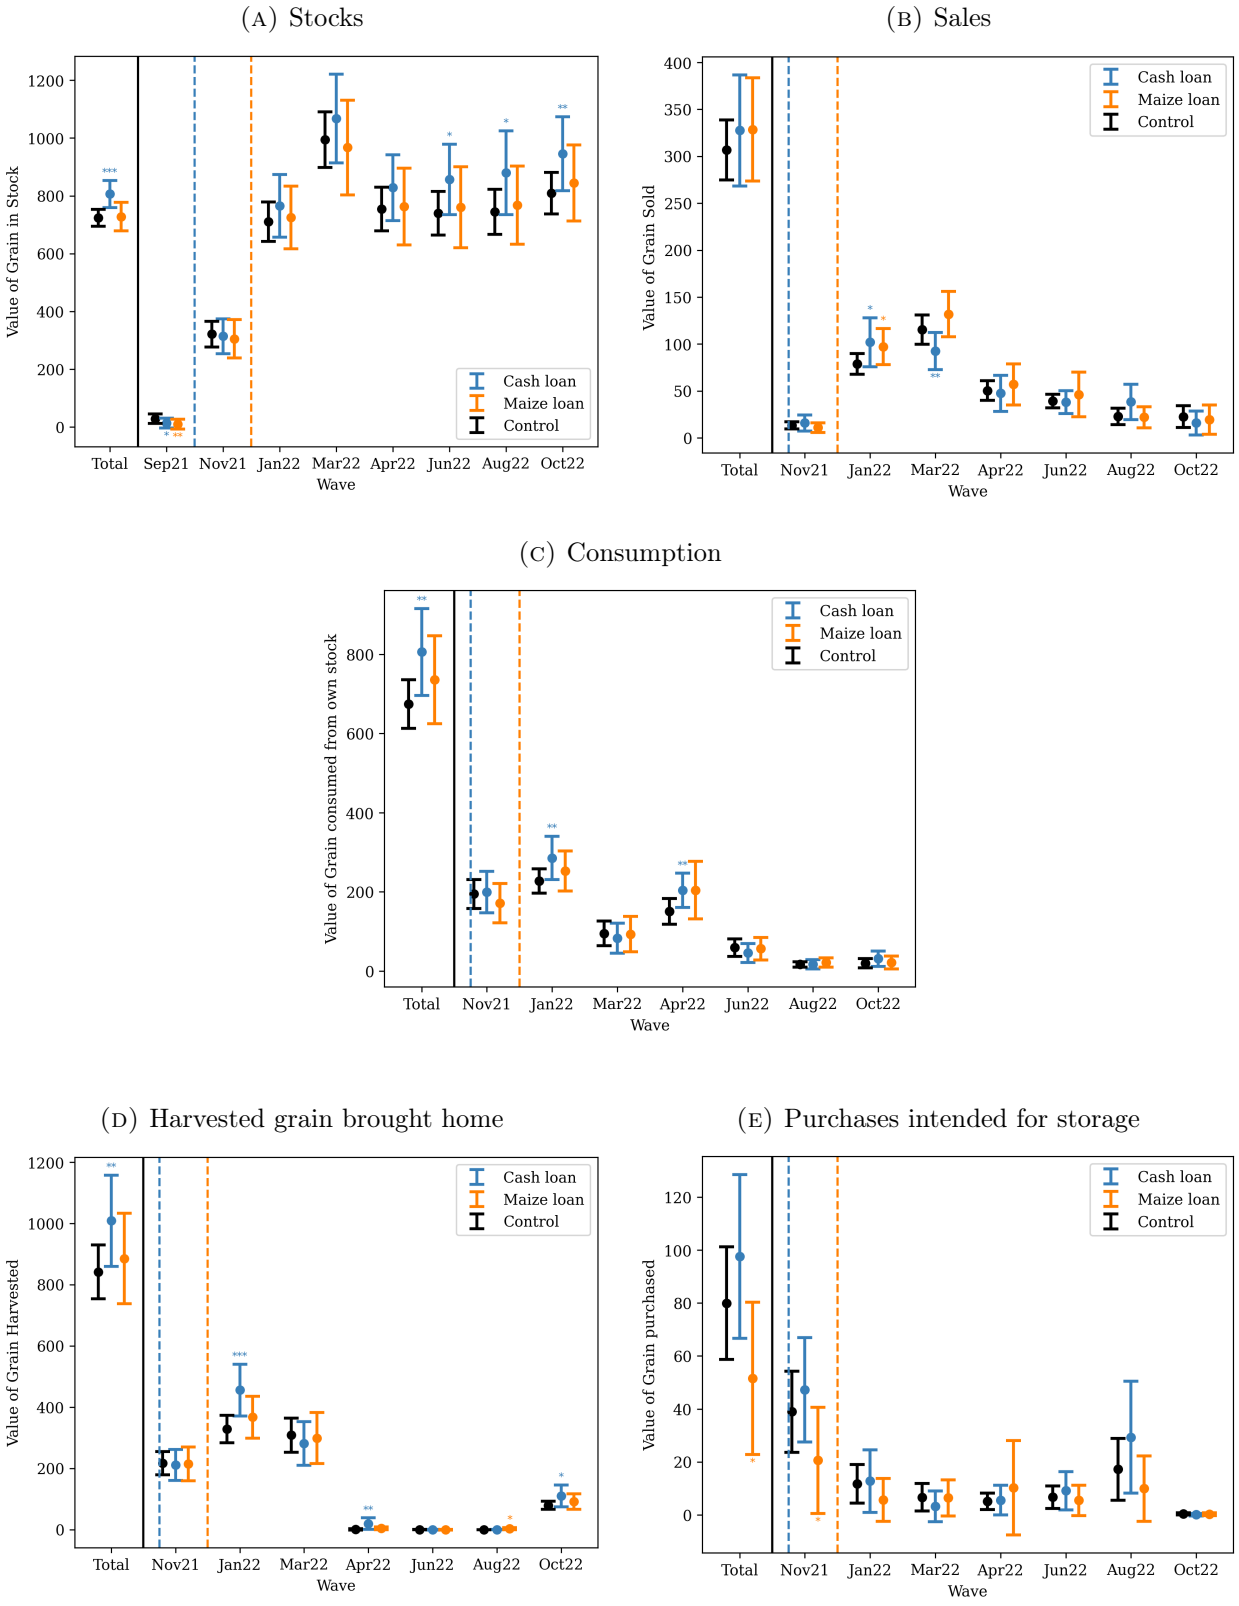


FIGURE 8. Treatment effects on value of grain stocks and flows by wave

Notes: Each panel plots Athey–Imbens treatment-effect estimates and 95% confidence intervals by survey wave. Outcomes are valued in Naira at contemporaneous local prices.

of seasonal hunger for both arms. Mapping expenditure into welfare, we see null effects on the average estimated MUE over the season, but two spikes for both treatment arms in April and at endline (higher values corresponding to lower welfare). The endline results are consistent with treated households reducing consumption when repaying loans or forfeiting assets after defaulting.[21]

Taken together, there's little evidence of higher or smoother consumption expenditures throughout the season. We can rule out increases of 2.6% (7.5%) in consumption expenditure and decreases of 2.1% (3.2%) in the MUE from the cash (maize) treatment. The contrast between expenditure-based and MUE-based measures is instructive: the MUE, which captures the marginal value of resources rather than just their level, detects welfare changes—particularly around loan delivery—that consumption expenditures are too noisy to resolve. The loan's main welfare benefit operates through reduced risk exposure rather than large improvements in consumption levels. But the significantly lower estimated welfare at endline suggests that default—for households that had taken the loan and didn't see high returns—may have been costly.

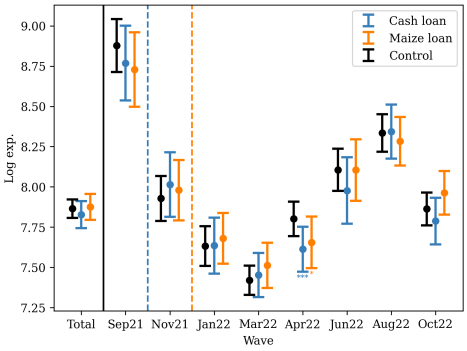
5.5. **Effects on Investment.** Other pre-specified outcomes include farm investment and profits from the 2022 planting season, non-agricultural business expenditures, borrowing and lending, and durable purchases.

We find that households in the cash arm significantly increased their investment in non-farm businesses (Figure 10) and livestock (Table 5.5), much as they did with their grain stocks. We see positive but noisy effects on acreage planted and agricultural expenditure for the 2022 season (Table 5.5). Since endline occurred early in the 2022 harvest season—when mostly millet and beans were being harvested—we find noisy positive effects on harvest values but large and significant effects on expected total harvest. Households in the cash arm also have significantly more chickens and goats than control households at endline; they own similar levels to baseline, while other households disinvested. The maize loan, by contrast, shows a null effect on each of the investment outcomes.

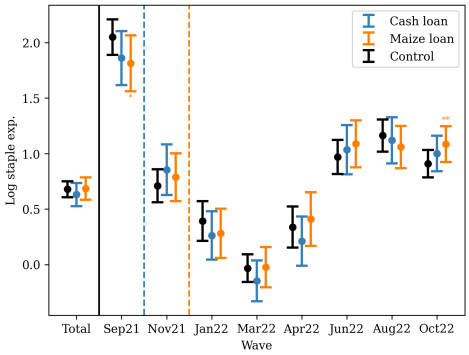
5.6. **Robustness.** We subject our main results to several pre-specified and other supplementary robustness checks. Appendix E.3 reports estimates (i) dropping clusters unbalanced by group-level attrition, (ii) using regressions with strata fixed effects and clustered standard errors following De Chaisemartin and Ramirez-Cuellar (2024), (iii) using Double Post-LASSO control selection (Belloni et al. 2012), and (iv) controlling for baseline outcomes in an ANCOVA specification (McKenzie 2012). Appendix C.3 adds Lee (2009) bounds on selective

[21]The positive effect on non-storable expenditure for the maize arm may seem at odds with the positive (welfare-reducing) effect on $\log \lambda$. This can arise because $\log \lambda$ reflects the *marginal utility* of expenditure, which depends on the full consumption bundle—including storable goods—not just the level of non-storable spending.

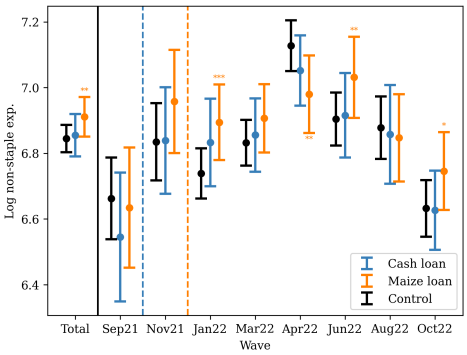
(A) Log food expenditure



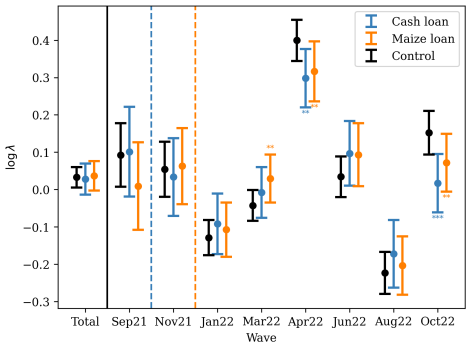
(B) Log storable expenditure



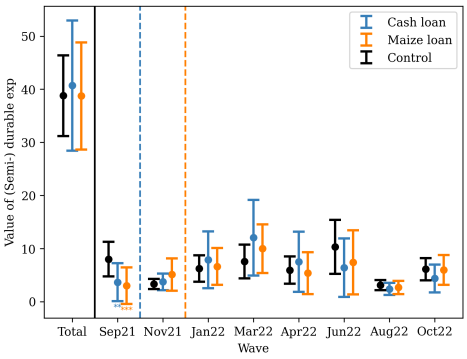
(C) Log non-storable expenditure



(D) MUE ($\log \lambda$)



(E) (Semi-)durable expenditure



(F) RCSI

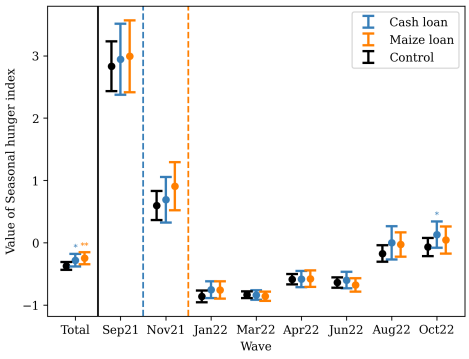


FIGURE 9. Treatment Effects on Consumption and Welfare, by Wave

Notes: Each panel plots Athey–Imbens treatment-effect estimates and 95% confidence intervals by survey wave.

attrition and Appendix E.4 reports a specification replacing strata fixed effects with the Taimaka group credit score, enumerator, and wave fixed effects. A complete concordance mapping our pre-analysis plan to the paper appears in Appendix E.

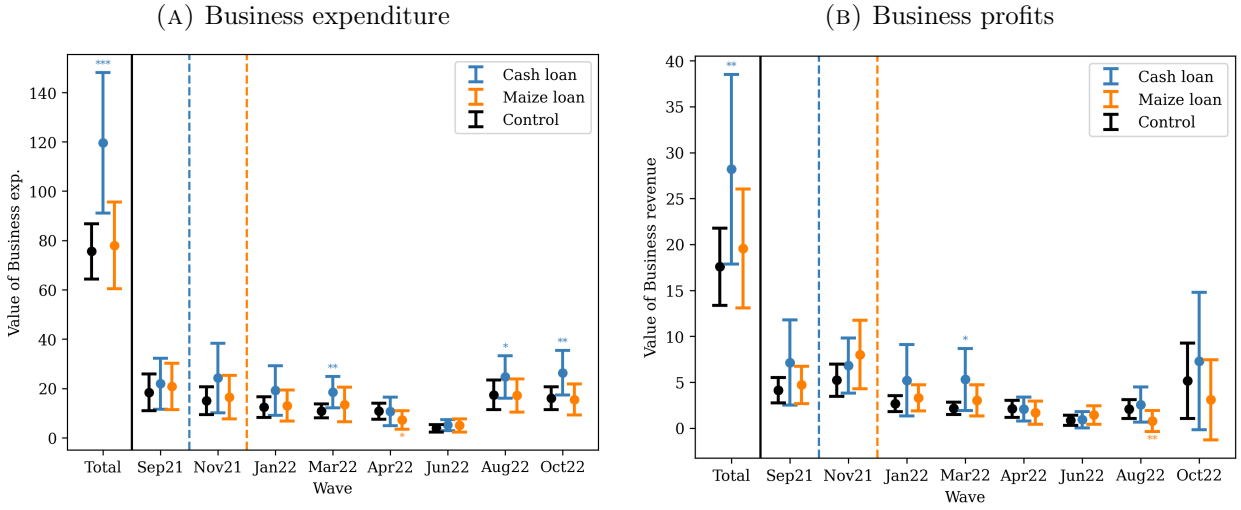


FIGURE 10. Treatment effects on non-farm business, by wave

Notes: Each panel plots Athey–Imbens treatment-effect estimates and 95% confidence intervals by survey wave.

TABLE 4. Effects of treatment on agricultural investment and output

Coefficient	Planted area (ha)	Ag. exp. ('000 Naira)	Dry season area (ha)	Dry season ag. exp.	Harvest value	Exp. harv value
Cash Loan	0.082	15.665	0.000	0.108	37.621*	184.961**
	(0.466)	(16.031)	(0.022)	(0.840)	(22.604)	(89.188)
Maize Loan	−0.334	4.677	−0.011	−0.101	23.778	130.060
	(0.394)	(14.005)	(0.019)	(0.884)	(17.361)	(110.670)
Any Loan	−0.137	8.882	−0.006	−0.021	27.543*	145.054**
	(0.387)	(12.792)	(0.017)	(0.721)	(15.956)	(73.524)
Control mean	5.482	144.122	0.024	0.894	92.447	784.949
N	829	829	829	829	829	829

This table contains estimates of the treatment effects of cash and maize loans on measures of agricultural investment in the 2022 planting season. The outcome in the first column is total area planted in hectares during the main (rainy) season of 2022, in hectares. The second column is total farm expenditure, including land, labor, fertilizer, seed, and equipment during the 2022 rainy season in 000's of Naira. The next two columns repeat the same outcomes for the 2022 dry season. The fifth column reports the value of crops harvested from the 2022 season by endline in 000's of Naira (at current market prices), while the final column combines actual and anticipated harvests as not all households had completed harvesting at endline. Results obtained using the estimator proposed by Athey and Imbens (2017) for tuple-wise randomized experiments.

None of these specifications alter the paper's main conclusions. Point estimates are qualitatively similar throughout; the regression-based specification tends to produce wider confidence intervals, consistent with having only about 30 clusters. The Lee bounds are tight, on the order of the 95% confidence intervals from the main specifications, reflecting the balanced attrition documented above.

TABLE 5. Effects of treatment on endline livestock holdings

	Chickens	Goats	Cows	Sheep	Donkeys
Cash Loan	1.747**	0.923**	0.519	0.448	−0.004
	(0.827)	(0.400)	(0.325)	(0.434)	(0.012)
Maize Loan	0.144	−0.354	−0.234	−0.494	−0.007
	(0.538)	(0.353)	(0.251)	(0.362)	(0.011)
Any Loan	0.807	0.260	0.133	−0.033	−0.005
	(0.548)	(0.315)	(0.249)	(0.364)	(0.010)
Control mean	2.262	3.843	1.528	2.472	0.010
N	829	829	829	829	829

This table contains estimates of the treatment effects of cash and maize loans on the number of different animals held at endline Results obtained using the estimator proposed by Athey and Imbens (2017) for tuple-wise randomized experiments.

TABLE 6. Estimated elasticities of w with respect to different shocks.

Shocks	$\hat{\delta}$
Business Expenses	0.085***
$N = 749$	(0.026)
Business Revenue	0.144***
$N = 2189$	(0.038)

Regressions of $w \equiv -\log\lambda$ logs of business expenses and revenues, estimated separately. Zeros are dropped. Robust standard errors are used.

6. TEST OF FULL INSURANCE

In the next section we consider idiosyncratic variation in IMRSs and test whether it's orthogonal to returns on different assets. That can be thought of as a test of whether portfolios are efficiently constructed, given some set of available assets.

Here we conduct a stronger test: we suppose that asset markets are complete, so that within Gombe households can perfectly smooth out idiosyncratic variation in $w = -\log\lambda$. This is essentially the well-known two-way fixed effects test of full insurance (Deaton 1992; Townsend 1994), implemented using our constructed MUEs—as theory suggests—rather than aggregate consumption expenditures. The estimating equation, explored by Ligon (2023), takes the form

$$-\log\lambda = \alpha_i + \eta_t + \delta\text{Shock}_{it} + \epsilon_{it},$$

where α_i is a household fixed effect, η_t is a time effect, and ϵ_{it} is idiosyncratic variation in MUEs. The question is whether this idiosyncratic variation is related to "shocks"—answered by testing whether the estimated δ is significantly different from zero.

We consider two related sources of idiosyncratic shocks: "business" revenues and expenses, which explicitly exclude agricultural operations. Not all farmers report any revenue or expenses, and few report these in every period. As one would expect, there's considerable seasonality in both expenses (usually early in the season) and the receipt of revenues (usually after harvest). The two-way fixed effects regression eliminates seasonality in both w and in the log of revenues and expenses, and eliminates also any cross-sectional correlation between w and these flows. Table 6 reports the estimated elasticities. Both are positive and significant. An idiosyncratic one percent increase in revenue is associated with roughly a 0.14% increase in w, while the corresponding expense elasticity is about 0.09%.

The sign on expenses may seem a puzzle. But of course expenses are endogenous: farmers who decide to spend more on their operations presumably expect a positive return. The point isn't that we're claiming a positive causal impact of expenses (or revenues) on w—it's that idiosyncratic variation in these flows is systematically predictive of idiosyncratic variation in w, allowing us to reject the (strong!) null hypothesis of full insurance in this environment.

Rejecting full insurance tells us that idiosyncratic risk is not perfectly shared in Gombe—some households bear shocks that others don't. But this alone does not tell us whether households are making the best of the assets available to them. In the next section we test this weaker condition: not whether all risk is shared, but whether portfolios are efficiently constructed given the assets households can actually trade. The distinction matters, because if portfolios are efficient but full insurance fails, then the binding constraint is the set of available assets—that is, poor market integration—not households' failure to optimize locally.

7. Test of Euler Equation

In Section 5 we established that while the cash treatment may have helped households smooth consumption, neither treatment led to significant increases in average arbitrage profits. But *could* households have made themselves better off *ex ante* by holding more maize? The realized price changes in 2021–22 weren't large by historical standards, but a purchase of maize in April 2022 followed by a sale in August would have returned about 10% over four months—an annualized yield of 33%. It would have been profitable to finance this using Taimaka's loans (15% interest) or commercial banks. But these are *ex post* returns. We'd like to say something about the *expected* welfare effects of investing in grain given the uncertainty over seasonal price increases.

Recall from Section 2 that the central prediction of the model with well-functioning local financial markets is that realizations of $y^j_{it+1} = (m_{it+1} - \bar{m}^j_{t+1})R^j_{t+1}\delta^j_{it}$ should be uncorrelated with any variable z^j_{it} in the time t information set. Three variables are of particular interest. First, lagged IMRS: if credit constraints were important for some households, their IMRSs

would be lower than for others, and the constraint would alter their portfolios. Second, treatment assignment: does giving a farmer a bag of maize make that farmer better off when prices rise? Third, lagged stocks—there's likely considerable persistence in stocks held, so if a farmer held more maize last period we might expect its IMRS to respond differently to maize price changes than the average farmer's.

TABLE 7. Tests of the Euler asset pricing equations

	Treatment	L^2m x loan	q	All
χ^2	2.20	4.36	4.93	16.15
df	11	11	5	29
p-value	1.00	0.96	0.42	0.97

Different columns involve tests of the orthogonality of errors to different sets of variables. Column 1 uses a dummy for assignment to either the cash or maize treatment arms. L^2 denotes the second lag of the IMRS, m_{it-1}, which is interacted with the treatment dummy. Column 3 uses the elicited bond price q and Column 4 tests these restrictions jointly.

We can construct estimated IMRSs m_{it+1} using ratios of time $t+1$ to time t MUEs for each household i, multiplied by some β. But our use of contemporaneous expenditure data to estimate $\log \lambda_{it}$ identifies the MUE only up to an unknown scalar σ. Thus,

$$m_{it+1} = \beta \left(\frac{\lambda_{it+1}}{\lambda_{it}} \right)^{\sigma}.$$

Using these, along with stock-holdings data, lets us construct $\bar{m}^h_{t+1}$, and so

$$\mathbb{E}_t \left(\left(\frac{\lambda_{it+1}}{\lambda_{it}} \right)^{\sigma} - \sum_{k=1}^{N} \left(\frac{\lambda_{kt+1}}{\lambda_{kt}} \right)^{\sigma} \frac{\delta^j_{kt}}{\sum_{k=1}^{N} \delta^j_{kt}} \right) R^j_{t+1} \delta^j_{it} = 0. \tag{7}$$

Note that the unknown discount factor β cancels, leaving only σ to estimate.

We estimate σ and test (7) using the continuously-updated GMM estimator of Hansen, Heaton, and Yaron (1996). Our preferred estimate of σ is 0.80, and Table 7 reports the test results. We're interested in whether forecast errors y_{it+1} are orthogonal to variables in the time t information set. Column 1 asks whether treatment status is correlated with these errors; there is no evidence at all of such a correlation. Column 2 asks whether twice-lagged m_{it-1}, interacted with receiving a loan, is orthogonal—one might think having more or fewer resources in the past would help predict the forecast error, but it does not. Column 3 considers our experimentally elicited IMRS, lagged one period. There's some suggestive evidence of a correlation, but not enough to reject the null. Finally, Column 4 puts all these variables together; unsurprisingly, we again fail to reject.

Note that we test (7) using returns on grain, not on the PHL itself. Since all but three farmers borrow the maximum amount, the loan ceiling binds and the Euler equation doesn't apply to that asset—the PHL is a corner solution, not an interior portfolio choice. What the tests *do* tell us is whether, given the assets farmers actually trade, there are unexploited local arbitrage opportunities.

The conclusion is straightforward: the low and seasonally variable levels of consumption expenditure in Gombe are an aggregate phenomenon. Offering PHLs to a subset of poor households is unlikely to address it.

8. Conclusion

Post harvest loans have gained attention as a way to help farm households earn additional income and smooth consumption expenditures across the season. But this theory of change rests on grain prices rising enough to cover the loans' interest rate—a highly risky proposition in sub-Saharan Africa. The implications for expected returns are ambiguous: beyond the possibility of low returns and high default when prices don't rise *ex post*, loans also shift *ex ante* risk from states with extremely high prices to those with low prices. Demand for the loans is therefore influenced by how the household expects its consumption expenditures to covary with grain prices. Whether defaults in our data are driven by genuine inability to repay when prices disappoint or by strategic considerations—as in Giné and Yang (2009)'s finding that limited liability reduces demand for insurance bundled with credit—is a question our data cannot resolve, but the distinction matters for evaluating the welfare consequences of risk-shifting through PHLs.

Our results from an RCT in Gombe State, Nigeria—in a year when prices didn't rise—show that the cash loan induced households to store more grain in their household for longer, but the in-kind maize loan did not. Neither treatment led to large increases in consumption expenditures, but the estimated MUE—a more sensitive welfare measure—detects significant improvements around delivery, and both treatments significantly reduced the experimentally elicited shadow rate and self-reported hunger. Over the course of the season, we also see small increases in business investment and livestock holdings in the cash loan group. This is consistent with credit constraints being relaxed—allowing households to store more and invest—but these investments didn't pay off. Throughout, two patterns recur. Seasonal swings in consumption, welfare, and stocks are far larger than any treatment effect—consistent with the lean season being an aggregate phenomenon, not a consequence of misallocation. And where treatment effects are significant, cash and maize loans look alike. The specific theory of change behind PHLs—that encouraging grain storage is what matters—doesn't find support in these data.

We use a simple model of intertemporal arbitrage to say something broader about the *expected* returns to storing grain, and in particular whether households could be better off *ex ante* by holding different portfolios. We do so by testing whether variation in households' realized IMRSs is correlated with variables in their prior-period information set. We fail to reject the null that treatment assignment, lagged estimated IMRS, stocks, and experimentally elicited IMRS are uncorrelated with this variation, though we come close for the elicited IMRS. This suggests that while Gombe's integration with the broader economy is poor, farmers aren't failing to respond optimally to local prices and returns. It doesn't rule out significant market failures—but it does suggest they aren't of the kind that PHLs are designed to fix.

Taken together, the uncertainty of prices, the lack of evidence against *ex ante* portfolio optimality, and the potential for adverse selection suggest that PHLs are not the right tool for improving seasonal welfare. Crop-specific credit programs like PHLs are particularly fragile: they concentrate risk in a single asset whose seasonal price increase may or may not materialise. One might wonder whether untied credit—allowing households to arbitrage across crops with stronger lean-season returns—would fare better. Our Euler-equation tests suggest not, at least within Gombe: we find no evidence that *any* locally available asset was systematically mispriced, so the scope for profitable reallocation across the local asset menu appears limited. More generally, what *cannot* fix seasonality is a "credit" product like PHLs (or many other microcredit products) that is functionally a fixed transfer. For a PHL to function as credit rather than a transfer, farmers would have to be allowed to borrow or save amounts the supply of loanable funds would have to be large enough that borrowers don't all hit a corner. That supply of resources can't come from within Gombe, where the risk-free rate itself varies with the season. It would have to flow in from outside, which is precisely the kind of market integration we argue would lead to real welfare improvements.

So what *could* fix the lean season? Better market integration, which could take two forms. The first form is integration in goods markets. If a Gombe farmer could buy and sell unlimited amounts of maize at external prices, the resulting seasonal inflows and outflows would tend to mitigate price variability directly. We find little evidence against the efficiency of allocation and production within Gombe *given local prices*—but local maize prices are only weakly correlated with world commodity prices, so local supply shocks affect prices more than they would in a better-integrated economy. What frictions presently prevent this integration? We don't know, but some candidates include transportation costs, poor contract enforcement, or simply a lack of information about prevailing prices in differerent places.

The second form is financial. If Gombe's farmers could borrow and lend at rates set by the broader economy, and the aggregate risk-free rate had less (or different) seasonal variation,

then this alone would tame seasonality in local IMRSs. This is the sense in which credit *can* fix the lean season. But as with the goods market, some friction prevents it. Commitment problems and default risk? Transaction costs like the actual difficulty of transferring funds to and contracting with individual smallholders? Endogenous market incompleteness, perhaps because outside finance would threaten the profits of local incumbents? Our data can't distinguish these mechanisms, but any could explain why Gombe remains isolated despite the apparent gains from integration.

Either goods or financial market integration could go a long way toward reducing the *predictable* seasonality in this environment, yet neither form of integration would eliminate the *risk* that farmers face. For that farmers would need financial instruments beyond credit. The introduction of forward contracts or other contingent contracts for maize and agricultural goods might be a less expensive and more useful innovation than PHLs. Who could absorb the price risk on the other side of such contracts? The counterparty needn't be a farmer. Anecdotally, some of our own enumerators—educated, urban, and relatively well-to-do—were already engaged in speculative grain storage before becoming involved in this project. Agribusinesses, cooperatives, or urban traders with better access to capital and storage infrastructure are natural candidates, and their greater diversification across crops, locations, and seasons would make them better suited to bear aggregate price risk than individual smallholders. That such intermediaries don't already play a larger role in Gombe points back to the same integration failures we've documented: poor transport links, thin local markets, and weak contract enforcement make it difficult for outside capital—whether in the form of goods or financial instruments—to reach these farmers on terms that both sides find attractive.

Finally, we note that the framework developed here extends beyond the evaluation of any single policy. The large cross-crop and cross-year variation in seasonal prices we document is likely characteristic of many agrarian economies in sub-Saharan Africa. A model of optimal portfolio choice given borrowing constraints, storage technology, and the stochastic structure of prices provides a disciplined way to compare policy designs (crop-specific credit, general liquidity, storage improvements, forward contracts) and to ask which margins each relaxes. With high-frequency price data increasingly available through FEWS NET and similar sources, and household panel data from the LSMS-ISA covering multiple countries, the moment conditions we derive can be taken to other settings. We hope this framework proves useful for studying seasonality and welfare more broadly.

References

Aggarwal, Shilpa, Eilin Francis, and Jonathan Robinson (2018). "Grain today, gain tomorrow: Evidence from a storage experiment with savings clubs in Kenya". In: *Journal of Development Economics* 134, pp. 1–15.

Aiken, Emily L et al. (2020). "Targeting Development Aid with Machine Learning and Mobile Phone Data: Evidence from an Anti-Poverty Intervention in Afghanistan". In: *Proceedings of the 3rd ACM SIGCAS Conference on Computing and Sustainable Societies*, pp. 310–311.

Allen, Treb and David Atkin (2022). "Volatility and the Gains from Trade". In: *Econometrica* 90.5, pp. 2053–2092.

Athey, Susan and Guido W Imbens (2017). "The econometrics of randomized experiments". In: *Handbook of economic field experiments*. Vol. 1. Elsevier, pp. 73–140.

Augenblick, Ned et al. (2023). *Retrieval Failures and Consumption Smoothing: A Field Experiment on Seasonal Poverty*. Tech. rep. Technical Report, University of California, Berkeley.

Barrett, Christopher B and Paul A Dorosh (1996). "Farmers' welfare and changing food prices: Nonparametric evidence from rice in Madagascar". In: *American Journal of Agricultural Economics* 78.3, pp. 656–669.

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

Belloni, Alexandre et al. (2012). "Sparse models and methods for optimal instruments with an application to eminent domain". In: *Econometrica* 80.6, pp. 2369–2429.

Bergquist, Lauren Falcao, Craig McIntosh, and Meredith Startz (2021). *Search cost, intermediation, and trade: Experimental evidence from Ugandan agricultural markets*. eScholarship, University of California.

Bruhn, Miriam and David McKenzie (2009). "In pursuit of balance: Randomization in practice in development field experiments". In: *American economic journal: applied economics* 1.4, pp. 200–232.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak (2014). "Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh". In: *Econometrica* 82.5, pp. 1671–1748.

Burgess, Robin and Dave Donaldson (2010). "Can openness mitigate the effects of weather shocks? Evidence from India's famine era". In: *American Economic Review* 100.2, pp. 449–453.

Burke, Marshall, Lauren Falcao Bergquist, and Edward Miguel (2019). "Sell low and buy high: arbitrage and local price effects in Kenyan markets". In: *The Quarterly Journal of Economics* 134.2, pp. 785–842.

Cardell, Lila and Hope Michelson (2022). "Price risk and small farmer maize storage in Sub-Saharan Africa: New insights into a long-standing puzzle". In: *American Journal of Agricultural Economics*.

Caselli, Francesco et al. (2020). "Diversification through trade". In: *The Quarterly Journal of Economics* 135.1, pp. 449–502.

Channa, Hira et al. (2022). "Overcoming smallholder farmers' post-harvest constraints through harvest loans and storage technology: Insights from a randomized controlled trial in Tanzania". In: *Journal of Development Economics* 157, p. 102851.

Christian, Paul and Brian Dillon (2018). "Growing and learning when consumption is seasonal: long-term evidence from Tanzania". In: *Demography* 55.3, pp. 1091–1118.

De Chaisemartin, Clément and Jaime Ramirez-Cuellar (2024). "At what level should one cluster standard errors in paired and small-strata experiments?" In: *American Economic Journal: Applied Economics* 16.1, pp. 193–212.

Deaton, Angus (1992). "Household Saving in LDCs: Credit Markets, Insurance, and Welfare". In: *Scandinavian Journal of Economics* 94, pp. 253–273.

Delavallade, Clara and Susan Godlonton (2023). "Locking crops to unlock investment: Experimental evidence on warrantage in Burkina Faso". In: *Journal of Development Economics* 160, pp. 1029–1059.

Dillon, Brian (2021). "Selling crops early to pay for school: A large-scale natural experiment in Malawi". In: *Journal of Human Resources* 56.4, pp. 1296–1325.

Dupas, Pascaline and Jonathan Robinson (2013). "Savings constraints and microenterprise development: Evidence from a field experiment in Kenya". In: *American Economic Journal: Applied Economics* 5.1, pp. 163–192.

Fafchamps, Marcel (1992). "Cash crop production, food price volatility, and rural market integration in the third world". In: *American journal of agricultural economics* 74.1, pp. 90–99.

Fink, Günther, B Kelsey Jack, and Felix Masiye (2020). "Seasonal liquidity, rural labor markets, and agricultural production". In: *American Economic Review* 110.11, pp. 3351–3392.

Giné, Xavier and Dean Yang (2009). "Insurance, credit, and technology adoption: Field experimental evidence from Malawi". In: *Journal of Development Economics* 89, pp. 1–18. DOI: 10.1016/j.jdeveco.2008.09.007.

Hansen, Lars Peter, John Heaton, and Amir Yaron (1996). "Finite-sample properties of some alternative GMM estimators". In: *Journal of Business & Economic Statistics* 14.3, pp. 262–280.

Lee, David S (2009). "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects". In: *The Review of Economic Studies* 76.3, pp. 1071–1102.

Ligon, Ethan (2020). "Estimating household welfare from disaggregate expenditures". In: *CUDARE Working Paper*.

– (2023). "Risk sharing tests and covariate shocks". URL: `https://escholarship.org/uc/item/2zr503fq`.

Maxwell, Daniel, Bapu Vaitla, and Jennifer Coates (2014). "How do indicators of household food insecurity measure up? An empirical comparison from Ethiopia". In: *Food policy* 47, pp. 107–116.

McKenzie, David (2012). "Beyond baseline and follow-up: The case for more T in experiments". In: *Journal of development Economics* 99.2, pp. 210–221.

Meghir, Costas et al. (2022). "Migration and informal insurance: Evidence from a randomized controlled trial and a structural model". In: *The Review of Economic Studies* 89.1, pp. 452–480.

Merfeld, Joshua D and Jonathan Morduch (2023). "Poverty at higher frequency". In: *KDI School of Pub Policy & Management Paper No. DP23-03*.

Omotilewa, Oluwatoba J et al. (2018). "Does improved storage technology promote modern input use and food security? Evidence from a randomized trial in Uganda". In: *Journal of Development Economics* 135. Publisher: Elsevier, pp. 176–198.

Porteous, Obie (2019). "High trade costs and their consequences: An estimated dynamic model of African agricultural storage and trade". In: *American Economic Journal: Applied Economics* 11.4, pp. 327–366.

Rosenzweig, Mark R and Christopher Udry (2020). "External validity in a stochastic world: Evidence from low-income countries". In: *The Review of Economic Studies* 87.1, pp. 343–381.

Samphantharak, Krislert and Robert M. Townsend (2018). "Risk and return in village economies". In: *American Economic Journal: Microeconomics* 10.1, pp. 1–40.

Stephens, Emma C and Christopher B Barrett (2011). "Incomplete credit markets and commodity marketing behaviour". In: *Journal of agricultural economics* 62.1, pp. 1–24.

Townsend, Robert M. (1994). "Risk and Insurance in Village India". In: *Econometrica* 62.3, pp. 539–591.

Appendix A. Loan Terms

This appendix describes the terms of the two loan products offered by the Taimaka Project in the 2021–22 season. Both treatments required joint liability in groups of five, with the promise of increased loan size in future years conditional on full group repayment.

A.1. **Cash Loan.**

Eligibility: Applicants must reside in the study area, grow maize, apply in a group of five, and consent to monthly data collection.

Loan size: Determined by the number of bags of maize a farmer chose to commit to Taimaka. Committing a bag meant tagging it, promising to hold it until July 15, 2022, and selling it then to repay the loan. The schedule was:

- 4 bags: 50,000 Naira (47,600 cash + 4 PICS bags at 600 each)
- 3 bags: 40,000 Naira (38,200 cash + 3 PICS bags)
- 2 bags: 30,000 Naira (28,800 cash + 2 PICS bags)

All but three farmers chose to commit the maximum of 4 bags.

Repayment: Total amount owed = loan amount + 15% administrative fee (e.g., 57,500 Naira for a 50,000 Naira loan). Minimum monthly payments of 3,000 Naira beginning December 2021, with full repayment due by July 15, 2022.

Delivery: September 16 to approximately October 10, 2021.

A.2. **Maize Loan.**

Eligibility: Same as cash loan.

Loan size: Every recipient received a flat allocation of 3 bags of maize plus 3 PICS bags, regardless of own stocks.[22]

Repayment: Total amount owed = (3 × purchase price per bag + 3 × 600 Naira for PICS bags + transport cost) + 15% administrative fee. Because maize purchase prices and transport costs varied across sites, total obligations ranged from 62,000 to 63,000 Naira.[23] Minimum monthly payments of 3,000 Naira beginning January 2022, with full repayment due by July 15, 2022.

Delivery: November 15 to November 29, 2021.

A.3. **Notes on Enforcement.** PICS bags were provided with the expectation that farmers would re-bag their committed (cash arm) or received (maize arm) grain into them. In practice, enforcement was uneven: some farmers re-bagged, others used the PICS bags for other purposes. Similarly, the commitment to hold bags until the due date was not

[22]An earlier design would have allowed farmers to choose the number of bags (up to their verified stock), mirroring the cash arm. This was simplified to a flat 3 bags after loan officers raised concerns about palatability of larger maize loans.

[23]One site (Gama Dadi) was contracted at 60,000 Naira due to an administrative error that excluded the PICS bag cost; this could not be revised after contracts were signed.

systematically monitored, though loan officers used it as leverage with borrowers who sold early without making progress on repayment.

Appendix B. Additional Tables and Figures

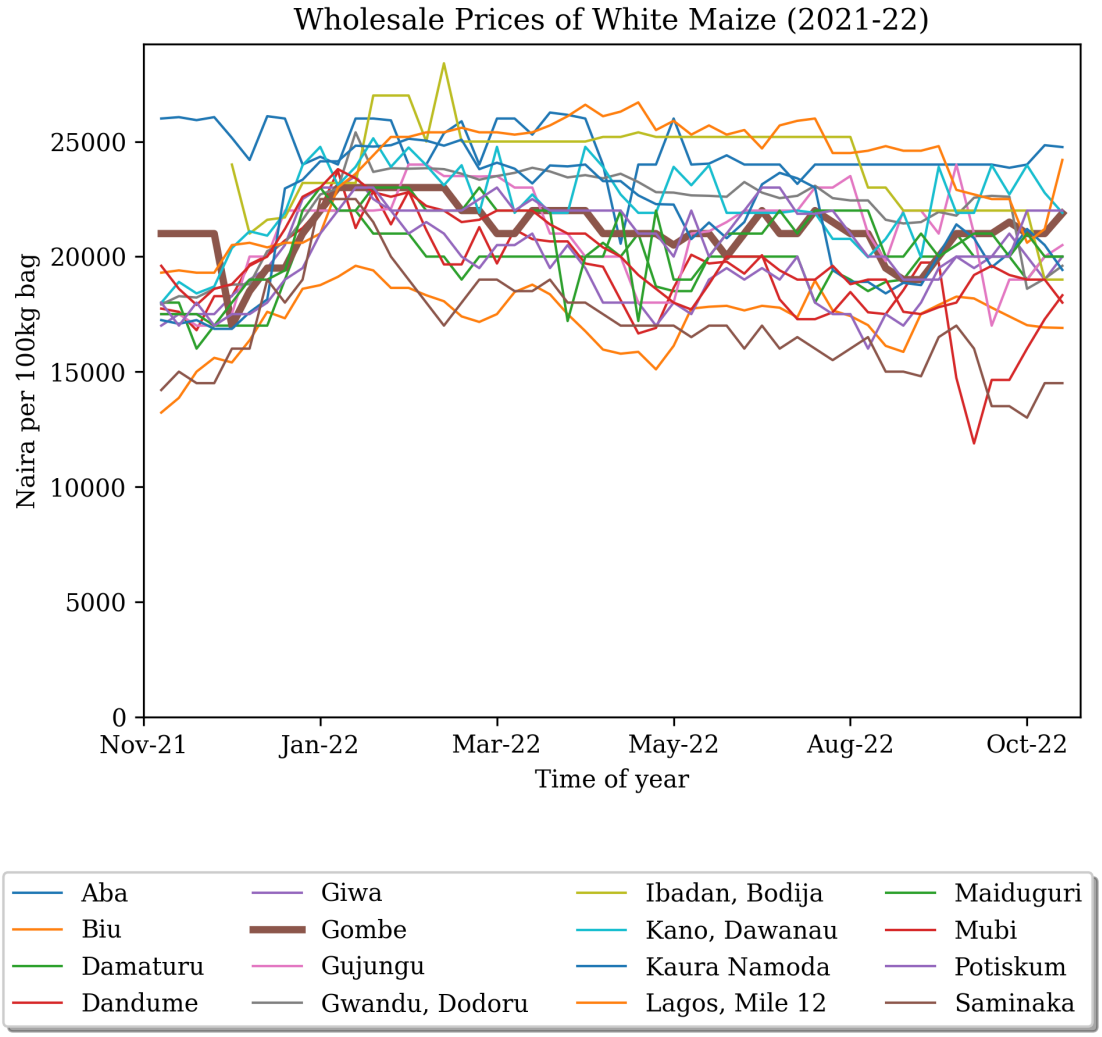


Figure A1. Maize prices across Nigeria during study period

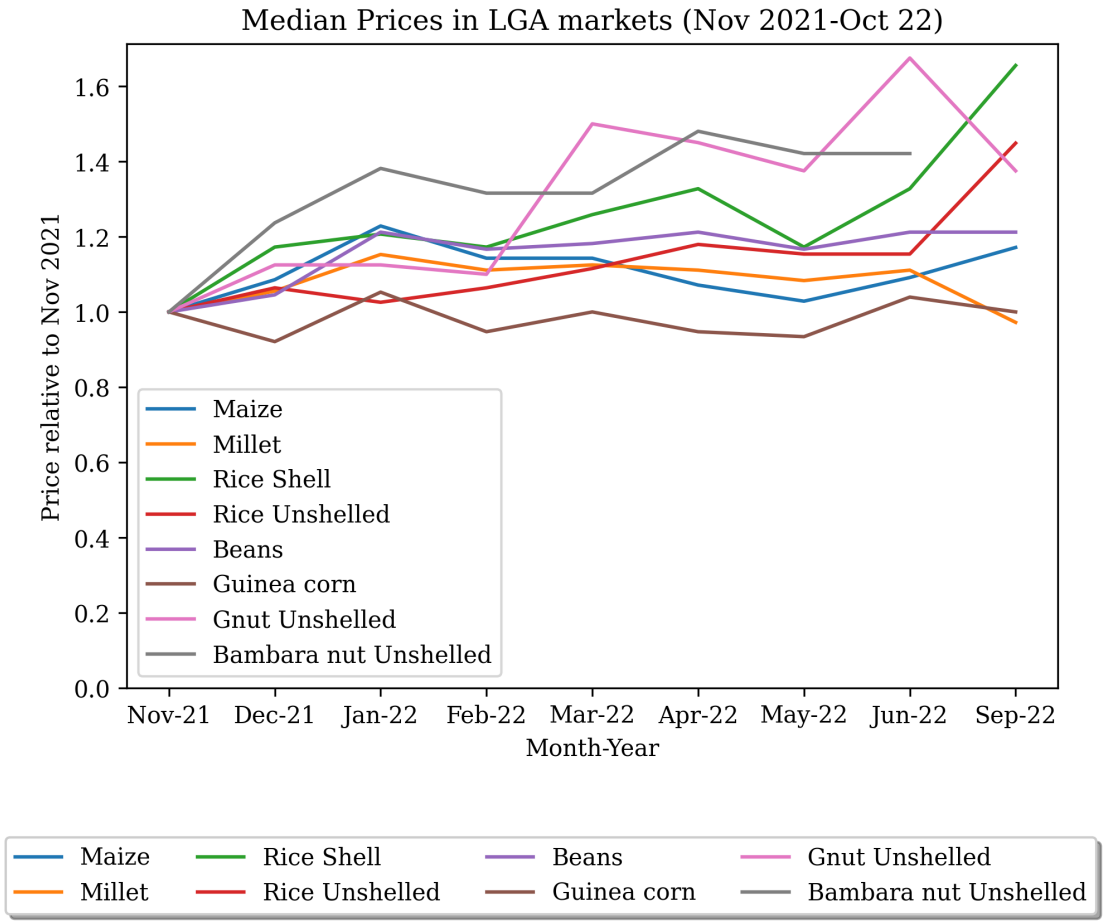


FIGURE A2. Prices at local markets within Gombe collected during study period

TABLE A1. Summary statistics on maize sale prices

Statistic	Value
Total maize sales	847
Fraction sold after January 2022	0.782
Fraction sold below 18,000 Naira/bag	0.151
Fraction sold below 20,000 Naira/bag	0.434
Fraction sold below break-even (15% above Nov)	0.163
Mean sale price (Naira/bag)	19,550
Median sale price (Naira/bag)	20,000

Sale price computed as revenue divided by quantity sold (Naira per kg), converted to price per 100kg bag. Observations filtered to sale prices between 10 and 5,000 Naira/kg. Break-even threshold assumes November 2021 Taimaka median maize price of 16,500 Naira/bag plus 15\Source: `data/stock_vals.parquet`.

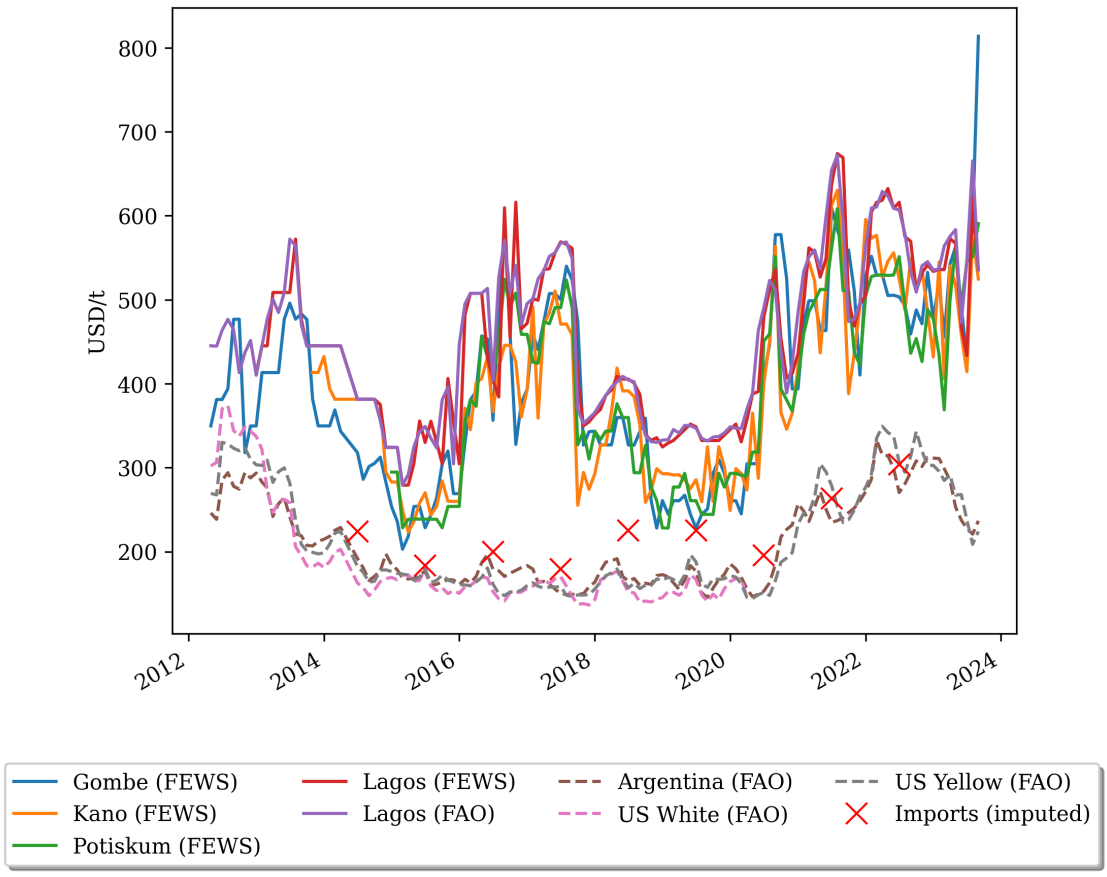


FIGURE A3. Series of local Nigerian (solid) and international (dashed) maize prices in USD

TABLE A2. Effects of treatment on total value of crops sold (000's Naira).

	Maize	Beans	Guinea Corn	Rice	Millet	All
Cash Loan	11.099	4.480	−9.028	5.043	19.788	30.651
	(9.140)	(13.698)	(6.722)	(5.631)	(21.880)	(33.701)
Maize Loan	11.832	−4.062	−4.258	3.638	3.864	17.981
	(8.866)	(14.798)	(6.877)	(5.973)	(8.778)	(29.876)
Any Loan	11.811	0.418	−6.556	4.103	11.486	24.145
	(7.482)	(11.897)	(6.731)	(4.757)	(12.450)	(25.881)
Control Mean	55.377	154.165	15.187	14.325	75.744	335.174
N	922	922	922	922	922	922

$^{***}p < .01, ^{**}p < .05, ^{*}p < 0.10$

This table contains estimates of the treatment effects of cash and maize loans on the total reported sales of each crop, using the estimator proposed by Athey and Imbens (2017) for tuple-wise randomized experiments.

TABLE A3. Maize sale prices and timing

Panel A: Sale prices		
Price threshold (Naira/bag)	% sold (by weight)	N transactions
Below 16,000	6.2%	36
Below 18,000	31.3%	231
Below 20,000	82.7%	677
Below 22,000	94.3%	794
Total transactions		848
Panel B: Timing of sales		
Survey period	% sold (by weight)	N transactions
Oct–Dec 2021	4.4%	35
Dec 2021–Feb 2022	23.0%	164
Feb–Mar 2022	31.0%	229
Apr–May 2022	15.0%	146
Jun–Jul 2022	11.3%	137
Aug–Sep 2022	9.8%	103
Sep–Oct 2022	5.6%	34

Panel A reports the cumulative share of maize sold (by weight) at prices below the indicated threshold. Panel B reports the share sold in each survey wave. Prices are per 100kg bag. Source: household survey data, 2021–22 season.

TABLE A4. Bond price elicitation quality by survey wave

Panel A: Refusal rates			
Wave	Refusal rate	N surveyed	N refused
Nov 2021	33.0%	540	178
Jan 2022	24.2%	653	158
Mar 2022	18.1%	818	148
Apr 2022	24.2%	790	191
Jun 2022	13.7%	629	86
Aug 2022	19.8%	758	150
Oct 2022	19.8%	799	158
Panel B: WARP violation rates			
Wave	Violation rate	N valid	N violations
Nov 2021	37.8%	291	110
Jan 2022	2.9%	245	7
Mar 2022	2.9%	381	11
Apr 2022	7.2%	388	28
Jun 2022	2.7%	445	12
Aug 2022	2.4%	494	12
Oct 2022	0.8%	516	4

Panel A reports the share of surveyed households that refused to participate in the bond price elicitation exercise. Panel B reports the share of non-refused responses that violated the Weak Axiom of Revealed Preference (selecting non-contiguous interest rates). Source: household survey data, 2021–22.

TABLE A5. Effects of treatment on total value of consumption of different crops.

	Maize	Beans	Guinea Corn	Rice	Millet	All
Cash Loan	33.977	50.036***	4.725	−13.794	49.465***	131.759**
	(30.880)	(13.561)	(9.955)	(11.323)	(15.585)	(55.981)
Maize Loan	27.384	34.526***	0.513	−14.051	15.689	61.458
	(32.208)	(13.277)	(8.854)	(11.281)	(14.680)	(56.657)
Any Loan	23.360	42.508***	2.820	−14.789	32.956**	88.372**
	(24.585)	(10.659)	(8.385)	(9.858)	(12.614)	(44.954)
Control Mean	303.658	97.344	47.173	78.589	139.847	674.250
N	922	922	922	922	922	922

$^{***}p < .01,^{**}p < .05,^{*}p < 0.10$

This table contains estimates of the treatment effects of cash and maize loans on the consumption of own grain stocks by crop. Results using the estimator proposed by Athey and Imbens (2017) for tuple-wise randomized experiments.

TABLE A6. Treatment effects on household consumption expenditures and welfare

	log exp.	log perishable exp.	−RCSI	$\log \lambda$	Shadow rate $1/q$
Cash Loan	−0.037	0.005	0.091**	−0.010	0.001
	(0.032)	(0.025)	(0.041)	(0.016)	(0.006)
Maize Loan	0.014	0.059**	0.085**	0.003	−0.013**
	(0.031)	(0.024)	(0.038)	(0.015)	(0.006)
Any Loan	−0.017	0.030	0.094***	−0.010	−0.007
	(0.027)	(0.021)	(0.032)	(0.013)	(0.006)
Fixed Effects	Strata	Strata	Strata	Strata	Strata
Control mean	7.865	6.854	0.377	−0.005	0.836
N	5733	5714	4021	5250	4724

Different outcomes related to consumption expenditures using the estimator proposed by Athey and Imbens (2017) for tuple-wise randomized experiments. Log perishable expenditures sums all non-durable consumption expenditures. Log expenditures adds in expenditures on storable grains. −RCSI is the (negative) of the "reduced coping strategies index" which attempts to measure hunger. The log of the marginal utility of expenditures is $\log \lambda$, while the shadow rate $1/q$ is the minimum gross return the household requires on a bond. The "cash loan" and "maize loan" treatments are as indicated in the text; the "any loan" treatment is a loan received in either cash or kind.

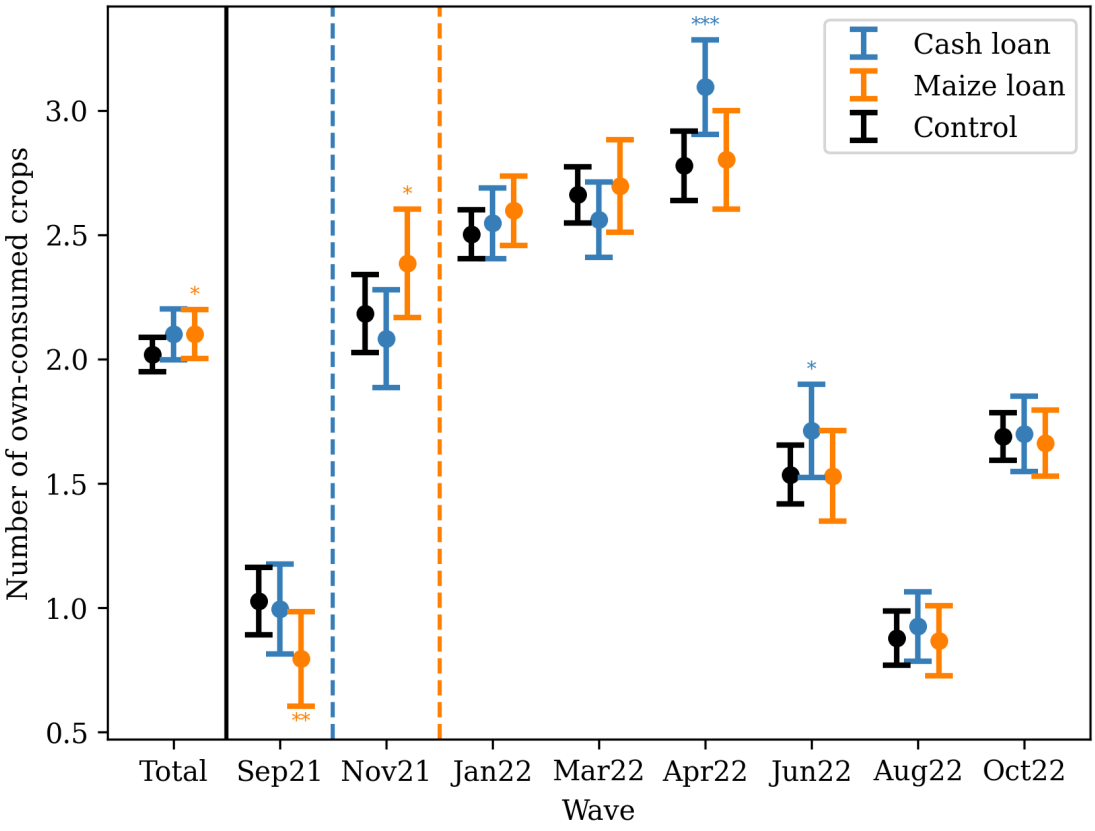


FIGURE A4. Treatment effects on number of crops consumed from own stock by wave

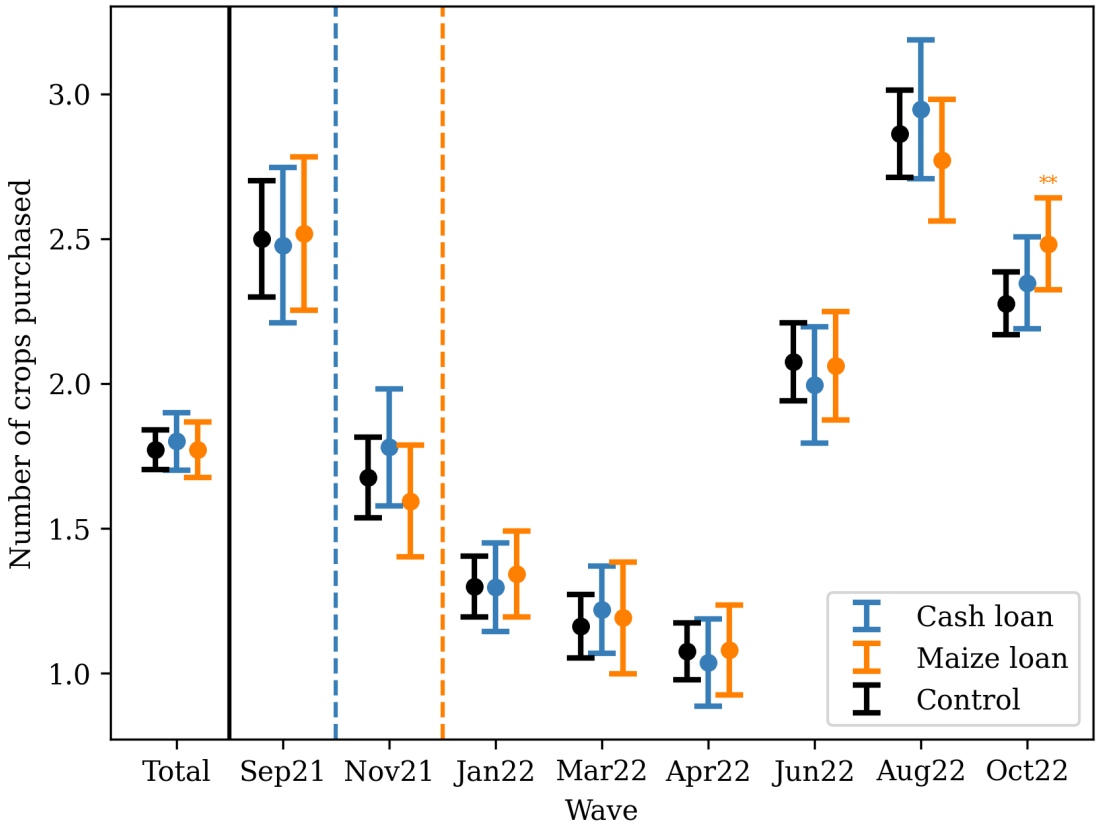


FIGURE A5. Treatment effects on number of crops purchased by wave

TABLE A7. Covariate balance between treatment arms at baseline.

	Control	Cash	Maize	p-val Cash vs. Control	p-val Maize vs. Control
Female HH head	0.076	0.056	0.059	0.320	0.413
HH Size	6.471	5.996	5.782	0.182	0.039
Age HH Head	39.165	38.641	37.047	0.621	0.028
No educ.	0.685	0.672	0.677	0.724	0.815
Some pri. school	0.083	0.096	0.076	0.583	0.755
Finished pri. school	0.116	0.103	0.092	0.631	0.345
Finished sec. school	0.109	0.070	0.086	0.100	0.332
Tertiary educ.	0.007	0.059	0.069	0.001	0.000
Cows	4.305	4.732	4.742	0.559	0.565
Goats	6.077	6.522	6.354	0.342	0.540
Sheep	5.419	6.068	6.124	0.260	0.182
Chickens	11.957	12.219	11.046	0.857	0.509
Has Business	0.169	0.186	0.194	0.598	0.416
Has loan	0.043	0.057	0.068	0.432	0.171
Cultivated area (ha)	5.867	5.698	6.621	0.663	0.367
Ag. exp.	81.532	76.019	87.146	0.462	0.498
Ag. revenue	106.790	92.667	84.520	0.392	0.175

B.1. **Covariate Balance.**

B.2. **Imputation of Stocks.** It is very clear that the raw data is not always an accurate accounting of farmers' inventory. Common inconsistencies include:

- Households will report selling or consuming goods that we don't see in previous stocks, purchases, or harvests
- Households will report purchasing or harvesting goods that we never see them dispose of, but at endline they report having less than the acquired amounts in stock
- In later rounds, households' answers to whether they have consumed a good from their own stock in the inventory module do not match their answers to whether they have acquired the good from their own production or purchased it.

To resolve these issues, we apply the following algorithm, which allows us to compute a number of different indicators of stocks, with higher indices corresponding to more imputation. We proceed until baseline stocks plus the cumulative sum of flows matches endline stocks, with the restriction that stocks never fall below 0 at any point in the period. We use responses from the consumption module as an additional source of information about whether households had positive stocks at a given period. The goal is to take flows as seriously as possible rather than reported stocks.

(1) Save reported stocks (which are only available in waves 1-3 and at baseline). Assign dummy value of 1 if positive
(2) Then compute stocks as cumulative sum of report flows (purchases + harvest - sales - consumption). Replace 0 dummy with 2 if this is positive. Add reported initial stocks (positive for only a few HH) and replace 0 dummy with 3 if this is positive
(3) Replace 0 dummy with 4 if the household reports consuming crop from own stock in consecutive waves
(4) For waves 1-3, impute consumption as the difference between the change in stocks between the previous and current period net of flows (harvest + purchases for storage - sales). Update flows accordingly (updating stocks is trivial)
(5) If the household has harvested a positive amount between waves 1-3, add the residual between reported stocks and harvests to the harvest. Assign dummy value of 6 when this turns stocks positive. Update flows and stocks accordingly. Note that we need to do this iteratively (because if a HH meets this condition for multiple waves but has activity in the interim, one iteration won't account for this.)
(6) Still have some people who sell/consume crops they've harvested after period 3 that are unaccounted for. For example, bg-11d claims to harvest and sell 400kg of beans in period 2 but then sells another 200 in period 4 and 100 in period 5. It seems more likely that these came from the harvest rather than purchases. Likewise for consumption that's reported as being own-produced. We are going to want to attribute these

to the latest feasible harvest (between waves 1-3). Assign dummy value of 7 when this turns stocks positive. Update stocks and flows again

(7) Where there are unaccounted for sales/consumption in later waves for crops that a household never harvests, attribute these to purchases (`arb`). Assign dummy value of 8

(8) When households have excess positive stocks in excess of the endline stocks they report, attribute the maximum amount that will not cause subsequent stocks to go negative at any point to current period consumption. This again needs to be done iteratively. Assign dummy value of 9 and update accordingly.

(9) When households don't report consuming these crops from own production, attribute to sales and assign dummy value of 10.

(10) Assume any remaining discrepancy is stock carried over to endline.

We end up with 5 stock variables:

- Stock 1 is raw sum of flows plus baseline stock
- Stock 2 Tries to ensure that stock1 never goes negative by attributing differences to previous harvests, if feasible based on timing and what the household grew
- Stock 3 Ensures that stocks don't go negative by offsetting excess outflows with previous purchases if they can't be attributed to harvests
- Stock 4 corrects for any excess stocks at endline (relative to reported) by attributing it to consumption such that stocks don't go negative
- Stock 5 Attributes remaining differences to unreported sales

Appendix C. Attrition

Our sample faces two distinct sources of attrition: the loss of entire farmer groups from the program, and the failure to survey individual farmers in particular rounds. Below, we show that the main results of the paper are robust to both sources.

C.1. **Group-Level Attrition.** Seven groups were removed from the sample after treatment assignment but before loan disbursement. One was dropped because a traditional leader in the group had not disclosed his status during enrollment; he and his group were reassigned to the cash arm outside the study. Two were dropped due to verification failures discovered during delivery. The remaining four withdrew voluntarily—traditional leaders reversed their endorsement in two cases, and the groups themselves chose to withdraw in the other two. Only a handful of these individuals appear even in the baseline survey.

A further seven groups were lost after the first post-baseline round: five from the maize arm and two from the control arm. Three opted out of the program, two were women's groups whose husbands forbade their continued participation, and two were dropped for other reasons.

Because treatment was assigned at the group level within randomization clusters (strata of six adjacent groups), losing groups from a cluster can disturb balance. Our preferred approach is to drop the seven affected clusters after baseline and five additional clusters after round 1. This yields a conditional average treatment effect, averaging over the remaining clusters. We show in Appendix C.1.1 that the main results are robust to this restriction.

C.1.1. *Dropping Unbalanced Clusters.*

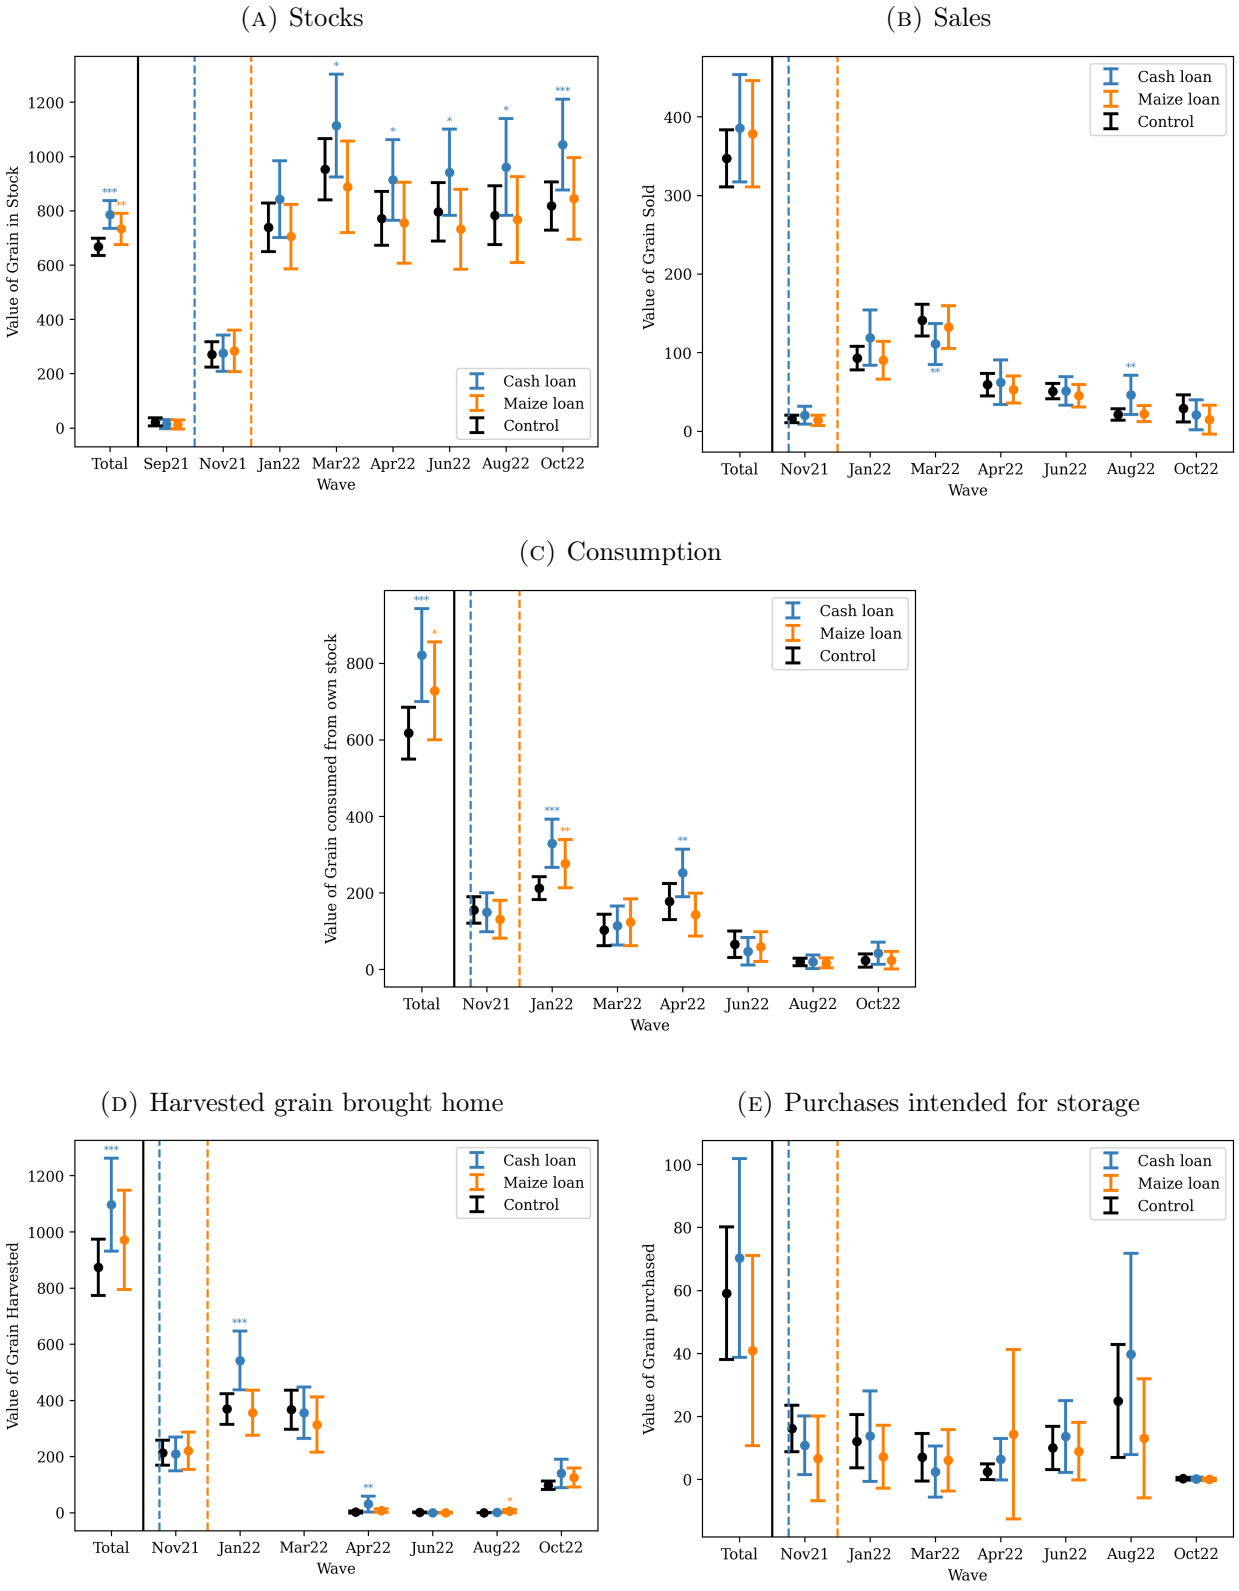
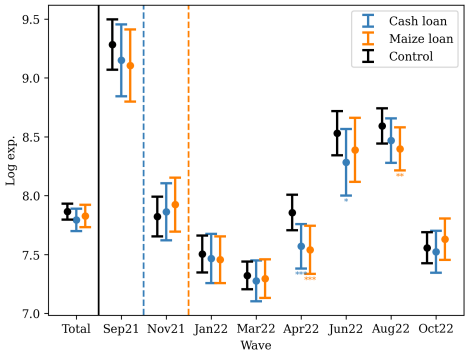


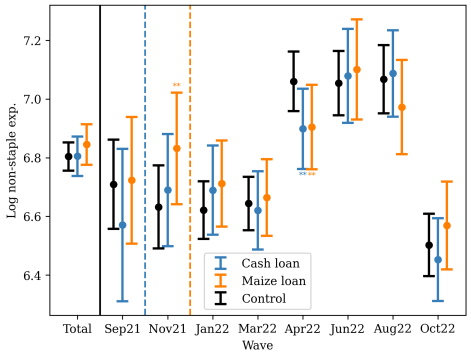
FIGURE A6. Treatment Effects on Stocks, by Wave, dropping unbalanced clusters due to attrition

Notes: Replicates Figure 8 dropping strata unbalanced by group-level attrition.

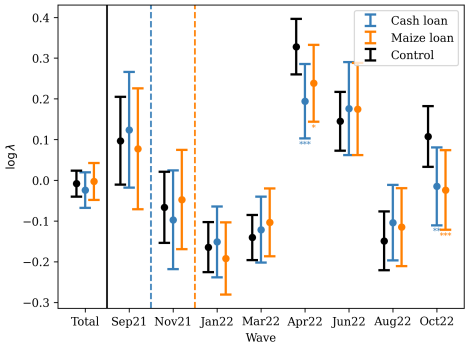
(A) Log food expenditure



(B) Log non-storable expenditure



(C) $\log \lambda$



(D) (Semi-)durable expenditure

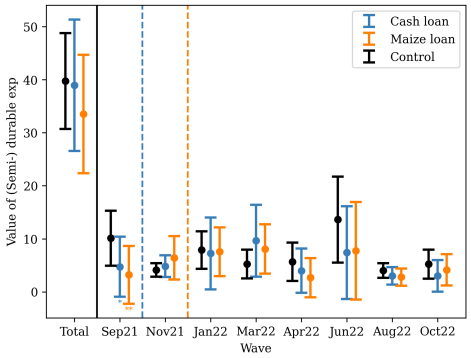
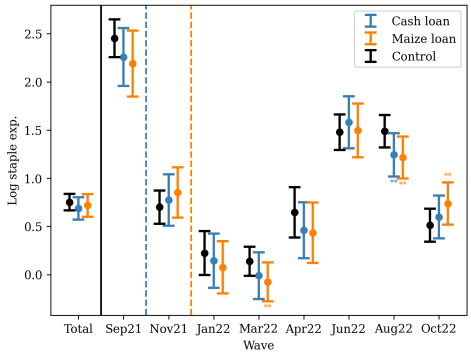


FIGURE A7. Treatment Effects on Expenditures, by wave, dropping unbalanced clusters due to attrition

(E) Log storable expenditure



(F) RCSI

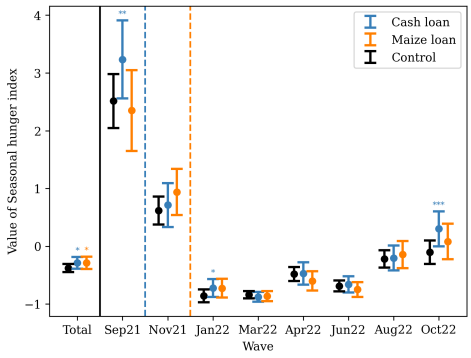


FIGURE A7. Treatment Effects on Expenditures, by wave, dropping unbalanced clusters due to attrition (continued)

Notes: Replicates Figure 9 dropping strata unbalanced by group-level attrition.

TABLE A8. Percent of target sample surveyed

Wave	Cash Loan	Maize Loan	Control
0	0.644	0.600	0.612
1	0.803	0.921	0.868
2	0.961	0.897	0.911
3	0.950	0.932	0.940
4	0.889	0.900	0.924
5	0.886	0.806	0.879
6	0.921	0.906	0.940
7	0.921	0.919	0.908
All (post-treatment)	0.903	0.898	.910

Probability of being surveyed across survey waves by treatment arm. Note that baseline was cut short due to time constraints

C.2. **Individual-Level Attrition.** Table A8 reports the share of the target sample successfully surveyed in each wave, by treatment arm. The baseline was cut short due to logistical constraints, and response rates are correspondingly lower (60–64%) and approximately balanced across arms. In subsequent rounds, response rates are consistently high (80–96%), with the cash arm typically 1–3 percentage points more likely to be surveyed than control. The in-kind (maize) arm tracks control closely on average, though individual-round differences fluctuate in sign.

Table A9 reports the results of regressing an indicator for being surveyed on baseline covariates and wave fixed effects, with standard errors clustered at the group level. Only baseline log food expenditure is a statistically significant predictor of being surveyed ($p < 0.01$); no other covariate—including household size, livestock holdings, farm area, or education—significantly predicts attrition.

Individual cases of attrition include four deaths, one refusal (a farmer who did not receive maize on time), and a small number of farmers who appear in only a subset of rounds.

C.3. **Lee Bounds.** Lee (2009) bounds address potential bias from differential attrition by trimming observations from the arm with higher response rates within each stratum-wave cell. We adapt the procedure to the tuple-wise Athey and Imbens (2017) estimator by trimming within each stratum–arm cell to match the baseline response rate before computing the usual cluster-weighted difference-in-means. Horizontal bars show the upper and lower bounds; circular markers show the point estimates. The bounds are tight, consistent with the balanced attrition documented above.

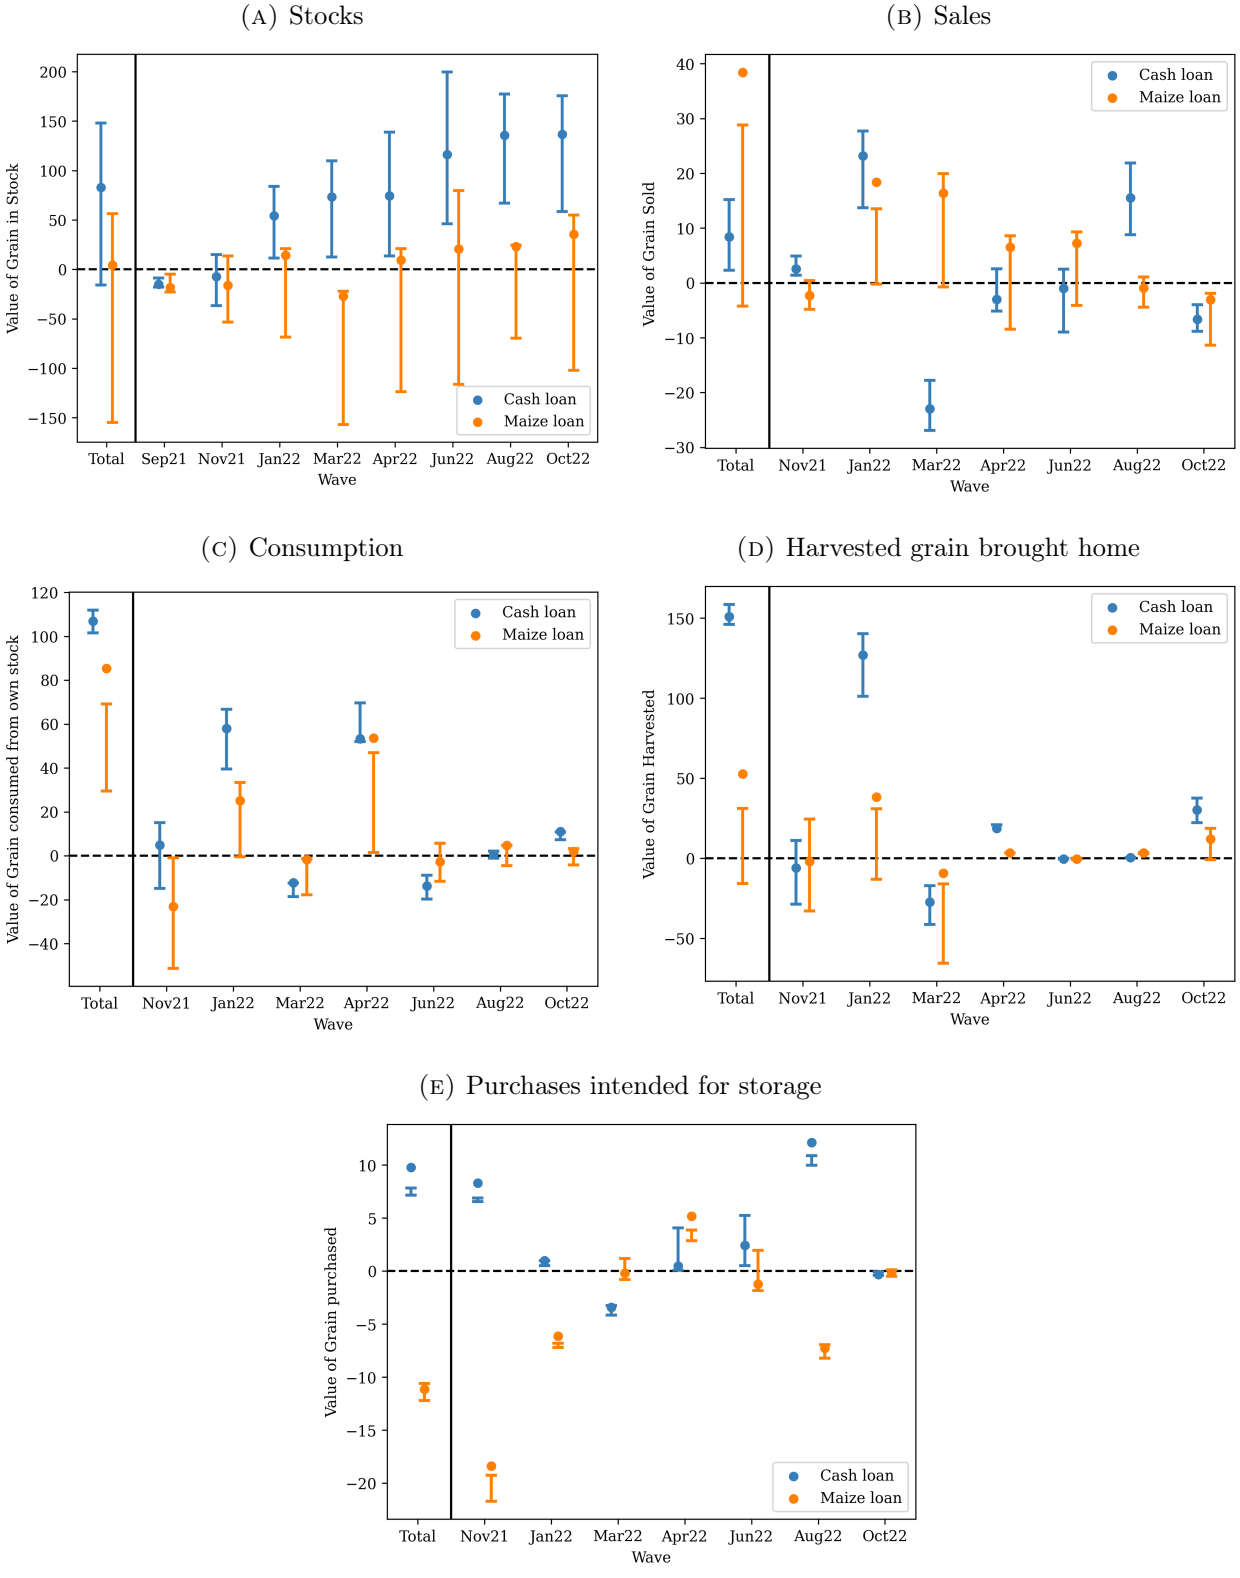


FIGURE A8. Lee Bounds on Treatment Effects on Stocks (Athey–Imbens Estimator)

TABLE A9. Predictors of attrition

$\mathbf{1}(\text{Surveyed}_{it})$			
hhsize	0.003	rainy rev	0.000
	(0.002)		(0.000)
age head	0.000	lexp	0.021***
	(0.001)		(0.004)
cows	-0.002	loglambda	-0.001
	(0.002)		(0.011)
goats	0.003	biz exp	0.000
	(0.001)		(0.000)
sheep	0.001	biz rev	0.000
	(0.002)		(0.000)
chickens	0.000	large item exp	-0.000
	(0.001)		(0.000)
any biz	-0.004	comppri	-0.018
	(0.022)		(0.038)
any borrow	-0.017	compsec	-0.042
	(0.019)		(0.046)
area ha	-0.003	none	-0.047
	(0.002)		(0.039)
ag exp	-0.000	somepri	-0.009
	(0.000)		(0.038)
female	0.019	N=	5368
	(0.067)		

Coefficients from a regression of a dummy for appearing in the sample on baseline covariates with wave fixed effects. Standard errors clustered at the farmer group level.

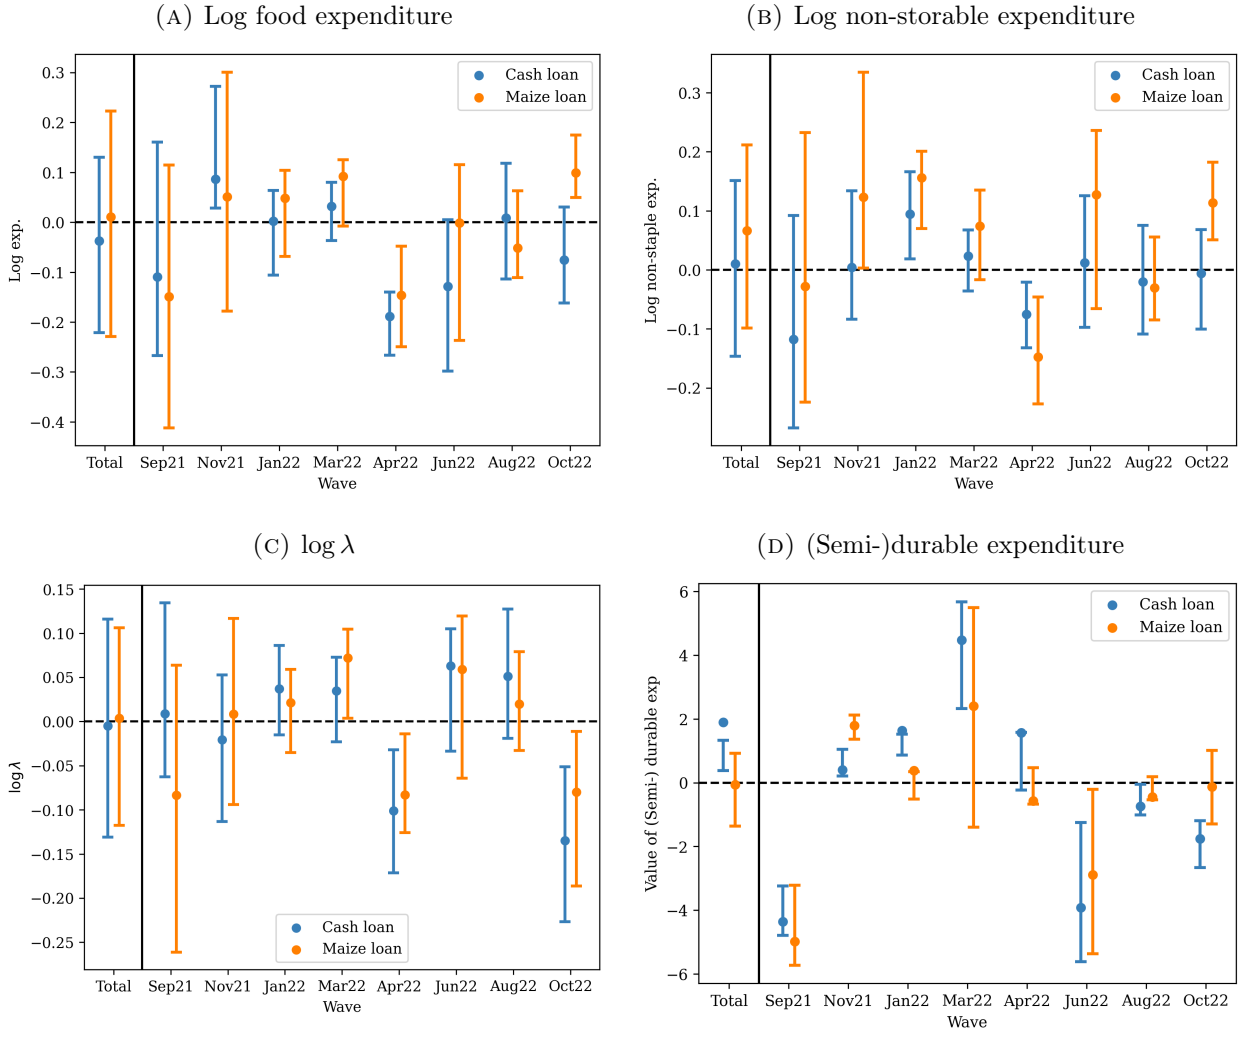
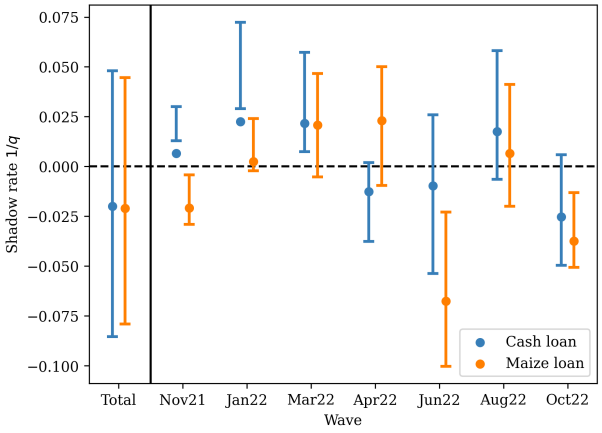
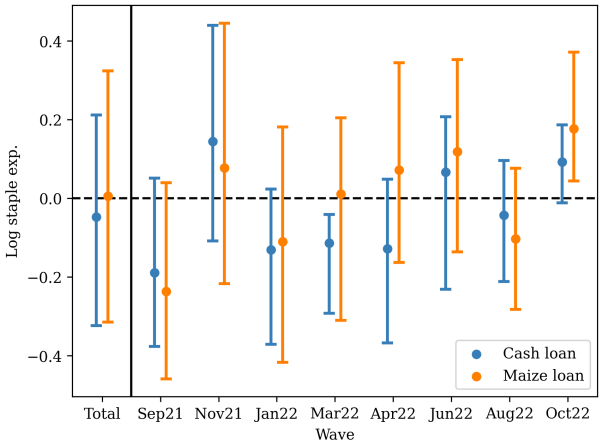


FIGURE A9. Lee Bounds on Treatment Effects on Expenditures (Athey–Imbens Estimator)

(E) Shadow rate



(F) Log storable expenditure



(G) RCSI

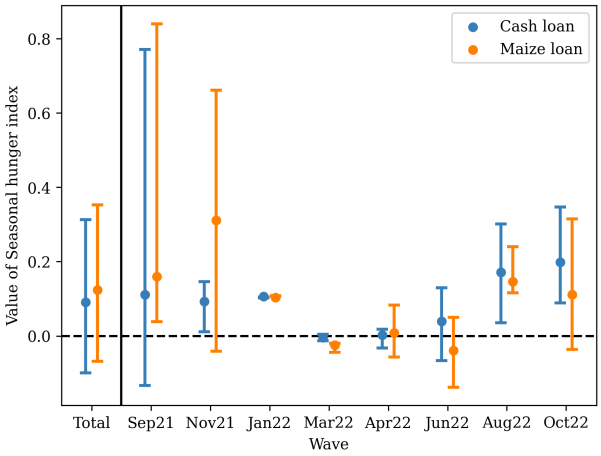
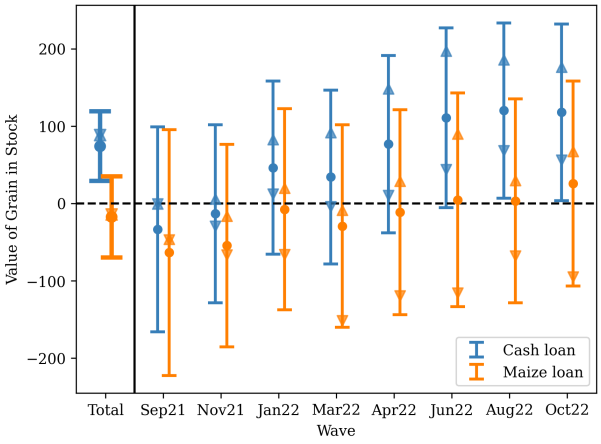


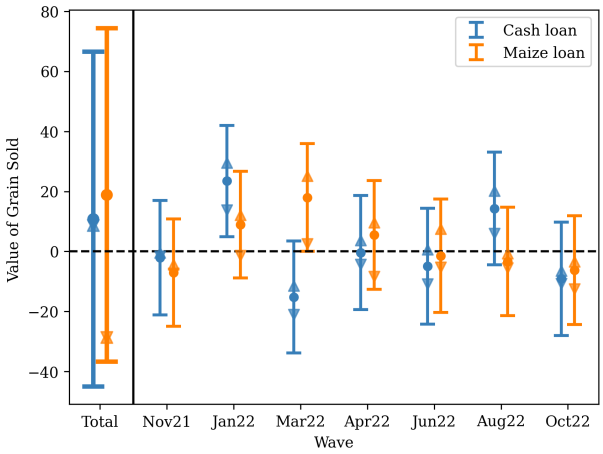
FIGURE A9. Lee Bounds on Treatment Effects on Expenditures (Athey–Imbens Estimator, continued)

Following Lee (2009), we also compute Lee bounds using the regression specification with strata fixed effects. Triangular markers show the upper and lower bounds; circular markers show the point estimates from Appendix E.3.1.

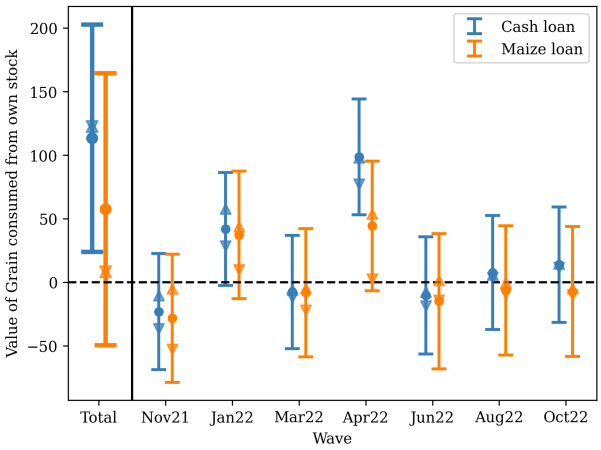
(A) Stocks



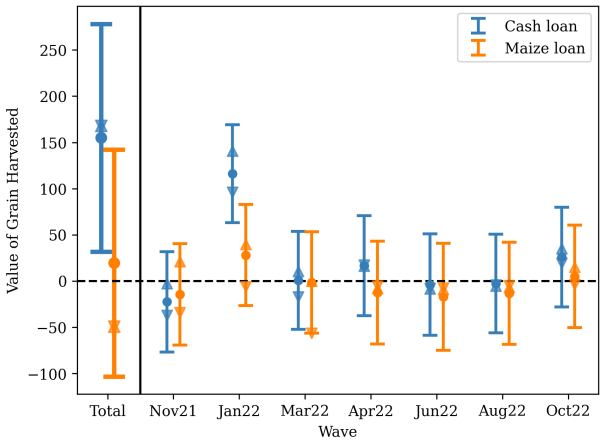
(B) Sales



(C) Consumption



(D) Harvested grain brought home



(E) Purchases intended for storage

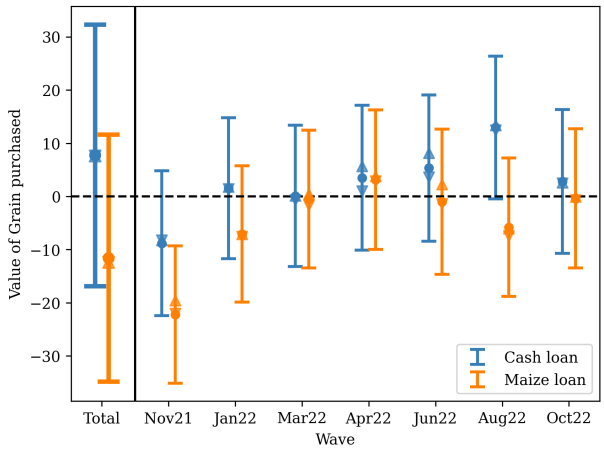
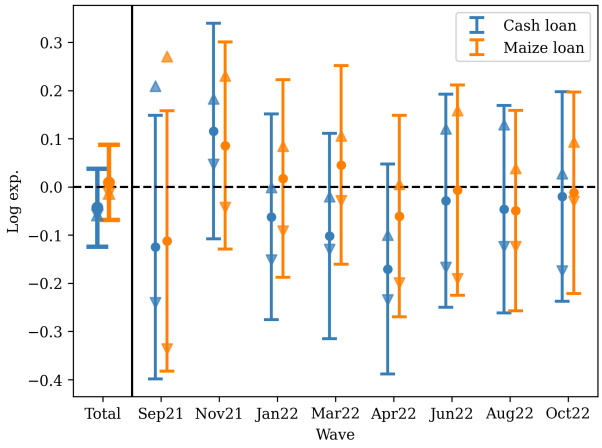
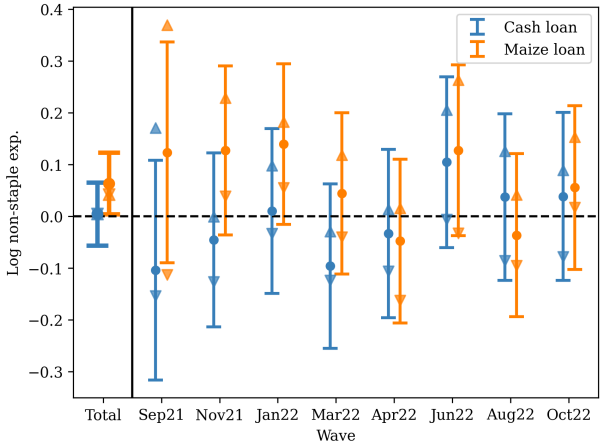


FIGURE A10. Lee Bounds on Treatment Effects on Stocks (Regression with Strata FE)

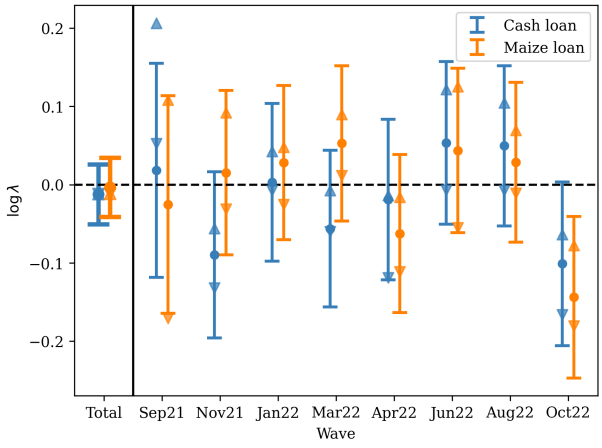
(A) Log food expenditure



(B) Log non-storable expenditure



(C) $\log \lambda$



(D) (Semi-)durable expenditure

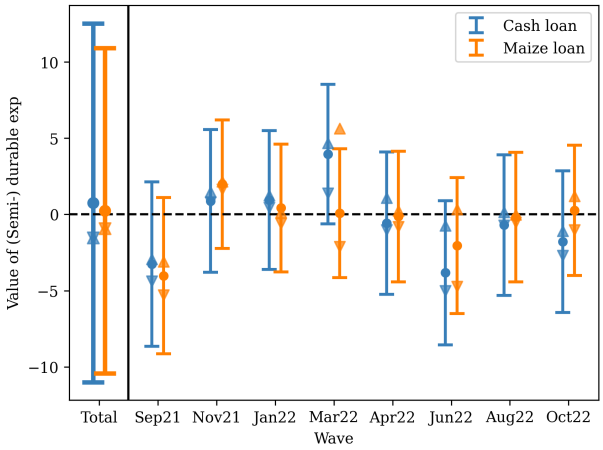
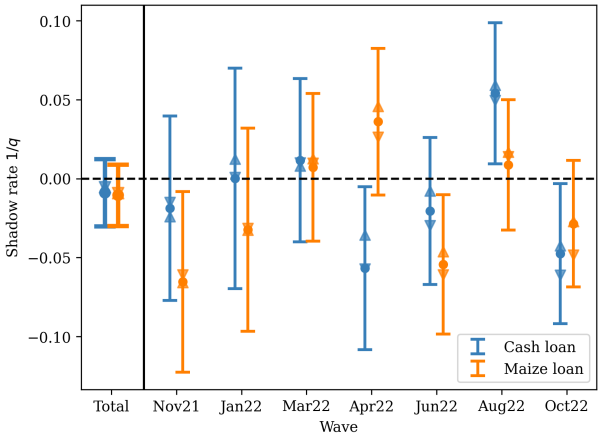
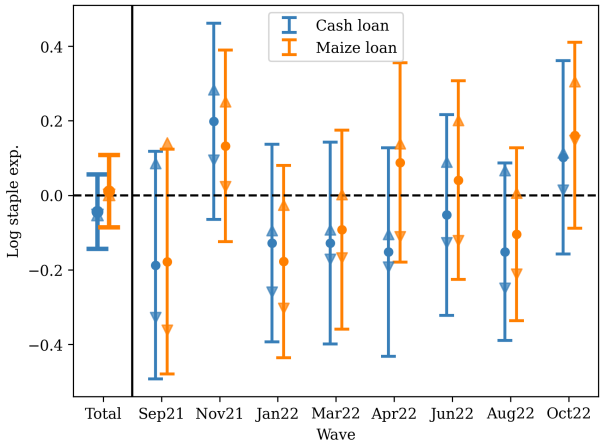


FIGURE A11. Lee Bounds on Treatment Effects on Expenditures (Regression with Strata FE)

(E) Shadow rate



(F) Log storable expenditure



(G) RCSI

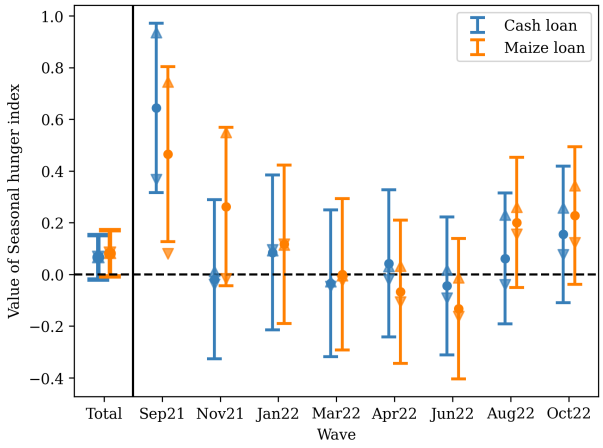


FIGURE A11. Lee Bounds on Treatment Effects on Expenditures (Regression with Strata FE, continued)

Appendix D. Descriptive Evidence

D.1. **Spatial Correlation in Grain Prices.** A central argument of the paper is that lean-season welfare losses in Gombe stem more from poor spatial market integration than from failures of local financial markets. This section provides direct evidence on the degree of price integration between Gombe and other markets, at both national and local scales.

D.1.1. *National Price Co-movement.* We use weekly wholesale price data from the Famine Early Warning Systems Network (FEWS NET), which covers 16 markets across Nigeria, to measure the correlation between log price changes in Gombe and every other monitored market. Figure A12 maps these correlations for white maize.

Price co-movement between Gombe and the rest of Nigeria is weak. The highest correlations are with nearby northeastern markets—Damaturu (0.37), Biu (0.22), and Mubi (0.22)—but even these are modest. Correlations with major commercial centers such as Lagos (0.07), Kano/Dawanau (0.09), and Ibadan (0.14) are close to zero. These patterns are consistent across other staple grains (millet, sorghum, cowpeas) and hold at both weekly and monthly frequencies.

The weakness of these price linkages implies that local supply shocks in Gombe—above all, the timing and size of the harvest—translate directly into local price movements with little dampening from arbitrage with distant markets. This is precisely the environment in which seasonal price variation is large and post-harvest lending is risky: the seasonal price rise that would make grain storage profitable is not guaranteed, because it depends on local conditions rather than on a predictable national or global price cycle.

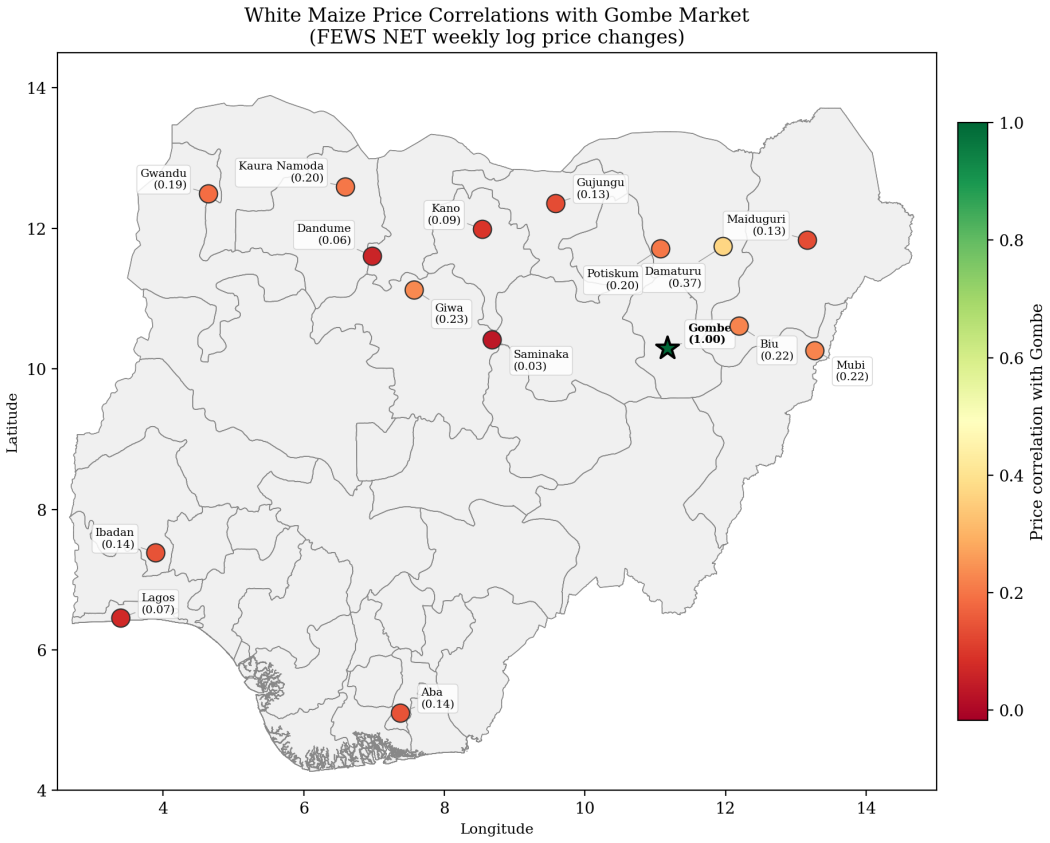


FIGURE A12. White maize price correlations with Gombe market (FEWS NET weekly log price changes)

D.1.2. *Local Price Co-movement.* We also examine whether markets *within* Gombe State move together, using price data collected in our own survey from eight rural markets alongside the FEWS NET Gombe reference price. Figure A13 maps the share of periods in which local and FEWS Gombe maize prices move in the same direction.

Co-movement within Gombe is higher than between Gombe and distant markets, but still imperfect. Most local markets agree with the FEWS Gombe price direction 50–70% of the time—better than the near-zero national correlations, but well short of the tight co-movement one would expect in a well-integrated local market. This suggests that even within Gombe State, spatial frictions—poor roads, transport costs, thin markets—limit price arbitrage.

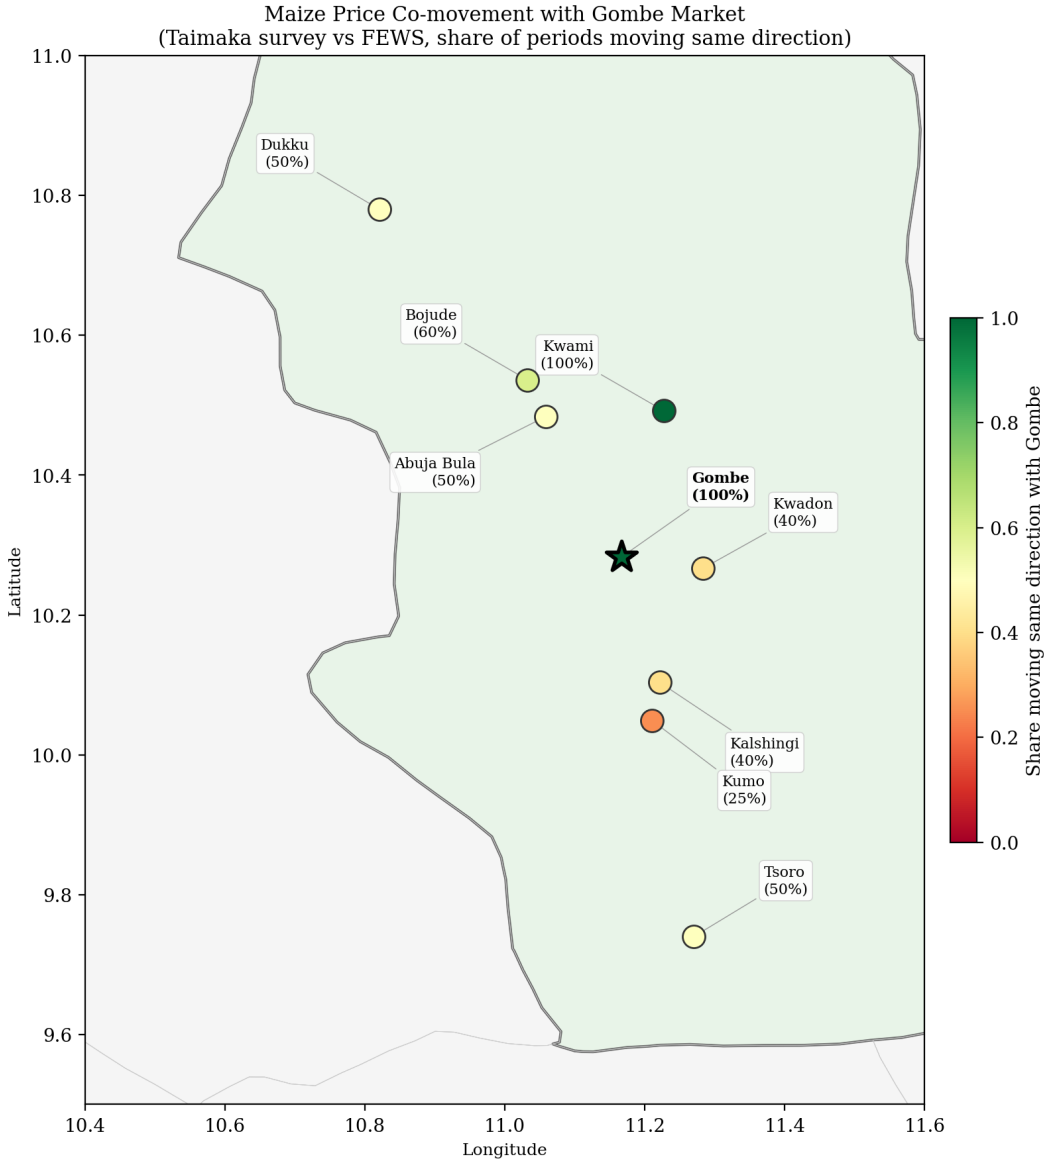


FIGURE A13. Maize price co-movement with Gombe market (share of periods moving in same direction)

D.2. **Household Socioeconomic Narrative.** :header-args:python: :exports results

(1) Comparison Data Sources

To contextualize our sample, we benchmark baseline characteristics against rural agricultural households in Northeast Nigeria from two sources: the 2018-19 Living Standards Measurement Study (LSMS, N=673) and the 2021 Nigeria Living Panel Survey (NLPS, N=320). The LSMS is a comprehensive face-to-face household survey that captures detailed monetary variables including agricultural expenditure, revenue, consumption, and livestock counts. Despite its detail, the LSMS data predate our baseline by three years. The NLPS is a phone-based panel survey conducted by the World Bank across multiple rounds during 2020-2023. Our NLPS extract draws on several rounds, not all of which fall within the Taimaka study period (August 2021 to November 2022). Demographics come from Round 1 (December 2021) and credit/borrowing from Round 4 (November 2022), both within our study window. Agriculture variables—main crop, cultivated area, fertilizer, herbicide, and hired labor—come from Round 9 (April 2023), but the agricultural questions refer to the "last completed agricultural season," i.e., the 2022 rainy season, which overlaps with our endline period. Education and livestock ownership come from Round 7 (June 2023); these are time-invariant or slow-moving characteristics unlikely to differ meaningfully from the study period. Food security measures (HDDS, FIES) also come from Round 9 and reflect conditions at the April 2023 survey date. Non-farm business data come from Round 11 (late 2023), which is further removed from our study period. As a phone survey, the NLPS also has important limitations in variable coverage: it does not collect monetary agricultural expenditure, revenue, or consumption data; livestock ownership is recorded as binary (owns any) rather than counts, precluding herd size comparisons; and crop participation captures only the respondent's main crop, making participation rates lower bounds on true cultivation. Education in the NLPS refers to the farmer respondent, who is not necessarily the household head. Despite these caveats, the NLPS offers a useful second benchmark—particularly for household demographics, crop choice, input use, and livestock ownership prevalence—with core demographic and financial variables drawn from within our study period and agricultural variables that retrospectively cover the relevant cropping season. Not all variables are measured identically across surveys (e.g., differences in recall periods, question wording, or crop year definitions). The comparisons we present with our sample are therefore suggestive rather than exact.

D.2.1. *Demographics and Household Composition.* Households in our sample are relatively large, reflecting extended family structures common in rural Nigeria. We collected detailed

roster information including age, sex, and relationship to household head for up to 20 members.

(1) Household Size and Structure

Table A10 summarizes the demographic characteristics of our sample at baseline, alongside means from the LSMS (2018–19) and NLPS (2021) for rural Northeast Nigeria. Household heads in our sample are younger (38 vs. 46 years in the LSMS and 49 years in the NLPS) and less formally educated: 64% have no schooling compared to 42% in the LSMS and 25% in the NLPS (though NLPS education refers to the farmer respondent, not necessarily the household head). Household size is comparable across the three samples (9.3 in Taimaka vs. 7.7 in the LSMS and 8.4 in the NLPS), and female headship is rare in all three (6–8%).

TABLE A10. Household Demographics at Baseline

	N	Mean	Std. Dev.	LSMS NE	NLPS NE
Household size	935	6.091	4.231	7.722	8.353
Age of HH head	908	38.272	12.43	46.385	48.914
Female HH head	935	0.062	0.241	0.082	0.061
No education	935	0.640	0.480	0.415	0.253
Some primary	935	0.079	0.270	0.049	0.012
Finished primary	935	0.097	0.297	0.22	0.134
Finished secondary	935	0.084	0.278	0.204	0.244
Tertiary education	935	0.042	0.200	0.051	0.194
Dependency ratio	900	1.035	1.071		

LSMS NE = rural agricultural households in NE Nigeria (LSMS 2018–19, N=673). NLPS NE = rural agricultural households in NE Nigeria (NLPS 2021, N=320). NLPS education refers to the farmer respondent, not necessarily the household head.

(2) Seasonal Changes in Household Composition

Household composition is not static—members join and leave throughout the year due to seasonal migration, marriage, fostering of children, and other life events. The monthly survey records household size (`hhsize`) and the number of members who joined (`hh_join`) or left (`hh_leave`) since the previous round. These flows are important because household size affects both labor supply and consumption needs, and seasonal migration is a common coping strategy in the Nigerian semi-arid zone.

Table A11 tracks mean household size, the share of households experiencing any joining or leaving events, and the mean net change in membership across rounds. Persistent churning—with both joiners and leavers present in most rounds—is common, even when mean household size remains relatively stable.

TABLE A11. Household Size and Composition Changes by Round

Round	Mean HH size	SD	Joiners	Leavers	Net change	N
Baseline (Aug-Sep 21)	6.1	4	0.000	0.000	0	600
Oct-Dec 21	6.3	4.1	0.051	0.023	-0.01	839
Dec 21-Feb 22	7.3	5	0.041	0.005	0.04	861
Feb-Apr 22	7.3	5.1	0.059	0.007	0.07	851
Apr-May 22	7.4	5.1	0.040	0.009	0.05	819
May-Jul 22	7.4	5	0.037	0.008	0.06	775
Jul-Sep 22	7.5	5.1	0.059	0.023	0.08	835
Endline (Sep-Nov 22)	7.5	5.1	0.019	0.004	0.02	829

D.2.2. *Agricultural Production.* The primary economic activity for households in our sample is rainfed agriculture. The main growing season runs from May to October (rainy season), with some households also engaging in dry season irrigation.

(1) Baseline Agricultural Characteristics

Table A12 summarizes crop portfolio, farm-level outcomes, and input use at baseline. Crop portfolios differ from the broader population: our sample grows more millet (56% vs. 29% in the LSMS and at least 13% in the NLPS) and Bambara nut (18% vs. 2% in the LSMS; the NLPS does not report Bambara nut), traditional crops well-suited to the semi-arid conditions of Gombe State. They grow less rice (8% vs. 26% in the LSMS and at least 11% in the NLPS) and groundnut (13% vs. 42% in the LSMS and at least 11% in the NLPS). Maize cultivation is similar across all three samples (71% Taimaka vs. 62% LSMS vs. at least 50% NLPS), consistent with its role as a regional staple. NLPS crop participation rates should be interpreted as lower bounds, since the survey captures only the respondent's main crop.

Our sample farms substantially larger holdings than the broader rural Northeast Nigeria population—mean cultivated area is roughly three times the LSMS mean (2.2 ha) and the NLPS mean (2.2 ha; notably identical across the two surveys). Agricultural expenditure and revenue are correspondingly higher: mean expenditure is more than double the LSMS average, and mean revenue roughly twice as high, reflecting both larger farm sizes and greater commercialization. The NLPS does not collect monetary agricultural data, but reports that only 16% of households sold any harvest versus 56% in the LSMS, consistent with lower commercialization in the broader population. Dry season cultivation is rare—about 1% of households report any irrigated production.

Fertilizer adoption in our sample is high and closer to the LSMS rate for Northeast Nigeria (62%) than to the NLPS rate (15%). Hired labor use in our sample also exceeds the NLPS rate (32%) but falls well below the LSMS rate (81%, mean 42.6

person-days, costing 49,500 Naira among users). LSMS data additionally show that only 8% of plots use improved seed, with 85% relying on saved seed from the previous season.

TABLE A12. Agricultural Characteristics at Baseline

	Mean	SD	LSMS	NLPS
Crop portfolio (share growing)				
Maize	0.711		0.618	0.500
Millet	0.564		0.285	0.131
Guinea Corn	0.353		0.565	0.119
Rice	0.079		0.257	0.106
Beans	0.586		0.421	0.006
Bambara nut	0.176		0.024	—
Groundnut	0.126		0.415	0.109
Farm-level outcomes				
Land owned (ha)	5.3	4.6		
Cultivated area (ha)	5.8	4.4	2.211	2.179
No. crops grown	3.5	1.3		
Ag. expenditure (N)	82646.1	97129.2	3.242e+04	
Ag. revenue (N)	103158.9	179856.3	4.438e+04	
Input use (share using)				
Fertilizer	0.895			0.153
Agrochemicals	0.662			0.634
Purchased seed	0.641			
Hired labor	0.471			0.322

NLPS crop rates are lower bounds (main crop only). LSMS/NLPS = rural agricultural HH in NE Nigeria.

(2) Crop-Level Outcomes

Table A13 reports crop-level harvest outcomes aggregated across the three post-harvest survey rounds (October 2021 to April 2022), which together span the main harvest reporting window. For each household, we sum reported harvests, sales, and revenues across these rounds to approximate full-season totals. Maize and millet are the dominant crops by harvest volume, while beans generate the highest per-seller revenue, reflecting their high market value.

D.2.3. *Livestock.* Table A14 summarizes livestock ownership at baseline.

D.2.4. *Grain Arbitrage.* The monthly survey asks households about grain purchased explicitly for storage, separately from grain purchased for immediate consumption, recording the quantity (in kg) and expenditure for each crop. This distinction is intended to capture arbitrage purchases—buying grain when prices are low (typically post-harvest) and storing it for later consumption or sale when prices rise. In practice, some households may not have clearly distinguished storage purchases from consumption purchases, so these figures should

TABLE A13. Crop-Level Harvest Outcomes (summed over post-harvest rounds, among growers)

Crop	Growers	Harvest (kg)	Sellers	Sold (kg)	Revenue (N)
Maize	601	2072	317	448	101472
Millet	479	1684	354	517	118244
Guinea Corn	385	375	89	164	44387
Rice	147	1851	70	533	79329
Beans	580	673	497	332	142870
Groundnut	51	740	20	272	28909
Bambara nut	113	814	80	491	94496

TABLE A14. Livestock Ownership at Baseline

	Own any	Mean	Mean (own.)	Median	LSMS	NLPS
Chickens	0.543	6.4	11.8	2	0.5	0.256
Goats	0.766	4.8	6.3	4	1.1	0.541
Sheep	0.517	3	5.9	1	0.3	0.300
Cows	0.489	2.2	4.6	0	0.3	0.359
Donkeys	0.014	0	1.2	0		

LSMS reports mean counts; NLPS reports binary ownership (% owning any).

be interpreted as an upper bound on deliberate arbitrage activity. Questions on grain flows were not asked at baseline; the data below cover rounds 1–7 only.

(1) Arbitrage Activity and Expenditure

Few households engage in grain storage purchases in any given round. Table A15 shows the share of households making any such purchase by round, along with the number of distinct crops purchased and total expenditure among purchasers. Table A16 breaks down activity by crop, pooled across all household-months.

TABLE A15. Grain Storage Purchases by Round

Round	Share purch.	N crops	Total exp. (N)	N
Oct-Dec 21	0.099	1.9	16729	839
Dec 21-Feb 22	0.071	1.4	85038	861
Feb-Apr 22	0.047	1.3	113062	851
Apr-May 22	0.037	1.4	195367	819
May-Jul 22	0.054	1.7	137414	775
Jul-Sep 22	0.184	1.5	50133	835
Endline (Sep-Nov 22)	0.000			829
Ever	0.351			922

Excludes baseline (grain flow questions not asked). N crops and total expenditure conditional on any purchase.

TABLE A16. Grain Storage Purchases by Crop (rounds 1–7)

Crop	HH-month share	Mean qty (kg)	Med. qty (kg)	Mean exp. (N)	N purch.
Maize	0.055	585	200	42119	317
Millet	0.009	495	300	88958	51
Guinea Corn	0.008	277	150	33776	44
Rice	0.017	1032	950	20587	94
Beans	0.021	245	200	72931	114
Groundnut	0.002	422	300	48688	9
Bambara nut	0.002	600	300	209570	9
Cassava	0.001	238	100	29125	4

D.2.5. *Income and Non-Farm Activities.*

(1) Seasonal Income Patterns

The overall income source composition reinforces the seasonal labor allocation story. The share of households reporting no off-farm income falls from 45% at baseline to 18% at before the 2022 harvest as households take on casual labor during the lean season. Likewise, casual labor participation rises from 15% at baseline to 48% in the peak pre-harvest period, while own-business income mirrors the dedicated business participation pattern.

TABLE A17. Income Sources by Round (share of HH)

Round	No off-farm	Casual labor	Own business	Okada	Farm labor	N
Baseline (Aug-Sep 21)	0.447	0.148	0.178	0.080	0.012	600
Oct-Dec 21	0.426	0.154	0.212	0.091	0.029	839
Dec 21-Feb 22	0.418	0.232	0.135	0.106	0.015	861
Feb-Apr 22	0.323	0.370	0.137	0.088	0.047	851
Apr-May 22	0.283	0.403	0.104	0.079	0.022	819
May-Jul 22	0.323	0.270	0.086	0.068	0.004	775
Jul-Sep 22	0.178	0.477	0.104	0.054	0.127	835
Endline (Sep-Nov 22)	0.344	0.368	0.138	0.053	0.183	829

(2) Off-Farm Income Sources

Each month, households report all sources of off-farm income, including casual labor, own-business activity, motorcycle taxi (okada) services, salaried employment, and remittances. Table A18 shows the prevalence of each income source pooled across all household-months. Casual labor is the most common off-farm activity. A large fraction of household-months report no off-farm income at all, reflecting the agricultural focus of this population.

TABLE A18. Off-Farm Income Sources (share of HH-months)

Source	Share of HH-months
No off-farm income	0.340
Casual labor	0.308
Own business	0.136
Motorcycle taxi (Okada)	0.078
Farm labor	0.057
Salaried employment	0.011
Informal wage work	0.017
Remittances	0.009
Other	0.110

(3) Seasonal Business Participation

Non-farm business participation declines over the agricultural season (Table A19), consistent with households reallocating labor toward farming during the growing season. At baseline, 18% of households report operating a non-farm business—substantially less than the 42% rate in both the LSMS and NLPS for rural Northeast Nigeria, suggesting these households are more specialized in agriculture and less diversified into non-farm activities. This rises slightly to 20% in Oct-Dec 2021 (the immediate post-harvest period, when cash is available to invest), then declines steadily to 6% by the peak farming season (May-July). Participation recovers modestly to 9-10% in the subsequent rounds. The most common business types are tailoring, carpentry, trading, and motorcycle taxi (okada) services.

TABLE A19. Non-Farm Business Participation by Round

Round	Share with business	N
Baseline (Aug-Sep 21)	0.177	600
Oct-Dec 21	0.197	839
Dec 21-Feb 22	0.122	861
Feb-Apr 22	0.130	851
Apr-May 22	0.094	819
May-Jul 22	0.063	775
Jul-Sep 22	0.089	835
Endline (Sep-Nov 22)	0.095	829

D.2.6. *Consumption and Food Security.*

(1) Food Acquisition: Own Production vs. Purchases

Households acquire staple foods through three channels: own production, market purchases, and gifts. Table A21 shows the share consuming, acquisition method,

TABLE A20. Business Revenue and Expenditure by Round (Naira, among operators)

Round	Mean rev.	Med. rev.	Mean exp.	Med. exp.	N
Baseline (Aug-Sep 21)	21539	10000	46113	11000	106
Oct-Dec 21	25693	12000	59241	15000	165
Dec 21-Feb 22	20603	10000	59543	20000	105
Feb-Apr 22	19171	10000	55181	20000	111
Apr-May 22	16160	10000	45591	20000	77
May-Jul 22	10822	4000	9043	0	49
Jul-Sep 22	18427	10000	63946	10000	74
Endline (Sep-Nov 22)	46987	15000	37620	0	79

and purchase expenditure by crop and round. For maize, the dominant staple, only 37% of consumers source from own production at baseline (just before harvest), with 60% purchasing. As stocks accumulate post-harvest, own production dominates—by February-April, over 90% of maize is home-sourced. During the lean season, households draw down reserves and return to market dependence, with only 38% consuming own produced maize by July-Sep 2022. Millet and Guinea corn follow similar patterns. Rice, which few sample households grow (Table A12), is almost entirely purchased year-round. Expenditure among purchasers mirrors these seasonal patterns: maize and millet purchase spending is highest at baseline and during the lean season when market reliance peaks, and lowest in the post-harvest months when own production covers most consumption. Consumption prevalence also varies seasonally: maize and millet are consumed by a large majority of households in most months, while Guinea corn and beans are concentrated in the post-harvest period when stocks are available.

(2) Food Expenditure

Total food expenditure captures spending on all purchased foods (Table A22). At baseline, median food spending is 5,000 Naira, reflecting high market reliance just before harvest. This drops sharply as households consume own production—median expenditure falls to 1,600 Naira in February-April. It then rises as stocks deplete, peaking at 3,550 Naira at the peak of the lean season.

(3) Large Non-Food Expenditures

Expenditure on large items (purchases exceeding 10,000 Naira, such as building materials, clothing, or durable goods) is widespread: 50-64% of households report such expenditure in any given round. The median amount is relatively stable at 10,000-15,000 Naira, with a slight upward drift toward the endline.

(4) Reduced Coping Strategies Index

TABLE A21. Food Acquisition and Expenditure by Crop and Round

Crop	Round	Cons.	Own	Purch.	Mean exp.
Maize	Baseline (Aug-Sep 21)	0.950	0.372	0.596	12016
	Oct-Dec 21	0.825	0.744	0.237	9377
	Dec 21-Feb 22	0.787	0.879	0.102	14402
	Feb-Apr 22	0.825	0.923	0.073	5811
	Apr-May 22	0.951	0.899	0.095	7888
	May-Jul 22	0.876	0.710	0.283	17004
	Jul-Sep 22	0.910	0.387	0.604	4786
	Endline (Sep-Nov 22)	0.807	0.565	0.425	5457
Millet	Baseline (Aug-Sep 21)	0.675	0.383	0.560	8412
	Oct-Dec 21	0.697	0.697	0.270	3035
	Dec 21-Feb 22	0.652	0.738	0.232	2320
	Feb-Apr 22	0.656	0.819	0.158	878
	Apr-May 22	0.786	0.789	0.191	1424
	May-Jul 22	0.582	0.543	0.412	1897
	Jul-Sep 22	0.502	0.399	0.566	1975
	Endline (Sep-Nov 22)	0.776	0.672	0.323	2538
Guinea Corn	Baseline (Aug-Sep 21)	0.485	0.254	0.680	9133
	Oct-Dec 21	0.354	0.697	0.259	3394
	Dec 21-Feb 22	0.574	0.901	0.061	1311
	Feb-Apr 22	0.537	0.858	0.081	1164
	Apr-May 22	0.556	0.897	0.095	2090
	May-Jul 22	0.601	0.519	0.442	4914
	Jul-Sep 22	0.408	0.238	0.724	2427
	Endline (Sep-Nov 22)	0.404	0.403	0.588	3768
Rice	Baseline (Aug-Sep 21)	0.508	0.105	0.879	3697
	Oct-Dec 21	0.734	0.231	0.758	2764
	Dec 21-Feb 22	0.833	0.215	0.775	2378
	Feb-Apr 22	0.803	0.286	0.698	1243
	Apr-May 22	0.668	0.298	0.691	2485
	May-Jul 22	0.674	0.125	0.862	2049
	Jul-Sep 22	0.668	0.027	0.964	2470
	Endline (Sep-Nov 22)	0.721	0.132	0.860	2813
Beans	Baseline (Aug-Sep 21)	0.670	0.172	0.803	4378
	Oct-Dec 21	0.754	0.556	0.425	1803
	Dec 21-Feb 22	0.786	0.746	0.241	1071
	Feb-Apr 22	0.800	0.742	0.247	465
	Apr-May 22	0.755	0.843	0.149	2256
	May-Jul 22	0.692	0.364	0.604	846
	Jul-Sep 22	0.851	0.224	0.764	1594
	Endline (Sep-Nov 22)	0.799	0.346	0.651	1518

Expenditure in Naira, conditional on purchasing. Acquisition shares among consumers of each crop in each round.

The survey includes the Reduced Coping Strategies Index (rCSI), which measures food insecurity through five coping behaviors in the past 7 days: relying on less preferred foods, borrowing food, limiting portion sizes, restricting adult consumption

TABLE A22. Total Food Expenditure by Round (Naira)

Round	Mean	Median	Std. Dev.	N
Baseline (Aug-Sep 21)	18381	3815	32586	600
Oct-Dec 21	6143	2365	13229	839
Dec 21-Feb 22	4410	1730	10644	861
Feb-Apr 22	2515	1600	5094	851
Apr-May 22	3926	1800	10551	819
May-Jul 22	8867	2350	25155	775
Jul-Sep 22	8041	3500	11801	835
Endline (Sep-Nov 22)	7241	2350	11592	829

TABLE A23. Food Expenditure by Purchased Item (Naira, among purchasers)

Item	N	Mean	Median
Maize	1632	8877	3000
Millet	1355	3162	900
Guinea Corn	1037	4431	1500
Rice	3650	2422	1000
Beans	2315	1789	600
Groundnut	979	520	200
Chicken	97	1699	1500
Beef	2049	799	500
Fish (fresh)	237	457	400
Fish (dried)	100	593	500
Eggs	299	379	250
Onions	3927	173	100
Tomatoes	2876	210	150
Greens	1281	160	100
Bread	4176	327	200
Sugar	5404	215	100
Milk	2227	208	100

for children, and reducing the number of meals. These are weighted according to the standard rCSI protocol (weights 1, 2, 1, 3, 1 respectively).

Some enumerators recorded zeros for all rCSI components in certain rounds, rendering those observations unreliable. Following the approach in our data notes, we drop these enumerators when computing rCSI statistics. Table A26 shows the mean of each component by round.

(5) Food Security in the NLPS

TABLE A24. Large Non-Food Expenditure by Round (Naira, among those with expenditure)

Round	Share	Mean	Median	N
Baseline (Aug-Sep 21)	0.642	24355	10000	385
Oct-Dec 21	0.561	17904	10000	471
Dec 21-Feb 22	0.510	25432	10000	439
Feb-Apr 22	0.550	29537	10000	468
Apr-May 22	0.569	19576	10000	466
May-Jul 22	0.563	21236	10000	436
Jul-Sep 22	0.545	29290	12500	455
Endline (Sep-Nov 22)	0.595	33373	15000	493

TABLE A25. Large Non-Food Expenditure by Category (pooled across rounds)

Category	N	Share	Mean (N)	Median (N)
Construction/building	719	0.112	17189	10000
Clothing	467	0.073	10821	8000
Farm equipment	844	0.132	37522	20000
Livestock purchase	238	0.037	34204	15000
Household items	47	0.007	17004	10000
Festival/ceremony	421	0.066	49645	20000
Other	887	0.138	13797	8000

Category based on primary (first-listed) large item type. Expenditure in Naira among those reporting the category.

TABLE A26. rCSI Components by Round (mean days in past 7, excluding bad enumerators)

Round	Less pref.	Borrow	Limit port.	Restrict	Fewer meals	N
Baseline (Aug-Sep 21)	10.09	4.45	7.39	8.58	4.18	600
Oct-Dec 21	4.98	2.19	2.43	2.55	1.2	839
Dec 21-Feb 22	0.47	0.28	0.35	0.36	0.12	861
Feb-Apr 22	0.54	0.19	0.18	0.09	0.02	851
Apr-May 22	1.33	0.61	0.67	0.44	0.17	819
May-Jul 22	1.34	0.72	0.62	0.42	0.17	775
Jul-Sep 22	2.79	1.94	1.63	1.36	0.68	835
Endline (Sep-Nov 22)	3.17	1.77	1.65	1.64	0.62	829

The NLPS uniquely captures two food security measures not available in the Taimaka or LSMS surveys. The Household Dietary Diversity Score (HDDS), measured on a 0-9 scale, averages 7.2 (N=169) among rural Northeast agricultural households, indicating consumption from roughly seven of nine food groups. The Food

Insecurity Experience Scale (FIES), measured on a 0-8 scale, averages 3.9 (N=151), and 82% of respondents report at least some food insecurity. These figures provide important context for the levels of food stress in the region. Although the Taimaka survey does not include the HDDS or FIES, it captures related information through the Reduced Coping Strategies Index (rCSI), which measures behavioral responses to food insecurity (Table A26). The high prevalence of food insecurity in the NLPS is consistent with the elevated rCSI scores observed in our sample.

D.2.7. *Transfers and Gifts.* Informal transfers play a negligible role for these households. Table A27 shows the share of households receiving food gifts or remittances in each round. Both are rare: fewer than 5% of households receive any food as a gift in a typical month, and remittance receipt is similarly uncommon throughout the study period. This is consistent with the limited urban connections and low commercialization of this rural population, and underscores the absence of informal safety nets that the Taimaka program aims to substitute for.

TABLE A27. Transfers Received by Round (share of HH)

Round	Any food gift	Remittances	N
Baseline (Aug-Sep 21)	0.062	0.003	600
Oct-Dec 21	0.067	0.006	839
Dec 21-Feb 22	0.058	0.009	861
Feb-Apr 22	0.079	0.012	851
Apr-May 22	0.033	0.018	819
May-Jul 22	0.081	0.003	775
Jul-Sep 22	0.049	0.008	835
Endline (Sep-Nov 22)	0.033	0.011	829

D.2.8. *Credit and Savings.* The monthly survey tracks borrowing through several variables: a binary indicator for whether the household currently has an outstanding loan (`borrowing`), the loan amount (`borrow_amt` at baseline/endline, `borrow_val` in monthly rounds), and the amount repaid during the recall period (`borrow_repay`). Formal borrowing is limited: only 11% of LSMS Northeast households report having a loan, and an even lower rate in our sample. The NLPS reports a substantially higher borrowing prevalence of 49%, though this likely reflects differences in question framing—the NLPS may capture informal credit arrangements (e.g., borrowing from family or neighbors) that the LSMS and Taimaka surveys do not. The low rate in our sample is consistent with the well-documented shallow penetration of formal credit in rural Nigeria, and helps motivate the Taimaka intervention itself.

(1) Borrowing and Repayment by Round

Borrowing prevalence follows a seasonal pattern (Table A28). Borrowing is highest at baseline, when some households carry pre-existing loans into the study period, and lowest during the post-harvest months when crop sales provide cash. Savings are also rare: fewer than 10% of households report any savings at baseline. Repayment activity mirrors this seasonal cycle, declining during the lean months when cash is scarce and picking up again post-harvest.

TABLE A28. Borrowing and Repayment by Round

Round	Share borrow	N borrow	Mean loan	Share repay	N
Baseline (Aug-Sep 21)	0.272	160	23118	0.267	600
Oct-Dec 21	0.191	160	28609	0.191	839
Dec 21-Feb 22	0.064	55	28535	0.064	861
Feb-Apr 22	0.066	56	11777	0.066	851
Apr-May 22	0.074	60	11387	0.061	819
May-Jul 22	0.052	40	7942	0.050	775
Jul-Sep 22	0.127	105	18171	0.126	835
Endline (Sep-Nov 22)	0.124	102	42643	0.123	829
Ever	0.514	481	24070		935

Implied interest rate (total repaid / total borrowed − 1, among 480 HHs with both): mean 2.6%, median 0.0%.

Appendix E. Supplementary Pre-Specified Analyses

This section reports results from all pre-specified analyses that do not appear in the main text. We begin with a concordance table mapping the pre-analysis plan to the paper, then present each supplementary analysis in turn.

E.1. **Pre-Analysis Plan Concordance.**

E.1.1. *Treatment Effects — Main Specifications.* The tables below map each pre-specified outcome family to the corresponding results.

(1) Consumption and Welfare

PAP item	Main paper location
Log marginal utility of exp.	Figure 7 (panel), Table B, Figure 9
Log consumption expenditure	Table B, Figure 9
Log non-storable expenditure	Table B, Figure 9
Log storable expenditure	Figure 9
Elicited IMRS	Table B, Figure 9
RCSI (hunger index)	Table B, Figure 9
Large purchases	Figure 9
Business exp./revenue	Figure 10

(2) Grain Stocks and Flows

PAP item	Main paper location
Stocks (cumulative flows)	Figure 8 panel (a)
Sales	Figure 8 panel (b), Table B, Figure A25
Consumption from own stock	Figure 8 panel (c), Table B
Harvests (grain brought home)	Figure 8 panel (d)
Purchases for storage (arbitrage)	Figure 8 panel (e)
No. crops consumed from own stock	Figure A4
No. crops purchased	Figure A5

(3) Other Investments

PAP item	Main paper location
Agricultural inputs	Section 5, Agricultural Outcomes
Farm output / profit	Section 5, Agricultural Outcomes
Dry-season farming	Section 5, Agricultural Outcomes
Non-agricultural business	Figure 10

Continued on next page

Continued from previous page

PAP item	Main paper location
Amounts borrowed	Figure 10

E.1.2. *Treatment Effects — Robustness.* The PAP specified three control strategies for the main treatment-effect regressions:

(1) **No controls (strata fixed effects only).** This is the main specification, implemented via the Athey and Imbens (2017) estimator (Figures 8, 9).
(2) **Taimaka credit score as control.** Implemented as a supplementary check in Section E.4. Because the randomization was stratified on the Taimaka score, the preferred Athey and Imbens (2017) estimator absorbs this variation by construction; the score-controlled specification uses enumerator and wave fixed effects instead and confirms that results are robust.
(3) **Double Post-LASSO** (Belloni et al. 2012). Implemented in Figures A16 and A17.

The paper also reports robustness checks not pre-specified in the PAP:

- **ANCOVA** (controlling for baseline outcome): Figure A18.
- **Dropping unbalanced clusters**: Figures A6, A7.
- **Regression with strata fixed effects**: Figures A14, A15.

E.1.3. *Heterogeneity.* The PAP specified interaction terms with two variables:

PAP interaction	Consumption outcomes	Grain sales outcomes
Gender of household head	Table A36	Table A37
Baseline wealth index	Table A38	Table A39

These are reported in Appendix B of the main paper.

E.1.4. *Balance and Attrition.*

PAP item	Main paper location
Covariate balance	Table A7
Attrition rates	Table A8
Predictors of attrition	Table A9
Lee bounds	Section C.3

E.1.5. *Structural Estimation.*

PAP item	Main paper location

Continued on next page

Continued from previous page

PAP item	Main paper location
Euler equation / asset pricing test	Table 7 (Section 7)
Full insurance test	Section 6
Counterfactual policy simulations	Not implemented (see Section E.1.6)

The Euler equation test uses the continuously-updated GMM estimator of Hansen, Heaton, and Yaron (1996) to test conditional moment restrictions from the household's first-order conditions. The full insurance test is presented in Section 6.

E.1.6. *Deviations from the Pre-Analysis Plan.* The following items were pre-specified but are not reported in the main text. We describe the status of each.

PAP item	Status
Variance of outcomes	Not reported. Dominated by measurement error in high-frequency consumption data; not informative about treatment effects on risk.
Amounts lent	Not reported. Not collected in the survey instrument.
Counterfactual simulations	Not reported. Noted in the PAP as a longer-term objective (Research Question 3). The Euler equation results in Section 7 provide the key structural evidence; full counterfactual simulations under alternative price distributions are left as future work.

TABLE A36. Heterogeneous effects on consumption outcomes, by gender

	Log expenditure	Log perishable exp.	−RCSI	$\log \lambda$	Shadow rate $1/q$
Cash Loan	-0.041	0.008	0.094*	-0.014	-0.011
	(0.067)	(0.042)	(0.051)	(0.030)	(0.008)
Cash Loan ×Female Head	-0.062	-0.146	-0.213*	-0.068	0.043
	(0.119)	(0.208)	(0.110)	(0.082)	(0.027)
Maize Loan	-0.004	0.048	0.085	-0.003	-0.015*
	(0.061)	(0.043)	(0.053)	(0.028)	(0.008)
Maize Loan ×Female Head	0.063	0.081	0.002	-0.060	0.018
	(0.165)	(0.179)	(0.187)	(0.063)	(0.025)
Female Head	-0.277**	-0.305***	0.116	-0.061*	0.001
	(0.111)	(0.085)	(0.075)	(0.032)	(0.011)
Any Loan	-0.021	0.030	0.089*	-0.008	-0.013*
	(0.058)	(0.039)	(0.047)	(0.026)	(0.007)
Loan ×Female Head	-0.016	-0.056	-0.127	-0.067	0.031
	(0.120)	(0.141)	(0.128)	(0.057)	(0.022)
Female Head	-0.270**	-0.293***	0.124	-0.060*	-0.000
	(0.110)	(0.078)	(0.080)	(0.033)	(0.011)
Fixed Effects	Strata	Strata	Strata	Strata	Strata
Control mean	7.865	6.854	-0.378	0.041	1.268
N	5733	5714	4021	5431	2760

E.2. **Heterogeneous Effects.** Here we report the results on main consumption related outcomes and grain sales interacting the treatments with sex of the household head and a baseline asset index, as pre-specified.

TABLE A37. Heterogeneous effects on grain sales, by gender

	Maize	Beans	Guinea Corn	Rice	Millet	All
Cash Loan	14.68	-2.01	-2.96	-0.80	1.86	14.50
	(11.62)	(23.72)	(2.30)	(3.97)	(22.41)	(36.95)
Cash Loan ×Female Head	-28.07	60.74	-2.40	0.14	17.78	37.03
	(27.74)	(37.55)	(4.30)	(4.94)	(30.42)	(69.91)
Maize Loan	13.04	-7.08	1.40	2.41	1.17	21.51
	(12.02)	(25.63)	(1.52)	(6.29)	(15.35)	(49.57)
Maize Loan ×Female Head	-27.17*	-15.34	-6.91	-1.78	-20.25	-86.82
	(15.08)	(36.83)	(4.48)	(6.77)	(21.05)	(68.77)
Female Head	7.39	-30.44	0.79	-3.04	-5.22	-29.08
	(12.57)	(20.80)	(4.37)	(4.38)	(13.66)	(39.88)
Any Loan	13.79	-4.90	-0.59	0.95	1.42	18.12
	(10.46)	(21.27)	(1.67)	(4.65)	(13.97)	(35.76)
Loan ×Female Head	-27.52*	21.85	-4.96	-1.02	-1.76	-27.11
	(15.83)	(33.73)	(4.08)	(5.04)	(18.61)	(53.66)
Female Head	7.36	-31.90	0.84	-2.99	-5.90	-31.02
	(12.43)	(20.63)	(4.35)	(4.41)	(13.63)	(40.22)
Fixed Effects	Strata	Strata	Strata	Strata	Strata	Strata
Control Mean	52.486	145.566	6.988	12.557	68.729	295.292
N	930	930	930	930	930	930

TABLE A38. Heterogeneous effects on consumption outcomes, by baseline wealth

	Log expenditure	Log perishable exp.	−RCSI	$\log \lambda$	Shadow rate $1/q$
Cash Loan	-0.040	0.010	0.057	-0.022	-0.007
	(0.068)	(0.042)	(0.046)	(0.027)	(0.008)
Cash Loan ×Bl Asset Index	0.040	0.022	0.079*	-0.026	-0.002
	(0.049)	(0.033)	(0.043)	(0.016)	(0.005)
Maize Loan	0.016	0.069	0.081	-0.006	-0.016*
	(0.060)	(0.046)	(0.049)	(0.028)	(0.008)
Maize Loan ×Bl Asset Index	0.043	-0.020	0.105**	-0.008	0.006
	(0.034)	(0.032)	(0.040)	(0.017)	(0.004)
Bl Asset Index	-0.004	0.046	-0.067*	0.020	-0.004
	(0.039)	(0.031)	(0.033)	(0.015)	(0.004)
Any Loan	-0.010	0.040	0.070	-0.014	-0.012*
	(0.057)	(0.038)	(0.043)	(0.024)	(0.007)
Loan ×Bl Asset Index	0.044	0.000	0.094***	-0.016	0.002
	(0.039)	(0.031)	(0.031)	(0.016)	(0.004)
Bl Asset Index	-0.006	0.043	-0.067**	0.019	-0.003
	(0.039)	(0.030)	(0.033)	(0.015)	(0.004)
Fixed Effects	Strata	Strata	Strata	Strata	Strata
Control mean	7.865	6.854	-0.378	0.041	1.268
N	5733	5714	4021	5431	2760

TABLE A39. Heterogeneous effects on grain sales, by baseline wealth

	Maize	Beans	Guinea Corn	Rice	Millet	All
Cash Loan	10.35	2.19	-2.76	-0.93	-0.79	9.75
	(12.27)	(23.48)	(2.21)	(3.81)	(21.36)	(36.94)
Cash Loan ×Bl Asset Index	-4.66	-25.25	-0.26	-4.54	-1.34	-40.40
	(10.78)	(17.23)	(1.95)	(9.30)	(8.77)	(28.36)
Maize Loan	14.34	-0.66	1.59	1.84	2.13	29.09
	(11.18)	(24.27)	(1.30)	(5.88)	(14.14)	(46.31)
Maize Loan ×Bl Asset Index	-3.85	-17.98	0.11	-3.25	-12.71*	-42.85
	(9.83)	(19.40)	(1.92)	(9.73)	(7.15)	(29.63)
Bl Asset Index	3.48	28.84	0.16	2.61	20.79***	61.45**
	(10.05)	(18.79)	(2.10)	(9.49)	(6.41)	(28.74)
Any Loan	12.55	0.71	-0.36	0.61	0.68	20.36
	(10.24)	(20.65)	(1.55)	(4.42)	(12.97)	(34.64)
Loan ×Bl Asset Index	-4.13	-21.35	0.04	-3.78	-7.47	-41.31*
	(8.97)	(17.20)	(1.89)	(9.36)	(4.91)	(22.98)
Bl Asset Index	3.37	28.98	0.04	2.54	20.61***	60.87**
	(9.96)	(18.83)	(2.12)	(9.53)	(6.52)	(28.92)
Fixed Effects	Strata	Strata	Strata	Strata	Strata	Strata
Control Mean	52.486	145.566	6.988	12.557	68.729	295.292
N	930	930	930	930	930	930

E.3. **Robustness.** The following subsections reproduce all main treatment-effect figures under each alternative specification discussed in Section 5.6.

E.3.1. *Regression-Based Estimates.* In our pre-analysis plan, we pre-specified the tuple-wise randomization procedure described in Athey and Imbens (2017). The authors prescribe a simple procedure for estimation and inference in such settings — essentially a weighted average of within-cluster differences-in-means and variances. However, we also pre-specified a regression-based approach, controlling for strata fixed effects and clustering standard errors at the stratum level, following De Chaisemartin and Ramirez-Cuellar (2024). However, we prefer the former approach, given its direct link to the experimental design and the potential restrictiveness of the regression-based approach with only about 30 clusters. In this appendix, we reproduce our results using the regression-based approach. For the most part, the point estimates are qualitatively similar to those obtained using the preferred Athey and Imbens (2017) approach, but with wider confidence intervals.

E.3.2. *Double Post-LASSO.* As an alternative to the stratification-based estimator, we re-estimate treatment effects using the Double Post-LASSO (DPL) procedure of Belloni et al. (2012), which selects controls from a large candidate set via penalized regression. Figures A16 and A17 show that results are robust to this approach.

E.3.3. *ANCOVA.* We also check robustness to an ANCOVA specification that controls for the baseline (wave 1) value of each outcome. Figure A18 shows that the main results are qualitatively unchanged.

E.4. **Controlling for Taimaka Credit Score.** The PAP pre-specified the Taimaka group credit score as a control variable. Because the randomization was stratified on the score, the preferred Athey and Imbens (2017) estimator absorbs this variation by construction. As a supplementary check, Figures A19 and A20 replicate the main treatment-effect estimates in a regression specification that adds the score as an explicit control, with enumerator and wave fixed effects replacing strata fixed effects. The results are qualitatively unchanged.

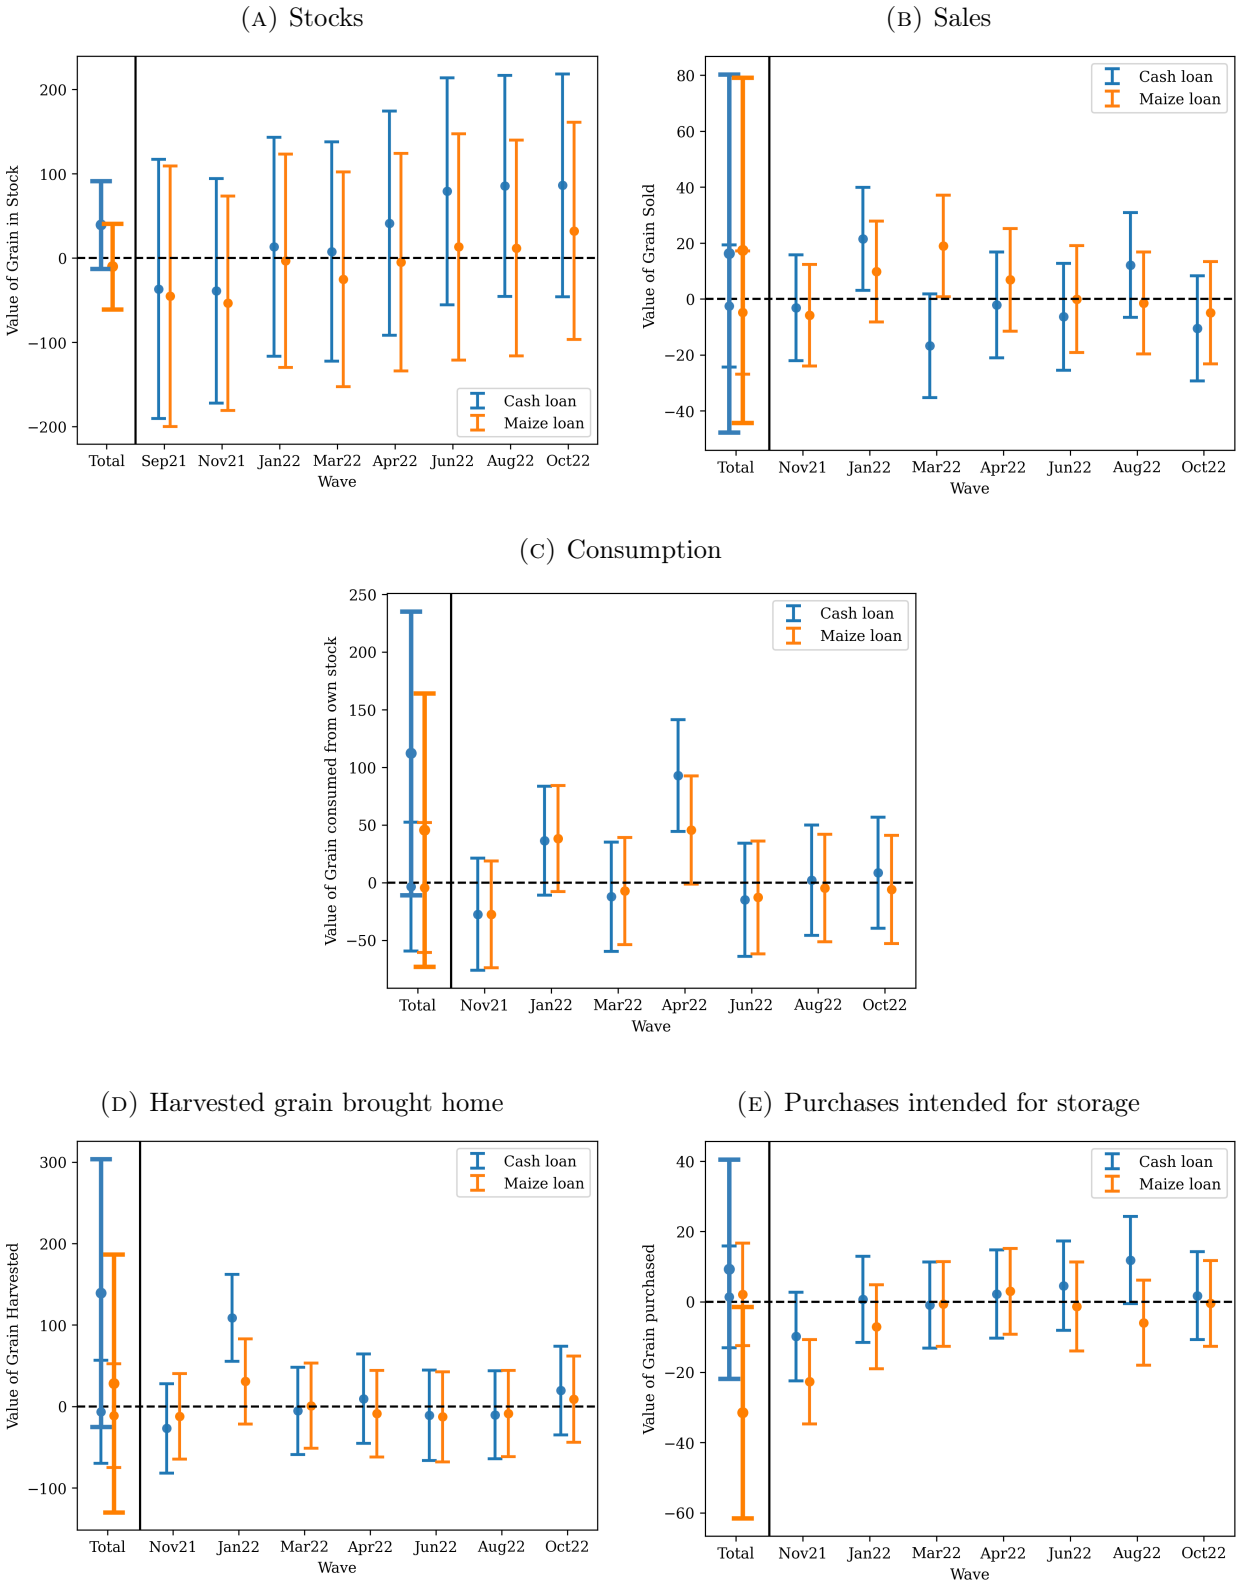
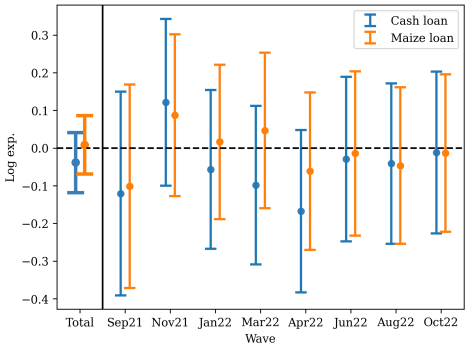


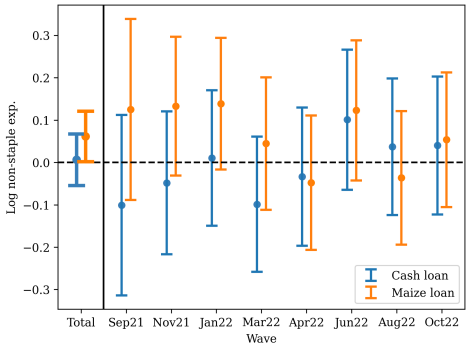
FIGURE A14. Treatment Effects on Stocks, by Wave, regressions with strata FEs

Notes: Replicates Figure 8 using regressions with strata fixed effects and clustered standard errors.

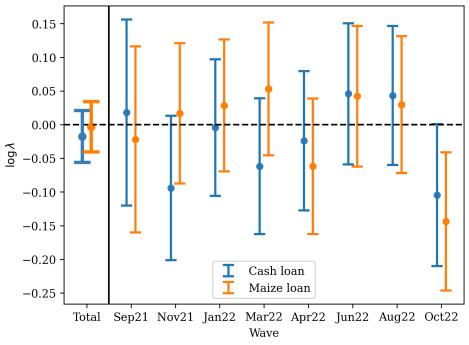
(A) Log food expenditure



(B) Log non-storable expenditure



(C) $\log \lambda$



(D) (Semi-)durable expenditure

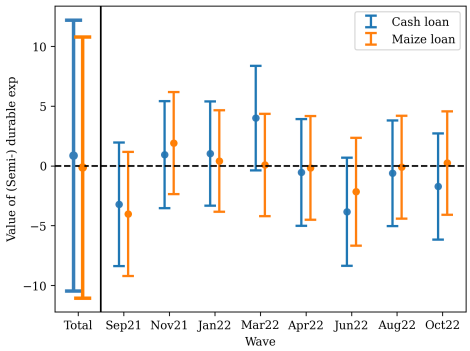
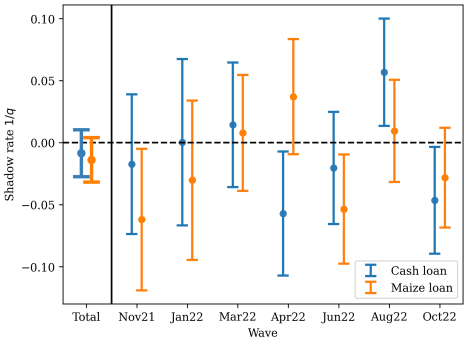
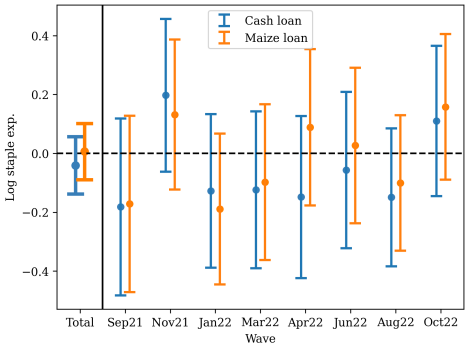


FIGURE A15. Treatment Effects on Expenditures, by Wave, regressions with strata FEs

(E) Shadow rate



(F) Log storable expenditure



(G) RCSI

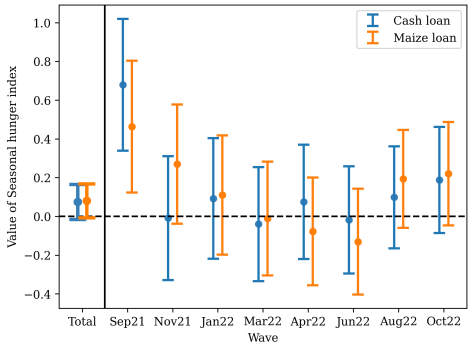


FIGURE A15. Treatment Effects on Expenditures, by Wave, regressions with strata FEs (continued)

Notes: Replicates Figure 9 using regressions with strata fixed effects and clustered standard errors.

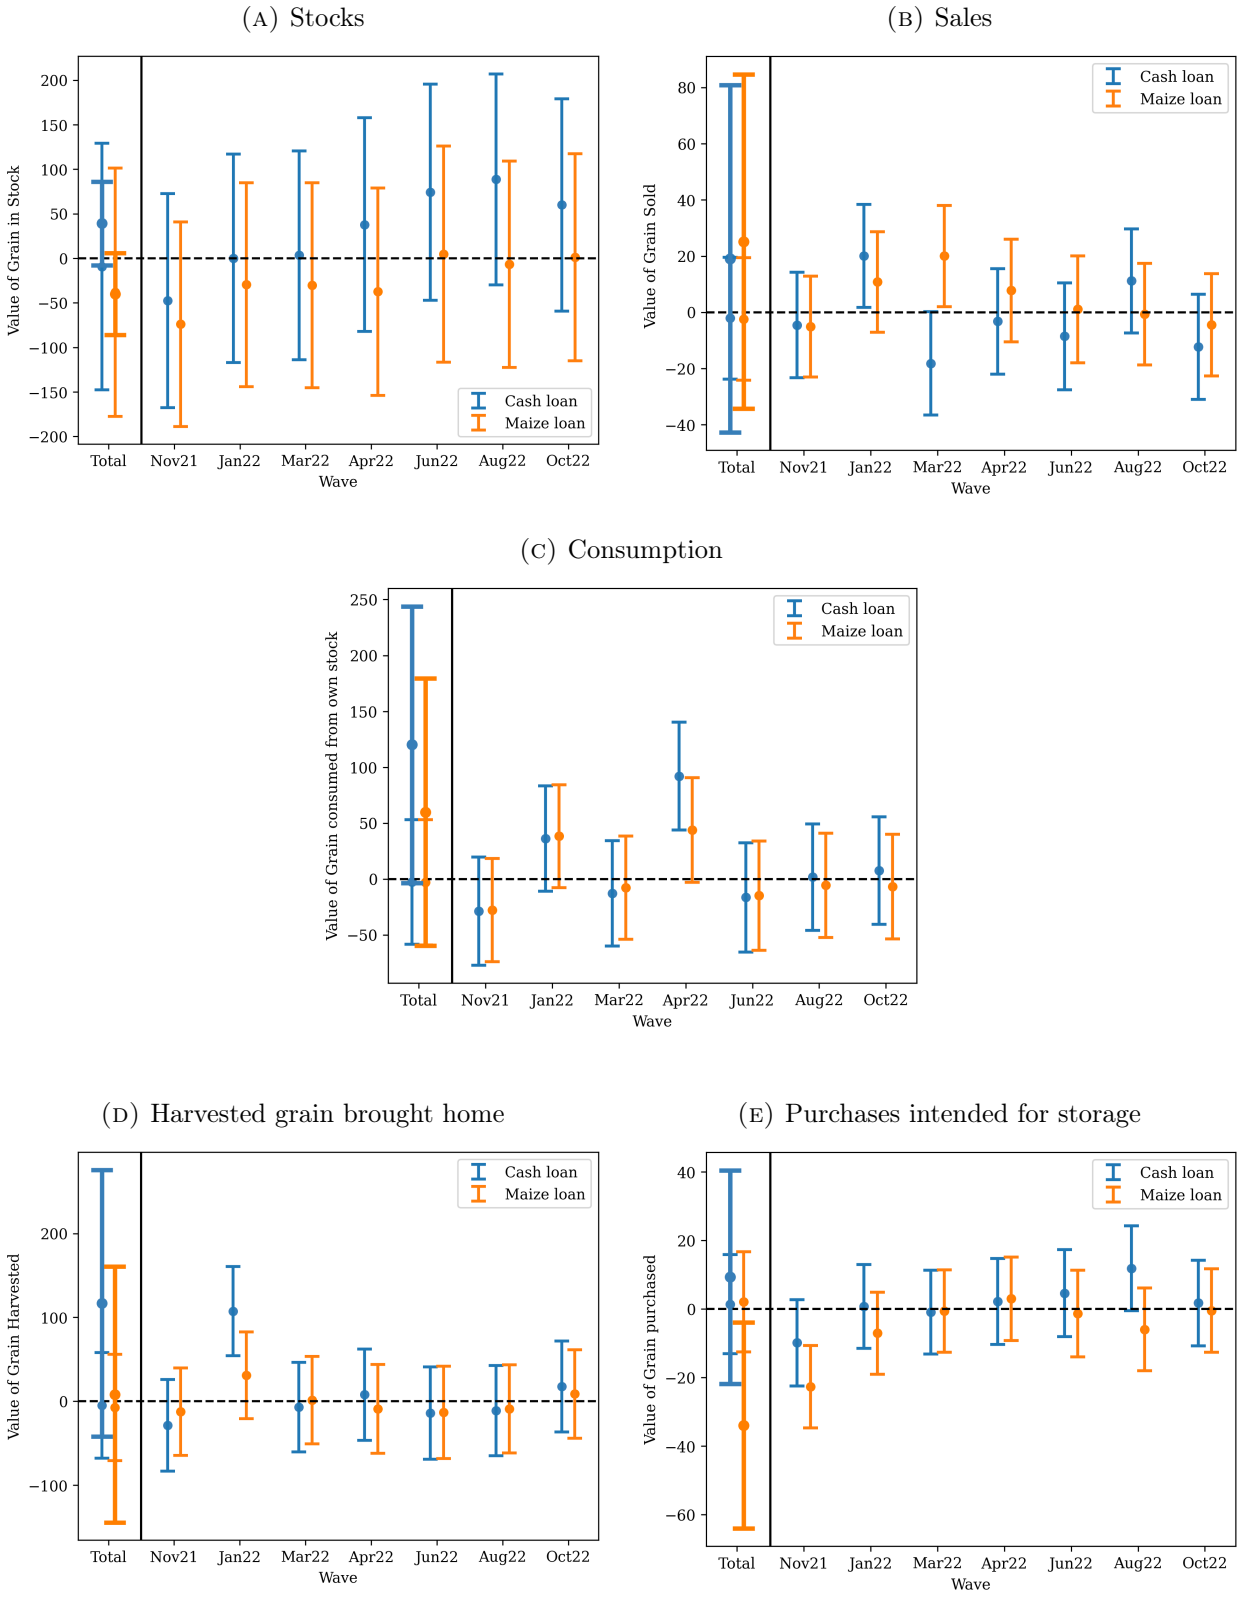
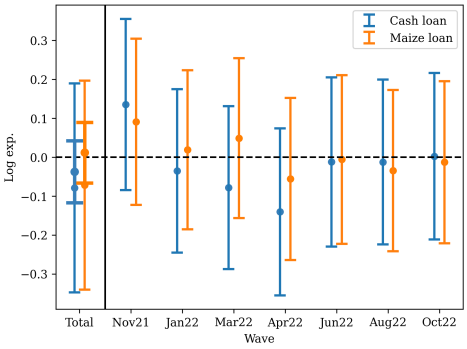


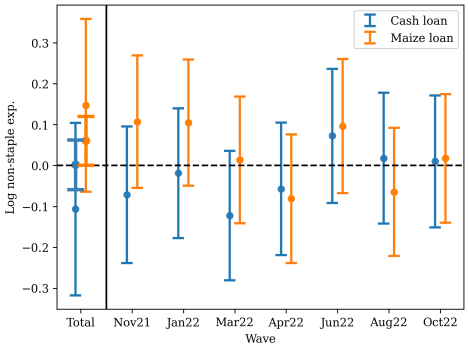
FIGURE A16. Treatment effects on value of grain stocks and flows by wave, Double Post LASSO

Notes: Replicates Figure 8 with Double Post-LASSO control selection.

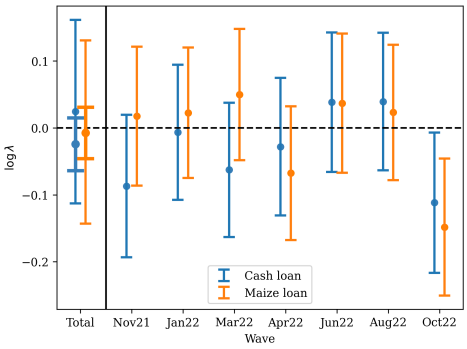
(A) Log food expenditure



(B) Log non-storable expenditure



(C) $\log \lambda$



(D) (Semi-)durable expenditure

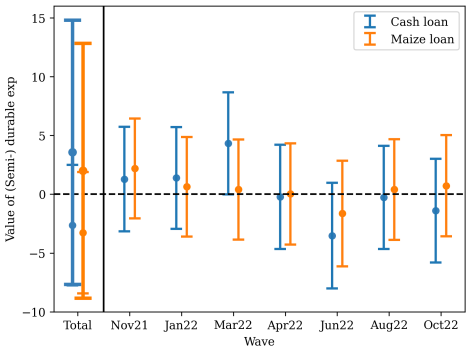
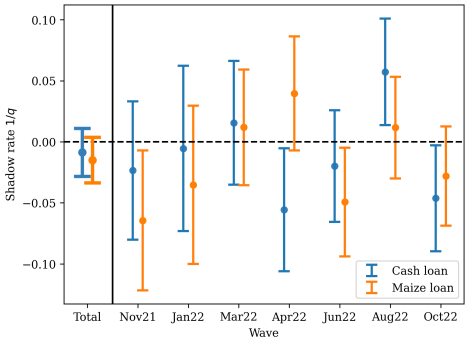
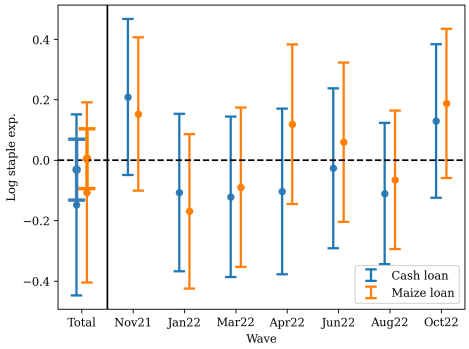


FIGURE A17. Treatment Effects on Expenditures, by Wave, Double Post LASSO

(E) Shadow rate



(F) Log storable expenditure



(G) RCSI

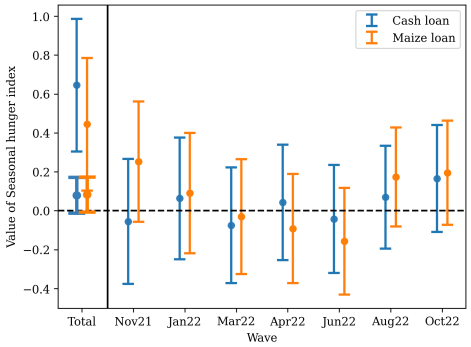
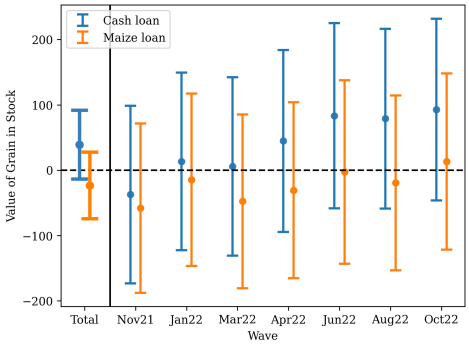


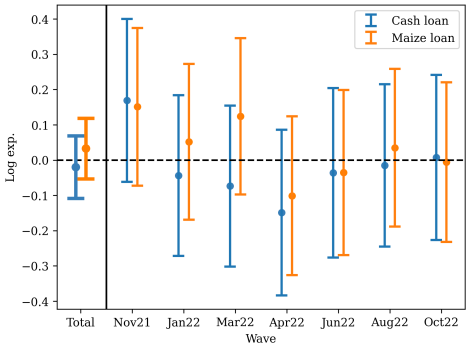
FIGURE A17. Treatment Effects on Expenditures, by Wave, Double Post LASSO (continued)

Notes: Replicates Figure 9 with Double Post-LASSO control selection.

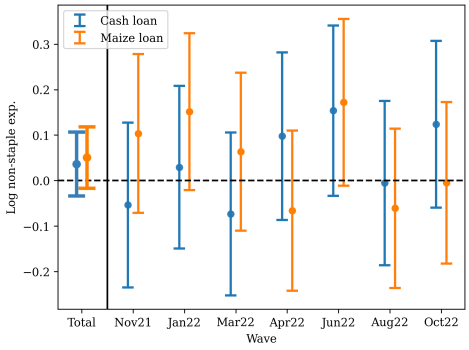
(A) Stock values



(B) Log food expenditure



(C) Log non-storable expenditure



(D) $\log \lambda$

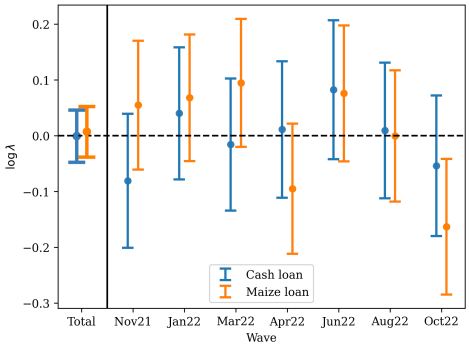
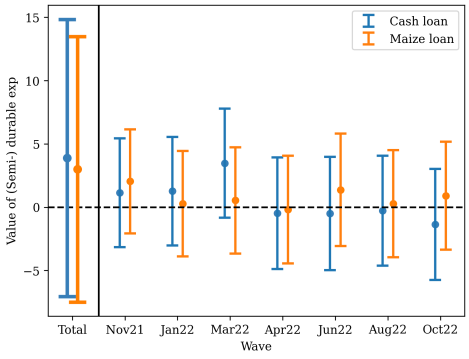


FIGURE A18. Treatment Effects on Stocks and Expenditures, controlling for baseline outcome

(E) (Semi-)durable expenditure



(F) Log storable expenditure

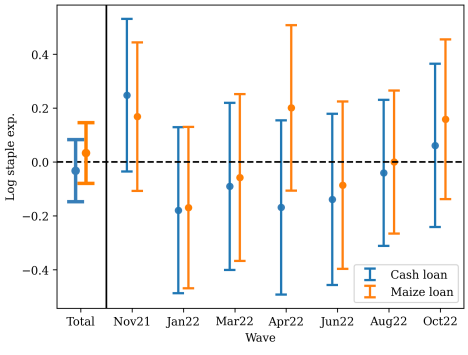
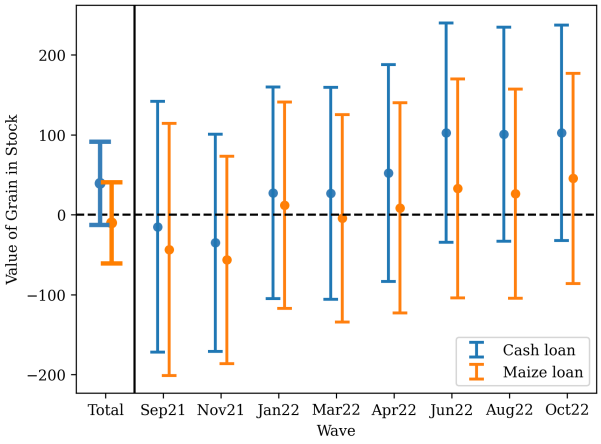


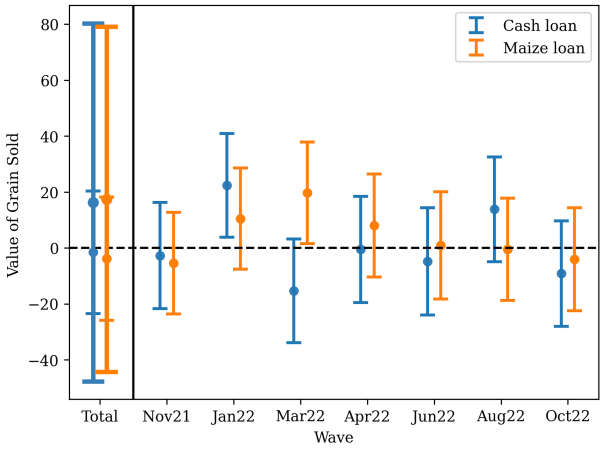
FIGURE A18. Treatment Effects on Stocks and Expenditures, controlling for baseline outcome (continued)

Notes: Replicates main results controlling for baseline outcome values (ANCOVA).

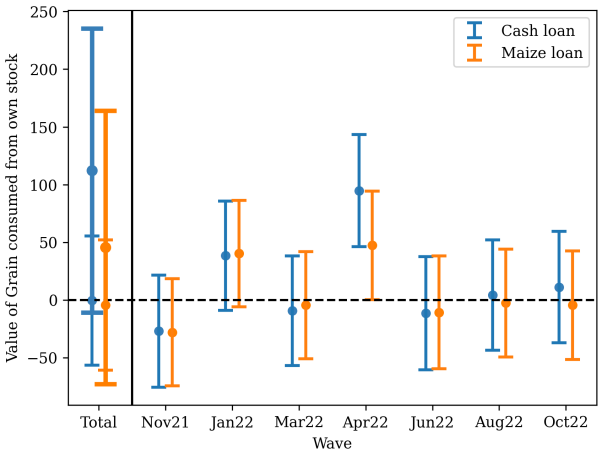
(A) Stocks



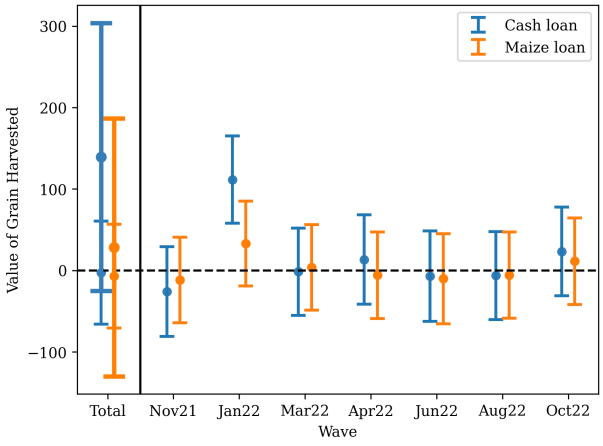
(B) Sales



(C) Consumption



(D) Harvested grain brought home



(E) Purchases intended for storage

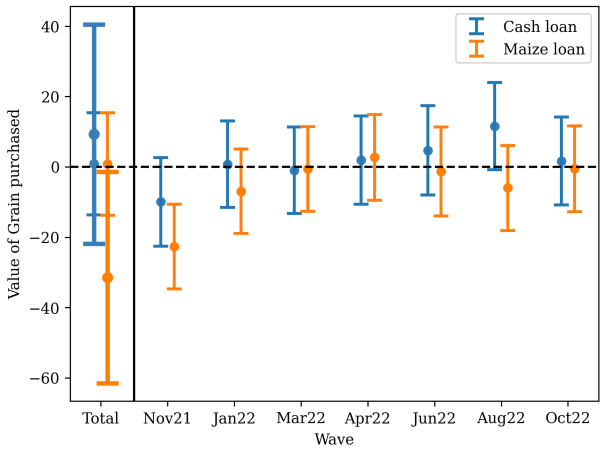


FIGURE A19. Treatment Effects on Stocks, Controlling for Credit Score

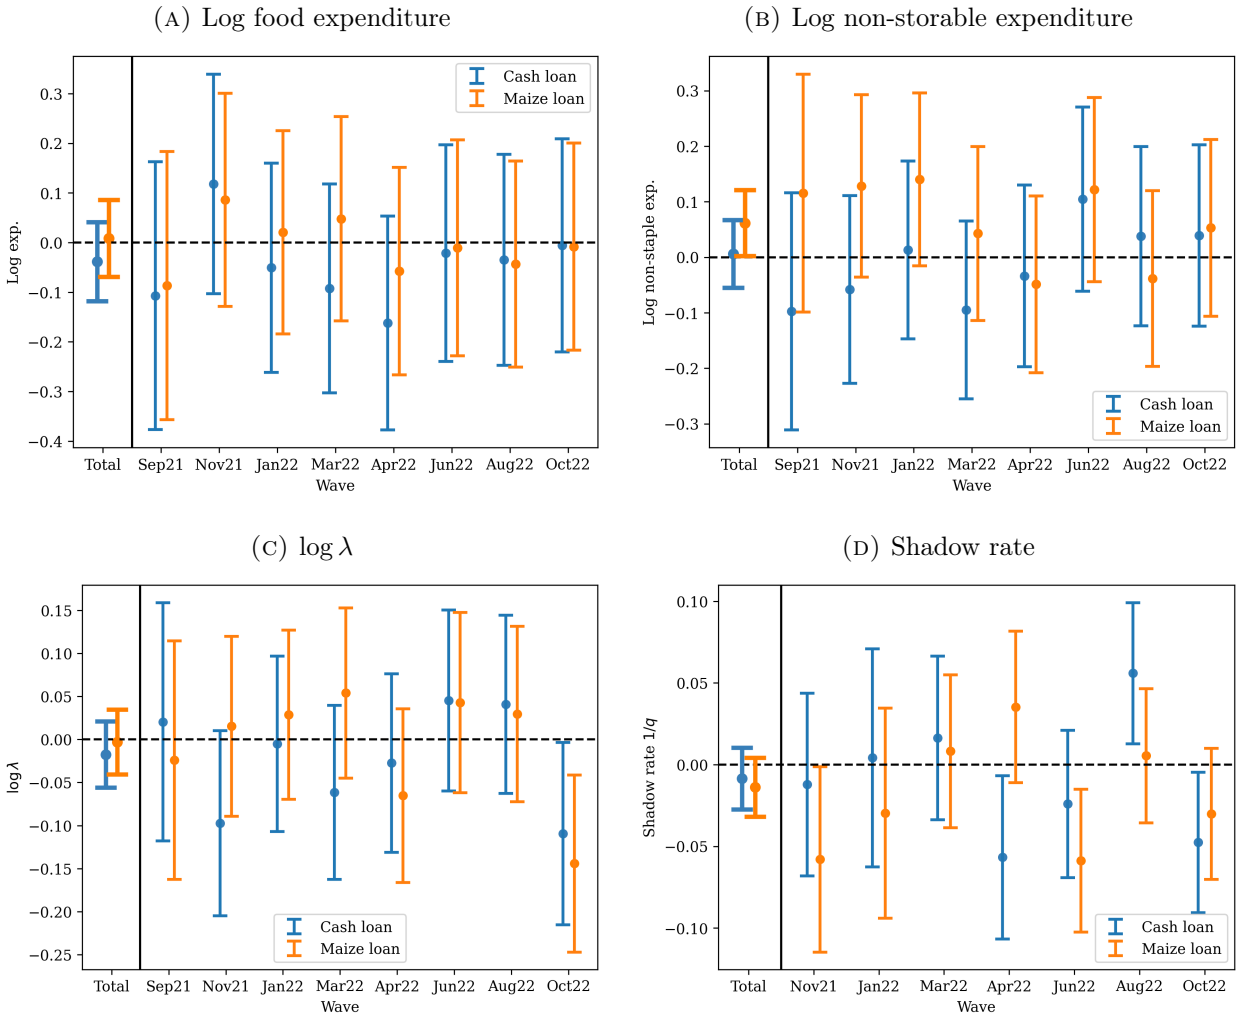
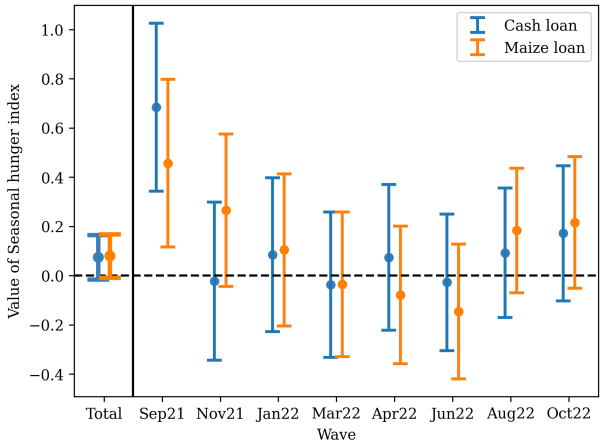


FIGURE A20. Treatment Effects on Expenditures, Controlling for Credit Score

(E) Seasonal hunger index



(F) Log storable expenditure

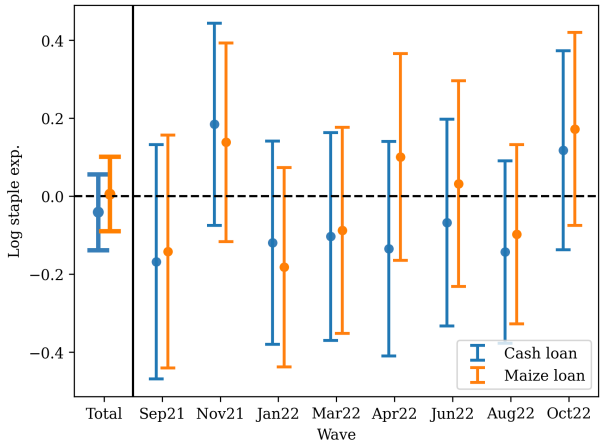
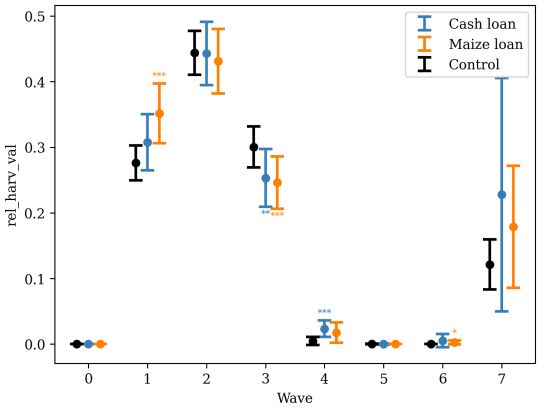


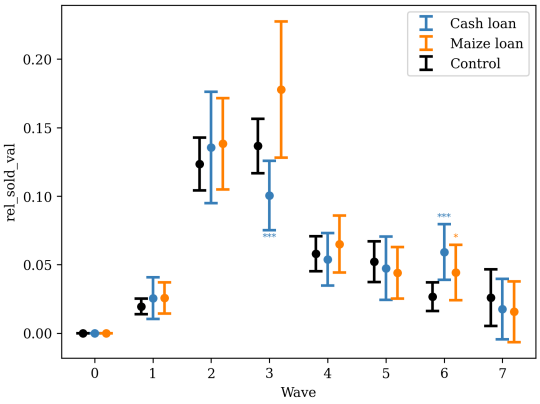
FIGURE A20. Treatment Effects on Expenditures, Controlling for Credit Score (continued)

E.5. **Stocks as Share of Harvest.** To assess treatment effects on the composition of grain management rather than the level, Figures A21 express each stock outcome as a share of total 2021 harvest value.

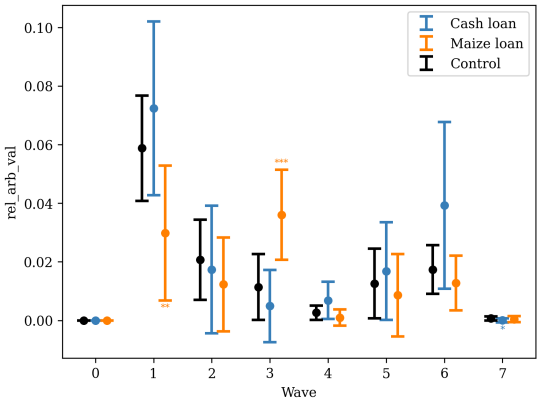
(A) Harvest



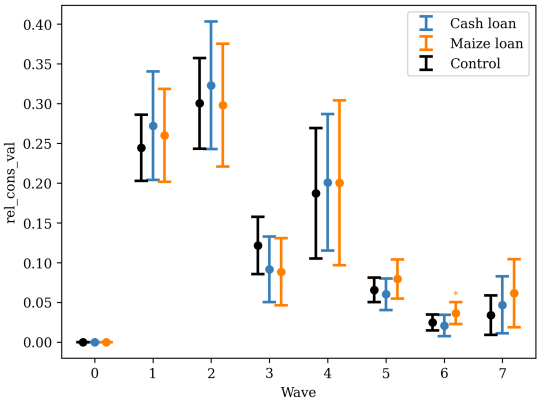
(B) Sales



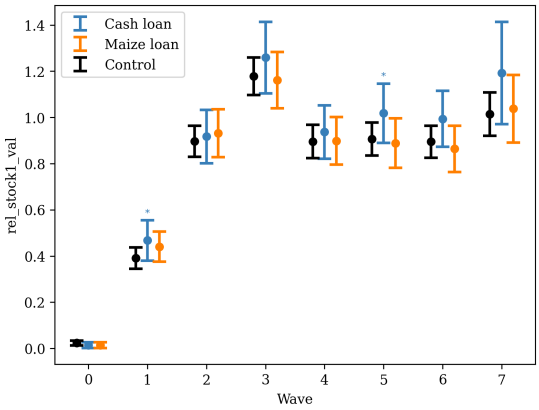
(C) Purchases for storage



(D) Own-stock consumption



(E) Stock value



(F) Stocks plus cumulative sales

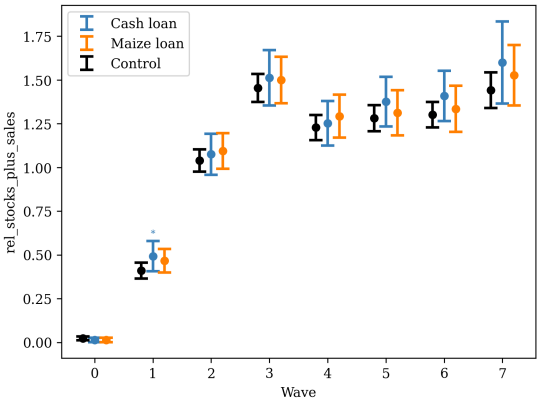


FIGURE A21. Treatment Effects on Stock Outcomes as Share of Total Harvest Value

E.6. **Changes in shadow rate.** Figure A22 shows treatment effects on the first difference of the elicited shadow rate, using only non-empty elicitation choices.

(A) Δ shadow rate (non-empty choices)

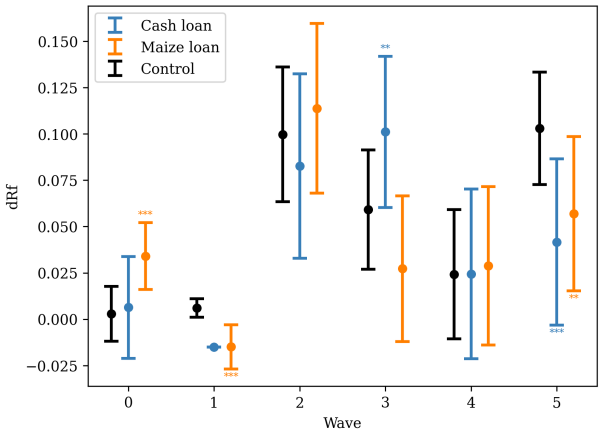
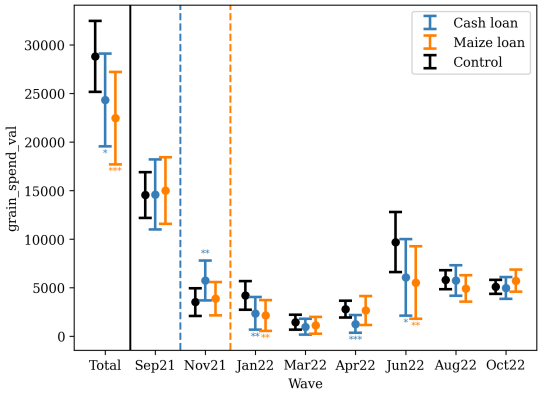


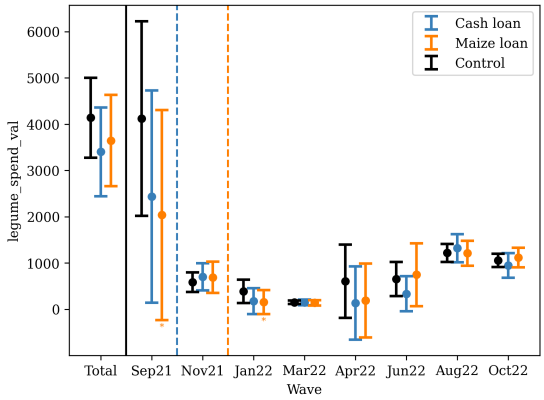
FIGURE A22. Treatment Effects on Changes in shadow rate

E.7. **Treatment Effects on Expenditure Categories.** Figures A23 decompose the treatment effect on food expenditure into six categories: grains, legumes, staples (grains plus legumes plus cassava), protein (meat, fish, and eggs), vegetables, and other food (milk, oil, bread, sugar).

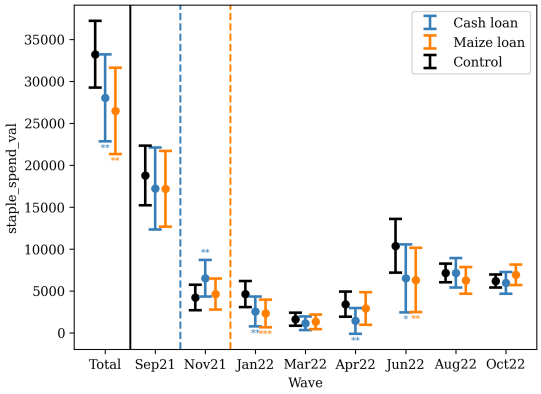
(A) Grain expenditure



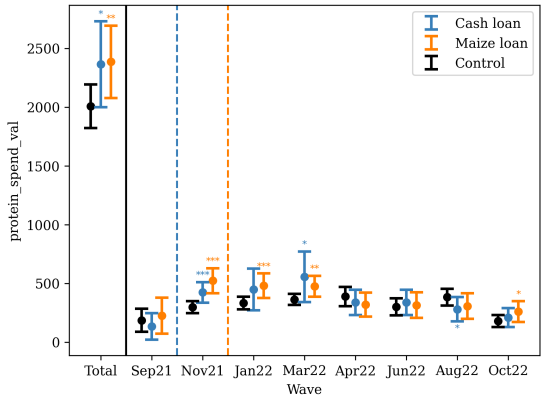
(B) Legume expenditure



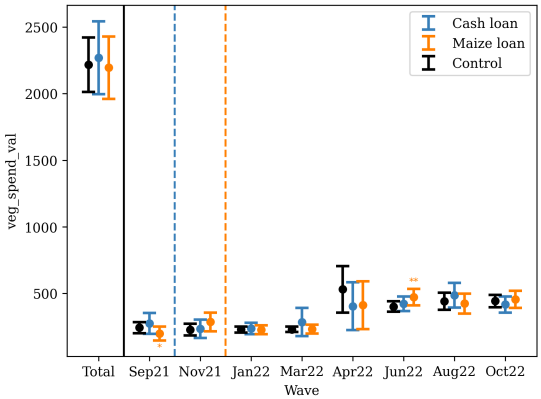
(C) Staple expenditure



(D) Protein expenditure



(E) Vegetable expenditure



(F) Other food expenditure

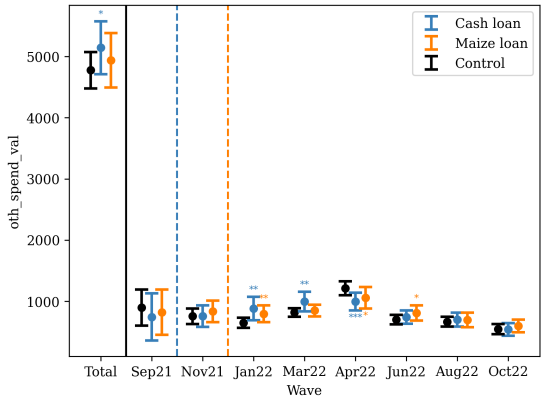


FIGURE A23. Treatment Effects on Food Expenditure by Category

E.8. **Own-Consumption of Produced Crops.** Figures A24 show treatment effects on an indicator for consuming own-produced grain, for each crop separately.

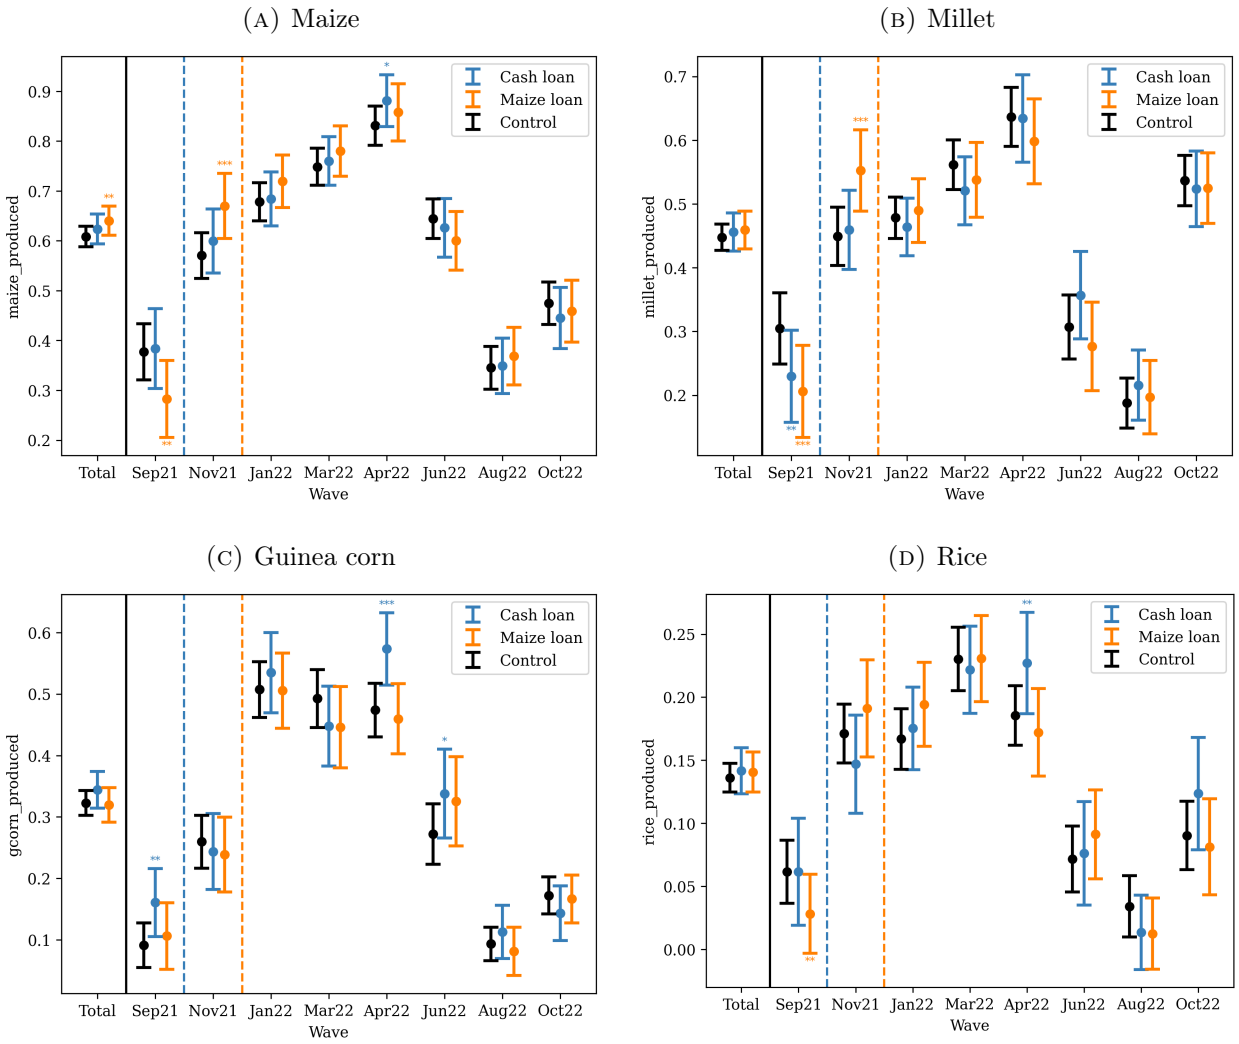
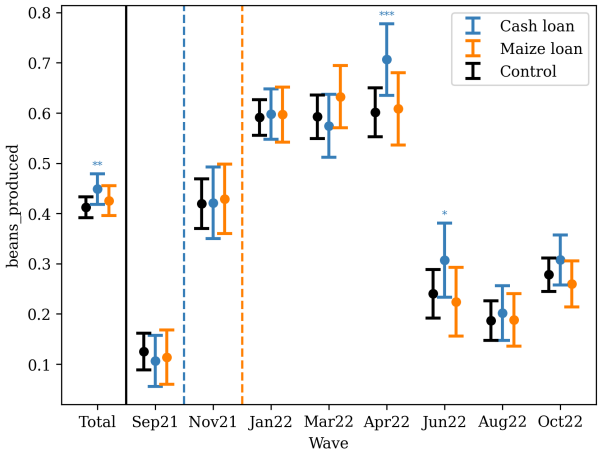
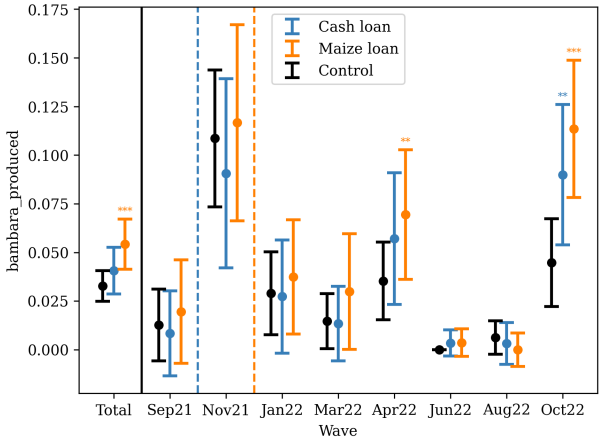


FIGURE A24. Treatment Effects on Own-Consumption of Produced Crops

(E) Beans



(F) Bambara nut



(G) Groundnut

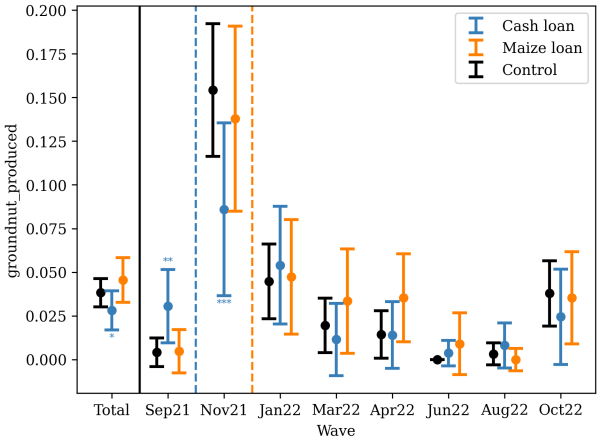
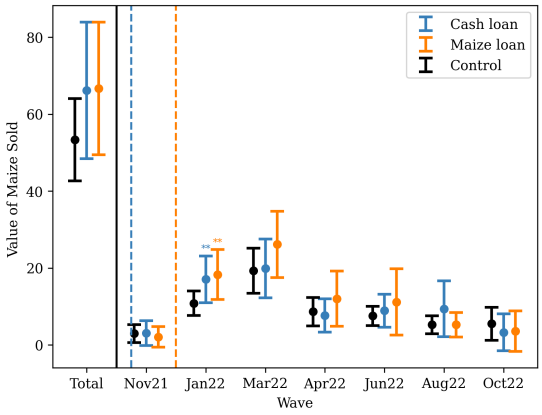


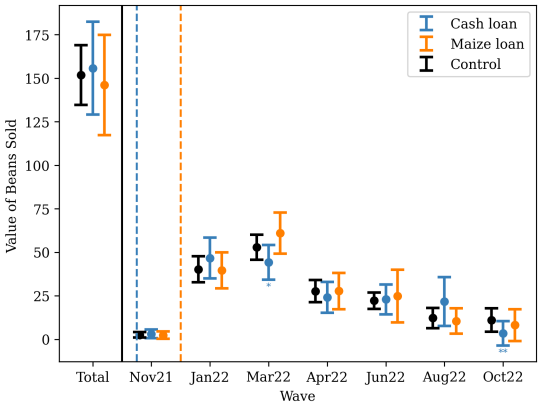
FIGURE A24. Treatment Effects on Own-Consumption of Produced Crops (continued)

E.9. **Crop-Level Sales.** Figure A25 decomposes the aggregate sales treatment effects from Table B by crop and survey wave.

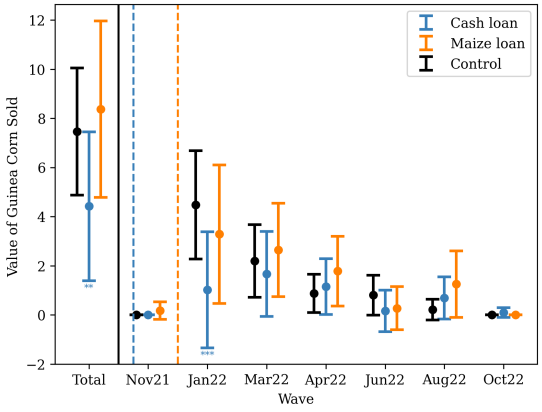
(A) Maize



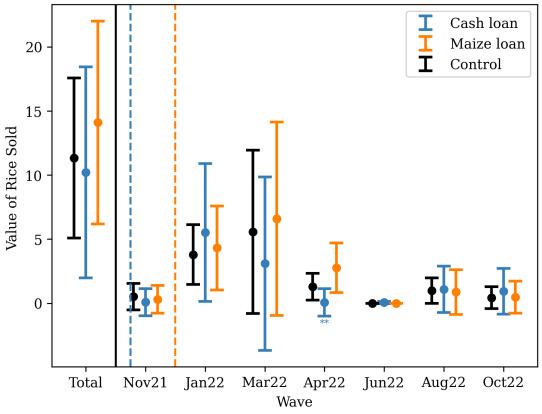
(B) Beans



(C) Guinea corn



(D) Rice



(E) Millet

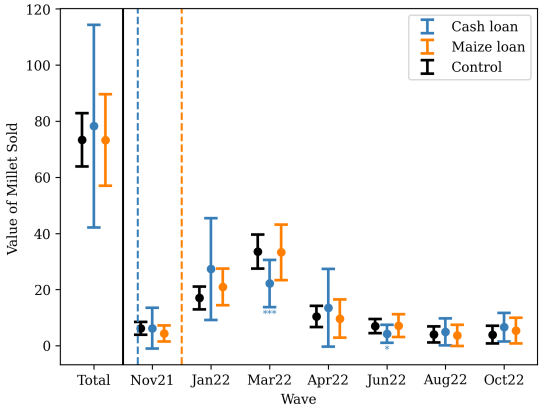


FIGURE A25. Treatment Effects on Crop Sales by Wave

E.10. **Stock Depletion Timing.** Figure A26 shows the mean wave at which households fully deplete their grain stocks, by crop and treatment arm. Depletion is defined as the first wave with zero end-of-period stocks and no further outflows (sales, consumption, or production from stored grain). Crops that were never held are excluded.

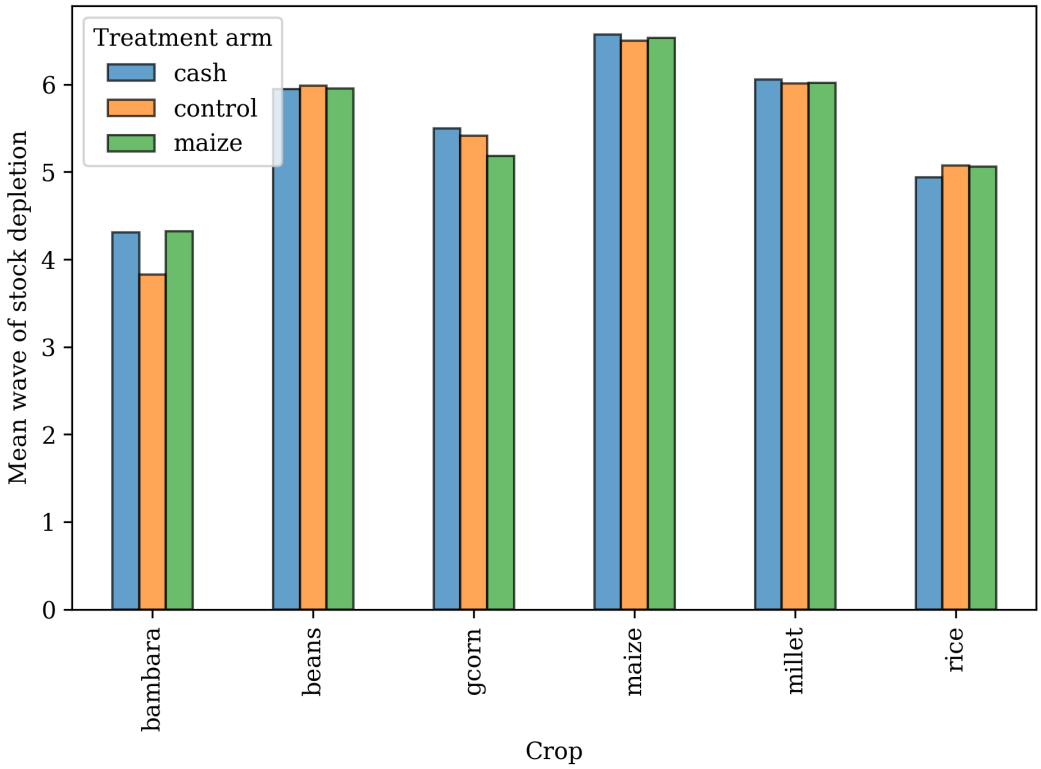


FIGURE A26. Mean Wave of Stock Depletion, by Crop and Treatment Arm

E.11. **Bond Reneging and Payouts.** The PAP pre-specified an analysis of reneging on accepted bond deals. In each survey wave, households were offered a deal to save a specified amount at a given interest rate; reneging is defined as accepting a deal but depositing zero. Figure A27 shows treatment effects on the probability of reneging, by wave. We find no evidence that treatment assignment affects reneging behavior.

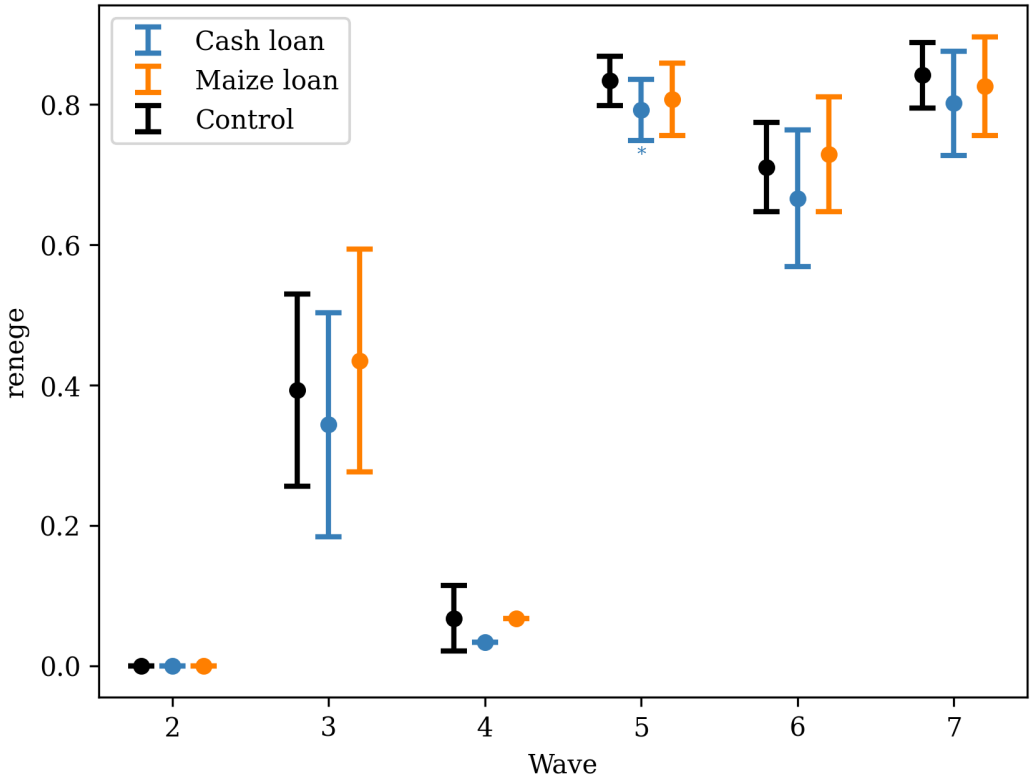


FIGURE A27. Treatment Effects on Probability of Reneging on Accepted Bond Deal

We also examine whether receiving a cash payout from the bond elicitation affects welfare. Regressing $\log \lambda$ on the lagged payout amount (in thousands of Naira) with household and wave fixed effects yields a coefficient of 0.193 (s.e.\ 0.105, $p = 0.066$), suggesting that bond payouts modestly reduce the marginal utility of expenditure, consistent with a relaxation of liquidity constraints.