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

August 2021

Tessa Bold,[1] Selene Ghisolfi,[2] Frances Nsonzi,[3] Jakob Svensson[4]

Abstract

Smallholder farming in many developing countries is characterized by low productivity and low quality produce. Low quality limits the price farmers can command and thus their potential income from farming. We conduct a series of measurement and field experiments among smallholder maize farmers in western Uganda to shed light on the barriers to quality upgrading at the farm level and to study its potential in raising productivity and rural incomes. First, we measure maize quality at the farm gate and in the lab and show that quality is low and at least partly observable at the farm gate. Second, we generate exogenous variation in the quality of the maize farmers sell to local markets by offering a randomly selected sub-set a post-harvest service package to improve the quality of their maize. The market return to this quality improvement is zero, suggesting that the market for quality maize is effectively missing. Third, we generate experimental variation in access to a market for premium quality maize, combined with training on agricultural best-practices. Over time, the majority of treatment farmers sold maize of high quality. Profit from maize farming in the treatment group increased by 40-80%; an effect driven both by increased productivity and higher prices for both premium and lower quality maize in treatment villages. Fourth, we separately assess the impact of the training intervention and show that extension service alone does not lead to a change in farmers' practices or agricultural outcomes. Our findings reveal the importance of demand-side constraints in limiting rural income and productivity growth.

---

* We gratefully acknowledge the editor Esther Duflo and two anonymous referees for many valuable comments and suggestions. We have also benefitted 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 is greatly appreciated.

[1] IIES, Stockholm University, Stockholm, Sweden; CEPR (tessa.bold@iies.su.se).
[2] Leap, Bocconi University, Italy (ghisolfi@unibocconi.it).
[3] Department of Food Technology and Nutrition, Makerere University, Kampala, Uganda; PNDK Solutions, Kampala, Uganda (nsonzif@gmail.com).
[4] IIES, Stockholm University, Stockholm, Sweden; BREAD; CEPR (jakob.svensson@iies.su.se).

1. Introduction

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 and to study its potential. We proceed in four steps. 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 all the way to the farmer. 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: traders pay the same price for high and low quality maize.

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 be the underlying reason for the lack of demand for high quality maize we observe in the local markets. In turn, buyers of high quality maize may not be 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. To ensure that farmers had up-to-date knowledge about pre- and post-harvest practices necessary for producing maize of sufficiently high quality, the market access intervention was combined with a learning-by-doing extension service component. Following farmers over four years, including four post-treatment seasons, we find that a majority of farmers increase the quality of their maize and sell to the high quality buyer when offered the opportunity. Treatment farmers' profit increases substantially – an effect driven both by increased farm productivity and higher prices for both premium and lower quality maize in treatment villages.

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, and revenue, expenses, and yield remain essentially the same in the treatment and the control group. Although the results do not rule out that an extension component was required for the access to market intervention to achieve its impact, the comparison of the two trials suggests that it is market access – not extension service – that is necessary (if possibly 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-differences approach, the higher price for sales to local buyers can account for 30% of the increase in average prices in the treatment relative to the control group. Despite the higher price for lower quality maize, the evidence suggest that the increase in output in the treatment relative to the control 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 et al. (2017), we exploit experimental variation in access to a market/buyer of quality products. Their intervention, which connects rugs producers in Egypt to foreign buyers paying a premium for higher quality rugs, led to large improvements in both quality and productivity. In contrast to this study, which focusses on established small firms, 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, and constrain the firm's actions in several ways so that only farmers in randomly selected villages have access to a local market for quality. 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 design the intervention to run for 4 seasons and thus assess the implications of having access to a market for quality over a longer period.

Our paper is also related to, and complements, Macchiavello and Miquel-Florensa (2019). They employ a difference-in-difference strategy to estimate the impact of a quality upgrading program in Colombia and find that eligible farmers increased quality and received higher farm gate prices. 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, including the costs of quality upgrading in response to getting access to an output market for quality maize, and to decompose the treatment effect on output into components attributable to changes in the inputs that we can measure and to increases in total factor productivity. 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.[5]

Knowledge about the pre-conditions and determinants for agricultural technology adoption has grown vastly over the last decade (for a review, see de Janvry, et al., 2017). The evidence – drawn primarily from randomized controlled trials in Sub-Saharan Africa and

[5] 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. Important theoretical contributions in this literature is Hausmann et al. (2007), who emphasize the link between specialization in exporting high quality and subsequent higher growth. Empirical results (largely non-experimental) are reviewed, for example, in De Loecker and Goldberg (2014). Ashraf, Gine and Karlan (2009) experimentally evaluated an intervention in Kenya that helped farmers to adopt and market export crops. The authors find small effects on adoption and income, but a 32% income gain for adopters. Our study also relates to the (mainly non-experimental) 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 (e.g. Bergquist and McIntosh, 2021).

South Asia – makes clear that there are productivity-enhancing supply side innovations available today that can increase technology adoption and productivity among smallholders. Measured effects on farmer income, however, have been much more limited.[6] We add to this literature by studying the impact of lifting both a demand (inclusion in a value chain) and a supply (knowledge) constraint. De Janvry and Sadoulet (2019) discuss the complexity of rigorously evaluating value chain inclusion, noting the double challenge of having to implement treatment at the community level – making the intervention costly – and the difficulty in finding implementing partners who are willing to expand their business in a way that is amenable to rigorous evaluation. We overcome these problems here, and document widespread increases in adoption (here mostly in terms of quality upgrading) over time and a significant increase in farm profit.

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 context can help explain why farmers raised yields predominantly by working more and producing a higher output per hour worked, rather than applying modern inputs, such as improved seeds and fertilizer to their fields. 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 had 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.[7] It does not, however, allow us to study the behavior of the layers of intermediaries that dominate the low-quality segment of the market.[8]

[6] For example, Karlan et al. (2014) find that farmers with insurance invest more in cultivation, but without any significant impacts on profit. Beaman, Karlan, Thuysbaer and Udry (2013) show that receiving the recommended dose of fertilizer increases harvest value by 30% but profits only by 0.5%. An exception is the evaluation of the One Acre Fund's small farmer program – a bundled program where participating farmers receive training on improved farming practices, input loans, and crop insurance. Deutschmann, et al. (2019) document an increase in profit of 8%-16% after one season.

[7] The decision to work with a vertically integrated buyer to purchase high quality maize grain is reminiscent of the work by Hansman et al. (2020), who study vertical integration as a response to a higher quality premium in the context of the Peruvian fishmeal industry, as well as a theoretical literature that emphasizes the relationship between the nature of output and the boundary of the firm (Baker et al, 2001).

[8] 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

2. Context

2.1. Introduction

Uganda remains highly dependent on agriculture. It is estimated that the sector contributes over 70 percent of export income and 65 percent of the population is active in the sector. 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.[9] Smallholder farmers account for roughly 75 percent of maize production and 70 percent of marketable surplus.

Maize has different end uses depending on the geographic regions of the producers (Daly, et al. 2016). In the US only 12% of maize produced is used for human consumption, with the remainder split between animal feed and ethanol fuel production (Ranum, 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 tend to be low. For example, Bold et al (2017), using four waves of LSMS data for Uganda, report an average yield (metric tons per hectare) for smallholders of 1.4. As a comparison, average maize yield based on data from farm demonstrations in Uganda (with recommended crop management and modern inputs) is over 4 tons per hectare (World Bank 2007) and average corn yield in the U.S. was close to 12 tons per hectare in 2017 (USDA, 2019). There is also concern about the quality of the maize (and other crops) bought and sold in Uganda, though there is little systematic measurement of quality.[10]

(2015) estimate in the context of Sierra Leonean cocoa producers that the local markets in which traders operate to buy farmer cocoa are highly competitive with a low differentiation parameter. Bergquist and McIntosh (2021) find small effects on trader prices and margins of providing farmers in areas with market surplus 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.

[9] According to official statistics, maize exports accounted for about 2% of the country's total exports (Uganda Bureau of Statistics, 2015). Based on interviews with stakeholders in the sector, Daly et al. (2016) estimate that 70-80% of maize that is bought and sold in Uganda is channeled through informal channels. To account for the size of the informal market, previous surveys have used multipliers of between 3-3.5 for formal trade data (Gates Foundation, 2014).

[10] For the low quality of maize and coffee in Uganda, see Daly et al. (2016), Gates Foundation (2014), and Morjaria and Sprott (2018), respectively.

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 0.80 USD per day (UBOS, 2019).[11]

2.2. Local markets

The local, or village, output market for smallholders can be described as a spot market.[12] The farmer and the buyer agree right before the sale, usually after a short visual inspection of the 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 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, 79% of the sales and 78% of the sales volume 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 fetches, on average, an 8% higher price than a sale to a local trader. On average, farmers sold 80% of what they produced, and kept 20% 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% of the smallholders sold to the same buyer in at least two out of the last five seasons and 31% sold to the same buyer in at least four out of five seasons. Most households sell once per season (79%) and accounting for multiple sales to the same buyer in a given season, 90% sell to only one trader per season.

We collected data on market prices from the nearest trading center for each village in the sample (altogether five trading centers). Prices tend to increase (slowly) between harvest seasons, while there is substantial idiosyncratic variation in market prices within harvest seasons. The trading center price follows a similar pattern as the prices in the main wholesale

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

[12] 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.

markets in Kampala and Nairobi, with the average trading center price 16% lower than the price in Kampala, which in turn is 30% below the price in Nairobi.

3. 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 post-harvest 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 or is market access necessary for improvements in agricultural practices, income and productivity?

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. Sample Frame 1 forms the basis for the market access intervention. We first randomly selected 20 communities (census enumeration areas; i.e., villages or parts of villages, to which we will refer to as "villages" henceforth), each at least 5 kilometers apart, from digital maps of Kakumiro district in western Uganda. For each of the selected villages, we completed a census and identified smallholder farmer households (with maize gardens of no more than 5 acres of land) who cultivated maize in the previous season.

To learn about complementarities between demand and supply side constraints, we compare the market-cum-extension experiment with the results from an additional trial focused solely on the impact of the extension service component. The extension service intervention was offered to a randomly selected set of villages in the neighboring district Kibaale (Sample Frame 2). Villages and households in this experiment were selected using the same approach as in Sample Frame 1.

Finally, we identified all villages adjacent to the control group villages in Sample Frame 1 and randomly selected 20 of these villages (Sample Frame 3). We completed a census and identified smallholder farmer households in the villages, and selected around 10 households in each village. We use this sample to measure the quality of maize in local markets and to infer the causal return to quality upgrading in local markets. As discussed in section 6.3.2, we also use a random subset of Sample Frame 3 villages as a quasi-control group for counterfactual maize quality in treatment villages in Sample Frame 1.

4. Maize Quality and Verifiability of Quality

4.1. Introduction

Maize is sold and handled in large quantities, with the smallest unit typically a 100-170 kilogram bag. A bag of maize is considered high quality if it contains sufficiently large and dry maize kernels of the right color and neither non-grain substances (e.g., stones, dirt, and insects) nor defective (e.g., broken, immature, damaged, rotten, or moldy) grain. More formally, maize quality in East Africa is classified 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 non-grain substances and defective grain: graded maize, undergrade maize and reject maize (the latter two, we label 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 potential economic and nutritional value and whether it is safe for human consumption. The presence of non-grain substances and defective grain adds to the weight of the bag without adding value and increases processing costs. Non-grain substances and defective grain are also indicators that 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 which inhabit the soil.[13] Contamination and infestation can spread quickly through the bag and while waste and defective maize kernels can be sorted and cleaned at a later stage in the value chain, there is an elevated risk that the remaining grain is (already) contaminated.

The moisture content influences expected maize quality through the same two channels: by increasing gross (but not dried) maize weight and by raising the risk of infestation. Dry

[13] 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).

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 ten-fold in just a few days if maize grain is not dried properly (Hell, at al., 2008).

An effective quality control strategy thus requires the prevention of defects and control of moisture in the maize as early as possibly in the value chain. Farmers can influence maize quality by harvesting at the right time, drying the maize quickly and thoroughly, shelling the cob without breaking or cracking the grains, not drying or storing cobs on the bare ground, and cleaning and storing the grain correctly.

The different components of maize quality can, in principle, be tested. As the full EAS testing protocol requires maize to be tested in the lab, it is rarely applied in the informal market. Instead, quality testing, when performed, relies on inspections of the bags of maize grains for the presence of various defects and measuring of moisture using either subjective tests (e.g. the crush test) or portable grain moisture meters.

4.2. Quality Measurement

To measure the quality of maize sold by farmers, and the extent to which quality can be verified at the farm gate, we performed four tests: visual inspection and moisture measurement at the farm gate, and laboratory testing and aflatoxin measurement in Kampala.[14] For the measurement experiment, we enrolled 100 farmers from nine villages in Sample Frame 3 that were about to harvest their maize and assigned half of them to a treatment group and half of them to a control group. This sample forms the basis for the returns to quality experiment discussed in the next section. Here, we focus on the subset of control households (50 households).

To measure farm gate maize quality, trained enumerators visited each farmer at the time of sale and recorded, based on visual inspections, the presence (or absence) of 10 types of defects in each bag put up for sale. We denote the mean number of those defects (in a given bag) as “Visually verifiable defects”.[15]

After inspection, one bag per farmer was randomly selected and bought from the household and transported to a laboratory in Kampala for testing.[16] Samples of 200g were

[14] Details on the test protocol are in online Appendix section B.
[15] Enumerators verified and recorded whether the maize in the bag was dirty, included cobs, stones, dust, insects (live or dead), and broken, immature, damaged, rotten, and mold-infested grain.
[16] The lab testing protocol followed the EAS approved objective test methods for defects. 44 bags were tested in the lab (see online Appendix section B).

drawn from each bag, and the weight of all non-grain substances and defective grain recorded. The total weight of the defects, expressed as percent of the sample weight, is denoted “Lab verified defects”.

To test for aflatoxin, we randomly sampled an additional 30 households from six villages in Sample Frame 3 over two consecutive seasons. In this sample, at the time of sale, one bag per farmer was randomly selected and bought from the household. The purchased bags were brought to Kampala for lab measurement of defects, as described above, and also tested for aflatoxin.[17]

In both samples, the field enumerators 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 1 for maize bags with a moisture content above 13%.

Finally, we combined the lab verified defects and the moisture measure to classify all samples tested in the lab using the East African Quality Standard.

4.3. 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% defects; i.e., a quarter of the weight of maize sold consists of waste with defective grains constituting the majority of defects. 29% of the households sold maize with a moisture content higher than 13%.

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% of the bags contained maize of grade 2 quality and 9% contained maize of grade 3 quality. The remaining bags, 71%, contained under-graded 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 0 to 2, the percentage of waste in the lab sample increases from 10 to 15%. As the number of defects doubles from 2 to 4, the percentage of waste in the lab sample also doubles.

[17] 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 (or ng/g) and the US standard is 20 ppb (Sserumaga et al, 2020).

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 the limit imposed by the Uganda National Bureau of Standards. 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 (> 10 ppb).

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.

5. Returns to quality experiment

5.1. Introduction

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 agriculture 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 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:

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

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.

5.2. 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 post-harvest (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.

5.3. 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 post-planting survey was administered. Online Appendix G Table 3 compares pre-harvest outcomes between treatment and control groups in the experiment. None of the collected covariates show statistically significant differences across assignment groups and a joint balance test fails to reject the null hypothesis that the pre-harvest 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. Compliance was 100%; i.e., all treatment households 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 UGX 10,000 (approximately USD 3) if they did so. Farmers in the comparison group were also visited and offered a (larger) monetary reward (UGX 30,000; approximately USD 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 (or absence) of defects in all bags the farmer was planning to sell (see section 4.2. and online Appendix B for a detailed description of the measurement). The enumerators also weighed the bags and tested the moisture content using a mobile moisture meter. Altogether 622 bags were visually inspected. In addition, one bag, drawn at random, was bought from each farmer for further quality analysis in the lab.[18] After selling their maize, all farmers were visited for a third and final time and asked about sales volume and prices. One farmer decided not to sell any maize in the season under study and of the remaining 98 farmers, four could not be reached. In total, we collected data on 116 sales from 94 households.

[18] As discussed in online Appendix C, one farmer did not sell a bag for testing and 16 bags could not be tested due to administrative constraints. In total 84 bags were tested in the lab.

5.4. Results: returns to quality

The free service package successfully raised quality in the treatment group. As illustrated in Figure 2 Panel A, the average sale in the control group contained 2.5 defects per bag (maximum 10 defects), ranging from 0 to 7 defects. 86% 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% of bags contained one or more defects, yielding a mean difference in defects of 2.2 (p =.000); see Table 2, column 1. Observable quality differences between treatment and control groups were equally stark when focusing on the randomly selected bag sent for lab testing and when assessing quality using laboratory testing. One third of the content in the average bag brought in for testing in the control group consisted of defective maize grain and waste. In the treatment group, the corresponding number was 6.5%. The difference in means across groups is highly significant; see Table 2, column 2. None of the bags purchased in the treatment group contained excessive moisture while 14% 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 4.2). 95% of the lab-tested bags in the treatment group were graded maize, of which 80% were grade 1 or grade 2 maize. In the control group, 58% were classified as ungraded maize, and the remainder as either grade 2 (30%) or grade 3 quality (12%).

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 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 = .84). The non-parametric 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 (13% of households sold more than once during the season). The treatment effect is essentially zero; i.e., there is no evidence that higher maize quality – equivalent to fewer defects counted in the bag – systematically yields a higher price. The coefficient is also tightly estimated, with the 95% CI spanning a reduction in price by 2.3% to an increase in price of 3.1% in the treatment group relative the control group.

Why do traders not pay higher prices for better quality maize? A first possible explanation for the absence of a relationship between farm gate quality and price is that traders simply cannot infer the true quality of maize from observable defects and hence do not adjust prices. In Section 4.3, however, we showed 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. Based on what farmers report, it is also common that traders perform some inspections before buying.

A second possible explanation for the results is that adjustments for quality differences are not done through prices, but through deductions in weight.[19] 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% and in one out of four sales more than 5% of the weight is deducted), the extent of deductions is similar across groups.[20] 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 kilo 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.[21]

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 (Antras, 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 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,

---

[19] Here we pool together data on deductions and cheating. Deductions refer to a transparent process whereby the buyer weighs the bag, reports the correct weight to the farmer, and then states they will deduct a share of the content of the bag 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%), 2.9 pp are deductions. The remaining discrepancy (0.8pp) is (possibly) due to cheating.
[20] The Kolmogorov-Smirnov D statistic for the test of equality of the treatment and control distributions for deductions is 0.17 [p = .39].
[21] 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.

our results still suggest that the farmer would not be rewarded in the first season they upgrade quality, which lowers the return to upgrading.

6. Market for quality experiment

6.1. Introduction

The results from the returns to quality experiment in section 5 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 4.

A possible interpretation of these findings is that the prevailing market equilibrium is primarily caused by a lack of demand for quality maize. 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. That is, maize that exceeded a certain quality threshold $\bar{q}$ was paid a premium ω over and above the prevailing market price $\tilde{p}$, resulting in an inverse demand schedule $p(q) = \tilde{p}(1 + I_{q \geq \bar{q}}\omega)$, where $I_{q \geq \bar{q}}$ is an indicator function equal to 1 if quality exceeds the minimum threshold and zero otherwise.

The intervention also aimed to improve farmers' knowledge on how to produce higher quality maize (as well as increasing their general knowledge of best-practice pre- and post-harvest agronomic activities) through extension services. 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.

6.2. Interventions

6.2.1. Market intervention

The intervention was designed to provide treated farmers with access to a market for high quality maize. To this end, we collaborated with a Ugandan vertically integrated agro trading company which committed to buy quality maize in treatment villages. The company used agents, overseen by a manager, to run their buying operation. Agents contacted all predetermined households in person or by phone before buying commenced and were present in the villages throughout the buying-seasons.[22,23] When a household was ready to sell quality maize, 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 reject bags that included waste or defective maize, as well as maize with excessive moisture and only buy maize bags that were of sufficient quality at a predetermined price. If the farmer was selling several maize bags of different (observed) quality, the agents bought the bags of sufficiently high quality. The households were informed why a maize bag had been rejected. The company then processed the maize bought in the treatment villages, and sold quality flour to customers in Kampala willing to pay a premium.

The research team randomly selected which villages the company should be active in and randomly selected the households, in each village, invited to participate in the trial.

The research team also determined the premium for quality maize with the aim of reproducing a market equilibrium. Since such a market does not exist in the villages, we used a simple model predicting the 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 framework we use (see online Appendix D) farmers can produce and sell either low (denoted by subscript L) or high (denoted by subscript H) quality maize, with

[22] Most farmers (70%) sell their maize during a two-month period, between mid-January and mid-March in the spring season and mid-July and mid-September in the fall season. Some farmers, likely for financial reasons, sell their maize early, while some sell outside the main selling seasons, when prices usually, but not always, are higher. In the first follow-up season, the company was active buying in the treatment villages for one month. In the remaining three follow-up seasons, the company was active for the full season (8-10 weeks).

[23] 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 non-sampled neighbor and sell it as if it was the farmer's own maize. The company did not buy from villages in the control group, or from villages in Sample frame 3, or in the villages involved in the extension-services-only trial. The research team strictly enforced this constraint.

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 turn low quality maize into high quality maize by sorting away defects and waste and, in the benchmark model, that doing so is costless. We then solve for the minimum premium, $(p_H - p_L)$, which the buyer needs to pay for high quality maize, which is simply the share of defective kernels and waste in low quality maize, valued at the premium price. Based on pre-treatment data, we predicted that maize with no visually verifiable defects, and a moisture level below 13%, would constitute maize of essentially grade 3 quality (using the EAS grading system); i.e., that the company would be able to buy grain containing 14 p.p. less waste and defects compared with the average quality on the market (see section 4.3). We further assume, again based on pre-treatment data, that local prices, on average, are 10% lower than prices in the trading centers. These assumptions yielded a target premium relative to trading center prices of 5% and an estimated premium relative to local (or village) prices of approximately 15%.

The quality premium we chose should be viewed as a lower bound of a market-based premium. First, in the model, the premium is set such that the farmers are indifferent between producing and selling high and low quality; i.e. the farmer's participation constraint binds. Second, the premium increases if we relax the assumption that the cost of sorting and cleaning away waste and defective kernels is zero. 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 in section 6.4.2, our estimate of the share of defects in the maize produced by treatment farmers in the intervention turned out to be too high.

6.2.2. Extension service intervention

To ensure that farmers had up-to-date knowledge about the pre- and post-harvest practices considered necessary for producing maize of sufficiently high quality, the agro trading company also 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 post-harvest tasks. All treatment households were invited to the demonstrations, and close to 70% of the invitees attended the meetings in the market experiment, while 78% 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.

6.3. Experimental design, data, and power

6.3.1. Trial design: Market experiment

We chose a clustered repeated measurement design for the experiment. Specifically, we restricted the number of clusters (20) but expanded on the number of waves, or seasons (7). The 20 clusters were randomly assigned to two groups: 12 to the buying group and 8 to the control group.

This 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 buying 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 rather than 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) because of large aggregate variations, impacts may vary dramatically from season to season (see Rosenzweig and Udry, 2000). Subject to these considerations, the final combination of clusters and waves was then 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.

6.3.1. Trial design: Extension service experiment

The trial design for the extension service experiment followed closely the market experiment design, but with 18 clusters followed over six seasons (see online Appendix H Figure 1).[24] 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.

6.3.2. 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.

[24] The initial design had 20 clusters. A land conflict broke out in two villages (one control village and one treatment village) in the end of 2019. This resulted in two changes of 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.

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 and labor use and sales data. In order to ensure data quality, GPS data from the previous season was pre-loaded 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, sales data were collected for each sale. The survey also collected detailed expense data, including on chemical use, seed varieties, and various pre-harvest and post-harvest 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% 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 (or the) 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 (which mechanically influences the income from farming). 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 that is very similar to the agricultural trading company's activities: the company was present in treatment villages during the selling period; it assessed quality before buying; and it paid a premium (albeit for quality maize). Because of this overlap in activities between quality measurement and treatment, we deemed the risk arising from potential Hawthorne effects (e.g., farmers may have perceived that any maize brought for sale was externally monitored for quality defects and changed behavior accordingly) and disruption of normal trading activities in the villages (which in turn could have influenced prices for sales to other traders) as too large. In sum, in the market access trial, the measurement of quality and the buying intervention may interact in ways (positively or negatively), which effectively mean that it would not have been possible to net out any direct effects of quality measurement by simply comparing treatment and control quality outcomes.

To minimize these potential problems, data on maize quality was collected using a complementary approach. In the treatment group quality testing was weaved into the buying operation; i.e., randomly selected bags from all farmers selling maize 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 sub-set 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.[25]

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 treatment effect by treating the problem of missing data in the treatment group as one of non-random one-sided attrition. In online Appendix G Table 10 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 significant differences in mean outcomes across households in the control group in the market access experiment and households in the quasi-control group. The joint balance tests fail to reject the null hypothesis that household characteristics and farm enterprise outcomes do not predict assignment to the control group ($p = .771$ and $p = .223$, respectively).

[25] 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 (post-harvest 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.

6.3.3. Estimator and power calculations

Our benchmark ANCOVA specification uses only follow-up data for the dependent variable and regresses outcome Y_{ijt} , where sub-script i denotes individual, j denotes cluster, and t wave or season, on a treatment indicator, D_{jt}, which takes on the value 1 in treatment clusters and zero in control clusters, a full set of season dummies and a lag-dependent variable, i.e., the value of the outcome pre-treatment $\bar{Y}_{ij,PRE}$:

$$Y_{ijt} = \gamma D_{jt} + \sum_{n-3}^{n} \delta_t + \theta \bar{Y}_{ij,PRE} + \varepsilon_{ijt} \ , \tag{3}$$

where n denotes the number of rounds and where ε_{ijt} is an idiosyncratic error. The coefficient of interest, γ, gives the average causal effect over the four follow-up rounds. We report point estimates and p-values, with the latter estimated both based on clustered-by-village standard errors and computed using randomization inference with permutations of treatment done at the village level.

We estimated the power of the designs before the trial began using data from a small pilot. Online Appendix G Table 5 updates these power calculations for the ANCOVA estimator using data from the control group. Expressed as share of the control group mean, the MDEs for the market access experiment vary from approximately 10% (for price and yield) to 35% (for expenses and harvest). The power pattern for the extension service experiment is similar even though the design includes one less follow-up round; a result driven by a lower intraclass correlation in this sample.

6.3.3. 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 10 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 5 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 pre-intervention 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 6).

Follow-up in the market access experiment lasted for four seasons. As reported in Figure 4 and online Appendix Table 6, less than 5% (9 households) of the 189 households in the final baseline sample attritted. The attrition rates were similar across assignment groups (see online Appendix G Table 7). Of the non-attritters (180 households), 86% were re-surveyed 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. The re-survey rates; i.e., the share of the 180 households that were surveyed per season, were similar across assignment groups (see online Appendix G, Table 7). 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 seven seasons.

The sample for the extension service experiment included 498 household-by-season observations from 173 households in 18 clusters (see Figure 4 and online Appendix G Table 6). Follow-up lasted for three seasons. 9 households (5%) attritted. The attrition rates were similar across assignment groups (see online Appendix G, Table 7). Of the non-attritters, 84% were re-surveyed 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. The re-survey rate was similar across assignment groups. The complete baseline and follow-up panel consists of 931 household-by-season observations from 164 households in 18 clusters over six seasons.

Table 3 reports summary statistics and mean comparisons between the treatment and control groups across a broad set of outcomes in the market access experiment. Panel A focuses on household characteristics and Panel B, and online Appendix G Table 9 Panel A, present baseline farm enterprise outcomes. There are large variations across seasons. For example, the price of maize was more than 50% higher in the first season than in the third and profit in season two was roughly 50% higher than in season three.

Summary statistics for household characteristics and farm enterprise outcomes (the latter pooled across the three baseline seasons) by assignment group in the market access experiment are reported in columns (4)-(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 grouping jointly predict treatment assignment. Our joint balance tests fail to reject the null hypothesis that neither household characteristics ($p = .998$) nor farm enterprise outcomes ($p = .617$; $p = .601$) predict assignment to treatment.

Summary statistics and mean comparisons between the treatment and control group in the extension service experiment are reported in online Appendix G Table 8 and Table 9 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 fail to reject the null hypothesis that the household characteristics (p=.237) and farm enterprise outcomes (p=.141; p=.121) do not predict assignment to treatment.

6.4. Results

6.4.1. Summary

We begin by summarizing how the market access intervention affected the main outcome of interest: profits. Figure 5 Panel A1 shows that the cumulative distribution function (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 = .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 \$63-\$98 or 36%-80% (depending on how own and family labor is priced). In contrast, the treatment and control profit functions in the extension service experiment, where farmers were provided with hands-on training of pre- and post-harvest best practices but not with access to a market for quality maize, lie on top of each other (Panel B1) and the Kolmogorov-Smirnov test fails to reject the null that the two distributions are equal (D statistic = 0.062, p = .787).

6.4.2. Quality upgrading

To measure the extent of quality upgrading, we examine two types of data: (i) the share of farmers who sold to the high quality buyer coupled with administrative data on how the company enforced quality standards; (ii) laboratory measured quality in bags.

In each post-treatment season, the agricultural trading company offered to buy maize of sufficiently high quality from pre-selected households in the treatment villages. Averaging across the four post-treatment seasons, 40% 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%). 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 Panel C-D.

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.[26] Approximately four out of 10 of the households did not sell (quality) maize to the company. One-third of the households sold all they wanted to sell. For 15% of the farmers, the company first refused to buy (and required the maize to be sorted, cleaned and/or dried further), but then bought at least a subset of bags once the farmer had upgraded the quality. For one in ten households, the company refused to buy because the quality was low.

Further evidence that farmers learned over time (or were willing to spend effort and money) to produce high quality maize is evident when considering farmers who attempted to sell in season 6 but were refused by the company at first. Half of those approached the company also in season 7, but this time with sufficiently high quality maize, implying the company bought all the households wanted to sell directly. For another 40%, the company bought once additional post-harvest processing had been successfully completed by the farmers. The remaining ten percent of households decided to sell to other buyers.

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. Or in other words, that many farmers were eager to sell to the company, and learned over time how to produce maize of sufficiently high quality.

In addition, we measured maize quality in the lab for randomly selected bags from all farmers selling to the agricultural trading company in season 7. Column (1) in Panel A of Table 4 reports the results based on the EAS classifications for maize quality. Pooling the data at the household level, 89% 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.

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

To assess the impact of the intervention on lab measured 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% of the maize bags tested were classified as 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 59pp higher or almost three times as high as in the quasi-control villages (p = 0.000).

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 graded maize across the two groups could be explained by non-random 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 non-random missing data, we assume that the 35% farmers in treatment villages who do not sell to the high quality trader in season 7 produce maize quality that is equivalent to the distribution of maize quality in control up to the 35th percentile. Specifically, we estimate the Horowitz-Manski lower bound (Horowitz and Manski, 2000), which imputes unobserved maize quality in the treatment group as ungraded, and the Lee lower bound (Lee, 2009), which trims the excess share of observations in the control group, dropping those with the lowest quality.

Even allowing for this extreme form of selection, we estimate a treatment effect of 28pp or 93% (p = 0.094) and 43pp or 143% (p = 0.092) in the two bounding exercises (see Table 4, Panel B). Comparing these estimates to the quality difference of 59pp between maize sold to the high quality buyer and maize sold in quasi-control, we can conclude that non-random missingness explains at most half of the observed difference in maize quality between treatment and control. By implication, at least half or more of the measured difference in maize quality between treatment and quasi-control is driven by treatment farmers upgrading the quality of their maize.

Thus both types of evidence point in the same direction: there is a large difference in quality between treatment and control and this was driven by farmers actively learning how to upgrade the quality of their maize over time.

To put these results in perspective, assume a buyer is aiming to procure/produce 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 receive maize of quality we observe in the treatment group. Using data from the laboratory tests, buying from the treatment group, the buyer would need to sort away on average 0.8% of the maize bought. Buying from the quasi-control, the buyer will need to sort away on average 20.4%. Assume the marginal cost of sorting is 0. 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 $\tilde{p}/(1-.204)$. 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%; i.e., almost 10 pp 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% of the bags had live insects in the bag tested in the lab, the maximum premium the buyer would be willing to pay would be even higher.

6.4.3. Investments and productivity

Market access may encourage farmers to invest more via two channels: (i) the intervention offered farmers in the treatment group 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 pre-harvest activities (see Table 5, Panel A), investments that primarily – though not exclusively – affect how much maize is produced. Specifically, farmers spent an additional \$2.3 or 74% (p =.048, control mean \$3.1) 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, pesticides and herbicides, increased by \$4.3 or 30% (p =.06, control mean \$13.3). Although these treatment effects represent large relative increases, in absolute terms, modern input use is low: 3% of control farmers used inorganic fertilizer, 13% used improved seeds, and input expenses amounted to 13% of all expenses. Second, farmers spent \$16 or 30% (p = .275, control mean \$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.

Farmers also invested more post-harvest, 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, 13% sorted their maize and one fifth winnowed it (columns (5)-(7) in Table 5, Panel A). With access to a market for high quality maize the share of farmers who engaged in these practices nearly doubled: 60% dried their maize properly (a difference of 24 percentage points, p = .001), 27% sorted the maize (a difference of 14pp, p = .001) and 34% winnowed it (a difference of 15pp, p = .036). Consistent with this, spending on harvest and post-harvest activities rose by 20% (p = .255, control mean $30), an increase mainly driven by higher expenses on hired labor (a difference of 40%, p = .144, control mean $15.6).

Summing across all items of cultivation expenditure, farmers in treatment villages invested $18 more than those in control villages (Table 6, Panel A, column (6)), an increase of 17% (p = .305, control mean $107). The area under cultivation did not change in treatment (Table 6, Panel A, column (2)).

Farmers in treatment villages increased their total maize harvest as well as their yields. Figure 5, Panel A2, 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 112 kg or 15% (p = .036, control mean 793 kg), and total harvest by 239 kg or 13% (p = .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 remain the same, but the effects are more precisely estimated (see online Appendix G Table 11 Panel A).[27]

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; i.e., a production function, and add an additional assumption; namely that observable inputs are independent of unobservable inputs, or total factor productivity (TFP), given treatment status. The results suggest that one-third of the treatment effect on harvest comes from increases in measured inputs, with hired labor being the most important factor. The lion share of the output increase is thus accounted for by improvements in TFP.

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 B2, strongly suggest the answer is no. With the exception of winnowing, which increased by 8pp or 57% (p = .050, control

[27] The share sold is similar in the two assignment groups, implying both the amount consumed at home and the amount sold increased in the treatment group relative to the control group.

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 post-harvest processes was put to use with market access.[28] But the results from the extension service trial make it unlikely that this kind of supply intervention by itself would significantly and sustainably change how farmers operate.

6.4.4. Prices and income

Farmers with access to a high quality market 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 A3). The Kolmogorov-Smirnov D statistic is 0.34 (p = .000). The regression equivalent is presented in Table 6, Panel A, column (1): on average, farmers earned \$2.40 or 11% more per bag of maize (140 kg) they sold (p = .001, control mean \$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 \$78.7 or almost 30% per season (p = .079, control mean \$286.7).

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 $\bar{Y}^d$ and $\bar{p}^d$, and Δx the treatment effect on outcome x, the treatment effect on harvest value (pY) can be decomposed into a pure price/quality-effect ($\bar{Y}^0 \Delta p$), a pure quantity effect ($\bar{p}^0 \Delta Y$) and an interaction:

$$\Delta pY = \bar{Y}^0 \Delta p + \bar{p}^1 \Delta Y + \Delta p \Delta Y \ . \tag{4}$$

Given the treatment effects on price and harvested amount, the quantity effect accounts for 46% of the increase in harvest value, the quality effect accounts for 48%, 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

[28] 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 et al., 2007; Hanna et al., 2012; and Islam and Beg, 2020). For recent reviews of the literature on extension service, see Macours (2019), Magruder (2018), and Takahashi, et al. (2020). Bernard (2017) et al. show that adoption behavior may depend on demand side conditions: they find that Senegalese onion producers adopted quality-enhancing inputs when they expected market structure to change from rewarding volume to rewarding quality.

farmers' harvest value, we find that farmers in the treatment group on average earned \$63 or 36% more per season than farmers in control villages (p = .062, control mean \$178); Table 6, column (7). The treatment effect becomes even larger and more precisely estimated when trimming the top and bottom 1% 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 75 hours per season or 16% (p =.106, control mean 449 hours). On the other hand, and in line with the overall increase in expenses, treatment farmers hire substantially more labor.[29] Farmers in treatment villages increase their spending on hired labor by \$26 or 36% (p = .322, control mean \$71), equivalent to an additional 121 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, 2020). Summing the two sources of labor, we estimate a net increase in labor of 46 hours per season in the treatment group.

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.[30]

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

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

[29] There is an active market for hired labor in the study villages, with farmers in control villages spending on average \$71 on hired labor equivalent to 70% of all monetary expenses and roughly 315 hours per season at the hourly wage.

[30] 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 \$98 higher in treatment than in control villages (p = 0.029, control mean \$120). Trimming reduces the effect size (70% increase in treatment compared with control; p = .020). For $\varphi = 2/3$, which corresponds to our guesstimate of the relative productivity of family and hired labor (see footnote 30), the treatment effect on profits is \$88, or 55% (p = .040, control mean \$165).

These effects represent large absolute increases in the context of our study, where most people live on less than 1 dollar 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% increase in average annual income relative to a typical family in the region (UBOS 2019).

In the extension service experiment, all such effects on prices, revenue and income are absent. As seen in Panel B of Table 6, the estimated treatment effects are close to zero and insignificant.

6.4.5. 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 (inasmuch as they have any); or it may improve farmers' bargaining position relative to the trader simply by improving their outside option. Our trial, which induced variation in exposure to the new buyer across clusters, was designed to (partly) capture such spillover effects.[31]

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 price received and the type of buyer; i.e., whether a sale went to the high quality buyer (denoted HT) or to other

[31] Full details of the estimation in this section are given in online Appendix F.

traders (OT), where other traders are composed of local traders (LT) or commercial (CT) traders (see section 2.2). With this data, we estimate the average proportion of sales across all follow-up seasons $\bar{s}_{d,j}$ from households with treatment status $d = \{0,1\}$ to traders of type $j = \{OT, CT, LT\}$, as well as their difference, $\Delta\bar{s}_j$, between treatment and control (see online Appendix F for details). We also calculate the average price $\bar{p}_{1,j}$ paid by traders of type j to treatment households, as well as the difference between $\bar{p}_{1,j}$ and the average control group price $\Delta\bar{p}_{j,0} = \bar{p}_{1,j} - \bar{p}_0$.

Entry of the high quality buyer decreased the average market share of other traders by almost 40% ($\Delta\bar{s}_{OT}$ = -0.396; p = .000); see Table 7 Panel A. Of these, local traders' share of sales fell from 80% in control to 47% in treatment ($\Delta\bar{s}_{LT}$ = -0.325; p = .000) and commercial traders' share of sales decreased from 20% in control to 13% in treatment ($\Delta\bar{s}_{CT}$ = -0.071, p = .267). 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 ($\Delta\bar{p}_{OT,0}$ = 0.045; p = .107). This price increase is mainly driven by prices for sales to local traders, which increased by 6.2 percent relative to the control group (p = .028). 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 = .687).

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% (p =.056). 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%), 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. Conversely, sales to the high quality buyer at follow-up come from farmers who earned 4 percent more than others at baseline.[32]

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 30% of the average price increase in treatment relative to the control is driven by market spillover effects.

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.

[32] The difference in difference estimation for sales can be further extended to distinguish which trader type $j = LT, CT$ the farmer's sales in follow-up went to. As reported in Table 7, the entry of the high quality buyer raised prices for sales to local traders by 7.8 percent (p=.039) and commercial traders by 2.1% (p=.683). 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 the increase in output and productivity in the treatment relative to the control group can be (partly) explained by farmers not selling 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.

7. 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, however, also highlights challenges. First, while farmers increased their use of modern inputs – one pathway to increased productivity – adoption of these technologies remained low. A number of promising, and potentially interlinked constraints, focusing on the supply side, have recently been studied in the literature (see review in de Janvry, et al., 2017), including missing markets for risk (Karlan et al., 2014) and behavioral constraints (Duflo et al., 2011). Assessing the impact of coupling such supply interventions with market access or value chain inclusion and understanding complementarities between them is, we believe, an important area for future research.

The second challenge is operational. The market access intervention was implemented by an integrated domestic buyer to circumvent many of the potential agency- and information problems that plague the market for (lower quality) maize.[33] After factoring out all evaluation costs, the agro trading company broke even in two of the four buying seasons.[34] Adding

[33] As a reference, in high income countries, agricultural outputs are produced in closely aligned segments of the value chain by actors exploiting economies of scale. In Uganda and most of Sub-Saharan Africa, on the other hand, markets for agricultural outputs are segmented, and the low quality segment is characterized by a multitude of layers of small actors. For example, maize in Uganda often passes through several sets of traders, or aggregators, before reaching mills located in urban centers (Daly et al., 2016).

[34] The company did not break even in the first season of operation, partly because farmers were not easily convinced to sell to a new buyer and partly because the company itself first had to establish a customer base for high quality maize. Hence in the first season the amount of maize bought and flour sold and the premium for flour were all relatively low. As both buying maize and selling flour are associated with fixed cost, the average cost was high. After that, the amount of maize the company purchased and the amount of flour it sold as well as the premium increased with every season. This growth allowed the company to break even already from season 2. The same would likely have been true in season 4 had it not been for Covid-19 school closures (private boarding schools in Kampala were the company's main customer base). That is, the company had already secured flour orders for the lion share of the maize they had purchased, but when schools were closed demand for their product collapsed. The growth trend of the company was thus upward until the time of the Covid crisis

farmer profit, joint surplus in these two seasons was therefore strictly positive. Three structural features of the product and the economy constrained the company's ability to increase revenues. First, as quality is more difficult to determine once the 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, though 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 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 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 here, and the revenue and costs 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.

and we think it is reasonable to assume that this would have continued in the absence of the crisis. Note that by focusing on the financial surplus, we rule out any potential health benefits of higher quality maize.

References

ATAI. 2018. “Evidence for transformation: Framing a research agenda in agriculture for development.” CEGA, UC Berkeley.

Antras, Pol. (2015). Global production: Firms, contracts, and trade structure. 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 vol. 91 (4): 973-990.

Atkin, David, Amit K. Khandelwal, and Adam Osman. 2017. “Exporting and firm performance: Evidence from a randomized experiment.” The Quarterly Journal of Economics 132(2):551-615.

Baker, George, Robert Gibbons and Kevin J. Murphy. 2001. “Bringing the Market Inside the Firm?” American Economic Review, 91(2): 212-218.

Beaman, Lori, Dean Karlan, Bram Thuysbaert, and Christopher Udry. 2014. “Profitability of Fertilizer: Experimental Evidence from Female Rice Farmers in Mali.” American Economic Review vol. 103 (3): 381-86.

Bergquist, Lauren Falcao and Craig McIntosh. 2021. Search Cost, Intermediation, and Trade: Experimental Evidence from Ugandan Agricultural Markets. Working Paper, University of Michigan.

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 vol. 99 (4): 1096-1115.

Blouin, Arthur and Rocco Macchiavello.2019. Strategic default in the international coffee market. The Quarterly Journal of Economics, 134(2):895-951.

Burlig, Fiona, Louis Preonas, and Matt Woerman. 2020. “Panel data and experimental design.” Journal of Development Economics, vol. 144: 102458.

Casaburi, Lorenzo and Tristan Reed. 2017. “Competition in agricultural markets: An experimental approach.” CEPR Discussion Papers 11985.

Casaburi, Lorenzo and Tristan Reed. 2020. “Interlinked Transactions and Competition: Experimental Evidence from Cocoa Markets.” Working Paper.

Daly, Jack, Danny Hamrick, Gary Gereffi, and Andrew Guinn. 2016. Maize value chains in East Africa. Report F-38202-RWA-1, IGC, London.

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 vol. 9, 13158.

de Janvry Alain and Elisabeth Sadoulet. 2019. “Transforming developing country agriculture: Removing adoption constraints and promoting inclusive value chain development.” FERDI Working paper P253.

de Janvry, Alain, Elisabeth Sadoulet, and Tavneet Suri. 2017. "Field Experiments in Developing Country Agriculture". Handbook of Economic Field Experiments. Vol. 2. 427-466.

De Loecker, Jan and Pinelopi Koujianou Goldberg. 2014. "Firm Performance in a Global Market." Annual Review of Economics, 6(1): 201–227.

Deutschmann, Joshua W., Maya Duru, Kim Siegal, and Emilia Tjernström. 2019. "Can Smallholder Extension Transform African Agriculture?" NBER Working Paper No. 26054.

Dillon, Brian and Chelsey Dambro. 2017. "How Competitive Are Crop Markets in Sub-Saharan Africa?" American Journal of Agricultural Economics 99: 1344-1361.

Donaldson, Dave. 2015. "The Gains From Market Integration." The Annual Review of Economics 7: 619-647.

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), 2350–2390.

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

Hansman, Christopher, Jonas Hjort, Gianmarco León-Ciliotta, and Matthieu Teachout. 2020. "Vertical Integration, Supplier Behavior, and Quality Upgrading among Exporters." Journal of Political Economy, 128(9): 3570-3625.

Hausmann, Ricardo, Jason Hwang, and Dani Rodrik. 2007. "What You Export Matters." Journal of Economic Growth 12 (1): 1-25.

Hell, Kerstin, Pascal Fandohan, Ranajit Bandyopadhyay, Sebastian Kiewnick, Richard Sikora, Peter J. Cotty. 2008. "Pre- and postharvest management of aflatoxin in maize: an African perspective." Mycotoxins: detection methods, management, public health and agricultural trade. Leslie, J. F., Bandyopadhyay, R., Visconti, A. (eds). CAB International.

Heß, Simon, "Randomization inference with Stata: A guide and software" Stata Journal 17(3) pp. 630-651

Hoffmann, Vivian, Samuel Mutiga, Jagger Harvey, Michael Milgroom, and Rebecca Nelson. 2013. "A Market for lemons: Maize in Kenya." Working paper, University of Maryland.

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 (forthcoming).

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

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

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

Macchiavello, Rocco and Josepa Miquel-Florensa. 2019. “Buyer-Driven Upgrading in GVCs: The Sustainable Quality Program in Colombia.” CEPR Working paper No. 13935.

Macours, Karen. 2019. “Farmers’ Demand and the Traits and Diffusion of Agricultural Innovations in Developing Countries.” Annual Review of Resource Economics vol. 11:1: 483-499.

Magruder, Jeremy R. 2018. “An Assessment of Experimental Evidence on Agricultural Technology Adoption in Developing Countries.” Annual Review of Resource Economics vol. 10: 299-316.

McKenzie, David. 2012. “Beyond baseline and follow-up: The case for more T in experiments.” Journal of Development Economics vol. 99(2): 210-221.

Morjaria, Ameet and Martin Sprott. 2018. Ugandan Arabica coffee value chain opportunities. IGC Report F-43410-UGA-1.

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: 105-112.

Rosenzweig, Mark R. and Christopher Udry. 2020. “External Validity in a Stochastic World: Evidence from Low-Income Countries.” The Review of Economic Studies, vol. 87 (1): 343-381.

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 vol. 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 vol. 51: 31-45.

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

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

Table 1. Maize quality: Summary statistics

Quality measure:	Mean	Std	Min	Max	Obs
Visually verifiable quality (#)	2.5	1.5	0	7.0	355
Lab verified quality (%)	26.2	34.1	4.3	100	102
Wet maize (%)	29.4	-	0	1	109
Grade 1: EAS (%)	0	-	-	-	102
Grade 2: EAS (%)	20.6	-	-	-	102
Grade 3: EAS (%)	8.9	-	-	-	102
Ungraded: EAS (%)	70.6	-	-	-	102

Note. Unit of observation is a bag. Visually verifiable quality is the number of defects out of 10 detected in a bag of maize. Lab verified quality is grams of defects per 200g maize sample drawn from bags bought in the field (in %). Wet maize is a binary indicator for maize with a moisture level above 13%. Grade 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).

Table 2. Returns to quality: treatment effects

Specification	(1)	(2)	(3)	(4)	(5)
Outcome variable:	Defects (visually verifiable)	Defects (lab)	Price	Deductions	Net-price
Treatment	-2.16	-0.20	2.95	0.18	0.98
	(.212)	(.054)	(9.87)	(.63)	(10.6)
	[.000]	[.000]	[.766]	[.781]	[.927]
Observations	622	82	116	116	116
Households	99	82	94	94	94
R-squared	0.67	0.36	0.91	0.22	0.89
Mean control	2.46	0.32	731.2	3.73	704.6

Note. 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 10) verified by trained enumerators in maize bags for sale. Defects (lab) is grams of defects per 200g maize. Price is in UGX 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 UGX per kilogram. Specifications (1)-(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.

Table 3. Market access experiment: summary statistics and balance at baseline

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Sample		*Means*		*Difference in means*	
Variable	Mean	Std.	Obs.	T	C	Coeff.	*p*
				***Panel A*. Household characteristics**			
Main decision maker: female	0.18	0.39	189	0.19	0.16	0.026	.776
Main decision maker: completed primary school	0.41	0.49	189	0.42	0.39	0.026	.765
Number of household members	6.15	2.57	189	6.15	6.15	-0.006	.987
Distance to district capital (km)	29.8	9.53	189	29.5	30.2	-0.703	.873
Distance to main road (mins)	12.1	8.30	189	12.2	11.9	0.325	.872
				***Panel B*. Farm enterprise characteristics**			
Maize acreage	2.16	1.60	544	2.17	2.13	0.047	.900
Expenses (USD)	146.3	158.0	363	147.7	144.4	3.383	.929
Harvest (ton)	2.10	1.94	499	2.14	2.05	0.096	.844
Yield (ton/hectare)	2.09	1.04	499	2.13	2.04	0.092	.651
Share sold	0.82	0.24	499	0.82	0.83	-0.012	.699
Price kilogram (USD)	0.20	0.06	470	0.20	0.20	0.002	.682
Harvest value (USD)	433.6	426.2	498	451.0	409.3	40.20	.698
Profit I (USD)	264.0	292.5	362	274.1	250.0	23.27	.734
Joint balance test I							.998
Joint balance test II							.617
Joint balance test III							.601

Note. 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 2-3; sample size 336).

Table 4. Impact on maize quality

	(1)	(2)	(3)	(4)
	Mean		Difference	Obs.
	Treatment	Quasi-control		
***Panel A*. Quality per bag**				
Graded maize	0.89	0.30	0.59 [.002]	86
Grade 1 maize	0.07	0		
Grade 2 maize	0.52	0.20		
Grade 3 maize	0.30	0.10		
***Panel B*. Bounds**				
Horowitz-Manski lower bound			0.28 [.094]	116
Lee lower bound			0.43 [.092]	116

Note. Graded maize is grade 1-3 maize based on the East African Quality Standard (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 bought by the agro-trading company in season 7 (212 bags and 56 households), pooled at the household level. Data from the quasi-control group: one randomly selected bag bought from each of 30 randomly sampled households from six quasi-control villages in season 7 (30 bags). Panel A: unit of observation is household either 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 permutations of the treatment assignment in bracket. Panel B: unit of observation is a household. Sample is composed of all surveyed household in treatment and quasi-control selling maize in season 7. Horowitz-Manski lower bound with p-value from Fisher-permutations test based on 10,000 permutations of the treatment assignment in bracket. Lee lower bound, with bootstrapped standard errors (10,000 replications) adjusting for clustering and p-value in bracket. Number of selected observations (Lee lower bound) is 86.

Table 5. Impact on investment

Specification	(1) Expenses: seeds and fertilizer	(2) Expenses: all inputs	(3) Any improved seeds	(4) Any fertilizer	(5) Proper drying	(6) Sorting	(7) Winn-owing	(8) Pre-harvest expenses	(9) Post-harvest expenses	(10) Post-harvest expenses (labor)
					Panel A. Market access experiment					
Access to a market for quality maize	2.34	4.33	0.04	0.03	0.24	0.14	0.15	16.3	5.94	5.87
	(.048)	(.060)	(.246)	(.144)	(.001)	(.001)	(.036)	(.275)	(.255)	(.144)
	[.051]	[.077]	[.270]	[.159]	[.002]	[.001]	[.051]	[.296]	[.272]	[.153]
Observations	658	658	657	657	640	464	464	464	464	464
R-squared	0.31	0.33	0.08	0.02	0.18	0.03	0.05	0.20	0.26	0.22
Mean control	3.75	13.25	0.13	0.03	0.35	0.13	0.19	53.76	30.37	15.63
					Panel B. Extension service experiment					
Extension Service	-0.59	-1.02	0.02	0.02	0.03	0.04	0.08	-1.14	-1.10	-1.62
	(.643)	(.613)	(.678)	(.433)	(.759)	(.388)	(.050)	(.874)	(.794)	(.618)
	[.650]	[.634]	[.741]	[.481]	[.773]	[.426]	[.077]	[.880]	[.808]	[.626]
Observations	445	445	445	445	150	296	297	297	297	297
R-squared	0.20	0.40	0.03	0.01	0.01	0.00	0.06	0.39	0.25	0.05
Mean control	1.93	10.27	0.11	0.04	0.68	0.15	0.14	53.94	25.56	13.17

Note. ANCOVA specification. Clustered-by-village standard errors with *p*-values in parenthesis. *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, OPV or recycled hybrid seeds and fertilizer. Expenses on all inputs are expenses on modern seeds and fertilizer, booster and chemicals. Proper drying is a dummy equal to 1 if the maize was dried on a tarpaulin. Pre-harvest expenses are expenses on hired labor for pre-planting (ploughing and weeding), planting, weeding after planting and spraying. Post-harvest 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.

Table 6. Impact on productivity and income

Specification	(1) Price	(2) Maize acreage	(3) Harvest	(4) Yield	(5) Harvest value	(6) Monetary expenses	(7) Profit (monetary expenses)	(8) Profit (incl. own hours)
***Panel A*. Market access experiment**								
Access to a market for quality maize	0.02	0.05	239.3	111.7	78.7	18.3	62.9	97.5
	(.001)	(.829)	(.308)	(.036)	(.079)	(.305)	(.062)	(.029)
	[.004]	[.838]	[.350]	[.044]	[.103]	[.321]	[.079]	[.031]
Observations	617	677	658	658	658	640	640	464
R-squared	0.68	0.27	0.29	0.09	0.32	0.33	0.22	0.18
Mean for control	0.15	2.29	1887.5	793.1	286.7	106.5	177.6	120.3
***Panel B.* Extension service experiment**								
Extension service	0.01	-0.04	-110.0	33.7	-6.80	-4.40	2.81	13.4
	(.404)	(.800)	(.609)	(.589)	(.882)	(.657)	(.942)	(.727)
	[.416]	[.817]	[.631]	[.621]	[.893]	[.668]	[.946]	[.743]
Observations	420	458	445	445	445	448	443	443
R-squared	0.34	0.29	0.39	0.16	0.30	0.39	0.14	0.16
Mean control	.17	1.91	1678.4	872.9	303.01	86.7	218.3	168.1

Note. ANCOVA specification. Clustered-by-village standard errors with *p*-values in parenthesis. *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.

Table 7. Impact on trader prices and market shares

	(1) Other traders	(2) Local traders	(3) Commercial traders
Panel A. Difference in market shares and prices			
Difference in market shares	-0.396	-0.325	-0.071
	[.000]	[.000]	[.267]
Difference in prices vs. control	0.045	0.062	-0.016
	[.107]	[.028]	[.687]
Panel B. Difference in prices adjusting for selection			
Difference in prices vs. control	0.066	0.078	0.021
	[.056]	[.039]	[.683]

Note. See section 6.4.5 for details. 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, $\tilde{p}_{j,k,t} = (p_{j,k,t} - \bar{p}_{0,t})/\bar{p}_{0,t}$ where $p_{i,k,t}$ is the price for sale j by farmer k in period t and $\bar{p}_{0,t}$ 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.

Figure 1. Verifiability of maize quality

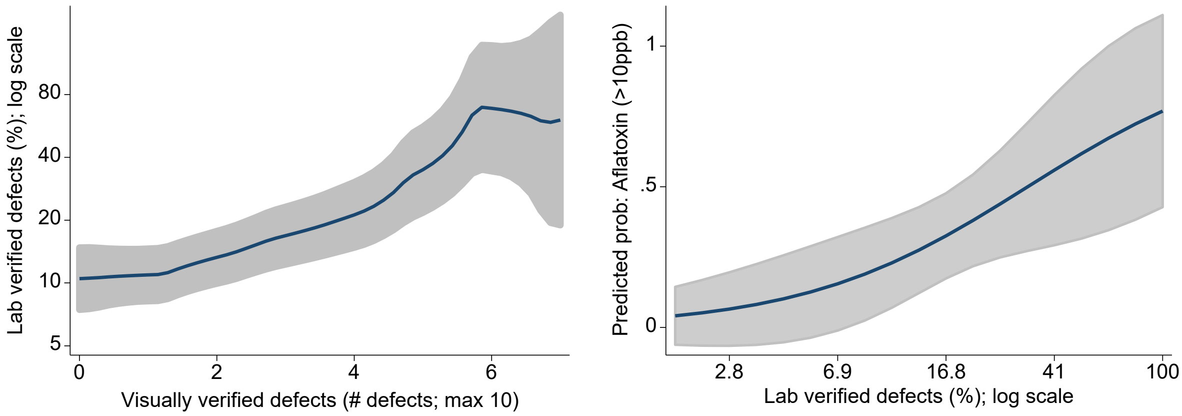

Note. Panel A: Smoothed values from a local polynomial regression of lab verified defects on visually verifiable defects. Grey shaded area represents the 95% confidence interval. Panel B: Predicted probability of aflatoxin above the UNBS cut-off as a function of share of defects in the bag (lab verified defect), from specification (3), in online Appendix G, Table 2, with the grey shaded area representing the 95% confidence interval. The unit of observation is a bag. See online Appendix B for details on the tests and samples used.

Figure 2. Quality outcomes in treatment and control

Panel A. Visually verifiable quality ***Panel B.*** EAS quality grades

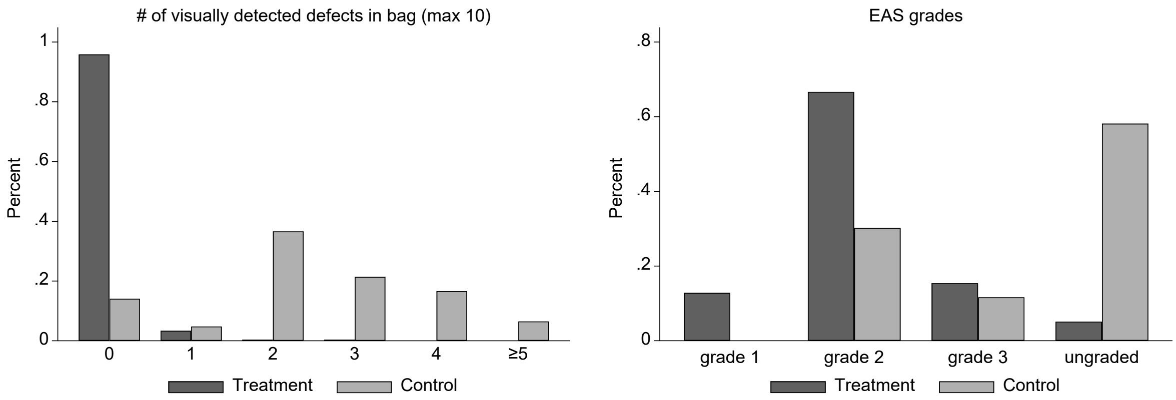

Note: Panel A depicts the average across all bags in visually verifiable quality; i.e., the number of visually detected defects (max 10) in maize bags. Panel B reports grades based on the East African Quality Standard (EAS) classification, using data from laboratory measurement of purchased maize bags.

Figure 3. Price and deductions in treatment and control

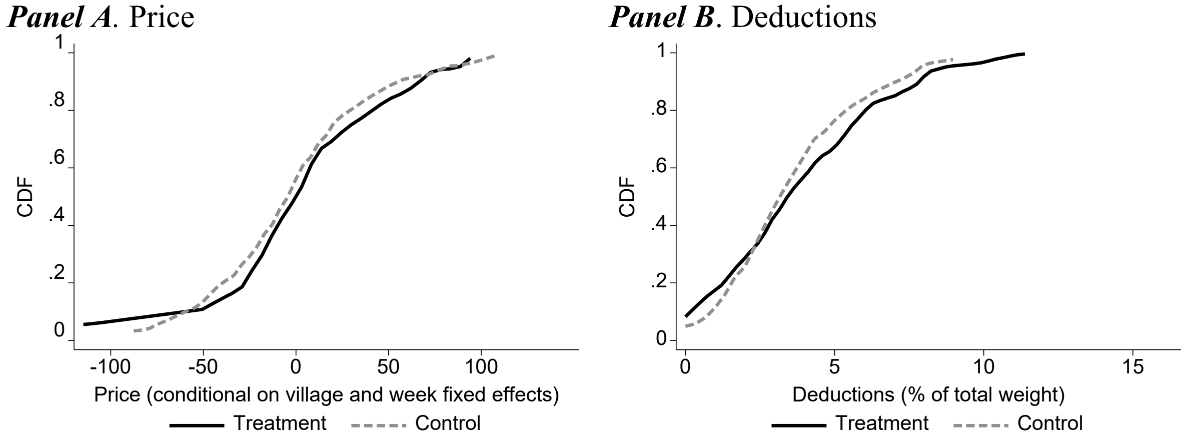

Note. Panel A: cumulative distribution functions of price (conditional on village and week fixed effects) in the assignment groups. Panel B: cumulative distribution functions of deductions in the assignment groups, with deduction defined as defined as $(y - z)/y$, where y is the weight of maize sold as measured by enumerators and z is the agreed upon sales volume. The Kolmogorov-Smirnov D statistic on the test of equality of the treatment and control distributions is 0.11 [p =.84] in Panel A and 0.17 [p = .39] in Panel B.

Figure 4. CONSORT flow diagrams

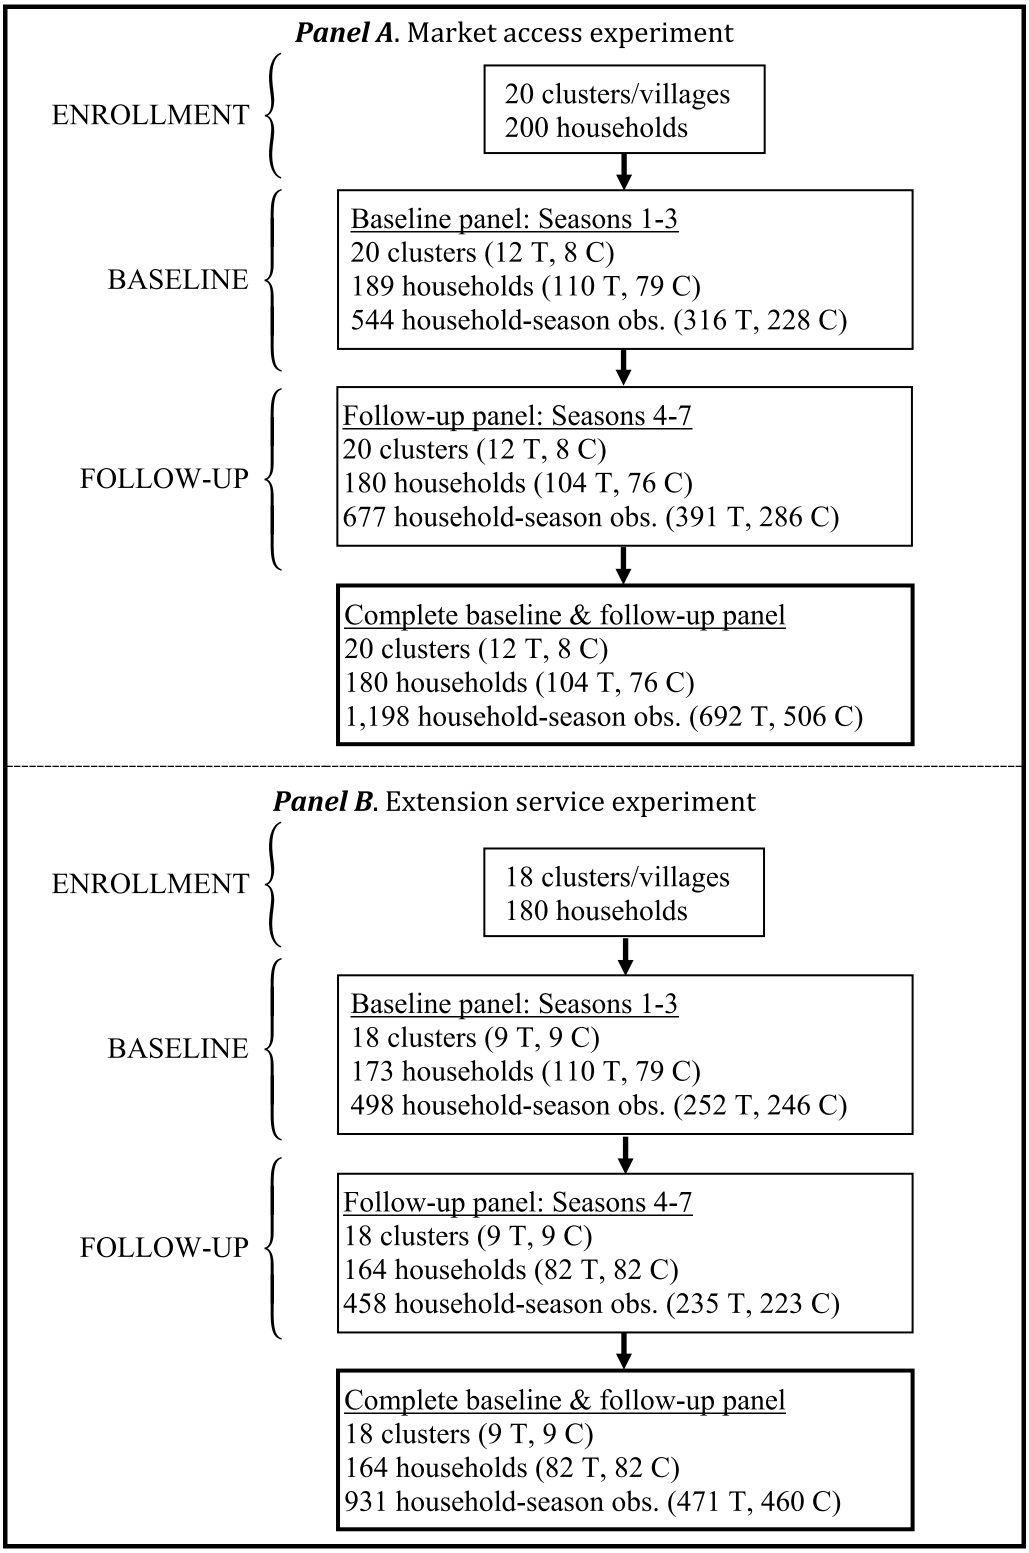

Figure 5. Effect of access to a market for quality maize on profit

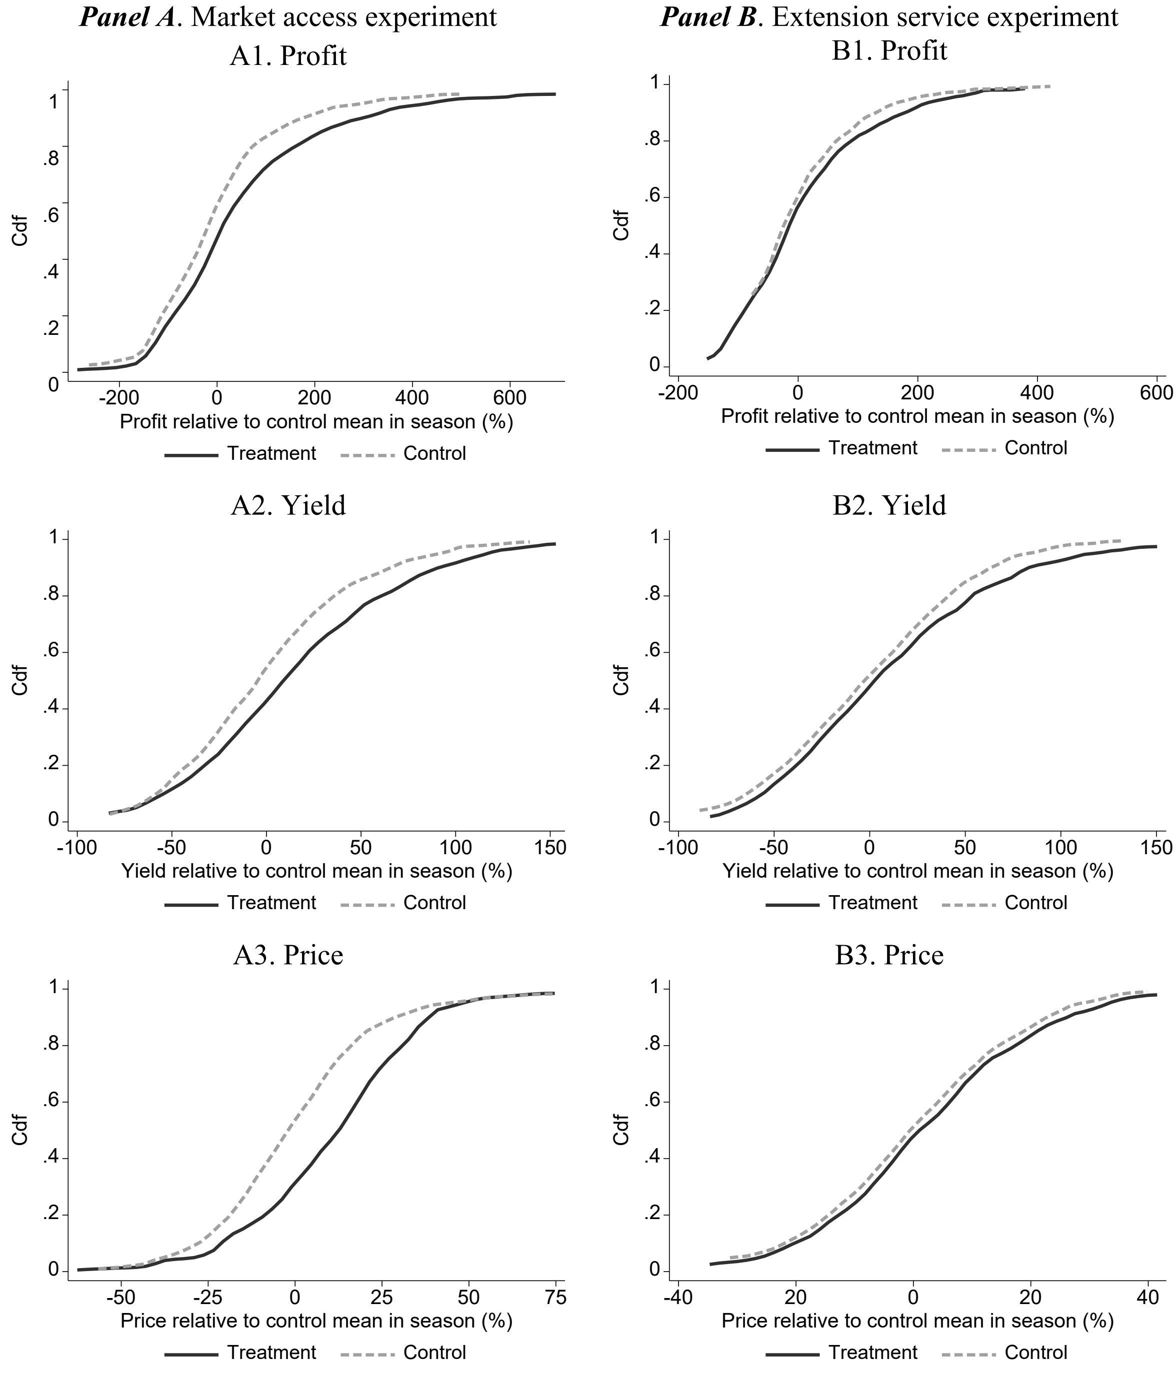

Note. All variables are expressed as difference from the control group mean in percent, by season. In Panel A, The Kolmogorov-Smirnov D statistic is 0.17 [p = .000] for profit, 0.14 [p = .003] for yield, and 0.34 [p = .000] for price. In Panel B, the Kolmogorov-Smirnov D statistic is 0.06 [p = 0.787] for profit, 0.09 [p = 0.314] for yield, and 0.08 [p =0.562] for price.