American Economic Review 2022, 112(8): 2518–2552
https://doi.org/10.1257/aer.20210122

Market Access and Quality Upgrading: Evidence from Four Field Experiments†

By Tessa Bold, Selene Ghisolfi, Frances Nsonzi, and Jakob Svensson*

Smallholder farming in many developing countries is characterized by low productivity and low-quality output. Low quality limits the price farmers can command and their potential income. We conduct a series of experiments among maize farmers in Uganda to shed light on the barriers to quality upgrading and to study its potential. We find that the causal return to quality is zero. Providing access to a market where quality is paid a market premium led to an increase in farm productivity and income from farming. Our findings reveal the importance of demand-side constraints in limiting rural income and productivity growth. (*JEL* C93, L14, L15, L22, O13, Q12, Q13)

Smallholder farmers in low income countries produce and sell output of low quality. Low quality limits the price farmers can command and can help explain why the returns to smallholder farming are low. At the same time many experts and policy makers argue that quality upgrading is key to raising income and productivity and the World Bank promises "double dividends" to poorer countries that participate in global value chains (World Bank 2020). Yet, few farmers upgrade the quality of their produce. This, in turn, suggests that either the link between quality upgrading and higher income is not as strong as hypothesized or constraints—possibly both on the demand and the supply side—trap farmers in a low quality-low productivity equilibrium.

This paper conducts four experiments among smallholder maize farmers in western Uganda to shed light on the impediments to quality upgrading at the farm level

* Bold: IIES, Stockholm University, Stockholm, Sweden (email: tessa.bold@iies.su.se); Ghisolfi: Catholic University of the Sacred Heart, Milan, Italy (email: selene.ghisolfi@unicatt.it); Nsonzi: Department of Food Technology and Nutrition, Makerere University, PNDK Solutions, Kampala, Uganda (email: nsonzif@gmail.com); Svensson: IIES, Stockholm University, Stockholm, Sweden (email: jakob.svensson@iies.su.se). Esther Duflo was the coeditor for this article. We gratefully acknowledge two anonymous referees for many valuable comments and suggestions. We have also benefited from comments and suggestions from seminar participants at University of Michigan, Bristol University, University of Zurich, IIES Stockholm University, CEPR-VDEV Webinar, and CEPR/LEAP Workshop in Development Economics. Financial support from the Swedish Research Council (2017-05665) and Riksbankens Jubileumsfond (M16:0345:1) is greatly appreciated. The trial was registered in the AEA RCT Registry as study AEARCTR-0002812 and received institutional review board approval from the Makerere University School of Social Sciences Research Ethics Committee (MAKSS REC), study protocol MAKSS REC 08.18.207, and the Uganda National Council for Science and Technology (UNCST), registration SS 4765. The data and code for the study are available through the AEA Data and Code Repository (Bold et al. 2022).

† Go to https://doi.org/10.1257/aer.20210122 to visit the article page for additional materials and author disclosure statements.

and to study its potential. In the first experiment, we measure the quality of maize sold at the farm gate and confirm that it is poor. The low quality of the final product sold to consumers can thus be traced back to farmers. We also show that maize quality measured through simple and quick tests at the farm gate strongly predicts maize quality measured by more elaborate laboratory tests. Maize quality is thus, at least partly, observable.

In the second experiment, we randomly assign farmers into two groups and offer treatment farmers a service package to improve the quality of the maize they sell. We find no return to selling high quality maize: maize of high and low quality earns the same price in local markets.

Being unable to access a market that rewards high quality may thus explain why farmers are loath to incur the extra cost and effort to increase quality. The low quality equilibrium we observe may also be caused by constraints on the farmers' side, however. Farmers may not be willing or able to produce maize of higher quality either because they are not aware of the required agricultural techniques or the quality standards, or because investing in quality is simply not profitable. These potential constraints, in turn, may help explain the lack of demand for high quality maize in the local markets. That is, buyers of high quality maize are not active in these markets simply because they do not expect to be able to procure maize of sufficient quality.

To test this hypothesis, we conduct a third experiment to investigate if and how farmers respond when offered access to a market where quality maize is paid a (market) premium. Following farmers over four years, including four posttreatment seasons, we find that a majority of farmers increase the quality of their maize and sell to the buyer of high quality when offered the opportunity. Treatment farmers' profit increases substantially—an effect driven both by increased farm productivity and higher prices.

The market access intervention was combined with a learning-by-doing extension service component to ensure that farmers had up-to-date knowledge about pre- and postharvest practices necessary for producing maize of sufficiently high quality. In a fourth experiment, we investigate the impact of the extension-service component alone. We find no evidence that farmers changed their farm practices as a result of this supply-side intervention: revenue, expenses, and yield remain essentially the same in the treatment and the control group. We cannot rule out that the extension component may have been required for the market access intervention to achieve its impact, and we even think it likely, but the comparison of the two trials shows that market access is necessary, and extension alone not sufficient, for agricultural transformation.

The market access intervention was designed to give treated households access to an output market for quality maize. To achieve this, we worked in close collaboration with a Ugandan vertically integrated agro trading company. The company committed to buy quality maize at a premium throughout the main buying seasons in treatment villages, with the premium determined by the difference in the amount of waste and defective grains in high versus low quality maize, valued at high quality prices. Because the market access intervention was randomized at the village level, we can also study how the intervention affected sampled farmers (in treated villages) who did not sell to the high quality buyer. We find that providing farmers with access to a buyer of high quality maize also resulted in higher prices for sales to other (local) traders. This effect raised prices for farmers continuing to sell average quality maize. Adjusting for selection using a difference-in-difference approach, the higher price for sales to local buyers can account for approximately 30 percent of the increase in average prices in the treatment group. Despite the higher price for lower quality maize, the evidence suggests that the increase in output in the treatment group can fully be accounted for by the subgroup of farmers selling premium quality maize to the high quality buyer.

Our results relate to a number of recent papers on the implications (for farmers or firms) of market (buyer) driven quality upgrading in a developing country setting. As in Atkin, Khandelwal, and Osman (2017), we exploit experimental variation in access to a market/buyer of quality products. Their intervention, which connects established small rugs producers in Egypt to foreign buyers paying a premium for higher quality rugs, led to large improvements in both quality and productivity. Here we focus on poor smallholders working with little or no physical capital. Further, we worked closely with an agro-trading company that makes direct outreach at the village level. Our village level intervention allows us not only to assess the direct effect of the intervention, but also to assess local market spillover effects. Finally, to allow farmers the time to upgrade quality and adjust their practices, and to build trust between the farmer and the buyer, we designed the intervention to run for four seasons and thus assess the implications of having access to a market for quality over a longer period.

Our paper is also related to Macchiavello and Miquel-Florensa (2019), who employ a difference-in-difference strategy to estimate the impact of a quality upgrading program in Colombia. While the intervention we exploit here also involves a vertically integrated domestic buyer—in our case a firm who buys quality maize at the farm gate and sells high quality maize flour in Kampala—the intervention, at the farm gate, was designed for research purposes. This enables us to directly measure and identify smallholders' choices and outcomes using experimental variation. By comparing outcomes from the market intervention with the extension-service only intervention, we can also, at least partly, unbundle the impact of demand and supply (extension service) factors.[1]

Knowledge about the preconditions and determinants for agricultural technology adoption has grown vastly over the last decade; for a review, see de Janvry, Sadoulet, and Suri (2017). The evidence—drawn primarily from randomized controlled trials in sub-Saharan Africa and South Asia—makes clear that there are productivity-enhancing supply side innovations that can increase technology adoption and productivity among smallholders. Effects on farmer income, however, have

[1] More broadly our findings relate to a macroeconomic and trade literature that studies effects of quality upgrading—often coupled with exporting—on productivity and growth. Empirical results (largely nonexperimental) are reviewed in, for example, De Loecker and Goldberg (2014). Ashraf, Giné, and Karlan (2009) experimentally evaluate an intervention in Kenya that helped farmers to adopt and market export crops. The authors find small average effects on adoption and income. Our study also relates to the (mainly nonexperimental) literature on the effects of market access and market integration reviewed in Donaldson (2015) as well as an emerging experimental literature that studies how attempts to improve the functioning of local markets (through market integration and increased competitiveness) affect prices and farmer welfare (see, e.g., Bergquist and McIntosh 2021).

been much more limited.[2] We add to this literature by studying the impact of lifting both a demand (inclusion in a value chain) and a supply (knowledge) constraint. While supply-side interventions, like the extension service component considered here, remain at the core of agricultural development programs throughout the developing world, our findings suggest that pre- and postharvest training on best agronomic practices alone has little impact on productivity or income without market access.

Our findings should be interpreted within the context of the study—a poor area of western Uganda. All farmers participating in the experiments are smallholders and farm largely with traditional methods. This can help explain why farmers raised yields predominantly by employing more labor and producing a higher output per hour worked, rather than by increasing the application of inputs. The setting may also have contributed to the large increase in farmers' income: by targeting the poorest farmers who are currently excluded from global value chains, the agricultural trading company provided market access to precisely those farmers who have the largest potential to benefit.

To implement the intervention and to circumvent potential agency and information problems, the collaborating company ran a vertically integrated operation. As quality upgrading is potentially a core motive for why firms change their organizational structure, the intervention provides a case study of the constraints of buying, processing, and selling quality maize for a vertically integrated domestic buyer. It does not, however, allow us to study the behavior of the layers of intermediaries that dominate the low quality segment of the market.[3]

I. Context

Uganda remains highly dependent on agriculture: the sector contributes over 70 percent of export income and 65 percent of the population is active in the sector.[4] As in most countries in the region, the agricultural sector is dominated by smallholder farmers, a majority of whom cultivate less than two hectares. Maize is the most important cereal crop and grown primarily as a cash crop. Smallholder farmers account for roughly 75 percent of maize production and 70 percent of marketable surplus (Daly et al. 2016).

Maize has different end uses depending on the geographic region of the producers (Daly et al. 2016). In the United States only 12 percent of maize produced is used for human consumption, with the remainder split between animal feed and

[2] Karlan et al. (2014) find that farmers with insurance invest more in cultivation, but without any significant impacts on profit. Beaman et al. (2013) show that receiving the recommended dose of fertilizer increases harvest value by 30 percent but profits only by 0.5 percent. Deutschmann et al. (2019), in an evaluation of a bundled program of training, input loans, and crop insurance, document an increase in profit of 8–16 percent after one season.

[3] There is a complementary literature examining the structure and competitiveness of intermediaries both downstream (as sellers) and upstream (as buyers). Regarding upstream competitiveness, Casaburi and Reed (2017) estimate that the local markets in which traders operate are highly competitive with a low differentiation parameter. Bergquist and McIntosh (2021) find small effects on trader prices and margins of providing farmers with a platform allowing them to sell to any buyer available. Dillon and Dambro (2017) conclude in their review of a wide variety of agricultural markets that producers tend to sell on fairly competitive markets.

[4] According to official statistics, maize exports accounted for about 2 percent of the country's total exports (UBOS 2015). Based on interviews with stakeholders in the sector, Daly et al. (2016) estimate that 70–80 percent of maize that is bought and sold in Uganda is channeled through informal channels.

ethanol fuel production (Ranum, Peña-Rosas, and Garcia-Casal 2014). In Africa, and especially in East Africa, maize is a staple food crop, accounting for nearly half of the calories and protein consumed (Macauley 2015). While maize grain of the lowest quality is also used for animal feed, the main feed ingredient sold in markets is maize bran—a byproduct of flour production or grits manufactured from maize grain.

Maize, as most other crops in Uganda, is produced using mainly traditional techniques and few farmers use modern inputs such as hybrid seeds and fertilizer. Yields also tend to be low. For example, Bold et al. (2017) report an average yield for smallholders of 1.4 metric tons per hectare As a comparison, average maize yield based on data from farm demonstrations in Uganda is over 4 tons per hectare (World Bank 2007) and average corn yield in the United States was close to 12 tons per hectare in 2017 (USDA 2019).

The research program is set in an area of western Uganda (Kakumiro and Kibaale districts), where smallholder maize farming is common. Rural Kibaale is poor, with an average consumption expenditure of US$0.80 per day (UBOS 2019).[5]

A. *Local Markets for Maize*

The local, or village, maize market for smallholders can be described as a spot market.[6] The farmer and the buyer agree right before the sale, usually after a short visual inspection of the maize bags by the buyer, about the amount and the price. The farmer is paid directly and the transaction takes place at the farm gate.

There are two types of buyers active in these local markets: (i) local traders or aggregators who often buy from a smaller set of farmers and resell to commercial traders/aggregators that are either passing through the village or located in a nearby trading center, and (ii) commercial buyers, who typically pass through the village with a truck, and buy directly from individual farmers (and local traders).

Over the five seasons for which we collected detailed sales data, 78 percent of the sales went to local traders (see online Appendix Table 1). Still, more than half the smallholders in the sample have sold to a commercial trader at least once during the last five seasons. A sale to a commercial trader is associated with, on average, an 8 percent higher price than a sale to a local trader. On average, farmers sold 82 percent of what they produced, and kept 18 percent of maize for home consumption (or to give away).

Farmers tend to know the local traders they sell to and repeated transactions across seasons are common: 98 percent of smallholders sold to the same buyer in at least two out of the last five seasons and 31 percent sold to the same buyer in at least four out of five seasons. Most households sell once per season (79 percent) and, accounting for multiple sales to the same buyer in a given season, 90 percent sell to only one trader per season.

[5] Kakumiro was created in 2016 from the split of Kibaale district and separate statistics are not available.

[6] We use data collected from the control group in the market experiments (sample frame 1) discussed below to describe the local market context. See online Appendix A for more details on the data used.

II. Sample Frames and Research Design: Overview

We combine field experiments with maize quality measurement using laboratory techniques and visual inspections to answer four questions. What is the quality of the maize produced and sold at the farm gate? Is quality rewarded in local markets? If offered access to a market where quality maize is paid a (market) premium, and provided with basic training on pre- and postharvest best practices, will farmers respond by producing higher quality, and if so what are the consequences for income and productivity? Can similar impacts be achieved from a training program alone?

Figure 1 in online Appendix H illustrates the design of the study, including an overview of the sample frames, and the timing of the surveys and interventions.

We draw on data from three sample frames. From sample frame 1, we randomly selected 20 communities, each at least five kilometers (km) apart, from Kakumiro district in western Uganda and assigned these to treatment and control villages for the market access plus extension experiment. From sample frame 2, we randomly selected 20 communities from neighboring Kibaale district, also at least five km apart, to participate in the extension-only experiment. For sample frame 3, we identified all villages adjacent to the control group villages in sample frame 1 and randomly selected 20 communities (sample frame 3).[7] In this sample, we measure the quality of maize sold in local markets and infer the causal return to quality upgrading in local markets. We also use a random subset of the sample as a quasi-control group for counterfactual maize quality in the market access experiment.

For each of the selected villages, we completed a census and identified smallholder farmer households (with maize gardens of no more than five acres of land) who cultivated maize in the previous season. We then sampled about ten farmers per village.

III. Maize Quality and Verifiability of Quality

Maize is sold and handled in large quantities, with the smallest unit typically a 100–170 kilogram (kg) bag. A bag of maize is considered high quality if it contains sufficiently large and dry maize kernels of the right color and neither nongrain substances (e.g., stones, dirt, and insects) nor defective (e.g., broken, immature, damaged, rotten, or moldy) grain. Maize quality in East Africa is classified formally according to the East African Grading Standard (EAS) (East African Community 2011), which divides maize into three broad quality categories based on moisture level and amount of nongrain substances and defective grain: graded maize, undergrade maize, and reject maize (we label the latter two as ungraded maize). Graded maize (quality maize) is further categorized into three grades: grades 1, 2, and 3, with grade 1 having the most stringent thresholds for defects.

The quality of maize determines its economic and nutritional value and whether it is safe for human consumption. The presence of nongrain substances and defective grain adds to the weight of the bag without adding value and increases transport and processing costs. Nongrain substances and defective grain are also indicators that

[7] Specifically, we sampled census enumeration areas, i.e., villages or parts of villages, to which we will refer to as "villages" henceforth, from digital maps of Kakumiro and Kibaale districts.

the maize has not been properly handled and may therefore be unsafe for consumption. For example, stones and dirt in the bag indicate that the farmer has stored or dried the maize directly on the ground, raising the risk that grains are contaminated by microorganisms. Insect parts or insect waste, pest damaged, rotten, diseased, and discolored grains are direct indicators of (acute) infestation. A particular concern is contamination with aflatoxin; poisonous carcinogens that are produced by certain mold species that inhabit the soil.[8] Contamination and infestation can spread quickly through the bag and may still be present even if waste and visibly defective kernels are sorted away at a later stage.

The maize's moisture content influences quality through the same two channels: by increasing gross (but not dried) maize weight and by raising the risk of infestation. Dry grains keep longer, are attacked by insects less often, and usually do not support mold growth. In wet grains, on the other hand, fungal growth and release of mycotoxins can occur quickly, especially during storage. Aflatoxin contamination can increase tenfold in just a few days if maize grain is not dried properly (Hell et al. 2008).

Both economic and safety considerations dictate the prevention of defects and control of moisture in the maize as early as possible in the value chain. Importantly, farmers' practices crucially affect maize quality: by harvesting at the right time, drying the maize quickly and thoroughly, shelling the cob without breaking or cracking the grains, drying or storing grain on tarpaulins rather than the bare ground, and cleaning and storing the grain correctly, they can ensure that the maize is free of waste, defects, and infestation.

A. *Quality Measurement*

We conducted four different measurements to assess maize quality and its observability: visual inspection and moisture measurement at the farm gate, and laboratory testing and aflatoxin measurement in Kampala (details on the test protocol are in online Appendix B).

For the measurement experiments, we enrolled 100 farmers who were about to harvest their maize from nine villages in sample frame 3 for one season. We assigned half of them to the treatment group in the returns to quality experiment discussed in the next section and half of them to a control group. Here, we focus on the subset of control households (50 households). In addition, we sampled 30 households from a further six villages in sample frame 3 over two consecutive seasons. In both samples, maize moisture was measured at the farm gate and maize quality was tested in the lab. Visual inspections were only conducted in the first sample and aflatoxin levels were only measured in the second.

To inspect maize quality at the farm gate and relate it to lab measured quality, trained enumerators visited each farmer at the time of sale and checked each bag put up for sale for the presence of ten types of defects. Specifically, the enumerators

[8] There is a large literature testing for the presence of aflatoxin in crops like maize. For a recent discussion of the literature on aflatoxin and health, see de Almeida et al. (2019). For a recent summary of evidence of aflatoxin measurement in Uganda, see Sserumaga et al. (2020). For research on the consequence of the unobservability of aflatoxin, see Hoffmann et al. (2013).

recorded whether the maize in the bag was dirty, included cobs, stones, dust, or insects (live or dead), and whether the grain was broken, immature, damaged, rotten, or mold-infested. The enumerators then bought one randomly selected bag per farmer, which was transported to Kampala for laboratory testing. In the lab, samples of 200 grams (g) were drawn from each bag, and the weight of all nongrain substances and defective grain recorded.[9] We denote the mean number of types of defects found in the bags put up for sale as "visually verifiable defects" and the share of the 200 g sample that was defective as "lab verified defects."

To diagnose aflatoxin levels and relate them to visible defects measured in the lab, we randomly sampled an additional 30 households from six sample frame 3 villages over two consecutive seasons. Enumerators purchased one randomly selected bag per farmer and sent it to the laboratory in Kampala. Again, the laboratory measured the weight of all nongrain substances and defective grain and tested whether aflatoxin levels exceed the limit imposed by the Uganda National Bureau of Standards (ten ppb).[10]

The field enumerators also measured moisture levels in the bags destined for the lab using a mobile moisture meter. We generate a binary indicator labeled "wet maize," taking the value one for maize bags with a moisture content above 13 percent.

Finally, we combined the lab verified defects and the moisture indicator to classify all samples tested in the lab using the EAS.

B. *Results: Quality of Maize at the Farm Gate*

The quality of the maize farmers sell is low (see Table 1). The average bag inspected at the farm gate contained 2.5 defects (out of 10 possible) and the maize samples tested in the lab contained on average 26.2 percent visible defects; i.e., a quarter of the weight of maize sold consists of waste, with defective grains constituting the majority of defects. Additionally, 29.4 percent of the households sold maize with a moisture content higher than 13 percent.

The results for grading the lab samples according to the EAS classification are reported in the last four rows of Table 1. None of the bags contained grade 1 grain (the highest grade), 20.6 percent of the bags contained maize of grade 2 quality and 8.8 percent contained maize of grade 3 quality. The remaining bags, 70.6 percent, contained undergraded or rejected maize.

Quality measured at the farm gate predicts quality measured in the lab, as shown in Figure 1 panel A (the corresponding regression is reported in online Appendix G Table 2), especially at higher levels of defects. When the number of defects found in the bag increases from zero to two, the percentage of waste in the lab sample increases from 10 to 15 percent. As the number of defects doubles from two to four, the percentage of waste in the lab sample also doubles.

[9] The lab testing protocol followed the EAS approved objective test methods for defects. Forty-four bags were tested in the lab (see online Appendix B).

[10] AflaCheck test kit (VICAM) was used to detect the presence of aflatoxin. The test strips can detect aflatoxin at two different cutoff levels depending on the protocol followed. We used the 10 ppb (parts per billion) cutoff level, which is the limit imposed by the Uganda National Bureau of Standards (UNBS). As a reference, the European Union standard is 4 ppb and the US standard is 20 ppb (Sserumaga et al.2020).

Table 1—Maize Quality: Summary Statistics

	Mean (1)	St. dev. (2)	Min (3)	Max (4)	Observations (5)
Visually verifiable quality (number)	2.5	1.5	0	7	355
Lab verified quality (%)	26.2	34.1	4	100	102
Wet maize (%)	29.4		0	100	109
Grade 1: EAS (%)	0				102
Grade 2: EAS (%)	20.6				102
Grade 3: EAS (%)	8.8				102
Ungraded: EAS (%)	70.6				102

Notes: This table presents summary statistics on maize quality as described in Section III. Unit of observation is a bag. Visually verifiable quality is the number of defects out of ten detected in a bag of maize. Lab verified quality is grams of defects per 200 g maize sample drawn from bags bought in the field (in percent). Wet maize is a binary indicator for maize with a moisture level above 13 percent. Grades 1–3 and ungraded (undergraded/reject) maize are East African Quality Standard (EAS) classifications for maize quality, with grade 1 having the most stringent thresholds for defects (see online Appendix B for details).

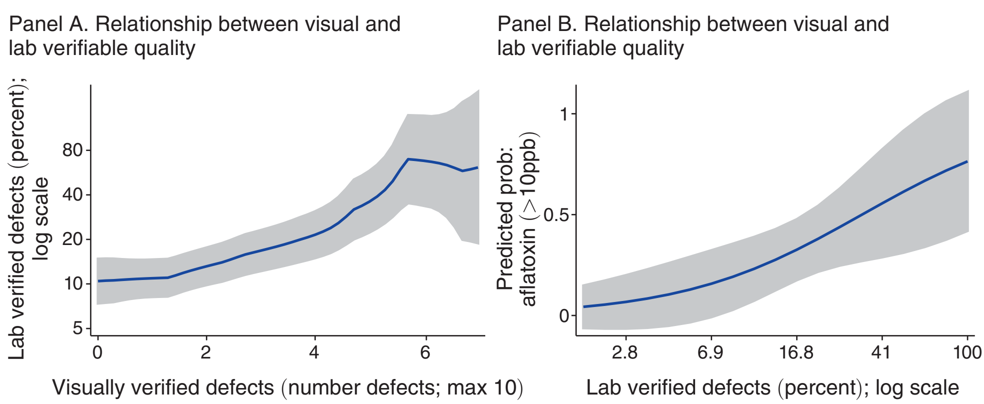

Figure 1. Verifiability of Maize Quality

Notes: This figure shows the smoothed values from a local polynomial regression of lab verified defects on visually verifiable defects (panel A) and the predicted probability of aflatoxin above the UNBS cutoff as a function of share of lab verified defect (panel B). Gray shaded area represents 95 percent confidence intervals. The predicted probability in panel B is from specification (3), in online Appendix G, Table 2. The unit of observation is a bag. See online Appendix B for details on the tests and samples used.

Lab measured defects, in turn, predict whether the sample contains dangerous levels of aflatoxin. Figure 1 panel B (specification (3) in online Appendix G, Table 2) plots the predicted probability of aflatoxin levels exceeding ten ppb as a function of lab verified defects. The relationship is roughly log-linear: as the share of the sample that is lost to waste and defects doubles, so does the predicted probability that aflatoxin levels are too high.

In sum, smallholders tend to sell maize of low and possibly unsafe quality. While testing for quality in the lab is costly, visual assessments that are easy and quick to execute at the farm gate can provide a proxy measure.

IV. Returns to Quality Experiment

If the economic value of maize depends on its quality, why is the quality of maize sold by farmers so low? A starting point to answer this question is the neoclassical agricultural household model. In this model, with complete markets, the production decision is separable from the consumption decision. Thus, a utility maximizing farmer chooses a vector of inputs to maximize profit. Consider a version with two inputs, x and z, with unit costs c_x and c_z, where input x primarily affects the quantity of output while input z primarily affects the quality. The farmer's decision problem can be stated as

$$\max_{x,z} \Pi = p(q(z))F(x) - c_x x - c_z z, \tag{1}$$

where $p(q)$ is the price as a function of quality q, $F(x)$ is output (assuming that land is a fixed factor). The farmer's choice of inputs is given by two first-order conditions:

$$p(\cdot)F'(x) - c_x \leq 0 \quad \text{and} \quad p'(\cdot)q'(z)F(x) - c_z \leq 0. \tag{2}$$

That is, the farmer will set the intensity of use of any particular input until its marginal value product equals its marginal cost. Thus price, or more precisely, the responsiveness of price to quality, is a key driver of the decision to produce high (or low) quality maize. But does the (local) market reward quality? To answer this question, we designed an experiment to measure the returns to quality.

A. *Intervention*

Low quality at the farm gate is determined by a number of factors, several of which the farmer can influence through good agronomic practices in harvesting, drying, decobbing, cleaning and storing the grain. To create random variation in the quality of maize farmers sell, we therefore developed a service package, which included assistance with several key harvest and postharvest (drying, winnowing, and sorting) activities. The services offered were implemented by agricultural workers with access to portable agricultural machinery and were managed by staff from the research team.

B. *Experimental Design and Data*

We attempted to enroll 100 maize farming households that had not yet begun harvesting, 99 of whom gave consent. The households were randomly assigned into treatment (49 households) and control (50 households) groups. At enrollment, a short postplanting survey was administered. Online Appendix G Table 3 compares preharvest outcomes between treatment and control groups in the experiment. None of the covariates show statistically significant differences across assignment and a joint balance test fails to reject the null hypothesis that the preharvest outcomes do not predict the assignment to treatment.

Before harvest, farmers in both groups were visited by staff from the research team. In the treatment group, households were offered the free service package. The offers were presented as a service from the research team and all treatment households

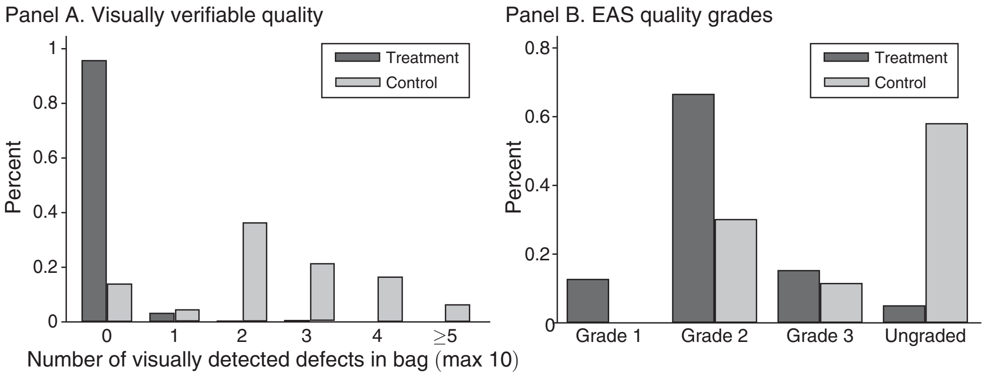

Figure 2. Quality Outcomes in Treatment and Control

Notes: This figure shows the distribution of visually verifiable quality, i.e., the number of visually detected defects (max ten) in maize bags (panel A) and the distribution of quality grades based on the EAS classification, using data from laboratory measurement of purchased maize bags (panel B).

accepted the offer. The households were also asked to contact the research team at the time of bagging the maize just before selling it, and were promised a reward of UGX10,000 (approximately US$3) if they did so. Farmers in the comparison group were also visited and offered a (larger) monetary reward (UGX30,000; approximately US$9), if they contacted the research team before selling their maize.

When the farmer was ready to sell, they were visited again, this time by trained enumerators who visually verified the presence of ten types of defects in all bags the farmer was planning to sell. The enumerators also weighed the bags and tested the moisture content using a mobile moisture meter. Altogether 622 bags were visually inspected (see online Appendix G Table 4). In addition, one bag, drawn at random, was bought from each farmer for further quality analysis in the lab. After selling their maize, all farmers were visited for a third and final time and asked about sales volume and prices. In total, we collected data on 116 sales from 94 households.[11]

C. *Results: Returns to Quality*

The service package successfully raised maize quality in the treatment group. As illustrated in Figure 2 panel A, the average sale in the control group contained 2.5 types of defects (out of a possible 10) per bag, ranging from 0 to 7 defects. Eighty-six percent of the bags contained at least one defect. In the treatment group, the average sale contained less than 0.05 defects per bag and only 4 percent of the bags contained one or more types of defects, yielding a mean difference in defects of 2.2 ($p = 0.000$); see Table 2, column 1. Quality differences between treatment and control groups were equally stark when assessing quality using laboratory testing. One-third of the content in the average bag brought in for testing in the control

[11] One farmer decided not to sell any maize in the season under study and of the remaining 98 farmers, 4 could not be reached.

Table 2—Returns to Quality: Treatment Effects

	Defects (visually verifiable) (1)	Defects (lab) (2)	Price (3)	Deductions (4)	Net price (5)
Treatment	−2.16	−0.20	2.95	0.18	0.98
	(0.212)	(0.054)	(9.870)	(0.633)	(10.621)
	[0.000]	[0.000]	[0.766]	[0.780]	[0.927]
Observations	622	82	116	116	116
Households	99	82	94	94	94
R^2	0.67	0.36	0.91	0.22	0.89
Mean control	2.46	0.32	731.25	3.73	704.63

Notes: This table presents estimates of the treatment effects in the returns to quality experiment. Unit of observation is a bag in column 1, a random sample from one randomly selected bag per household in column 2, and maize sale in columns 3–5. Defects (visually verifiable) is the number defects (out of ten) verified by trained enumerators in maize bags for sale. Defects (lab) is grams of defects per 200 g maize. Price is in Ugandan shillings per kilogram. Deductions are defined as $(y - z)/y$, where y is the weight of maize sold as measured by enumerators and z is the agreed upon sales volume, expressed in percent. Net price is price net of any weight deductions in Ugandan shillings per kilogram. Specifications (1) and (2) include villages fixed effect. All sales specifications include village fixed effects and week fixed effects. Robust standard errors in column 2 and clustered by household standard errors in parentheses in columns 1, 3–5; p-values in brackets.

group consisted of defective maize grain and waste. The mean in the treatment group is 20 percentage points lower. (see Table 2, column 2). Finally, only one of the bags (2 percent of the bags) purchased in the treatment group contained more than 13 percent moisture, while 12 percent of the bags in the control group did.

In Figure 2 panel B we compare maize quality in treatment and control on the basis of the EAS standard-maize grain classification system (see Section IIIA). Ninety-five percent of the lab-tested bags in the treatment group were graded maize, of which 79 percent were grade 1 or grade 2 maize. In the control group, 58 percent were classified as ungraded maize, and the remainder as either grade 2 (30 percent) or grade 3 quality (12 percent).

Despite the large differences in both visually and lab verified quality, buyers did not pay higher prices to farmers who had received the service package. Figure 3 panel A plots the cumulative distribution functions (CDFs) of price in the two assignment groups. The two CDFs lie effectively on top of each other and the Kolmogorov-Smirnov test fails to reject the null that the two distributions are equal (D statistic $= 0.11$, $p = 0.84$). The nonparametric results are confirmed by regression analysis. Table 2 column 3 regresses price on the treatment indicator, controlling for village and week of sale fixed effects. The unit of observation here is a sale. The treatment effect is essentially zero; i.e., there is no evidence that higher maize quality—equivalent to fewer defects detected in the bag—systematically yields a higher price. The coefficient is also tightly estimated, with the 95 percent confidence interval spanning a price change of between 3.1 percent to −2.2 percent relative to the control group.

Why do traders not pay higher prices for better quality maize? A first possible explanation is that traders cannot assess maize quality and therefore do not adjust prices. However, traders do conduct visual inspections of maize bags, and we showed in Section IIIB that there is a strong relationship between visually and lab verified maize quality. Thus "true" quality is—at least partly—observable at the farm gate.

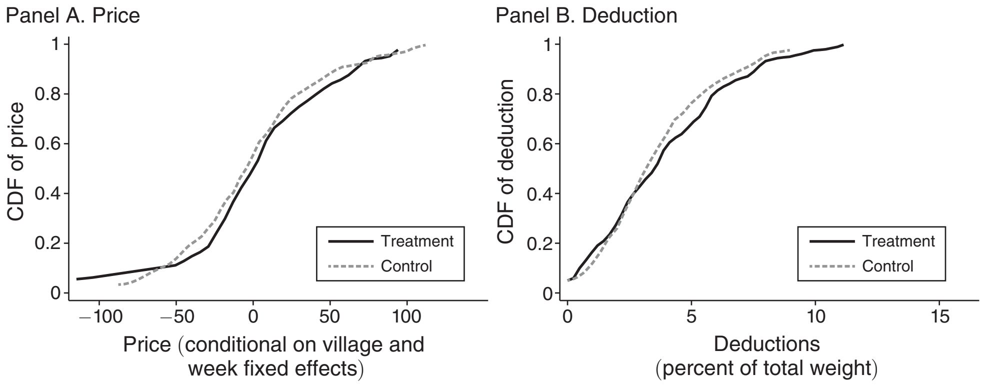

Figure 3. Price and Deduction in Treatment and Control

Notes: This figure shows the CDFs of price conditional on village and week fixed effects (in panel A) and deduction (in panel B) in the two assignment groups. The graph caps the bottom 1 percent observations. Deduction is defined as $(y - z)/y$, where y is the weight of maize sold as measured by enumerators and z is the agreed upon sales weight. The Kolmogorov-Smirnov D statistic on the test of equality of the treatment and control distributions is 0.114 [$p = 0.844$] in panel A and 0.168 [$p = 0.388$] in panel B.

A second possible explanation for the absence of a relationship between farm gate quality and price is that traders respond to higher maize quality by reducing weight deductions rather than increasing prices.[12] To test for this possibility, Figure 3 panel B, plots the CDFs of deductions in the two assignment groups, with deduction defined as $(y - z)/y$, where y is the weight of maize sold as measured by enumerators and z is the agreed upon (or buyer stated) weight. While deductions are common (the mean is 3.7 percent and in one out of five sales more than 5 percent of the weight is deducted), the extent of deductions is similar across groups. Table 2, specification (4), regresses deductions on the treatment indicator and specification (5) shows the treatment effect on the net sales price, pz/y, which is the per kilogram price scaled by the ratio of the buyer stated weight to the enumerator measured weight. We find no evidence of a systematic relationship between quality and the net price and conclude that quality is not rewarded by lower weight deductions.[13]

There are other possible explanations for the absence of a quality-price relationship that our experiment cannot speak to. A recent literature has emphasized limited contract enforcement and informational asymmetries in both input supply chains and output value chains (Antrás 2015; Blouin and Macchiavello 2019; Bold et al. 2017), which could limit traders', and, in turn, farmers', incentives for quality upgrading. A more direct, but also complementary, reason is that the local traders

[12] Data on deductions and cheating are pooled together. Deductions refer to a transparent process whereby the buyer weighs the bag, reports the correct weight to the farmer, and then states the share of the content deducted before paying. Cheating refers to the trader using a rigged scale and reporting a weight below the true weight before (possibly) applying deductions. Of the total difference between the measured weight and the agreed upon quantity sold (3.7 percent), 2.9 percentage points are deductions. The remaining discrepancy (0.8 percentage points) is (possibly) due to cheating.

[13] Casaburi and Reed (2020) find, in the context of cocoa production in Sierra Leone, that traders may extend trade credit rather than adjust prices. We do not observe such interlinked contracts in the village markets.

are active in a segment of the value chain where the final product is low quality, and possibly even hazardous, maize flour, and therefore they do not place any additional value on premium quality.

Our experimental results do not rule out that buyers would reward quality over time if the seller sold higher quality maize repeatedly. That is, buyers may offer a price based on expected quality not actual quality. Over time, a seller may be able to acquire a reputation for high quality maize and buyers may be willing to pay for it. Even if this mechanism is at play, our results still suggest that the farmer would not be rewarded in the first season they upgrade quality, which lowers the return to upgrading.

V. Market for Quality Experiment

The results from the returns to quality experiment in Section IV show that farmers face weak incentives to invest in high quality. We would therefore expect them to invest little in enhancing maize quality and the market to be characterized by trade in low quality maize, as we document in Section III.

A possible interpretation of these findings is that the prevailing market equilibrium is primarily caused by a lack of demand for quality maize at the local level. Yet low quality may also be caused by constraints on the supply side. That is, farmers may not be willing or able to produce maize of higher quality either because they are not aware of the required agricultural techniques and quality standards or because investing in quality is simply not profitable. These supply side constraints, in turn, may be the underlying reason for the lack of demand for high quality maize: buyers of high quality maize are not active in local markets because they do not expect to be able to procure maize of sufficient quality. The findings from the returns to quality experiment thus give rise to two important follow-up questions: Will farmers produce higher quality if the market values it? What are the implications for farmer profit and productivity of quality upgrading?

To answer these questions we designed an intervention offering farmers (or rather villages) access to a market for quality maize. In this market, maize above a quality threshold $\overline{q}$ is paid a premium ω over and above the prevailing market price $\tilde{p}$, resulting in an inverse demand schedule $p(q) = \tilde{p}(1 + I_{q \geq \overline{q}} \omega)$, where $I_{q \geq \overline{q}}$ is an indicator function equal to one if quality exceeds the minimum threshold and zero otherwise.

In addition, extension services were offered to improve farmers' knowledge on how to produce higher quality maize (as well as increasing their general knowledge of best-practice pre- and postharvest agronomic activities). To examine whether impacts of the market access program are driven only by its extension component, and to estimate the direct effects of hands-on training on cultivating and refining a well-known crop, a parallel trial with only the extension service component was also implemented.

Below we describe both interventions in detail. We also discuss the trial designs in the two experiments, and the data we collected, before presenting the main findings.

A. *Interventions*

Market Access Intervention.—To provide treated farmers with access to a market for high quality maize, we collaborated with a Ugandan vertically integrated agro-trading company, which committed to buy quality maize from randomly selected farmers in the treatment villages. To this end, company agents contacted all predetermined households in person or by phone before the buying season commenced and were present in the villages throughout the buying season.[14] When a household was ready to sell, the household and the agent agreed on a time and buying took place at the farm gate. Agents visually inspected the maize, weighed it, and measured moisture with mobile moisture meters. Agents were not allowed to make deductions or bargain about the price. Instead, they were instructed to only buy maize bags that were of sufficient quality at a predetermined price and explain to farmers why bags of insufficient quality were rejected. The company then milled the maize into quality flour and sold it at a premium to customers in Kampala.

Since there is no local market for high quality maize and thus no local market price, the research team calculated the minimum premium a trader wanting to buy high quality maize must pay farmers to induce them to sell high quality maize. To do so, we used a simple model predicting the required premium as a function of observable outcomes: the extent of defects and prices for (average quality) maize paid by commercial traders and in nearby trading centers. In the model (see online Appendix D) farmers can produce and sell either low or high quality maize, with low quality maize containing waste and defective kernels. Only high quality maize can be processed into quality flour. We further assume that a farmer can costlessly turn low quality maize into high quality maize by sorting away defects and waste. We then solve for the premium above the price for low quality maize that makes the farmer indifferent between selling high or low quality maize. This premium is simply the share of defective kernels and waste in low quality maize, valued at the premium price.

Based on the quality measurement experiment we estimated the reduction in waste and defects (and thus gross weight) the company would have to compensate the farmers for, if the company only bought maize with no visually verifiable defects, as 14 percentage points relative to the average quality on the market. If moisture levels in the maize are below 13 percent, this reduction would be equivalent to buying grade 3 maize. We further assume, again based on pretreatment data, that local prices are on average 10 percent lower than prices in the trading centers. These assumptions yielded a target premium relative to trading center prices of 5 percent and an estimated premium relative to local (or village) prices of approximately 15 percent.

The quality premium we chose should be viewed as the lower bound on a market-based premium that satisfies the participation constraint of the farmer. First,

[14] The company could also buy from other farmers in the treatment villages, conditional on the household selling quality maize. Allowing the company to buy from all households in treatment villages minimizes concerns that a sampled farmer in treatment would transport maize from a nonsampled neighbor and sell it as if it was the farmer's own maize. The company did not buy from villages in the control group, from villages in sample frame 3, or from the villages in the extension-services-only trial. The research team strictly enforced this constraint.

since the premium is set such that the farmer is indifferent between producing high or low quality, their participation constraint binds. Second, the premium increases if we allow that sorting and cleaning away waste and defective kernels is costly. Third, we did not factor in that more waste and defective kernels, and higher moisture levels, increase the risk that the maize becomes contaminated by various microorganisms, and therefore cannot, ex post, be sorted and cleaned into a higher quality product. Fourth, as reported below, the share of defects in maize sold to the company was lower than our ex ante estimate.

Extension Service Intervention.—To ensure that farmers had up-to-date knowledge about the pre- and postharvest practices necessary for producing maize of sufficiently high quality, the agro-trading company organized an extension service program in all treatment villages. A smaller plot was selected in each village and with the help of an extension service agent, a demonstration garden was set up. Throughout the first two seasons, five meetings were held at the demo garden, during which the extension service agent provided hands-on training on best agronomic practices for plot preparation, planting, weed and pest management, and harvest and postharvest tasks. All treatment households were invited to the demonstrations, and close to 70 percent of the invitees attended the meetings in the market experiment, while 78 percent of the invitees attended the meetings in the extension service only experiment. Other households in treatment villages could participate in the training as well, but few did.

B. *Experimental Design and Data*

Trial Design: Market Experiment.—We chose a clustered repeated measurement design for the experiment with 20 clusters surveyed over seven seasons. The 20 clusters were randomly assigned to two groups: 12 to the buying group and 8 to the control group.

The design was motivated by several features of the local market and market access intervention. First, we chose a clustered rather than an individual design because we anticipated that the market access intervention could also impact households in the treatment clusters who chose not to upgrade quality. Second, we chose to expand on the number of waves at the expense of the number of clusters for three reasons: (i) the intervention, essentially the creation of an integrated value-chain, was complex, and costly; (ii) it may take time for farmers to decide to upgrade and/or build up a relationship with the new buyer; (iii) impacts may vary substantially from season to season because of large aggregate variations (see Rosenzweig and Udry 2000). Subject to these considerations, the final combination of clusters and waves was chosen to have sufficient power to detect moderate treatment effects.

The trial design is illustrated in online Appendix H Figure 1. The first three seasons serve as baseline. The intervention began at the end of the third season and ran for four consecutive seasons.

Trial Design: Extension Service Experiment.—The trial design for the extension service experiment followed the market experiment design closely with 18 clusters followed over six seasons (see online Appendix H Figure 1).[15] The clusters were randomly assigned into two equal sized groups. The intervention began at the end of the third season and ran for three consecutive seasons.

Data.—The overall objective of the data collection was to measure the components of a farmer's profit function. For the market access trial, we also measured maize quality.

We measured the components of the profit function using household survey data. The household surveys were implemented at the end of the selling season when farmers had either planted or prepared the plot(s) for planting for the following season. The size of all maize plots that households had prepared for maize planting, or had already planted maize on, were collected using GPS trackers. To improve recall of revenues and expenses, households were provided with a form from the second season onwards, listing all maize plots in the current season, to be filled in with inputs, labor use and sales data. In order to ensure data quality, GPS data from the previous season were preloaded in the survey form, and farmers were shown satellite photos of their measured plots to confirm the plot sizes. All calculations were checked by the survey form and any discrepancy was immediately checked and corrected.

Data were collected on the amount harvested, amount sold, and the price and revenue received. For farmers that sold multiple times, data were collected for each sale. The survey also collected detailed expense data, including on chemical use, seed varieties, and various preharvest and postharvest practices, referring to the most recent season. Labor expenses and hours were collected for hired and family labor, respectively.

The data collected by the survey firm contained several observations with large positive values. We cannot rule out that these observations are correct (the outliers were rechecked for coding errors), and they therefore remain in our core sample. As these outliers may have an undue influence on the results, however, we also estimate treatment effects with outliers removed, trimming the top (and in the case of profits, which could take on both large positive and large negative values, also the bottom) 1 percent of observations.

Estimating the causal effect on maize quality requires measuring maize quality for all households in treatment and control. But doing so in an experiment where maize buying is an integral part of treatment is problematic. Specifically, the measurement of maize quality involves testing at the farm gate and the purchase of bags for laboratory testing. To get the farmer to agree to sell only part of their output (one bag), farmers also need to be paid a premium price. Moreover, to accurately measure the quality of maize available on the local market, it must be measured at the time when farmers are ready to sell. This, in turn, requires constant presence in the village throughout the selling period.

Hence, the measurement of quality requires a buying operation very similar to the agricultural trading company's activities. Because of this overlap in activities between quality measurement and treatment, we deemed the risk arising from

[15] The initial design had 20 clusters. A land conflict broke out in two villages (one control village and one treatment village) at the end of 2019. This resulted in two changes to the design. First, we initially delayed the start of the intervention in the extension service experiment for one season. Second, as the conflict remained active, we decided that the two clusters should be dropped from the trial.

potential Hawthorne effects and disruption of normal trading activities in the villages as too large. Thus, the quality measurement of maize destined for the existing local market may interact with the buying operation of the high quality buyer in ways (positively or negatively), which make it impossible to net out the direct effect of quality measurement by simply comparing treatment and control quality outcomes. To minimize these potential problems, data on maize quality were collected using a complementary approach. In the treatment group, quality testing was weaved into the buying operation; i.e., randomly selected bags from all sales to the agricultural trading company were brought into the lab for quality grading. The maize sold to other buyers was not tested for quality. Further, to avoid quality measurement-induced contamination in the control group, the research team organized maize buying in a random subset of villages and households in sample frame 3, i.e., a group of randomly selected villages adjacent to the control group villages. This sample of adjacent villages, therefore, serves as a quasi-control group to measure counterfactual maize quality.[16]

The chosen design does not provide a point estimate for the causal treatment effect of market access on quality upgrading because not all farmers in treatment sold to the agricultural trading company. Yet, as long as maize quality in the quasi-control villages provides a valid measure of status quo maize quality in the control group, we can estimate a lower bound on the causal effect by treating the problem of missing data in the treatment group as one of nonrandom one-sided attrition. In online Appendix G Table 9 we show that our assumption to use quality data from the quasi-control villages is warranted: in season 4, when comparable data were collected in sample 1 and sample 3, there are no systematic differences in mean outcomes across households in the control group in the market access experiment and households in the quasi-control group.

Estimator.—Our benchmark analysis of covariance (ANCOVA) specification uses only follow-up data for the dependent variable and regresses outcome Y_{ijt}, where subscript i denotes individual, j denotes cluster, and t wave or season, on a treatment indicator, D_j, which takes on the value one in treatment clusters and zero in control clusters, a season dummy δ_t and a lag-dependent variable, i.e., the value of the outcome pretreatment $\overline{Y}_{ij,PRE}$:

$$Y_{ijt} = \gamma D_j + \delta_t + \theta \overline{Y}_{ij,PRE} + \varepsilon_{ijt}. \tag{3}$$

The coefficient of interest, γ, gives the average causal effect over the four follow-up rounds. We report point estimates and p-values—both based on clustered-by-village standard errors and computed using randomization inference with permutations of treatment done at the village level.

[16]The potential problems associated with measuring quality also apply to the quasi-control villages, but the important difference is that the quasi-control villages do not serve as counterfactuals for the battery of agricultural outcomes (postharvest practices, income and profit) that we measure in the actual control villages. For similar reasons, these concerns are less problematic in the returns to quality experiment, where we were mainly interested in the effect of experimental variation in the quality produced on prices received from local traders. Moreover, since the returns to quality experiment was an individual-level trial, any other changes in market behavior at the village level would affect all trial participants in a similar way.

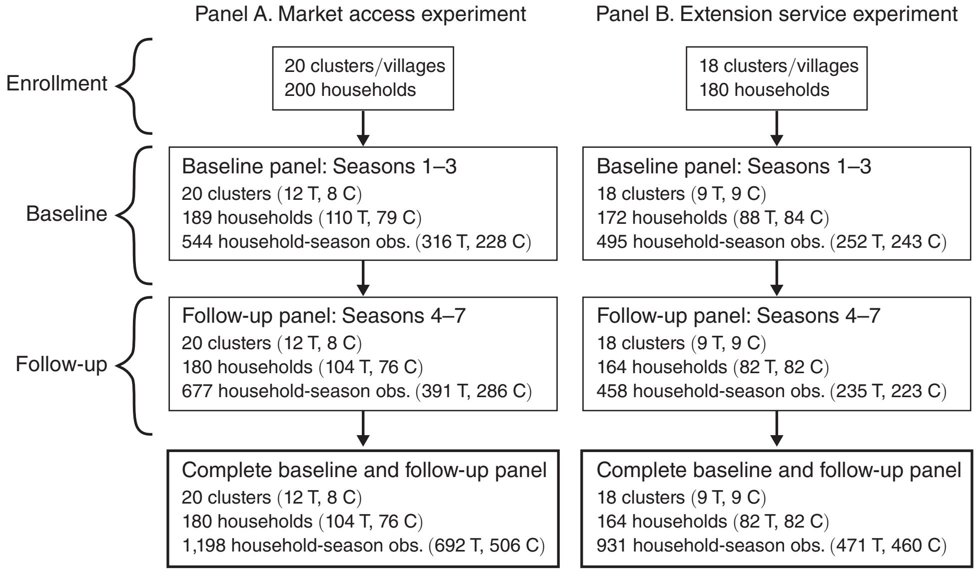

Figure 4. CONSORT Flow Diagrams

Note: Consolidated Standards of Reporting Trials (CONSORT) flow diagrams for the market access experiment (panel A) and the extension service experiment (panel B).

Assignment, *Attrition*, *and Baseline Balance*.—The sample population for both experiments consisted of smallholder maize farmers in farming communities in Kakumiro (market access experiment) and Kibaale (extension service experiment). In each trial cluster, we randomly selected ten households who had planted maize in the previous season, i.e., the season before the first baseline season. In addition, we randomly selected up to five replacement households in each cluster.

The first three seasons served as a baseline panel. After the first season, households that did not give consent to continue to participate, or that we could determine had moved, or that were involved in commercial maize trading, were replaced by households from the replacement list. No replacements were added after the first season.

At the end of the last preintervention season, the sample for the market access experiment included 544 household-by-season observations from 189 households in 20 clusters (see Figure 4 and online Appendix G Table 5). Follow-up lasted for four seasons. As reported in online Appendix G Table 6, less than 5 percent of the households attritted and attrition rates were similar across assignment groups. Of the nonattritters (180 households), 86 percent were resurveyed in each follow-up season, and the remainder were surveyed in some but not all seasons, yielding a follow-up sample of 677 household-by-season observations and an average resurvey rate of 0.94. Combining the baseline and follow-up data, we have a panel of households with baseline and follow-up data, with 1,198 household-by-season observations for 180 households in 20 clusters over 7 seasons (see Figure 4).

The sample for the extension service experiment included 495 household-by-season observations from 172 households in 18 clusters (see Figure 4 and online

TABLE 3—MARKET ACCESS EXPERIMENT: SUMMARY STATISTICS AND BALANCE AT BASELINE

	Sample			Means		Difference in means	
	Mean (1)	St. dev. (2)	Obs. (3)	T (4)	C (5)	Coeff. (6)	p (7)
Panel A. Household characteristics							
Main decision-maker (MDM): female	0.18	0.39	189	0.19	0.16	0.026	0.776
MDM: completed primary school	0.41	0.49	189	0.42	0.39	0.026	0.765
Number of household members	6.15	2.57	189	6.15	6.15	−0.006	0.987
Distance to district capital (km)	29.8	9.52	189	29.5	30.2	−0.703	0.873
Distance to main road (mins)	12.1	8.30	189	12.2	11.9	0.325	0.872
Panel B. Farm enterprise characteristics							
Maize acreage	2.16	1.60	544	2.18	2.13	0.047	0.900
Expenses (US dollars)	146.3	158.0	363	147.7	144.4	3.383	0.929
Harvest (ton)	2.10	1.94	499	2.14	2.05	0.096	0.844
Yield (ton/hectare)	2.09	1.04	499	2.13	2.04	0.092	0.651
Share sold	0.82	0.24	498	0.82	0.83	−0.012	0.699
Price per kilogram (US dollars)	0.20	0.057	470	0.20	0.20	0.001	0.797
Harvest value (US dollars)	433.2	425.9	499	450.1	409.3	39.463	0.703
Profit I (US dollars)	263.7	291.8	363	273.6	250.0	22.963	0.738
Joint balance test I							0.998
Joint balance test II							0.595
Joint balance test III							0.645

Notes: This table presents summary statistics and balance-at-baseline tests. T denotes the treatment group and C denotes the comparison group. Sample consists of households in the baseline panel sample. Panel A: measured at first baseline round. Panel B: pooled data over the three baseline rounds. Difference in means conditioning on season fixed effects in panel B. Standard errors are clustered at the village level. Expenses is expenses on inputs, equipment, transport, and hired labor. Data on hired labor was not collected in season 1. Harvest value includes own-produced consumption, valued at community-specific market value. Profit I is the difference between harvest value and expenses. The joint balance tests report p-values from testing whether the baseline outcomes predict enrollment into treatment, with profit dropped due to collinearity: all household characteristics in test I; farm enterprise outcomes except expenses in test II (seasons 1–3; sample size 470); all farm enterprise outcomes in test III (seasons 2 and 3; sample size 336).

Appendix G Table 5). Follow-up lasted for three seasons. The attrition rates were low and similar across assignment groups (see online Appendix G Table 6). Of the nonattritters (164 households), 84 percent were resurveyed in each follow-up season, and the remainder were surveyed in some but not all seasons, yielding a follow-up sample of 458 household-by-season observations and an average resurvey rate of 0.93. The complete baseline and follow-up panel consists of 931 household-by-season observations from 164 households in 18 clusters over 6 seasons.

Table 3 reports summary statistics and mean comparisons between the treatment and control groups in the market access experiment across a broad set of outcomes. Panel A shows household characteristics and panel B, and online Appendix G Table 8 panel A, present baseline farm enterprise outcomes. There are large variations across seasons. For example, the price of maize was more than 50 percent higher in the first season than in the third and profit in season 2 was roughly 50 percent higher than in season 3.

Summary statistics for household characteristics and farm enterprise outcomes (the latter pooled across the three baseline seasons) by assignment group are reported in columns 4 and 5 in Table 3 and a test of baseline balance is reported in columns 6 and 7. We find no evidence of differences in means among the household characteristics or the farm enterprise variables. The last three rows of Table 3 test whether the variables listed within each domain jointly predict treatment assignment. Our joint balance tests fail to reject the null hypothesis that neither household characteristics ($p = 0.998$) nor farm enterprise outcomes ($p = 0.595$; $p = 0.645$) predict assignment to treatment.

Summary statistics and mean comparisons between the treatment and control groups in the extension service experiment are reported in online Appendix G Table 7 and online Appendix G Table 8 panel B. Similar to the market access experiment, there is no evidence of differences in means among the household characteristics or the farm enterprise variables and the joint balance tests confirm this (for household characteristics ($p = 0.274$), for farm enterprise outcomes ($p = 0.125$; $p = 0.096$)).

C. *Results*

We begin by summarizing how the market access intervention affected the main outcome of interest: profits. The top graph in Figure 5 panel A shows that the CDF for profits is strongly shifted to the right for farmers who gained market access and we reject the hypothesis that the two distributions are equal (Kolmogorov-Smirnov D statistic is 0.17, $p = 0.000$). Several factors contributed to this profit increase: First, farmers who produced higher quality maize and sold to the agricultural trading company received higher prices. Second, and to a lesser degree, farmers who continued to sell to other traders in treatment villages also earned higher prices. Third, farmers in treatment villages grew more maize on a given plot of land. Together this produced a large increase in revenue. At the same time, farmers spent more on cultivating maize. The combined increase in revenue and expenses raised mean profits by US$63–98 or 36–81 percent (depending on how own and family labor is priced). In contrast, the treatment and control profit functions in the extension-only experiment lie on top of each other (top graph in panel B) and the Kolmogorov-Smirnov test fails to reject the null that the two distributions are equal (D statistic $= 0.07$, $p = 0.713$).

Quality Upgrading.—To measure the extent of quality upgrading, we examine two types of data: (i) data on farmers' interactions with the agro-trading company and administrative data on how the company enforced quality standards; and (ii) laboratory measured quality in bags.

In each posttreatment season, the agricultural trading company offered to buy maize of sufficiently high quality from preselected households in the treatment villages. Averaging across the four posttreatment seasons, 42 percent of farmers per season sold at least some bags of maize of sufficiently high quality to the company. The share of farmers who sold to the trading company increased with each additional season (see online Appendix H Figure 2 panel A). In the first season, about one in five households sold quality maize; in the fourth (and last) season, that ratio had more than tripled (to 65 percent). This upward trend suggests that it takes time for many households to make the necessary adjustment in their agricultural practices to produce maize of sufficient quality, but also that the switch to producing high quality maize is a permanent one. The distribution across villages in the share of farmers who sold to the premium quality buyer is reported in Appendix H Figure 2 panels C and D. For those farmers who sold high quality maize, on average 82 percent of sales went to the high quality buyer.

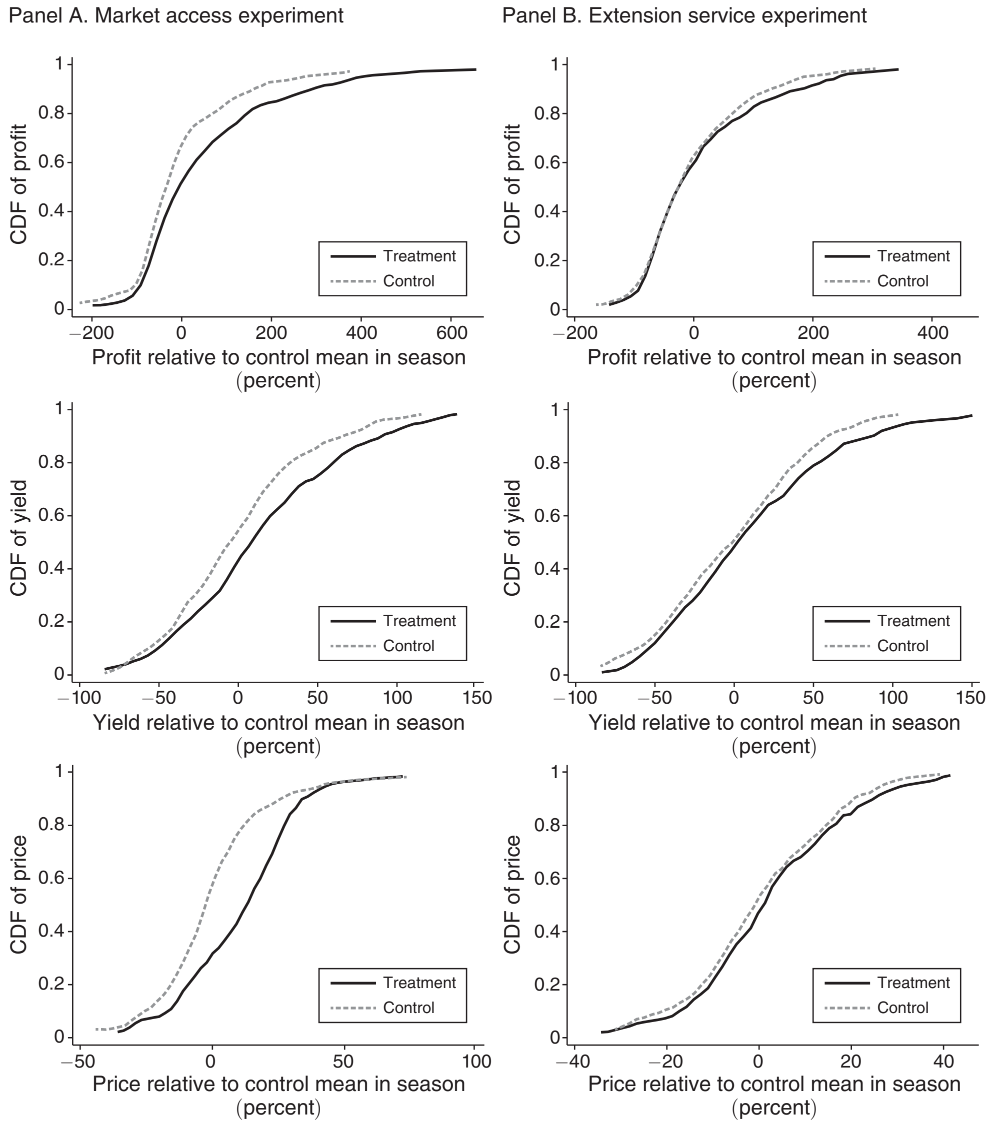

Figure 5. Effect of Market Access and Extension Service on Profit, Yield, and Price

Notes: This figure shows the CDFs of profit, yield, and price for the market access experiment (panel A) and the extension service experiment (panel B). All variables are expressed as difference from the control group mean in percent, by season. All graphs cap the top and bottom 1 percent observations. In panel A, The Kolmogorov-Smirnov D statistic is 0.174 [$p = 0.000$] for profit, 0.142 [$p = 0.003$] for yield, and 0.352 [$p = 0.000$] for price. In panel B, the Kolmogorov-Smirnov D statistic is 0.066 [$p = 0.713$] for profit, 0.091 [$p = 0.314$] for yield, and 0.077 [$p = 0.562$] for price.

Figure 2 panel B in online Appendix H provides detailed information on how the farmers and the company interacted, using data from the last two seasons.[17] Approximately four out of ten of the households did not sell (quality) maize to the

[17] The agro-trading company did not collect information on reasons for not buying in the first two buying seasons.

Table 4—Impact On Maize Quality

	Mean			
	Treatment (1)	Quasi-control (2)	Difference (3)	Observations (4)
Panel A. Maize quality				
Graded maize	0.89	0.30	0.593 [0.001]	86
Grade 1 maize	0.07	0.00		
Grade 2 maize	0.52	0.20		
Grade 3 maize	0.30	0.10		
Panel B. Bounds on average maize quality				
Horowitz-Manski lower bound			0.190 [0.205]	116
Lee lower bound			0.292 [0.027]	116

Notes: This table presents averages, differences in averages across assignment groups, and bounds on these differences, for maize quality. Graded maize is grade 1–3 maize based on the EAS classifications for maize quality, with grade 1 having the most stringent thresholds for defects (see online Appendix B). Data from the treatment group: randomly selected bags from all sales to the agro-trading company in season 7 (212 bags from 74 sales from 56 households), pooled at the household level. Data from the quasi-control group: 1 randomly selected bag bought from each of 30 randomly sampled households from 6 quasi-control villages in season 7 (30 bags). Panel A: unit of observation is a household selling maize to the high quality buyer (treatment group) or being randomly selected for testing (quasi-control group). Difference in means with p-value from Fisher-permutations test based on 10,000 clustered permutations of the treatment assignment in brackets. Panel B: unit of observation is a household and the outcome variable is the share of graded maize, out of total maize sold, by the household. The Horowitz-Manski lower bound p-value is calculated from Fisher-permutations tests with 10,000 clustered permutations of the treatment assignment in bracket. The Lee lower bound p-value is calculated from bootstrapped clustered standard errors (10,000 replications). Number of selected observations (Lee lower bound) is 86.

company, and another four out of ten sold all they wanted to sell. For 15 percent of the farmers, the company first refused to buy (and required the maize to be sorted, cleaned and/or dried further), but then bought once the farmer had upgraded the quality. For one in ten households, the company refused to buy because the quality was low.

These data suggest both that the company strictly enforced the quality standard, but also that the company did not cherry pick and buy only from farmers who already produced high quality to start with. In other words many farmers were eager to sell to the company, and learned over time how to produce maize of sufficiently high quality.

Maize quality was also measured in the lab for randomly selected bags from all farmers selling to the agricultural trading company in season 7. Pooling the data at the household level (column 1 in panel A of Table 4), 89 percent of the maize bought by the company was graded maize, with two-thirds of the graded maize classified as grade 1 or 2 maize—the highest two grades.

To assess the impact of the intervention on maize quality, we compare these results with maize quality data from the quasi-control group (Table 4, column 2). In the quasi-control group, 70 percent of the maize bags tested were ungraded maize. Of the remainder, two-thirds were grade 2 and one-third was grade 3. In sum, the share of graded maize sold to the company was 59 percentage points higher or almost three times as high as in the quasi-control villages ($p = 0.001$).

Because maize quality in treatment was only measured for farmers who sold to the agricultural trading company, part (or all) of the large difference in the share of graded maize across the two groups could be explained by nonrandom one-sided missing data. While we deem this unlikely based on the administrative data presented above, it is nevertheless possible that the distributions of maize quality in treatment and quasi-control are the same, and the treatment farmers who sell to the buyer of high quality maize simply come from the upper part of the quality distribution.

To test whether the observed quality difference could be generated by nonrandom missing data, we estimate the minimum share of graded maize produced in the treatment villages and the associated treatment effect. To do so, we exploit that each sale to the buyer of high quality was tested for quality. With this data, we calculate the (minimum) share of graded maize for households who sold to the high quality buyer as the share of graded maize sold to the high quality buyer times the share of total sales to the high quality buyer. This is equivalent to assuming that any maize they sold to other traders is ungraded. For treatment farmers who did not sell to the high quality buyer, the missing information on maize quality is addressed by conservative imputation or trimming.[18] For the imputation, we assume that all maize sold by treatment farmers who do not sell to the high quality buyer is ungraded and estimate the associated treatment effect, which is equivalent to the Horowitz and Manski (2000) lower bound estimator. Second, we also compute the Lee bound (Lee 2009), which trims the quasi-control sample to remove the farmers with maize quality in the zeroeth to thirty-fifth percentile of the quality distribution in quasi-control. Even allowing for this extreme form of selection, we estimate a treatment effect on average quality of 19 percentage points or 63 percent ($p = 0.205$) and 29 percentage points or 97 percent ($p = 0.027$) in the two bounding exercises (see Table 4, panel B).

To put these results in perspective, assume a buyer is aiming to procure at least grade 2 maize. The buyer has two options: buying maize at a market price $\tilde{p}$ per unit of maize and receiving the quality we observe in the quasi-control group, or buying at a premium $p^* = \tilde{p}(1 + \omega)$ and receiving maize of quality we observe in the treatment group. Using data from the laboratory tests, if buying from the treatment group the buyer would need to sort away on average 0.8 percent of the maize bought. If buying from the quasi-control, the buyer will need to sort away on average 20.4 percent. Assume the marginal cost of sorting is zero. The unit price for grade 2 maize when buying quality maize at a premium is then simply $p^*/(1 - 0.008)$ and the unit price for grade 2 maize when buying from quasi-control is $\tilde{p}/(1 - 0.204)$. By equalizing these prices we can solve for the maximum premium (ω) the buyer would be willing to pay for high quality maize, which is 24.6 percent, i.e., almost 10 percentage points more than the premium used in the experiment. Factoring in that every fourth bag tested in the quasi-control had too much moisture and thus required further drying (a task which further reduces the effective weight of the maize) and that 10 percent of the bags had live insects in the bag tested in the lab, the maximum premium the buyer would be willing to pay should be even higher.

[18] In the quasi-control group we measured quality for random bags (and random sales). We therefore assume the data from the random bags provide an unbiased estimate of the average quality of all sales. This assumption would hold, for example, if all farmers sell once (in the control group in the same season, over 80 percent sold once), or if the share sold is not correlated with quality.

Table 5—Impact on Investment

	Expenses: seeds and fertilizer (1)	Expenses: all inputs (2)	Proper drying (3)	Sorting (4)	Winn-owing (5)	Preharvest expenses (6)	Postharvest expenses (7)	Postharvest expenses (labor) (8)
Panel A. Market access experiment								
Access to a market for quality maize	2.37	4.04	0.24	0.14	0.15	16.2	5.92	5.86
	(0.045)	(0.075)	(0.000)	(0.002)	(0.033)	(0.275)	(0.256)	(0.144)
	[0.049]	[0.089]	[0.001]	[0.001]	[0.047]	[0.296]	[0.272]	[0.153]
Observations	658	658	640	464	464	464	464	464
R^2	0.31	0.32	0.21	0.03	0.04	0.20	0.26	0.22
Mean control	3.72	13.14	0.35	0.13	0.19	53.76	30.39	15.63
Panel B. Extension service experiment								
Extension service	−0.99	−0.54	0.029	0.037	0.071	−1.14	−1.18	−1.62
	(0.648)	(0.864)	(0.759)	(0.405)	(0.057)	(0.874)	(0.779)	(0.618)
	[0.677]	[0.883]	[0.770]	[0.423]	[0.069]	[0.881]	[0.788]	[0.621]
Observations	445	445	150	297	298	297	297	297
R^2	0.21	0.43	0.01	0.00	0.06	0.39	0.24	0.05
Mean control	4.21	12.55	0.68	0.15	0.14	53.94	26.22	13.17

Notes: ANCOVA regressions. Clustered-by-village standard errors with p-values in parentheses; p-values from Fisher-permutations test based on 10,000 permutations of the treatment assignment in brackets. Expenses on seeds and fertilizer are expenses on hybrid, open pollinated or recycled hybrid seeds and fertilizer. Expenses on all inputs are expenses on seeds and fertilizer, booster, and chemicals. Proper drying is a dummy equal to one if the maize was dried on a tarpaulin. Sorting is a dummy equal to one if the maize was sorted. Winnowing is a dummy equal to one if the maize was winnowed. Preharvest expenses are expenses on hired labor for preplanting (plowing and weeding), planting, weeding after planting and spraying. Postharvest expenses are expenses on hired labor for harvest and labor and equipment expenses for sorting, decobbing, winnowing, and bagging. All monetary values are in US dollars.

Investments and Productivity.—Market access may encourage farmers to invest more via two channels: (i) Farmers in the treatment group could earn higher prices conditional on producing high quality maize. It thus incentivized farmers to invest in upgrading quality. (ii) As farmers obtained higher prices for their crop, profit maximization predicts that they would use more inputs to increase the amount of (high quality) output to be produced. These predictions are borne out in the data: treatment farmers increased investments across a wide range of cultivation inputs and activities that improve both quality and productivity.

Farmers in treatment villages bought more inputs and hired more labor for preharvest activities (see Table 5 panel A), investments that primarily—though not exclusively—affect how much maize is produced. Specifically, farmers spent an additional US$2.4 or 64 percent ($p = 0.045$, control mean US$3.7) on hybrid and open pollinated seeds as well as inorganic fertilizer. The value of all agricultural input purchases, which also includes plant growth booster, animal manure, other seeds, pesticides and herbicides, increased by US$4.0 or 31 percent ($p = 0.075$, control mean US$13.1). Although these treatment effects represent large relative increases, in absolute terms, modern input use is low (see online Appendix G Table 10): 3 percent of control farmers used inorganic fertilizer, while 7 percent in treatment did (0.04 difference, $p = 0.054$), and 13 percent used improved seeds, compared to 17 percent in treatment (0.04 difference, $p = 0.250$).

Farmers did hire more labor. In the treatment group, farmers spent US$16 or 30 percent ($p = 0.275$, control mean US$54) more on hiring agricultural workers to prepare the land, plant maize seed, and weed and spray the crop, although the effect is not precisely estimated (column 6).

Farmers also invested more postharvest, which is viewed as crucial for maize quality. At baseline and in control villages, few farmers processed their crop properly: one-third dried their maize on a tarpaulin or in other ways that avoided direct contact with the soil (column 3), 13 percent sorted their maize (column 4) and one-fifth winnowed it (column 5). With access to a market for high quality maize the share of farmers who engaged in these practices nearly doubled: 59 percent dried their maize properly (a difference of 24 percentage points, $p = 0.000$), 27 percent sorted the maize (a difference of 14 percentage points, $p = 0.002$) and 34 percent winnowed it (a difference of 15 percentage points, $p = 0.033$). Consistent with this, spending on harvest and postharvest activities rose by 19 percent ($p = 0.256$, control mean US$30) and expenses on hired labor increased by 37 percent ($p = 0.144$, control mean US$15.6).

Summing across all items of cultivation expenditure, farmers in treatment villages invested US$18 more than those in control villages (Table 6 panel A column 6), an increase of 17 percent ($p = 0.305$, control mean US$106). The area under cultivation remained unchanged (Table 6 panel A column 2).

Farmers in treatment villages increased their total maize harvest as well as their yields. Figure 5 panel A shows that yield is higher in treatment than in control villages across the entire distribution. On average, yield (measured in kilogram per acre) rose by 109 kg or 14 percent ($p = 0.041$, control mean 792 kg), and total harvest by 239 kg or 13 percent ($p = 0.308$, control mean 1888 kg) as seen in Table 6 panel A columns 3 and 4. Trimming the data for outliers, the percentage increases for yield and harvest are larger and the effects are more precisely estimated (see online Appendix G Table 11 panel A).

In online Appendix E, we examine the relative importance of both measured and unmeasured inputs in explaining the increase in output and yield. To do so we need to specify the relationship between inputs and output and add an additional assumption; namely that observed inputs are independent of unobserved inputs, or total factor productivity (TFP), given treatment status. The results suggest that almost half of the treatment effect on harvest comes from increases in measured inputs, with hired labor and land being the most important factors. Thus, improvements in TFP account for the majority of the increase in output.

Can the increases we observe in both productivity and investments in the market access group be driven solely by the extension service component of the intervention? The results reported in Tables 5 and 6 panel B, and Figure 4 panel B, strongly suggest the answer is no. With the exception of winnowing, which increased by 7 percentage points or 50 percent ($p = 0.057$, control mean 0.14) in the treatment group in the extension trial, we find no significant impact of the extension intervention on input use or expenses.

These results do not rule out that the extension service program had an impact in the access to market intervention. One interpretation of the productivity increase we document, for example, and consistent with the findings reported in online Appendix E, is that better knowledge about best practice pre- and postharvest processes was put to use with market access.[19] But the results from the extension

[19] Previous research has documented positive, albeit small, impacts of providing accessible, tailored, and timely information through hands-on training on demonstration plots (see for example, Duflo, Kremer, and Robinson

Table 6—Impact on Productivity and Income

	Price (1)	Maize acreage (2)	Harvest (3)	Yield (4)	Harvest value (5)	Monetary expenses (6)	Profit (monetary expenses) (7)	Profit (incl. own hours) (8)
Panel A. Market access experiment								
Access to a market for quality maize	0.017	0.046	239.3	108.9	78.9	18.3	63.1	97.6
	(0.001)	(0.829)	(0.308)	(0.041)	(0.078)	(0.305)	(0.061)	(0.029)
	[0.004]	[0.838]	[0.350]	[0.049]	[0.102]	[0.321]	[0.079]	[0.031]
Observations	617	677	658	658	658	640	640	464
R^2	0.68	0.27	0.29	0.09	0.32	0.33	0.22	0.18
Mean control	0.15	2.29	1,887.52	791.97	286.17	106.38	177.23	120.25
Panel B. Extension service experiment								
Extension service	0.0040	−0.041	−110.0	33.7	−15.6	−4.10	0.063	23.0
	(0.404)	(0.800)	(0.609)	(0.589)	(0.726)	(0.693)	(0.999)	(0.509)
	[0.407]	[0.815]	[0.614]	[0.607]	[0.726]	[0.693]	[0.999]	[0.521]
Observations	420	458	445	445	445	445	443	443
R^2	0.34	0.29	0.39	0.16	0.36	0.38	0.16	0.16
Mean control	0.17	1.91	1,678.44	872.87	302.36	94.27	210.89	111.46

Notes: ANCOVA regressions. Clustered-by-village standard errors with p-values in parentheses; p-values from Fisher-permutations test based on 10,000 permutations of the treatment assignment in brackets. Price is the price per kilogram of maize. Yield is the harvest in kilogram per acre of cultivated land. Harvest value includes own-produced consumption, valued at community-specific market value. Expenses is expenses on input, equipment, transport, and hired labor. Profit (monetary expenses) is the difference between harvest value and expenses. Profit (incl. own hours) is harvest value minus monetary expenses and the cost of own labor valued at market wages. All monetary values are in US dollars.

service trial make it unlikely that this kind of supply intervention by itself would significantly and sustainably change how farmers operate.

Prices and Income.—Farmers with access to a market for high quality maize received significantly higher prices: the CDF of prices in treatment villages is strongly shifted to the right compared to the control group (Figure 5 panel A). The Kolmogorov-Smirnov D statistic is 0.35 ($p = 0.000$). The regression equivalent is presented in Table 6 panel A column 1: on average, farmers earned US\$2.38 or 11 percent more per bag of maize (140 kg) they sold ($p = 0.001$, control mean US\$21). This price increase together with the quantity increase translates into a significant and economically important increase of the value of farmers' harvest, column 5, which rose by US\$79 or 28 percent per season ($p = 0.078$, control mean US\$286).

How much do the price effect and the productivity effect each contribute to the increase in revenue? Denoting mean harvest and mean price in assignment group $d = \{1,0\}$ (treatment, control) by $\overline{Y}^d$ and $\overline{p}^d$, and Δx the treatment effect on outcome x, the treatment effect on harvest value (pY) can be decomposed into a pure price/quality-effect ($\overline{Y}^0\Delta p$), a pure quantity effect ($\overline{p}^0\Delta Y$) and an interaction:

$$\Delta pY = \overline{Y}^0\Delta p + \overline{p}^0\Delta Y + \Delta p\Delta Y. \tag{4}$$

2007; Hanna, Mullainathan, and Schwartzstein 2012; and Islam and Beg 2020). For recent reviews of the literature on extension service, see Macours (2019); Magruder (2018); and Takahashi, Muraoka, and Otsuka (2020). Bernard et al. (2017) show that adoption behavior may depend on demand-side conditions.

Given the treatment effects on price and harvested amount, the quantity effect accounts for 46 percent of the increase in harvest value, the quality effect accounts for 41 percent, and the remainder is explained by the interaction. Hence, the quality and quantity channels contribute in (almost) equal measure to the increase in harvest value.

The ultimate aim of linking farmers to a buyer of high quality maize is to increase farmer income and reduce rural poverty. After subtracting all monetary expenses from the farmers' harvest value, we find that farmers in the treatment group on average earned \$63 or 36 percent more per season than farmers in control villages ($p = 0.061$, control mean \$177); Table 6, column 7. The treatment effect becomes even larger and more precisely estimated when trimming the top and bottom 1 percent observations in each season (see online Appendix G Table 11 panel A). Access to a market for high quality maize thus presents a real opportunity to generate additional income in a context where such opportunities are few.

To give a complete picture of the profitability of quality upgrading, we also need to value farmers' own and family labor. Comparing the treatment effects on family and hired labor, we find that farmers in treatment villages reduce family labor hours by 112 hours per season or 23 percent ($p = 0.042$, control mean 486 hours). On the other hand, and in line with the overall increase in expenses, treatment farmers hire substantially more labor. Farmers in treatment villages increase their spending on hired labor by US\$26 or 38 percent ($p = 0.310$, control mean US\$69), equivalent to an additional 118 hours per season at the hourly wage. Again, this is an important effect in poor rural settings were employment opportunities are scarce and intermittent (de Janvry and Sadoulet 2019).

The total effect of this increase in labor hours and its changed composition on profitability depends on the relative productivity of family and hired labor and the relative cost of the two types of labor. Valuing family labor is challenging: one possible approach is to value family labor at the market wage, another is to put zero value on it, as no monetary costs are incurred. Family labor clearly has an opportunity cost, so valuing at zero is an extreme assumption. At the same time, family labor is most likely not a perfect substitute for hired labor: farmers typically hire labor for more difficult and physically demanding tasks, and even for the same task, hired labor tends to be adult labor while own/family labor is a mix of child and adult labor. Thus, costing family labor at market wage likely overestimates its value.[20]

In the end, we remain agnostic and let the value of own and family labor vary between zero and the market wage for hired labor. That is, we specify a profit function

$$\Pi = pY - cx - \varphi w L_F, \tag{5}$$

[20] We can compare the productivity of the two types of labor in the control group by relating the total amount of hours of hired labor per acre for a specific task (e.g., plot preparation, planting, weeding, harvesting) to the hours the average household member would take to perform the same task. These calculations suggest that family labor is about two-thirds as productive as hired labor.

where $pY - cx$ is harvest value minus monetary expenses (including hired labor), $\varphi \in [0,1]$, w is the hourly market wage for hired labor, and L_F measures hours of own and family labor.

Because treatment farmers reduced family labor in favor of hired labor, the treatment effect on profits increases with φ, the relative value of family labor. Valuing family labor at the market wage, farmer profits were US$98 higher in treatment than in control villages ($p = 0.029$, control mean US$120). Trimming reduces the effect size (77 percent increase in treatment compared with control; $p = 0.019$). For $\varphi = 2/3$, which corresponds to our guesstimate of the relative productivity of family and hired labor (see footnote 20), the treatment effect on profits is US$89, or 56 percent ($p = 0.042$, control mean US$157).

These effects represent large absolute increases in the context of our study, where most people live on less than US$1 a day. They also represent large increases relative to average annual income from all sources in the region: additional income from maize farming in the market access group represents a 16–24 percent increase in average annual income relative to a typical family in the region (UBOS 2019).

The increase in profit (focusing solely on the farmers) was also large relative to the cost of the extension program offered as part of the market access intervention. Specifically, the cost of the extension intervention, run over two seasons in a four season program, comes to US$15 per farmer and season, or a cost-to-income ratio of 15–24 percent.

In the extension-only experiment, all such effects on prices, revenue and income are absent. As seen in panel B of Table 6, the estimated treatment effects are small and insignificant.[21]

D. *Spillover Effects on Sales to the Local Market*

The entry of a high quality maize buyer in local (village) markets could affect (village) prices in two ways. First, and directly, households who successfully produced higher quality maize could sell it at a premium. Second, even in the case of differentiated products (higher or lower quality maize) the entry could affect the maize price also for households who did not sell to the high quality buyer. Such a spillover effect on local prices could come about through a number of channels: the external demand for high quality maize could reduce the supply of low quality maize and thus raise (local) prices; the entry of the high quality trader could reduce the market power of existing traders (if they have any); or it may improve farmers' bargaining position relative to the trader simply by improving their outside option. The trial, which induced variation in exposure to the new buyer across clusters, was designed to (partly) capture such spillover effects.

To measure how maize prices changed for households who chose not to sell to the high quality buyer, we use household-season-sale level data, which record the

[21] Although the samples were drawn from neighboring districts using the same sampling framework, there are some differences in the two samples. These differences raise the question whether the observed differences in treatment effects are (partly) driven by a difference in sample composition. Additional analysis, using entropy balancing and machine learning methods (see Hainmueller 2012; Crépon et al. 2021) to make the samples more comparable, strongly suggests that the difference in treatment effect is driven by a true difference in the causal effects of the two interventions (results are available upon request).

TABLE 7—IMPACT ON TRADER PRICES AND MARKET SHARES

	Other traders (1)	Local traders (2)	Commercial traders (3)
Panel A. Difference in market shares and prices			
Difference in market shares	−0.396	−0.325	−0.071
	[0.000]	[0.001]	[0.280]
Difference in prices versus control	0.045	0.061	−0.016
	[0.123]	[0.040]	[0.695]
Panel B. Difference in prices adjusting for selection			
Difference in prices versus control	0.066	0.078	0.021
	[0.071]	[0.052]	[0.680]

Notes: This table presents estimates on the impact of the market access intervention on prices and market share of maize traders in the trial villages. Panel A row 1 is estimated regressing market shares of other traders (column 1), local traders (column 2), and commercial traders (column 3) on the treatment assignment of the household who performs the sale. Row 2 is estimated regressing the normalized price, $\tilde{p}_{j,k,t} = (p_{j,k,t} - \overline{p}_{0,t})/\overline{p}_{0,t}$, where $p_{i,k,t}$ is the price for sale j by farmer k in period t and $\overline{p}_{0,t}$ is the average price in the control group in season t, on an indicator for the type of trader who bought the maize in the treatment group. Panel B row 1, reports the causal effect estimated regressing the difference between the (normalized) price at follow-up and the baseline normalized price on the indicator for trader type, within the treatment group. The unit of observation is household-sale (799 observations, a sale is included if the household has at least one registered sale in the last baseline season). Clustered-by-village standard errors with p-values in brackets.

price received and the type of buyer, i.e., whether a sale went to the high quality buyer (denoted HT) or to other traders (OT), where other traders are composed of local traders (LT) or commercial (CT) traders (see Section IIA). With this data, we estimate the average proportion of sales across all follow-up seasons $\overline{s}_{d,j}$ from households with treatment status $d = \{0,1\}$ to traders of type $j = \{OT, CT, LT\}$, as well as the difference, $\Delta\overline{s}_j$, between treatment and control (see online Appendix F for details). We also calculate the average price $\overline{p}_{1,j}$ paid by traders of type j to treatment households, as well as the difference between $\overline{p}_{1,j}$ and the average control group price $\Delta\overline{p}_{j,0} = \overline{p}_{1,j} - \overline{p}_0$.

Entry of the high quality buyer decreased the average market share of other traders by almost 40 percent ($\Delta\overline{s}_{OT} = -0.396$; $p = 0.000$); see Table 7 panel A. Of these, local traders' share of sales fell from 80 percent in control to 47 percent in treatment ($\Delta\overline{s}_{LT} = -0.325$; $p = 0.001$) and commercial traders' share of sales decreased from 20 percent in control to 13 percent in treatment ($\Delta\overline{s}_{CT} = -0.071$, $p = 0.280$). That is, farmers in the treatment group primarily switched from selling maize to local traders to selling quality maize to the agricultural trading company.

The (average) price for sales to other traders in the treatment group is 4.5 percent higher than the average price in the control group ($\Delta\overline{p}_{OT,0} = 0.045$; $p = 0.123$). This price increase is mainly driven by prices for sales to local traders, which increased by 6.1 percent relative to the control group ($p = 0.040$). Sales to commercial traders fetched 1.6 percent less than the average price in the control group, but we cannot reject the null hypothesis that prices are equal ($p = 0.695$).

The results in Table 7 panel A suggest that there is a positive spillover effect on prices for sales to other, mainly local, traders as a result of the opportunity to sell to the high quality buyer. But to determine whether this is truly the case, we need to disentangle the market equilibrium effect—a (reduced form) causal effect—from a possible selection effect. For example, the observed positive price difference between treatment and control villages might be driven by the fact that those who decide not to upgrade quality and sell to the high quality buyer would have earned higher prices even in the absence of the high quality buyer entering.

Recovering the causal spillover effect is challenging because we do not observe the average counterfactual price that a treatment farmer would have received had the premium quality buyer not entered. But with additional structure, and by exploiting the fact that we have both baseline and follow-up data on prices, we can make progress.

Specifically, we posit a selection model where farmers are heterogeneous with respect to a constant "ability" dimension, which determines both the price they demand on the existing market as well as their likelihood of upgrading quality and selling to the high quality buyer. The farmer's decision whether to sell to the high quality buyer or not in the treatment group is then as good as randomly assigned conditional on a farmer fixed effect and time trend.

The selection model embodies the parallel trends assumption needed for panel data to identify causal effects. Figure 3 in online Appendix H provides support for this assumption. Prior to the intervention, the prices for farmers who later on sell to the high quality buyer and farmers who continue to sell to other traders, follow parallel, and essentially time-invariant, trends. It should therefore be possible to estimate the local market spillover effect through a difference-in-difference procedure where differences relative to control and relative to baseline are used to purge the data of aggregate time-variation and farmer-specific time-invariant heterogeneity.

Table 7 panel B reports the results. We estimate that the entry of the high quality buyer raised the prices farmers received for sales to the local market by 6.6 percent ($p = 0.071$). Comparing the estimated causal effect in panel B to the average difference between prices for sales to other traders in treatment and control in panel A (4.5 percent) implies (given our assumptions) that those who do not sell to the buyer of premium quality are negatively selected in terms of price. On average, sales to the other traders in treatment villages come from farmers who earned 2 percent lower prices at baseline.[22]

How important, quantitatively, is the local market equilibrium effect in explaining the average price increase in treatment relative to control? Calculating the share of the average increase in price that comes from the increase in local market prices reveals that 32 percent of the average price increase in treatment relative to the control is driven by market spillover effects (see online Appendix F Section III for details).

In summary, the analysis suggests that treatment households who sold to the local market earned higher prices than they would have in the absence of the high quality buyer. The local spillover effect on prices is thus positive, which reduces the incentives for quality upgrading.

[22] The difference in difference estimation can be further extended to estimate spillover effects on prices for sales to local and commercial traders separately. As reported in Table 7, the entry of the high quality buyer raised prices for sales to local traders by 7.8 percent ($p = 0.052$) and commercial traders by 2.1 percent ($p = 0.680$). Hence, it appears that the spillover effect is stronger on prices for sales to local traders, while the response from commercial traders is more muted.

The finding that prices received by treatment farmers selling to other traders increased in response to the intervention raises the question whether higher output and productivity in the treatment group can be (partly) explained by farmers who do not sell to the high quality buyer. In online Appendix G Table 14, we estimate the causal effects on output of selling to the high quality buyer and other traders in response to the intervention, following the same identification strategy as we do for prices. While the point estimates are imprecise, the results suggest that the average causal effect on output (harvest) can be fully accounted for by farmers who sold to the high quality buyer.

VI. Discussion

We interpret our results as demonstrating a proof of concept: improving smallholders' access to markets where high quality produce is rewarded, and more generally linking farmers to value chains, has large potential. Our work highlights both new insights and challenges. On the former, while pre- and postharvest training on best agronomic practices remains a core pillar of most agricultural development programs, our findings suggest that such a supply-side interventions alone has little impact on productivity or income without market access. On the latter, while farmers increased their use of modern inputs—one pathway to increased productivity—adoption of these technologies remained low. This result, coupled with evidence from an increasing number of supply side interventions studied in the literature (see review in de Janvry, Sadoulet, and Suri 2017), including missing markets for risk (Karlan et al. 2014) and behavioral constraints (Duflo, Kremer, and Robinson 2011), points to an important area for future research—understanding complementarities between market access or value chain inclusion and supply constraints.

An additional challenge is related to the program's overall profitability. After factoring out all evaluation costs, the agro-trading company broke even in two of the four buying seasons.[23] Adding farmer profit, joint surplus in these two seasons was therefore strictly positive (even after factoring in the extension service costs). Three structural features of the product and the economy constrained the company's profits. First, as quality is essentially no longer observable once the maize grain is milled, customers need to learn about higher quality through consuming it. As a consequence, it takes time to build up a reputation for high quality maize flour and a domestic customer base willing to pay a premium for it. Second, the price elasticity of quality among large segments of consumers is low. This in turn is a consequence of low consumer awareness of the benefits of food safety and the inability of the government to publicize, test, or enforce quality standards at all stages of the value chain. Third, although quality maize can be exported at a (high) premium, a seller needs to incur large (fixed) costs (related to establishing contacts with international buyers and producing at the necessary scale) to enter the export market. These fixed costs help explain why the formal export market is dominated by a few large actors.

Other features of the business model raised costs. Specifically, the company's business model was not one of pure profit maximization. Unlike other vertically

[23] By focusing on the financial surplus, we rule out any potential health benefits of higher quality maize.

integrated firms on the market, it offered to buy maize from smallholders managing to sell maize of high quality, rather than identifying areas with largeholder farmers producing relatively high quality maize. While this strategy raised costs, and hence decreased company profits, it was likely crucial for achieving the large increases in farmer surplus we document. That is, the market intervention had such large positive impacts on income precisely because it provided market access to the poorest farmers who are currently excluded from global value chains.

Overall, the limited profitability of the vertically integrated model used by the company, and the revenue and cost constraints the company faced, provide important clues as to why market integration of large swathes of the rural population, and for many of the agricultural products they produce, is challenging—despite its potential. On the other hand, if the returns to sourcing high quality maize from smallholder farmers are not sufficiently high to attract private actors to enter this market, one could consider a subsidy to raise the returns. While we cannot properly compare the costs and benefits of market access to the various multifaceted programs to help the very poor (see, e.g., Bandiera et al. 2017; Banerjee et al. 2015), the effects on income we document suggest a market access program is a candidate worth investigating more closely.

REFERENCES

Antrás, Pol. 2015. *Global Production: Firms, Contracts, and Trade Structure*. Princeton, NJ: Princeton University Press.

Ashraf, Nava, Xavier Giné, and Dean Karlan. 2009. "Finding Missing Markets (and a Disturbing Epilogue): Evidence from an Export Crop Adoption and Marketing Intervention in Kenya." *American Journal of Agricultural Economics* 91 (4): 973–90.

Atkin, David, Amit K. Khandelwal, and Adam Osman. 2017. "Exporting and Firm Performance: Evidence from a Randomized Experiment." *Quarterly Journal of Economics* 132 (2): 551–615.

Bandiera, Oriana, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, Munshi Sulaiman. 2017. "Labor Markets and Poverty in Village Economies." *Quarterly Journal of Economics* 132 (2): 811–870.

Banerjee Abhijit V., Esther Duflo, Nathanael Goldberg, Dean Karlan, Robert Osei, William Pariente, Jeremy Shapiro, Bram Thuysbaert, Chris Udry. 2015. "A Multi-faceted Program Causes Lasting Progress for the Very Poor: Evidence from Six Countries." *Science* 348 (6236)

Beaman, Lori, Dean Karlan, Bram Thuysbaert, and Christopher Udry. 2013. "Profitability of Fertilizer: Experimental Evidence from Female Rice Farmers in Mali." *American Economic Review* 103 (3): 381–86.

Bergquist, Lauren Falcao, and Craig McIntosh. 2021. "Search Cost, Intermediation, and Trade: Experimental Evidence from Ugandan Agricultural Markets." Unpublished.

Bernard, Tanguy, Alain de Janvry, Samba Mbaye, and Elisabeth Sadoulet. 2017. "Expected Product Market Reforms and Technology Adoption by Senegalese Onion Producers." *American Journal of Agricultural Economics* 99 (4): 1096–1115.

Blouin, Arthur, and Rocco Macchiavello. 2019. "Strategic Default in the International Coffee Market." *Quarterly Journal of Economics* 134 (2): 895–951.

Bold, Tessa, Selene Ghisolfi, Frances Nsonzi, and Jakob Svensson. 2022. "Replication Data for: Market Access and Quality Upgrading: Evidence from Four Experiments." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. http://doi.org/10.3886/E158401V1.

Bold, Tessa, Kayuki C. Kaizzi, Jakob Svensson, and David Yanagizawa-Drott. 2017. "Lemon Technologies and Adoption: Measurement, Theory and Evidence from Agricultural Markets in Uganda." *Quarterly Journal of Economics* 132 (3): 1055–1100.

Casaburi, Lorenzo, and Tristan Reed. 2017. "Competition in Agricultural Markets: An Experimental Approach." CEPR Discussion Paper 11985.

Casaburi, Lorenzo, and Tristan Reed. 2020. "Interlinked Transactions and Competition: Experimental Evidence from Cocoa Markets." Unpublished.

Crépon, Bruno, Esther Duflo, Elise Huillery, William Pariente, Julien Seban, and Paul-Armand Veillon. 2021. "Cream Skimming and the Comparison between Social Interventions: Evidence from Entrepreneurship Programs for At-risk Youth in France." Unpublished.

Daly, Jack, Danny Hamrick, Gary Gereffi, and Andrew Guinn. 2016. *Maize Value Chains in East Africa*. London, UK: International Growth Centre.

de Almeida, Luis, Robert Williams, Dirce M. Soares, Harry Nesbitt, Graeme Wright, and William Erskine. 2019. "Aflatoxin Levels in Maize and Peanut and Blood in Women and Children: The Case of Timor-Leste." *Scientific Reports* 9 (1): 13158.

de Janvry, Alain, and Elisabeth Sadoulet. 2019. "Transforming Developing Country Agriculture: Removing Adoption Constraints and Promoting Inclusive Value Chain Development." Unpublished.

de Janvry, Alain, Elisabeth Sadoulet, and Tavneet Suri. 2017. "Field Experiments in Developing Country Agriculture". In *Handbook of Economic Field Experiments*, Vol. 2, edited by Abhijit Banerjee and Esther Duflo, 427–66. Amsterdam: Elsevier.

De Loecker, Jan, and Pinelopi Koujianou Goldberg. 2014. "Firm Performance in a Global Market." *Annual Review of Economics* 6 (1): 201–27.

Deutschmann, Joshua W., Maya Duru, Kim Siegal, and Emilia Tjernström. 2019. "Can Smallholder Extension Transform African Agriculture?" NBER Working Paper 26054.

Dillon, Brian, and Chelsey Dambro. 2017. "How Competitive Are Crop Markets in Sub-Saharan Africa?" *American Journal of Agricultural Economics* 99 (5): 1344–61.

Donaldson, Dave. 2015. "The Gains from Market Integration." *Annual Review of Economics* 7: 619–47.

Duflo, Esther, Michael Kremer, and Jonathan Robinson. 2007. "Why Don't Farmers Use Fertilizer? Experimental Evidence from Kenya." Unpublished.

Duflo, Esther, Michael Kremer, and Jonathan Robinson. 2011. "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya." *American Economic Review* 101 (6): 2350–90.

East African Community. 2011. East African Standard: Maize grains – Specification. East African Community, Arusha, Tanzania.

Hainmueller, Jens. 2012. "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies." *Political Analysis* 20 (1): 25–46.

Hanna, Rema, Sendhil Mullainathan, and Joshua Schwartzstein. 2012. "Learning through Noticing: Theory and Experimental Evidence in Farming." NBER Working Paper 18401.

Hell, Kerstin, Pascal Fandohan, Ranajit Bandyopadhyay, Sebastian Kiewnick, Richard Sikora, and Peter J. Cotty. 2008. "Pre- and Postharvest Management of Aflatoxin in Maize: An African Perspective." In *Mycotoxins: Detection Methods, Management, Public Health and Agricultural Trade*, edited by John F. Leslie, Ranajit Bandyopadhyay, and Angelo Visconti, 219–30. Wallingford, Oxfordshire, UK: Centre for Agriculture and Biosciences International.

Hoffmann, Vivian, Samuel Mutiga, Jagger Harvey, Michael Milgroom, and Rebecca Nelson. 2013. "A Market for Lemons: Maize in Kenya." Unpublished.

Horowitz, Joel, and Charles Manski. 2000. "Nonparametric Analysis of Randomized Experiments with Missing Covariate and Outcome Data." *Journal of the American Statistical Association* 95 (449): 77–84.

Islam, Mahnaz, and Sabrin Beg. 2020. "'Rule-of-Thumb' Instructions to Improve Fertilizer Management: Experimental Evidence from Bangladesh." *Economic Development and Cultural Change* 70 (1): 237–81.

Karlan, Dean, Robert Osei, Isaac Osei-Akoto, and Christopher Udry. 2014. "Agricultural Decisions after Relaxing Credit and Risk Constraints." *Quarterly Journal of Economics* 129 (2): 597–652.

Lee, David. 2009. "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects." *Review of Economic Studies* 76 (3): 1071–1102.

Macauley, Harold. 2015. *Background Paper: Cereal Crops: Rice, Maize, Millet, Sorghum, Wheat*. Abidjan, West Africa: African Development Bank.

Macchiavello, Rocco, and Josepa Miquel-Florensa. 2019. "Buyer-Driven Upgrading in GVCs: The Sustainable Quality Program in Colombia." CEPR Working Paper 13935.

Macours, Karen. 2019. "Farmers' Demand and the Traits and Diffusion of Agricultural Innovations in Developing Countries." *Annual Review of Resource Economics* 11: 483–99.

Magruder, Jeremy R. 2018. "An Assessment of Experimental Evidence on Agricultural Technology Adoption in Developing Countries." *Annual Review of Resource Economics* 10: 299–316.

Ranum, Peter, Juan Pablo Peña-Rosas, and Maria Nieves Garcia-Casal. 2014. "Global Maize Production, Utilization, and Consumption." *Annals of the New York Academy of Sciences* 1312 (1): 105–112.

Rosenzweig, Mark R., and Christopher Udry. 2020. "External Validity in a Stochastic World: Evidence from Low-Income Countries." *Review of Economic Studies* 87 (1): 343–81.

Sserumaga, Julius P., Alejandro Ortega-Beltran, John M. Wagacha, Charity K. Mutegi, and Ranajit Bandyopadhyay. 2020. "Aflatoxin-Producing Fungi Associated with Pre-harvest Maize Contamination in Uganda." *International Journal of Food Microbiology* 313: 108376.

Takahashi, Kazushi, Rie Muraoka, and Keijiro Otsuka. 2020. "Technology Adoption, Impact, and Extension in Developing Countries' Agriculture: A Review of the Recent Literature." *Agricultural Economics* 51 (1): 31–45.

Uganda Bureau of Statistics. 2015. Statistical Abstract. Retrieved from https://www.ubos.org/wp-content/uploads/publications/03_2018Stati stical_Abstract_2015.pdf.

Uganda Bureau of Statistics (UBOS). 2019. *Uganda National Household Survey 2016/2017 Report*. Kampala, Uganda: Uganda Bureau of Statistics.

US Department of Agriculture (USDA). 2019. United States Department of Agriculture, Data retrieved from National Agricultural Statistics Service.

World Bank. 2007. *World Development Report 2008: Agriculture for Development*. Washington, DC: World Bank.

World Bank. 2020. *World Development Report 2020: Trading for Development in the Age of Global Value Chains*. Washington, DC: World Bank.