

# Why do farmers sell immediately after harvest when prices are lowest? A pre-analysis plan

Joachim De Weerd<sup>\*</sup>, Brian Dillon<sup>†</sup>, Emmanuel Hami<sup>\*</sup>,  
Bjorn Van Campenhout<sup>‡</sup>, Leocardia Nabwire<sup>§</sup>

June 5, 2023

## Motivation

It is often observed that smallholder farmers sell most—if not all—of their marketable surplus or cash crops immediately after the harvest to itinerant traders at the farm gate. Selling immediately after the harvest is not optimal. Thin and poorly integrated markets mean that immediately post harvest, prices in excess supply areas drop. Later, during the lean season when some of the farmers run out of stock, prices have recovered, or even increase further since farmers start to buy back. This leads to the “sell low buy high” puzzle ([Stephens and Barrett, 2011](#); [Burke, Bergquist, and Miguel, 2018](#)). In addition to high supply immediately post harvest, agricultural commodities are often not yet in optimal condition. For instance, in the case of maize, fresh grains are generally not dry enough, requiring further processing and leading to increased risk of rot by the trader. Often, this is used by buyers as a reason to further drive down the price paid to the farmer.

There are many possible reasons why farmers choose to sell early at low prices instead of waiting a few months until prices recover. Farmers may simply not have the space and infrastructure available to safely store large quantities of maize for extended periods of time ([Omotilewa et al., 2018](#)). They may be in urgent need of cash after the lean season ([Burke, Bergquist, and Miguel, 2018](#); [Dillon, 2021](#)). Price movements may be unpredictable and farmers may be too risk averse to engage into intertemporal arbitrage ([Cardell and Michelson, 2020](#)). It may be that traders only visit villages immediately after harvest, and farmers do not have the means to transport maize to markets themselves. Furthermore,

---

<sup>\*</sup>Development Strategy and Governance Division, International Food Policy Research Institute, Lilongwe, Malawi

<sup>†</sup>Dyson School, Cornell University, Ithaca, United States

<sup>‡</sup>Development Strategy and Governance Division, International Food Policy Research Institute, Leuven, Belgium

<sup>§</sup>Development Strategy and Governance Division, International Food Policy Research Institute, Kampala, Uganda

issues related to social taxation may mean farmers convert maize to cash, which is easier to hide from friends and family.

Most of the explanations above focus on hard constraints to farmers' exploiting intertemporal arbitrage. In this study, we zoom in on two potential behavioural explanations why farmers seemingly sell at sub-optimal time. One potential explanation is situated at the household expenditure side, and assumes that households face challenges in accurately predictive future expenditures. Such budget neglect leads farmer to sell more early on and save too little for later in the year. A second potential explanation is situated at the household income side. Here the assumption is that farmers face cognitive challenges in making inter-temporal cost benefit calculations (Drexler, Fischer, and Schoar, 2014) and fail to commit to certain thresholds (Ashraf, Karlan, and Yin, 2006; Duflo, Kremer, and Robinson, 2011).

This document serves as a pre-analysis plan for the study that will be registered in a public repository. It provides background information, outlines hypotheses which will be tested, tools that will be used in the field, power calculations and sample size projections on which sampling is based, outcome variables that will be used to assess impact, and specification that will be estimated. As such, it will provide a useful reference in evaluating the final results of the study (Humphreys, Sanchez de la Sierra, and van der Windt, 2013).

## Literature

Why do farmers sell low and buy high? One of the most obvious neo-classical explanations is related to credit constraints. Using observational data, Stephens and Barrett (2011) find that to meet consumption needs later in the year, many farmers end up buying back grain from the market a few months after selling it, in effect using the maize market as a high-interest lender of last resort. Burke, Bergquist, and Miguel (2018) show that in a field experiment in Kenya, credit market imperfections limit farmers' abilities to move grain intertemporally. Providing timely access to credit allows farmers to buy at lower prices and sell at higher prices, increasing farm revenues and generating a return on investment of almost 30%. Dillon (2021) uses the fact that in Malawi, primary school began 3 months earlier in 2010 than in 2009, and notes that this prompted households with children to sell maize when prices are particularly low. To identify the impacts of liquidity during the lean season, Fink, Jack, and Masiye (2020) offered subsidized loans in randomly selected villages in rural Zambia and conclude that liquidity constraints contribute to inequality in rural economies. Channa et al. (2022) show that harvest loan recipients stored 29% more and sold 50% more maize in lean season on average (but found no impact of improved storage technology). While credit constraints thus seems to be an important reason why farmers sell immediately post harvest at low prices, we feel this is not the entire story. If farmers need urgent cash, it would make more sense to only sell part of the harvest to cover most urgent expenses. However, farmers generally sell all maize immediately post harvest at low prices.

Risk averse farmers may also fail to delay sales if there is considerable uncertainty about the price in the future. A recent article argues that the “sell low buy high” puzzle is not a puzzle at all, as price movements are insufficient for farmers to engage in inter-temporal arbitrage (Cardell and Michelson, 2020). However, their analysis use prices obtained from market centers, which may be a poor proxy for the farm gate prices that farmers face: prices in main markets are generally much better integrated in the wider national, regional and even global economy, and so will be less prone to extreme spikes and slumps. While we agree that uncertainty about prices is indeed an important reason to sell immediately post harvest for loss averse farmers (and indeed loss aversion lies at the core of one of our research hypotheses), we do feel that this is not a sufficient explanation in the face of large recurrent seasonal price movements.

A third reason that is often heard in the field is that farmers have nowhere to store, so they just sell. This could be a lack of space, as the average smallholder often harvest 10-20 bags of 100kg of maize. But there are also risk related to pests and diseases affecting the stored maize. If storage is the main reason why farmers do not engage in intertemporal arbitrage more, then providing storage technology should delay sales. Omotilewa et al. (2018) indeed find that households that received PICS bags stored maize for a longer period, reported a substantial drop in storage losses. Again, we feel storage is indeed part of the reason, but it does not explain everything. For instance Agricultural Commodities Exchange (ACE) in Malawi provides storage technology but still fails to fill its warehouses.

Another reason may be related to social taxation. If a farmer has a lot of maize stored in his house, this is visible for family and neighbours, and it will be very hard to deny if they come and ask for help. Therefore, farmers may choose to convert their harvest to money, which is easier to hide, even though this comes at a cost. Social taxation has been found important in a similar marketing decisions where household seem to forgo the benefits of buying in bulk (Dillon, De Weerd, and O’Donoghue, 2020).

## **Behavioural constraints to intertemporal arbitrage: Hypotheses and Interventions**

The first potential behavioural explanation is situated at the household expenditure side, and assumes that households face challenges in accurately predicting future expenditures. In other words, the first hypothesis assumes farmers suffer from budget neglect, which may lead to an overoptimistic view of the future. In particular, farmers may neglect some future expenditures when deciding on how much to sell immediately after the harvest. For example, immediately after harvest, they may budget for fresh seed from the agro-input dealer and for fertilizer, but they may forget that they also need pesticides and insecticides. Furthermore, farmers may underestimate the likelihood of, or simply forget to account for, unexpected events such illness within the family.

This hypothesis touches on cognitive limits of the household at the expenditure side. It is also related to the planning fallacy, where individuals typically underestimate the time it takes to complete a task, despite extensive experience with failure to complete the same task in a similar time frame in the past (Buehler, Griffin, and Peetz, 2010). Part of it may also be related to optimism bias if farmers neglect or underestimate the risk that adverse effects will happen to them (Sharot, 2011). For instance, farmers may not budget for pesticides or insecticides because they believe they will not be affected by pests or insects. Budget neglect is also found to be a main contributing factor to recurrent hungry seasons in Zambia.

To test the first hypothesis related to budget neglect, the focus will be on the expenditure side and we will design an intervention that takes the farmer through a detailed budgeting exercise. The budget exercise will involve three components. A first component uses recall to provide a first approximation of what will be necessary in the future. A second component consists of segmentation, which involves defining categories of expenditures for cognitive ease. Finally, we will look at a range of risks, which involve expenses that are not certain but may materialize. We try as much as possible to attach objective probabilities to these risks and also incorporate this in the budget.

This second hypothesis is also related to cognitive limitations when planning, but this time at the income side of the farm household. Farmers may have difficulties in making the intertemporal cost-benefit calculations necessary to determine the optimal reservation price and/or storage period. They often lack precise information about the fixed and variable costs involved, about the level and variability of the future stream of income from sales, or about the time frame of both cost and income (Van Campenhout, 2021). The fact that farmers are faced with uncertain prices and uncertain expenditures often means they abandon plans and engage in impulsive or distress sales.

To test the second hypothesis, we will develop, together with the farmer, a detailed plan of how much the farmer will sell over the coming year (per month or per quarter). For each sales event, the farmer will also be asked to commit to a minimum price. This will be done on a special form that farmers can then hang up in their house. Enumerators will be asked to take a picture of the plan. This is to 1) check if enumerators did their job 2) to signal to farmers that we will check if they keep to their commitments.

## Experimental design and power calculations

We propose parallel design with one common control group and two treatment arms (or a factorial design where the interaction cell is left empty). Kaur et al (personal communication) find that, in a similar budget neglect experiment, treated farmers enter the hungry season with 20 percent more maize (valued by current prices at 405 zambian kwacha instead of 335 zambian kwacha in the control group). If we assume that standard deviation is about 592 (1.6 times the mean of treatment and control means – the 1.6 is derived from maize production

data in Uganda), we get a sample size of 1123 in each sample. For one control group and two treatment arms we will thus need about 3400 farmers. As power is optimal when one allocates approximately 42% of the sample to control, and then equally (29% each) to each treatment arm (Muralidharan, Romero, and Wüthrich, 2019), will choose 13 households for the control and 9 households in each of the two treatment groups per village. This means we will need about 109 villages.

## Sampling

We use a multi-stage sampling procedure to create a self-weighting sample up to the village level and then just sample a fixed number of households per village. We then sample villages with the likelihood of a village being selected being proportionate to the number of people that live in this village (such that larger villages are more likely to end up in the sample). In particular, to get a nationally representative sampling frame of the smallholders farmers population Malawi, we rely on the list created by the Ministry of Agriculture for their Agricultural Input Programme (AIP). The AIP only targets smallholder farmers in the villages who mostly registered with the village chiefs.

We aggregate this list of households to the village level and remove villages that have less than 35 households (as we need at least 31 households per treatment arm in each village). We then sample 114 villages (109+5 to account for attrition), with the probability of a village being selected proportional to the number of households that reside in this village. We then randomly sample 31 households in each of these 114 villages

## Context and study area

The study focuses on the Central and Northern Region of Malawi (Kasungu, Mzimba (both North and South), Ntchisi, Rumphu, Dowa and Mchinji). In these areas, maize is generally regarded as the food crop for auto-consumption or to pay laborers in kind. Maize was also sometimes marketed, but mostly not as the most important one cash crop. The main cash crops in the area are soybean, ground nuts and tobacco. Prices of tobacco do are not seasonal. For the other crops, most farmers also mentioned significant seasonality similar to seasonal price movements of maize.

These areas are characterized by rained agriculture with a single season. The resulting seasonal price movements is illustrated in Figure 1 that shows maize price in kwacha per Kg in Rumpu over 2020. Planting of maize starts in December, and maize becomes increasingly scarce during the growing season. Harvesting starts around April 2020, which takes the pressure off the prices when farmers start consuming from their own maize. However, farm gate sales are still low as traders wait for maize to dry. This results in a relatively long period of low prices all the way to the start of the planting season towards the




Figure 1: Price of maize in Rumphi

end to the year. The aim of the study is to encourage farmers to wait just a few months longer before they sell.

Figure 2 shows that most sales happen only around August. So farmers do seem to hold on to their maize for reasonably long periods (suggesting some of the other explanations like lack of storage space or social taxation are less likely). Sales for other crops follows a similar pattern.

Taken together, the figures suggest that the best time for the interventions would be around April or May, immediately before farmer start to sell.

Farmers often indicate to have access to finance, but note that interest rates are prohibitively high (30-40%). There is also a strong cooperative movement in Malawi. Some of these cooperatives also provide access to warehousing and engage in collective marketing. Qualitative research suggests that it is pretty easy for farmers to sell even in the off-season. Traders operate in trading centers, writing prices on a blackboard. The trader we interviewed mentioned there were many others like him in the small trading center he was operating in. Traders also visit villages, often using ox carts. If they buy at farm gate, prices are discussed and depend on distance traveled. Traders buy from May to August and sell from December to February. Farmers are suspicious about scales used.

## Specifications

Instead of relying on a single endline, we will evaluate the interventions through multiple rounds of data collection, often using phone interviews. There are different reasons for this. First, when measuring noisy and relatively less auto-correlated outcomes such as amounts of commodities sold or household expenditure, one can increase power by taking multiple measurements at relatively short intervals to average out noise (McKenzie, 2012; Burke, Bergquist, and Miguel, 2018). Furthermore, it will allow us to assess the effect of the interventions at multiple points in time instead of just at endline.<sup>1</sup> The follow-up

<sup>1</sup>McKenzie (2012) also cautions about some of the downsides. For instance, it could be that repeated surveys affect people's reporting or even behavior.




Figure 2: Quantities of maize bought and sold

surveys will focus on tracking data on storage inventory of maize, groundnuts and soybean, marketing behavior of the three crops, consumption, and credit and savings behavior.

The expected low correlation of some of our key outcomes over time also means that we will focus mostly on simply comparing treatment and control post treatment (as opposed to difference-in-difference). Specification will also depend on the time horizon and relevance of pooled treatment effect. For instance, when flow variables like stocks held by farmers used as outcome variables, we expect that the difference between treatment and control will be largest in the early in the season. Therefore, when this outcome variable is used, we estimate a (separate treatment effect) for September 2022 and December 2022. This will be modeled with a simple OLS regression that also has an interaction dummy for treatment and the first two survey rounds (September 2022 and December 2022). For sales made by farmers, we expect that the treatment effect reverses over time, as early in the season, control farmers are likely to make sales, while later in the season (December 2022 and March 2023), treatment farmers are likely to make sales. This will be modeled with a simple OLS regression that also has an interaction dummy for treatment and the last two survey rounds (December 2022 and March 2023).

For other key outcomes (like for example prices received, proportion of transactions on which female co-head was consulted, proportion of transactions where sales were made to a trader, what proceeds were used for, etc) we will pool all post treatment observations and estimate ANCOVA models, where we also control for prices received during the 2021-2022 season (the latter being calculated as an average of transactions that were recalled by the farmer). For outcomes

further down the impact pathway on which no baseline data was collected (such as consumption expenditure of the last month), we will simply compare treatment and control households.

The exact model we estimate is the one that is used for factorial designs but without the interaction term (Muralidharan, Romero, and Wüthrich, 2019):

$$Y = \beta_0 + \mu_v + \beta_1 T_1 + \beta_2 T_2 + \theta Y_b + \varepsilon \quad (1)$$

Here  $Y$  is the outcome of interest post treatment, and  $Y_b$  is the outcome of interest measured at baseline.  $T_1$  is an indicator dummy for the first treatment and  $T_2$  for the second. We will focus on the within-village dimension of the data and so include village fixed effects  $\mu_v$  and an overall constant  $\beta_0$ . The parameters of interest are  $\beta_1$  as the average treatment effect for  $T_1$  and  $\beta_2$  as the average treatment effect for  $T_2$ .

Because we will test for treatment effects on a range of outcome measures, we are faced with issues related to multiple hypotheses testing. We will deal with this by means of two approaches. Firstly, we follow a method proposed by Anderson (2008) and aggregate different outcome measures within each domain into single summary indices. Each index is computed as a weighted mean of the standardized values of the outcome variables. The weights of this efficient generalized least squares estimator are calculated to maximize the amount of information captured in the index by giving less weight to outcomes that are highly correlated with each other. Combining outcomes in indices is a common strategy to guard against over-rejection of the null hypothesis due to multiple inference.

However, it may also be interesting to see the effect of the intervention on individual outcomes. An alternative strategy to deal with the multiple comparisons problem is to adjust the significance levels to control the Family Wise Error Rates (FWER). The simplest such method is the Bonferroni method. However, the Bonferroni adjustment assumes outcomes are independent, and so can be too conservative when outcomes are correlated. We therefore use a Bonferroni adjustment which adjusts for correlation (Aker et al., 2016; Sankoh, Huque, and Dubey, 1997)

## Data collection and endpoints

We will not organize a dedicated baseline survey, but rather ask a limited number of questions immediately prior to the interventions in May 2022. This information can then be used to demonstrate balance, to control for baseline outcomes for the primary outcome variables in an ANCOVA regression, and to explore heterogeneous treatment effects. Midline data will be collected in September-October 2022 and December 2022-January 2023, generally by mobile phone. To allow for farmers that do not have access to a mobile phone, we will make sure that in each village we identify someone with a phone that can be shared with the farmer or farmers that do not have a phone. A more elaborate in-person endline survey will be organized in March-April 2023.

## Baseline balance

To demonstrate baseline balance, we will construct a standard balance table consisting of the following variables household/demographic characteristics (inspired by balance tables in [Duflo, Kremer, and Robinson \(2011\)](#); [Karlan et al. \(2014\)](#)): household head is female (1=yes), household size (number of people), age of household head (years), number of years of education of the household head (years), material of roof (corrugated iron = 1), number of rooms in the house, cultivated acreage (maize+groundnuts+soybean), hired in agricultural labour (1=yes), distance to nearest all weather road (km), distance to nearest market (km).

We will report t-tests comparing treatment and control (unadjusted for multiple hypothesis testing) as well as a joint F-test from a regression of the treatment assignment on all variables in the balance table.

To explore heterogeneity in treatment effects, we will measure the following during baseline: Access to credit, access to storage facility, membership of (marketing related) cooperative, livestock asset ownership, whether the household already makes a budget. We will also assess balance on these characteristics at baseline.

During baseline we will also collect recall data on marketing of the three crops in the previous season. To explore some of the gender dimensions of the interventions, we will also ask for each transaction how decisions were made, and what expense categories were covered with the proceeds from the sale.

## Primary outcomes

Primary outcomes in this study include stocks of ground nuts, maize and soybean held by the farmers and how they evolve over time. As there is a particular focus on marketing behaviour, we will also collect detailed information on sales made, including quantities sold and prices received. We ask similar questions on for total amounts bought of each of the crops.

## Secondary outcomes

In addition to the primary outcomes above, we also look if the intervention affects who was sold to. In particular, we test if farmers increase the likelihood that they sold directly to the market, as this is often the place where they can get the best price for their crop ([Fafchamps and Hill, 2005](#)). We also test if there is a change related to decision making, in particular if the likelihood that decisions related to the sales of the commodity was taken jointly by the household co-heads. We also ask what were proceeds used for, and test if there is a change in the use of proceeds for expected expenditures (education) and unexpected expenditures (health).

Furthermore, we include a module on price expectations, which will be useful to see how expectations influence eventual prices obtained, and how interventions affect the relation between expectations and behaviour.

Budget neglect or failure to adhere to sales plans may also affect outcomes in the longer run, such as health of household members, resilience against shocks, and agricultural yield. At endline, we will therefore collect also detailed consumption data. Furthermore, we will collect information on farm investment, and also look at changes in the composition of livestock assets, which may have been used to smooth consumption (Balboni et al., 2021). In addition, poor planning may also affect labor allocations within the household, and so we will also investigate the effect of the interventions on, for instance, seasonal labour provided by household members (Fink, Jack, and Masiye, 2020).

At the time of endline data collection, households are normally harvesting, and if they did not start harvesting yet, they should have a good idea of the harvest they can expect. We will also look at the impact of the interventions on production and yield of the the three crops in April 2023.

## Variable construction

For continuous variables, 5 percent trimmed values will be used to reduce influence of outliers (2.5 percent trimming at each side of the distribution). Inverse hyperbolic sine transforms will be used if skewness exceeds 1.96. Trimming will always be done on end results. For instance, if the outcome is yield at the plot level, then production will first be divided by plot area, after which inverse hyperbolic sine is taken and the end result is trimmed. Outcomes for which 95 percent of observations have the same value within the relevant sample will be omitted from the analysis and will not be included in any indicators or hypothesis tests.

When we field our surveys, some respondents will not answer one or more questions that measure an outcome. We will handle missing variables from survey questions by checking whether item non-response is correlated with treatment status, and if it is, construct bounds for our treatment estimates that are robust to this. To be more precise, we will assess the relationship between missing outcomes and treatment assignment using a hypothesis test and report these results. If  $p < .05$  for the assessment of the relationship between treatment and missing outcomes, we will report an extreme value bounds analysis in which we set all of the missing outcomes for treatment to the (block) maximum and all missing outcomes for control to the (block) minimum. If  $p \geq 0.5$  for the assessment of the relationship between treatment and missing outcomes, we will impute the missing outcomes using the mean of the assignment-by-block subcategory.

## Ethical clearance

This research received clearance form the National Committee on Research in the Social Sciences and Humanities (P.01/22/615) as well as from IFPRI IRB (DSGD-22-0208).

## Transparency and replicability

To maximize transparency and allow for replicability, we use the following strategies:

- pre-analysis plan: the current document provides an ex-ante step-by-step plan setting out the hypothesis we will test, the intervention we will implement to test these hypothesis, the data that will be collected and specifications we will run to bring the hypotheses to the data. This pre-analysis plan will be pre-registered at the AEA RCT registry.
- revision control: the entire project will be under revision control (that is time stamped track changes) and committed regularly to a public repository (github).
- mock report: After baseline data is collected, a pre-registered report will be produced and added to the AEA RCT registry and GitHub. This report will differ from the pre-analysis plan in that it already has the tables filled with simulated data (drawn from the baseline). The idea is that after the endline, only minimal changes are necessary (basically connecting a different dataset) to obtain the final result, further reducing the opportunity of specification search.

## Epilogue

Right before endline data collection, it was decided to allow for (potential) follow-up in the longer run. This included re-treating some of the farmers. However, instead of simply repeating the treatment according to the original randomization, we also decided to include variation in treatment over time (year 1 versus year 2). In particular, we wanted to differentiate between farmers that were never treated ((CC)); farmers that were always treated ((T1T1) and (T2T2)); farmers that were treated only in the first year ((T1C) and (T2C)); and farmers that started receiving treatment in the second year ((CT1) and (CT2)).<sup>2</sup>

- In control households in year 1, we assign 28 percent to the first treatment (CT1) and 28 percent to the second treatment (CT2), while the remaining 44 percent remains in the control group in the second year (CC).
- Among households that received the first treatment in year 1, we assign 50 percent to the first treatment again in year 2 (T1T1). The remaining 50 percent did not receive the treatment again (T1C)

---

<sup>2</sup>In addition to measuring a dosage effect (once versus twice treated) the decision was also influenced by the fact that we found significant imbalance on a range of baseline characteristics for the second treatment. We hope that concerns related to this imbalance can be addressed by comparing differences between C and T2 in year one to differences in year 2 (comparing CC and CT2).

- Among households that received the second treatment in year 1, we assign 50 percent to the second treatment again in year 2 (T2T2). The remaining 50 percent did not receive the treatment again (T2C).

## References

- Aker, J. C., R. Boumnijel, A. McClelland, and N. Tierney. 2016. “Payment Mechanisms and Antipoverty Programs: Evidence from a Mobile Money Cash Transfer Experiment in Niger.” *Economic Development and Cultural Change* 65 (1): 1–37.
- Anderson, M. L. 2008. “Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects.” *Journal of the American statistical Association* 103 (484): 1481–1495.
- Ashraf, N., D. Karlan, and W. Yin. 2006. “Tying Odysseus to the Mast: Evidence From a Commitment Savings Product in the Philippines\*.” *The Quarterly Journal of Economics* 121 (2): 635–672.
- Balboni, C., O. Bandiera, R. Burgess, M. Ghatak, and A. Heil. 2021. “Why Do People Stay Poor?\*” *The Quarterly Journal of Economics* 137 (2): 785–844.
- Buehler, R., D. Griffin, and J. Peetz. 2010. “The planning fallacy: Cognitive, motivational, and social origins.” In “Advances in experimental social psychology,” vol. 43, 1–62. Elsevier.
- Burke, M., L. F. Bergquist, and E. Miguel. 2018. “Sell Low and Buy High: Arbitrage and Local Price Effects in Kenyan Markets\*.” *The Quarterly Journal of Economics* 134 (2): 785–842.
- Cardell, L. and H. Michelson. 2020. “*Sell Low, Buy High?*” - A New Explanation for a Persistent Puzzle. 2020 Annual Meeting, July 26-28, Kansas City, Missouri 304448, Agricultural and Applied Economics Association.
- Channa, H., J. Ricker-Gilbert, S. Feleke, and T. Abdoulaye. 2022. “Overcoming smallholder farmers’ post-harvest constraints through harvest loans and storage technology: Insights from a randomized controlled trial in Tanzania.” *Journal of Development Economics* 157: 102851.
- Dillon, B. 2021. “Selling Crops Early to Pay for School: A Large-Scale Natural Experiment in Malawi.” *Journal of Human Resources* 56 (4): 1296–1325.
- Dillon, B., J. De Weerd, and T. O’Donoghue. 2020. “Paying More for Less: Why Don’t Households in Tanzania Take Advantage of Bulk Discounts?” *The World Bank Economic Review* 35 (1): 148–179.

- Drexler, A., G. Fischer, and A. Schoar. 2014. “Keeping It Simple: Financial Literacy and Rules of Thumb.” *American Economic Journal: Applied Economics* 6 (2): 1–31.
- Duflo, E., M. Kremer, and J. Robinson. 2011. “Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya.” *American Economic Review* 101 (6): 2350–90.
- Fafchamps, M. and R. V. Hill. 2005. “Selling at the Farmgate or Traveling to Market.” *American Journal of Agricultural Economics* 87 (3): 717–734.
- Fink, G., B. K. Jack, and F. Masiye. 2020. “Seasonal Liquidity, Rural Labor Markets, and Agricultural Production.” *American Economic Review* 110 (11): 3351–92.
- Humphreys, M., R. Sanchez de la Sierra, and P. van der Windt. 2013. “Fishing, Commitment, and Communication: A Proposal for Comprehensive Nonbinding Research Registration.” *Political Analysis* 21 (1): 1–20.
- Karlan, D., R. Osei, I. Osei-Akoto, and C. Udry. 2014. “Agricultural Decisions after Relaxing Credit and Risk Constraints \*.” *The Quarterly Journal of Economics* 129 (2): 597–652.
- McKenzie, D. 2012. “Beyond baseline and follow-up: The case for more T in experiments.” *Journal of Development Economics* 99 (2): 210–221.
- Muralidharan, K., M. Romero, and K. Wüthrich. 2019. *Factorial designs, model selection, and (incorrect) inference in randomized experiments*. Tech. rep., National Bureau of Economic Research.
- Omotilewa, O. J., J. Ricker-Gilbert, J. H. Ainembabazi, and G. E. Shively. 2018. “Does improved storage technology promote modern input use and food security? Evidence from a randomized trial in Uganda.” *Journal of Development Economics* 135: 176–198.
- Sankoh, A. J., M. F. Huque, and S. D. Dubey. 1997. “Some comments on frequently used multiple endpoint adjustment methods in clinical trials.” *Statistics in medicine* 16 (22): 2529–2542.
- Sharot, T. 2011. “The optimism bias.” *Current biology* 21 (23): R941–R945.
- Stephens, E. C. and C. B. Barrett. 2011. “Incomplete Credit Markets and Commodity Marketing Behaviour.” *Journal of Agricultural Economics* 62 (1): 1–24.
- Van Campenhout, B. 2021. “The role of information in agricultural technology adoption: Experimental evidence from rice farmers in Uganda.” *Economic Development and Cultural Change* 69 (3): 1239–1272.