SEARCH COST, INTERMEDIATION, AND TRADE: EXPERIMENTAL EVIDENCE FROM UGANDAN AGRICULTURAL MARKETS

Lauren Falcao Bergquist¹, Craig McIntosh²

¹University of Michigan and NBER ²University of California, San Diego

July 29, 2021

PRELIMINARY DRAFT

Abstract

High search costs weaken market integration in developing country agricultural markets, harming both farmers and consumers. We present evidence from a large-scale experiment designed to reduce search costs in randomly selected subcounties in Uganda by introducing a mobile phone-based marketplace for agricultural commodities. The intervention drives increases in trade flows and reductions in price divergence across treated markets. Entry by traders into treated markets increases, and profits of incumbents decrease. However, small-scale farmers find it difficult to reach the scale necessary to find buyers on the platform; only the largest farmers use the platform. As a result, we are only able to detect significant increases in revenues among the farmers most likely to use the platform. Point estimates suggest effects that are meaningful in magnitude, but not statistically significant for the majority of farmers. Since farmers are so numerous and the cost per-farmer is low, these income gains per household aggregate to make the intervention strongly cost-beneficial from an overall welfare perspective.

JEL codes: O13, D51, Q12

Keywords: agriculture, market integration, randomized controlled trials

^{*}We gratefully acknowledge funding from USAID/BASIS, USAID/DIL, the Bill & Melinda Gates Foundation and the UK Foreign, Commonwealth & Development Office, awarded through the Agricultural Technology Adoption Initiative (ATAI), the International Growth Center, the Policy Design and Evaluation Lab, and an anonymous donor. We thank IPA, Stephanie Annyas and Laza Razafimbelo for excellent project management. This trial was registered as AEARCTR-0000749, and is covered by UCSD IRB #140744. All errors are our own.

1 Introduction

The integration of agricultural markets in developing economies is an issue of central welfare importance. On the production side, access to deep output markets is critical for farming households, for whom agricultural sales comprise the majority of their income. On the consumption side, well-functioning food markets are necessary to direct food to locations where it is most needed. Trading frictions that limit the movement of crops from relative surplus to relative deficit areas therefore have large welfare costs (Barrett, 2008; Rashid and Minot, 2010).

While transport costs are known to constitute a large fraction of these trading frictions, particularly in sub-Saharan Africa (Teravaninthorn and Raballand, 2009), growing attention has recently been paid to non-transport frictions. One of the most prominent of these is search costs: the frictions that prevent buyers and sellers from easily finding each other in a marketplace (Allen, 2014). These frictions can thwart otherwise profitable trades, leading to lower prices for suppliers and higher prices for consumers. Market-wide, they generate larger patterns of price dispersion across areas of relative surplus and relative deficit (Jensen, 2010). Search costs may also be a source of market power for intermediaries, as information frictions prevent traders from competing across larger geographical areas (Goyal, 2010; Antras and Costinot, 2011).

Against this backdrop, the introduction and rapid spread of mobile phones across sub-Saharan Africa has generated much excitement, offering the promise of dramatically reducing search costs. Indeed, the rollout of cell-phone towers in the early 2000s has been shown to have substantially reduced price dispersion in grain markets (Aker, 2010; Aker and Mbiti, 2010). Building off this success, recent efforts have attempted to move beyond the passive reduction in search costs facilitated by easier bilateral communication via mobile phones, and into more active facilitation of search on mobile platforms design for agriculture.

The first generation of these initiatives focused on the dissemination of price information to farmers via mobile phone. However, price information alone has had mixed results (Camacho and Conover, 2011; Fafchamps and Minten, 2012; Nakasone, 2014; Mitra et al., 2018; Hildebrant et al., 2020). The premise of these price-alert platforms is that farmers will be able to sell in better markets after receiving price information; however, in practice, farmers typically sell at farmgate

or in very local markets, perhaps because of limited access to transport small surpluses in a costeffective manner. This may limit the efficacy of such price alert systems. Moreover, most existing studies are randomized at the individual or village level, while prices are often determined in general equilibrium (GE) across a wider geographic area. GE effects can therefore drive violations of the Stable Unit Treatment Value Assumption (SUTVA), complicating inference.¹

A second generation of search technologies has emerged to offer more comprehensive, mobilebased marketplaces to farmers, intermediaries, and buyers of agricultural goods. These mobile trading platforms serve as clearinghouses, in which those buying and selling agricultural commodities can "match" on their phones. They have the potential to offer two advances over existing price-alert systems. Most directly, they may allow farmers to sell to a wider set of buyers at farmgate, as they no longer need to travel to far-away markets to reach additional buyers. And indirectly, because the platforms are open to traders as well, they may encourage wider movement and greater competition among intermediaries, which could also "trickle up" to benefit farmers, even if farmers do not directly themselves trade on the platform.

This paper presents the results from the first large-scale randomized control trial of such a mobile trading platform, designed to reduce search costs for agricultural commodities in Uganda. At the center of this platform is a novel mobile marketplace for food crops, which links potential buyers and sellers through a simple SMS-based platform. In-village support services are provided by a private-sector Ugandan brokerage firm. Finally, like other information services, the platform also gathers price data and broadcasts it back to farmers and traders using SMS. However, the information is drawn from a large set of national, regional, and local markets, providing a uniquely tailored information set to each user.

The introduction of the platform is randomized across 110 subcounties across Uganda, each of which contains a population of about 30,000 individuals. Our at-scale randomization enables us to measure impacts on local market prices, as well measure the impact on trade flows across treated

¹For example, Svensson and Yanagizawa (2009) find evidence that broadcasting prices via radio lead to higher farmgate prices in Uganda; however, a follow-up paper suggests that once accounting for general equilibrium effects, average farmer revenues impacts are minimal (Svensson and Yanagizawa-Drott, 2012). Hildebrant et al. (2020) also find evidence of spillovers from a price alert system in Ghana. In order to capture GE effects, this study employs "randomization at-scale," randomizing the intervention at the subcounty level. For more detail, see Section 2.

subcounties. To measure these impacts, we gathered data on market-level prices in 236 markets every two weeks for the three years in which the intervention ran. We also collect multiple survey rounds with a representative sample of traders in the study markets to analyze how the intervention drives their trading behavior, prices, and profits. Finally, we collect surveys of farming households to study the impacts of the platform on farmer revenues and welfare.

We find that the search platform increases trade on both the extensive and intensive margins, increasing the probability that treated markets trade with each other and the volumes traded. Prices increase in relative surplus areas that are treated and (weakly) decrease in relative deficit areas that are treated. As a result, price dispersion between treated market decreases. Increases in direct trade connections are concentrated in nearby markets, though price convergence effects may persist across further distances as multiple smaller intermediaries engage in transshipment.

We also find that the platform increases the number of traders who operate between treated markets, as reduced search costs make it easier for traders to find new markets. As a result of this increase in entry, profits of incumbent traders are significantly lower in treated markets. Evidence suggests that trader markups are squeezed, despite volumes increasing. However, pass-through is incomplete. Though traders' sales prices follow the path of market prices, going up in relative surplus areas and down in relative deficit areas, we see limited evidence of similar effects on the price at which they purchase from farmers.

We see that usage of the platform is concentrated among intermediaries; only the largest farmers find it profitable to use the system. These farmers see significant increases in maize revenues and quantities sold. The typical farmer in treated areas, however, only benefits from the general equilibrium effects on prices, and therefore she experiences revenue increases that are imprecisely measured and not statistically significant. Nevertheless, given the large number of people to whom these gains apply in general equilibrium and the relatively low cost per-person of running the platform, we find the platform is still cost-effective.

The rest of the paper is structured as follows: Section 2 discusses the setting and study design. Section 3 discusses the platform's effects on market integration, trade flows, and price convergence. Section 4 and 5 explore impacts on traders and farmers, respectively. Section 6 discusses findings from other sub-experiments that help to shed light on the mechanisms at play and on other related trading frictions. Section 7 explores the business case and the welfare case for the marketplace. Section 8 concludes.

2 Setting and Study Design

2.1 Study Setting

The data collected for the study include market-level outcomes as well as representative samples of traders and farmers. We identified all permanent trading centers (hereafter referred to as markets) within the 11 study districts; these 236 markets are located in 110 subcounties, which were the unit for random assignment of the intervention. Biweekly market surveys, as well as three rounds of trader surveys and two rounds of farmer surveys provide the data on which our analysis is based.

Market prices for maize, our core study crop, show strong variation both over space and over time (see Table A.1).² An East Africa-wide drought saw the price of maize rise from 19 cents per kilo in September of 2016 to almost 44 cents by the following June, and then fall again to 11 cents by the end of the study in September 2018. As a result, time fixed effects account for 83% of the variation in prices. However, we also see strong evidence of meaningful spatial heterogeneity in prices (see Table A.2). A major driver of this price dispersion observed across markets is transportation costs, which in Africa are the highest in the world (Teravaninthorn and Raballand, 2009). Transportation costs cannot, however, explain the full gap in prices observed across markets. Figure 1 presents the gap observed in prices across each pair of markets in our sample (black solid line). The dotted line presents the gap we would expect to see if the only factor driving this dispersion were transport costs, as predicted using self-reported transport costs from our trader surveys.³ While transport

³Traders reported the costs of traveling one-way along each of their five most commonly travelled routes, and

²Maize is the most commonly grown and consumed crop in our study area (and, as we will describe later, was the crop most commonly traded on the platform). Our market survey also follows beans, another non-perishable staple, and two perishable crops, tomatoes and bananas (green bananas are steamed and eaten as the most important staple starch crop in many parts of Uganda). Looking across crops, we see that maize and beans, storable crops with defined growing seasons, display moderate predictable seasonal variation across years (month-of-year R-squared of about .15 for maize and beans). In contrast, bananas and tomatoes do not have well-defined harvest seasons and display no seasonality in prices. Instead, the high transport costs associated with these perishable crops can be seen in the substantial explanatory power of trading center fixed effects. Bananas, being both perishable and heavy, display the strongest spatial variation in prices.

costs explain a majority of the observed price dispersion at longer distances, they do little to explain the substantial price gaps across nearby markets. The intervention studied in this paper seeks to work in the space between the transport-driven dispersion and the actual, much higher price gaps observed across the markets in our study.

Figure 1: **Price Differentials and Transport Costs.** The y-axis presents the absolute difference in prices across each market dyad (pair) in the sample. The solid black line presents the gap observed in prices across each pair of markets in our sample. The dotted line presents the gap we would expect to see if the only factor driving this dispersion were transport costs, as predicted using self-reported transport costs from our trader surveys. To generate this prediction, we asked surveyed traders to report the costs of traveling one-way along each of their five most commonly travelled routes and the vehicle size typically used. From this data, we construct an estimate of per kg transport costs, which we then estimate as a non-parametric function of the km traveled. The gray area represents the portion of price dispersion that cannot be explained by transportation costs.



The average market in our study has 11 traders, with a sharp distinction between "hubs" (the 19 regional or district trading centers, which have an average of 36 traders per market) versus "spokes" (the remaining 213 more rural TCs, with an average of nine). Churn among traders is low, with the median market seeing a little over one new and one exiting trader per year, though a few markets see a large amount of entry, such that the average number of traders per market increases from nine at baseline to over 12 at endline. Traders appear to work with large margins; at baseline traders bought maize at an average of 12.7 cents/kilo and sold at 16.4 cents/kilo, a nominal markup of

the vehicle size typically used. From this data, we construct an estimate of per kg transport costs, which we then estimate as a non-parametric function of the km traveled. This non-parametric estimate is presented in the dotted line in Figure 1.

29%. From baseline monthly revenues of \$2,243, traders report an average monthly profit of \$297.⁴ By comparison, average total monthly household expenditure in our farmer sample is \$65.

Given the complex price dynamics in these markets, traders and farmers operate in an informationhungry market environment. 84% of traders report at baseline that they would expand into new markets if they had the information to do so (placing it second behind credit, and ahead of personal connections, buyer contacts, and risk as a self-reported barrier to expansion). Mobile phone calls are the dominant form of price discovery. The average trader reports attempting 23 purchase transactions a week by phone, of which 16 are successful, but only four attempts (of which three are successful) to make a purchase by visiting a seller without prior information. Traders serve as aggregators and so the overall number of sales is lower than purchases, but sales are similarly dominated by transactions arranged over the phone: six successful phone-initiated sales per week versus one from traveling without prior information. Only 2% of our traders report no purchases initiated by phone calls. At baseline almost a half of traders were using radio broadcasts as a price discovery tool, and a tenth were using any kind of SMS service. Therefore, traders are partially informed, with access to some search technologies, but with demand and scope for additional market information and connections. Farmers are less well-informed than traders. In the endline control sample, only 7% of farmers report discovering prices through radio broadcasts, and 2% through SMS.

2.2 Intervention Design and Randomization

We conducted a cluster-randomized RCT that operated in the field for three years. We began the study selection process by identifying 11 districts of Uganda that our implementing partners selected as promising districts for the platform roll-out. These 11 districts are surplus producers of maize, have strong potential for commercialization, and yet are not immediately proximate to Kampala or the other major trading centers of the country (see Figure B.1 for a map from FEWS-NET of surplus maize areas in Uganda presented alongside a map of the 11 study districts).⁵

 $^{^{4}}$ Consistent with these figures, Bergquist and Dinerstein (2020) estimate that the median trader in their sample in Kenya retains 12% of total revenues in profits.

⁵These districts are Apac, Budaka, Butaleja, Dokolo, Hoima, Iganga, Kamwenge, Kasese, Masaka, Mubende, and Oyam.

We then listed all markets that were permanent (e.g. not meeting only on specific days of the week) and featured both buying and selling of maize (as opposed to wet markets where fruits and vegetables are only sold). This process identified 236 trading centers, hereafter referred to as markets. Markets were classified as "hubs" (major local commercial centers that are centers for aggregation and transshipment) and "spokes" (more remote local markets that typically trade with the outside world only through a hub). See Figure B.2 for a map of the hub-and-spoke structure of the study.

The multi-dimensional intervention used technology to provide farmers and traders with novel ways to obtain prices information and gain access to new markets. The heart of this was a platform called Kudu, developed at Kampala's Makerere University. The goal of Kudu is to provide a new means for producers to reach the market, reducing search costs and potentially circumnavigating intermediaries who may exploit search-based market power to depress farmgate prices. Users can post asks (sale offers) and bids (purchase offers) onto Kudu either using a smartphone or by registering their location and then using a basic feature phone to send messages to the platform via free-form text message or a USSD drop-down menu. A call-center also collects asks and bids by phone. Based on the price, quantity, and location of the buyer and seller, the system then matches supply to demand each day to find the Pareto-optimal set of sellers for each buyer.⁶

To deploy Kudu, we worked with AgriNet, a private sector agribusiness firm, to employ and train 210 Commission Agents (CAs) to serve as the on-the-ground agents of the project, promoting the mobile marketplace. Farmers and CAs could either post to Kudu through an AgriNet agent, or they could engage independently on the platform. In practice, almost all farmers who sold on the platform did so directly, rather than engaging in AgriNet brokered deals. Similarly, CAs, who were recruited from pools of existing traders in the area, operated almost exclusively as independent traders on Kudu. CAs were also not reliable promoters of Kudu, and the project ultimately hired salaried staff, not drawn from the local trader population, to promote the platform.

Once bids and asks were posted to the platform, there were two processes by which buyers

⁶The algorithm is designed for both buyers and sellers to post their "reservation price." However, in practice, qualitative interview suggest that buyers and sellers post prices that reflected strategic price offers, much in the way that one would typically make offers in more traditional, in-person negotiations.

and sellers could be matched. First was the Kudu algorithm that cleared the market each day, attempting to maximize the Pareto surplus from matches by crop, using a penalty function decreasing in the price difference between the bid and ask and increasing in distance.⁷ Second was a hand-matching process conducted by employees who could view a dashboard of the business on the platform and attempt to match trades manually.⁸ The hand-matching process proved dominant in the overall operation of the platform, accounting for 80% of all matches conducted on the platform. Hand-matched trades also had a higher success rate in translating matches into completed transactions, 9.2% versus 1.1% for the algorithm-matched bids and asks.

This core intervention (Kudu and AgriNet CAs) was randomly assigned at the subcounty level. In order to create a 2x2 design at the spoke level (is the spoke treated, is the hub treated) we blocked the design by whether the subcounty contains a hub (17%) or not (83%), and we stratified by a subcounty level price index (mean of the z-scores of the prices of each of our four crops at the trading centers in each subcounty). This generated a design in which half of the hubs are treated and half are not, but with random variation in the fraction of spokes for each hub that are treated. In total, this design results in 55 treated subcounties with 10 treated hubs and 115 treated spokes.

We also set up and ran a system to distribute high-frequency price information to both sides of the market for three years in treatment villages. Our "SMS Blast" system sent out market price information on the four crops study crops every two weeks to treatment traders, CAs, and farmers, as well as to all buyers registered on Kudu, regardless of location. All treatment traders and CAs were included in the Blast system, as well as a randomly selected two-thirds of the treatment farmer households in the study.⁹ Three core types of information were contained in the Blast system. First, a "Downstream Blast" gave each market participant price information for his or her respective local market, hub, and superhub. Second, a "Random Blast" randomly sampled five treatment markets each week and circulated price information on these markets to the entire treatment set of CAs,

⁷For more technical details on the Kudu platform, see Newman et al. (2018).

⁸Hand-matching was conducted by three AgriNet employees. When possible, they were asked to broker commissioned trade for AgriNet, and were explicitly permitted to favor AgriNet CAs and priority buyers to promote the commercial viability of the platform for AgriNet. However, it was often not possible to broker a commissioned trade for AgriNet, and therefore the majority of hand-matched trades were direct exchanges between the Kudu buyer and seller.

⁹This randomization was conducted at the household level, blocked on subcounty.

traders, and buyers. The purpose of this was to give a statistically high-powered estimate of whether prices in a given market change when traders all over the country know about that market in that week. Third, there was promotional information for Kudu; this included an advertisement and information on how to trade on the platform, either by registering directly on Kudu or by contacting their local CA, whose contact information was provided. The Blast system sent more than 25,000 SMS message a month and represents one of the largest experimental efforts to provide market price information; the farmer-level randomization allows us to understand the causal effect of the Blast system on the supply side of the market.

Finally, understanding that contractual risk was likely to be a major barrier to the introduction of impersonal, technologically-mediated trade, we randomly introduced a system of transport guarantees. These guarantees were aimed at mitigating the impact of contractual risk for buyers by compensating them for losses should transactions not be executed as promised on the Kudu platform. There were two levels of guarantee: the "Basic Guarantee," which covered the buyer against any shortfalls in quantity that occur when they arrive in the village to buy, and the "Comprehensive Guarantee," which additionally covered against shortfalls in quality or attempts by the seller to renegotiate price. Guarantees were randomized at two levels; first at the buyer level (some buyers were perpetually offered either the Basic or Comprehensive guarantee for any deals done on Kudu), and then, among those not offered a perpetual guarantee, at the match level.¹⁰

2.3 Survey Sampling and Timeline

The intervention ran for three years, starting in 2015 and concluding in 2018. This time period spans six major agricultural seasons. Figure B.3 presents a timeline for the project, and Figure B.4 provides a CONSORT diagram of study recruitment and attrition for each type of data.

We collect three core types of data for this project, using the 236 markets in our study as the primary sampling units. The first of these datasets is a high-frequency market survey. This survey gathered information in each market every two weeks by calling a key market informant, typically a trader whose store was based in the market. We collected data on the buying and

¹⁰For buyers in the guarantee control group, their matches were randomized into no guarantee, Basic guarantee, or Comprehensive guarantee at the time when he or she was being shown the optimized matches to a given bid.

selling price, availability, and average quality of four major food crops (maize, nambale beans, matooke bananas, and tomatoes). We also surveyed 20 hub markets in adjacent, non-study districts to provide an additional measure of potential spillover effects, as well as in the four 'super-hub' markets of Uganda.¹¹ The total number of markets reporting the biweekly Market Survey is thus 260, of which 236 form the core experimental sample. The market survey was collected for the three years during which the intervention ran.

The second dataset collected is a survey of traders in each study market. We first conducted a census of traders who were based in that market and who bought and sold at least one study crop. For markets that had fewer than 10 traders identified in the census, we surveyed all traders; for markets with more than this, we randomly sampled 10 traders. These traders were administered a baseline survey in 2015, prior to the initiation of any treatment, a midline survey in 2016 after one year of treatment, and an endline in 2018 after three years of treatment. The trader analysis is weighted to make it representative of all traders in study markets.

Finally, to understand the impact of the platform on farmers, we drew in a sample of agricultural households. We first listed all villages located in the subcounty.¹² We then selected the village containing the market (which are typically more urban) and randomly sampled one of the remaining villages within the same parish (which tend to be more rural). For these two villages, we then listed all the households based on administrative records held by the village chairperson, and randomly sampled households from these lists. We randomly sampled 8-9 farming households located within each village containing the market and another 4 in each rural village that does not contain the market. We imposed two eligibility criteria: (i) the household had to be engaged in agriculture, and (ii) the household had to have sold some quantity of any of the four crops included in the study in the previous year. Study households completed a baseline in 2015 and an endline survey in 2018 covering agricultural activities, farmgate prices, and marketed surpluses. Farmer analysis is weighted to make it representative of all farming households in sampled study LC1s.

¹¹These super hubs are the capital, Kampala, plus three border markets that trade grain with neighboring countries: Kabale on the border with Rwanda, Busia on the border with Kenya, and Arua which trades to the DRC and South Sudan.

¹²In Uganda, these villages are called Local Council 1, or LC1s.

2.4 Attrition and Balance

We now present the attrition and balance for each of the three types of data captured in the study: the market surveys, the trader surveys, and the household surveys. For the market survey, we have 88% of the attempted (market x survey round) observations, but 13% of markets in both the treatment and the control groups answer fewer than 75% of the market survey waves they were supposed to.¹³ For the trader midline, we were able to survey 1,358 of the 1,457 baseline traders (93.2%). For both the trader and household endlines, we ran standard panel tracking, and then conducted an intensive tracking exercise that attempted to follow up with a random sample of attritors. The trader endline originally located 1,248 traders (85.7%), after which we randomly sampled 20% of attritors (41 individuals) for intensive tracking, and successfully located 37 of these (92.7%). The weighted tracking rate in the trader endline is therefore 98.6%. The household endline originally located 2,744 of the 2,971 baseline respondents, and we then randomly sampled 17% or 39 households for intensive tracking. 31 of these households were successfully intensively tracked (79.5%), giving us a weighted household tracking rate of 98.7%.

Appendix Figure B.5 and Tables A.3 and A.4 present tests comparing attrition in the treatment to the control across the three data types. Among all the tests that we conduct only the intensive tracking rate in the trader survey appears differential, and given that this arises from finding 14 out of 14 control versus than 24 out of 27 treatment traders in the intensive tracking, this has relatively little influence on study-level effects. Overall, weighted attrition rates are very low and the overall unweighted attrition rate from the combination standard and intensive tracking is similar across treatment arms for all data types.

Table A.5 examines the balance of the market survey for the two main study crops (maize and beans) and the core variables in the market surveys (buying and selling price, number of traders, and quality measured on a 1-3 scale). Table A.6 uses the market survey data in dyadic form and examines the baseline balance of the experiment on price dispersion within dyads. The experiment is well balanced at the market level. For the trader and household analysis, balance is analyzed

¹³As a robustness check for attrition, we present appendix tables that show the main market survey data using interpolation; given the long panel (83 rounds of market surveys) and the highly interspersed nature of the missing observations this provides a reasonable check on the extent to which market survey attrition may influence our results.

using the sample still present at endline and is weighted using the attrition weights so as to mirror the structure of the outcome analysis. Table A.7 analyzes the baseline attributes of traders across seventeen different attributes and finds no evidence of baseline imbalance. Table A.8 conducts the same exercise for households, finding two out of seventeen outcomes significantly different at the 10% level and one at the 5% level, in line with what we would expect by random chance. We therefore proceed to the analysis section with confidence that the study is both representative and well-balanced.

2.5 Platform Usage

Over the three years that the Kudu platform was operational as a part of this project, it received 23,736 unique asks and 30,499 unique bids. Maize accounts for 67% of asks on the platform, though 19 total crops were successfully traded, with the next most common being soya, rice, and beans. Among those posting bids to buy on the Kudu system, 49% were study traders, 10% were AgriNet CAs, and 6% were study farmers. For those posting asks to sell, the corresponding percentages are 37% and 10% for study traders and CAs, and 7% of sellers are study farmers. 80% of treated traders and 26% of treated households posted to the platform at least once. Despite this heavy participation on the platform from study subjects, we still see 45% of bids and 54% of asks emanating from outside the study sample altogether, providing an initial suggestion that the study of new bids and asks posted on the platform per day, with supply climbing steadily through the first year to reach a steady maximum of about 200 tons per day, and demand following a similar time path to reach average levels somewhat less than twice supply.¹⁴ Figure B.7 shows the spatial distribution of asks, indicating study market centers across the country posting upwards of 1,000 asks each.

Subsequent to a Kudu match, the buyer was contacted by SMS and informed that the match had occurred, along with the contact information of the seller. The fully disintermediated version

¹⁴Standing up supply and demand simultaneously was an issue at inception of the project; an initial surge of asks from farmers in the first season overwhelmed demand, but then a drive to recruit buyers on to the platform was highly successful and for the remainder of the project the total demand on the platform exceeded supply.

of trade would then be that the buyer directly contacts the seller and arranges for a sale, which occurred very rarely in our study. More common was that a project employee would hand match a buyer and seller, and then reach out by phone to both to inform them about the match and gauge their interest in the deal. The manual matching process could also deal with failed matches in a flexible way (moving on immediately to the next counterparty), while the Kudu algorithm required them to go back into the matching pool for the next iteration.

Kudu instructs buyers to post their reservation bid prices and sellers to post their reservation ask prices. It was therefore assumed at the launch of the platform that there would be a sizable gap between the two, ideally substantial enough for the platform to broker trades with comfortable margins for all parties, such it could eventually charge commissions in order to make the platform self-sustaining financially. However, in practice, this rarely happened. Qualitative interview suggest that sellers often posted prices that reflected strategic price offers, fishing for higher prices, much in the way that one would typically make offers in more traditional, in-person negotiations. In fact, sellers often posted prices that were not only higher than their local market price, but even in excess of what was being paid in hubs or superhubs. As a result, ask prices were on average substantially above bid prices. ¹⁵ Buyers' average bid prices, on the other hand, track hub market prices very well. Figure B.8 plots these values over time for maize, and Figure B.9 provides a box-and-whisker plot of bid and ask prices within each season, in which we can see that the median bid price is typically at or below the 25th percentile of ask prices

Nonetheless, about 7,300 tons of grain were successfully transacted, worth about \$2.3 million USD. 22% of treated traders and 2% of treated households successfully traded on the platform. Figure B.10 shows the cumulative sales over the platform during the duration of the study.

3 Market Integration

We now turn the impacts of the platform. We first explore the effects on market integration and trade flows. Figure 2 presents impacts of the platform on several outcomes: whether any trade is

¹⁵In recognition of this, Kudu developed a feedback system that sent a message back to unmatched sellers stating "You would have had to ask for X price in order to match on Kudu." However, this failed to align prices, as average ask prices remained above bids the for the duration of the study.

occurring between subcounties, the number of traders engaged in trade between subcounties, the volume of trade flowing between subcounties, and price dispersion between markets. The first three of these outcomes is drawn from our panel survey of traders, in which we asked detailed questions about their trading behavior at the subcounty level (the level of randomization). The last is drawn from our market-level price surveys and is therefore at the market dyad level. Within these dyads, we analyze the experiment using indicator variables for dyads in which both markets are treated and dyads in which one market is treated, using no-treatment dyads as the control. Figure 2 presents Fan regressions of each outcome on the distance between the pair, estimated separately for our three treatment groups. Distance is measured as the road distance of the shortest route connecting the two.



Figure 2: Effects on trade linkages, number of traders, and trade volumes, and price dispersion by distance.

Before examining treatment effects, we first note some important patterns observed among our control-only pairs. In the upper left panel, we see that while the probability of any trade is high for nearby subcounties, this diminishes rapidly with distance. The probability of any trade occurring between the subcounties is close to zero beyond 200km distance. Consistent with this, the number of traders (upper right panel) and total trade volumes (lower left panel) also falls quickly with distance.¹⁶ These increasing trade costs with distance lead to notably higher price dispersion between markets located at further distances, as shown in the bottom right panel.

What is the effect of introducing a mobile clearinghouse? In Figure 2, we see increases in the probability of any trade occurring, the number of traders engaged in trade between subcounties, and the volume of trade flowing between subcounties. We also see a drop in price dispersion. These effects are more pronounced when both markets are treated than when just one market is treated.¹⁷

Table 1 presents these results in regression form, running the following specification:

$$Y_{dr} = \alpha + \beta_1 T 1_d + \beta_2 T 2_d + \beta_3 D_d + \varepsilon_{dr} \tag{1}$$

Here, Y_{dr} is the outcome of interest in subcounty or market dyad d in round r, pooling all post-treatment survey rounds in the same analysis. For the first three columns, the outcome of interest is whether any trade is reported, the number of traders trading, and the volume of trade flowing between subcounty dyads, as reported by traders in the traders midline and endline. In the fourth column, the outcome is the inverse hyperbolic sine transformation of the absolute value of the gap between prices across each possible market dyad (d) in round (r), which is two-week intervals as measured in the market survey. These outcomes are regressed on $T1_d$, a dummy for one of the markets being located in a treated subcounty, $T2_d$, a dummy for both being in a treated subcounty, and D_d , a measure of the shortest road distance between dyads. Standard errors are

¹⁶It is important to note that our study traders are those who *live* in study markets, and hence are likely involved in more localized forms of trade. It is possible that there is more long-distance trade directly connecting our study markets than we find among our traders, but it is being conducted by large-scale traders who live in major cities and hence were not sampled into our survey. Further, multiple short-distance traders engaged in sequential transshipment can still influence price dynamics over longer distances.

¹⁷Recall randomization occurs at the subcounty level, which is large unit than the market. We therefore cluster standard errors by subcounty.

clustered two-way by each subcounty (the unit of randomization).¹⁸

	Any Trade	Number Traders	Volume (tons)	Price Dispersion
One treated	$0.00 \\ (0.01)$	0.07^{*} (0.04)	2.44 (1.75)	-0.05^{*} (0.03)
Both treated	$0.02 \\ (0.01)$	$0.08 \\ (0.07)$	4.22^{*} (2.37)	-0.09 (0.07)
Dist (10km)	-0.00^{***} (0.00)	-0.02^{***} (0.00)	-0.53^{***} (0.13)	0.01^{***} (0.00)
Observations Mean DV	$\begin{array}{c} 11664 \\ 0.05 \end{array}$	$\begin{array}{c} 11236\\ 0.22 \end{array}$	$\begin{array}{c} 11664 \\ 4.90 \end{array}$	$\frac{1443397}{4.56}$

Table 1: Effects on trade linkages, number of traders, and trade volumes, and price dispersion

Again, we see increases in the probability of trade occurring (albeit not significantly so), increases in the number of traders operating between subcounties (significant for one-treated; not quite significant for both treated); increases in trade volumes (significant for both treated; not significant for one-treated), and reductions in price dispersion (significant for one-treated; not significant for both treated).¹⁹

Returning to Figure 2, we also note a striking pattern by distance. Treatment effects are strongly concentrated among nearby markets. In fact, we see almost no treatment effect beyond 200km, the point at which the probability of trade drops close to zero. The one exception to this is price dispersion, which continues to drop in our treated market pairs beyond 200km, albeit at a slower rate. This may be due either to treatment effects on large, long-distance traders not included in our sample as resident in study markets. Alternatively, we may observe transitive convergence (in which the platform connects Market A to Market C and Market B to Market C, such that Markets A and B see price convergence, despite no goods flowing between their markets) or to transshipment (in which the platform connects Market A to Market A to Market B and, subsequently, Market B to Market C. such that goods eventually flow from Market A to Market C, despite not having a

¹⁸Dyads in the same subcounty are dropped both here and in Figure 2, as they mechanically have the same treatment status.

¹⁹Note the number of pairs in which both members of the pair are treated is lower than those in which one is treated, which may explain the difference in power between the estimated treatment effects for each group.

direct connection).

Table 2 explores this further, estimating Equation 1 separately for market pairs above and below the median distance observed in our sample (242 km). We again see that treatment effects for all outcomes are concentrated in nearby subcounties and markets, with significant increases in the probability of trade occurring (among both treated), significant increases in the number of traders operating between subcounties, significant increases in trade volumes (among both treated), and significant reductions in price dispersion. In contrast, we see almost no effect of the platform on direct trading outcomes (extensive margin connections, number of traders, and trade volumes) for markets that are at above median distance, though the point estimates suggest that price convergence results may persist along a farther distance, albeit at a smaller rate.

Table 2: Effects on trade linkages, number of traders, and trade volumes, and price dispersion by distance

	Below Mean				Above Mean				
	Any Trade	Num Traders	Vol (tons)	Price Disp	Any Trade	Num Traders	Vol (tons)	Price Disp	
One treated	$0.01 \\ (0.01)$	0.18^{**} (0.09)	5.78 (4.28)	-0.08^{**} (0.04)	-0.00 (0.00)	-0.00 (0.00)	$0.04 \\ (0.14)$	-0.03 (0.03)	
Both treated	0.05^{**} (0.02)	0.23^{*} (0.13)	11.14^{*} (5.75)	-0.15^{*} (0.09)	$0.00 \\ (0.00)$	$0.00 \\ (0.01)$	-0.03 (0.12)	-0.05 (0.08)	
Dist (10km)	-0.02^{***} (0.00)	-0.10^{***} (0.01)	-2.26^{***} (0.57)	0.02^{***} (0.00)	-0.00^{*} (0.00)	-0.00^{*} (0.00)	-0.01^{*} (0.00)	0.01^{***} (0.00)	
Observations Mean DV	$4443 \\ 0.13$	$\begin{array}{c} 4270\\ 0.56\end{array}$	$4443 \\ 12.58$	$559270 \\ 4.31$	$\begin{array}{c} 7221 \\ 0.01 \end{array}$	$6966 \\ 0.01$	$7221 \\ 0.17$	884127 4.71	

The platform is therefore quite successful in generating additional short-distance trade. However, it falls to live up to the often-touted promise of such online marketplaces to *directly* connect remotely-located farmers and markets with urban consumers (though it is possible that it may indirectly achieve this goal through transshipment).

While perhaps initially surprising, this pattern of large effects over the shortest distances is consistent with Figure 1, which shows that very little of the existing price dispersion at short distances is explained by transport costs. Given the ubiquity of mobile phones even in the control group, it is likely that the very large price gaps necessary to motivate long-distance trade are already arbitraged away, meaning that the more marginal improvements in information revealed

Figure 3: Effects on price levels by relative surplus vs. deficit areas. The left panel shows the level of prices in treatment vs. control treatment centers, with respect to the average market surplus per farmer, as measure in tonnes at baseline. The right panel shows the difference between the two (the treatment effect), along with the 90% and 95% confidence intervals from a bootstrap estimation.



by our system typically only exceed the pecuniary costs of trade over shorter distances.

3.1 Unpacking price convergence

What is driving the observed price convergence? Figure 3 presents treatment effects on price levels in relative surplus vs. relative deficit areas, as measured by average marketed surplus per farmer at baseline. First, we note in the lefthand panel that, as expected, prices in the control group are higher in relative deficit areas and lower in relative surplus areas. However, we see a less steep relationship in the treatment group, as treatment lowers prices in relative deficit area and raises prices in relative surplus areas. The righthand panel presents this treatment effect, along with the 90% and 95% confidence intervals. We see that prices are weakly lower in deficit areas and statistically significantly higher in surplus areas.

Table 3 presents similar results in regression form. We see in Column 1 that the overall effect on price levels is a statistical zero. This is consistent with the netting out of two competing effects seen in the previous figure (the density in the right-hand panel of Figure 3 shows that for the median trading center, the average price effect is roughly zero). Column 2 presents heterogeneity by baseline average marketed surplus. We again see that prices are weakly lower in relative deficit areas (as evidenced by the negative coefficient on the treatment term) and higher in relative surplus areas (as evidenced by the significant and positive coefficient on the interaction term). With an average baseline marketed surplus of about 1 tonne, these effects almost exactly offset each other for the median market. Column 3 divides our sample into areas of relative surplus and deficit, with the cutoff defined by being above or below 1 tonne.²⁰. First, we note that prices are significantly lower in surplus areas overall, as expected. Second, we see that, with the introduction of Kudu, these surplus area experience significantly higher prices than they would have otherwise, while deficit areas experience weakly lower prices.

	(1)	(2)	(3)	(4)
Treatment	-4.801	-20.78	-11.58	
	(11.56)	(16.40)	(13.74)	
Baseline marketed surplus		-28.66^{***}		
		(7.716)		
Treat*Baseline marketed surplus		19.23^{*}		
		(9.995)		
Surplus dummy			-48.60***	-47.88***
			(9.868)	(9.834)
Treat*Surplus dummy			29.39^{*}	17.81^{*}
1			(15.10)	(9.369)
Treat*Deficit dummy				-10.26
v				(13.62)
Observations	15211	15161	15211	15211
Mean DV	831.6	831.6	831.6	831.6
Mean Baseline Marketed Surplus		0.907		
Percent Surplus			0.275	0.275
P-Val Treat*Deficit=Treat*Surplus				0.0640
R2	0.844	0.847	0.847	0.847

Table 3: Effects on price levels by relative surplus vs. deficit areas

In terms of impacts on average outcomes in study markets, our intervention has no detectable effects. Buying and selling prices, *within market* margins (defined as the buy minus sell price within

²⁰In addition to being roughly the mean surplus level, this is also the empirically-driven definition of surplus vs. deficit, as this is where treatment effects cross zero in Figure 3, suggesting that areas below this cutoff saw inflows and higher prices, while areas above saw outflows and lower prices.

the same market), and the number of traders who are permanently based in treatment markets all remain comparable to control markets. These results are presented in Table A.9. Hence, while the intervention had meaningful effects on price dispersion among nearby market dyads, it led to no average shift in market-level outcomes.

The fact that we see increased trade flows and reductions in one-treated-market dyads suggest that there are positive spillovers from treatment.²¹ What drives these increases in trade between one-treated-market dyads? One possibility is spillover of the technology use itself. We do see some use of Kudu within control communities; however, rates are fairly low (e.g. only 16% of control traders and 1% of control farmers use Kudu, compared to 80% and 24% of their respective counterparts in treated communities). Given the magnitude of these adoption rates, spillover of the technology itself is likely to low to explain the full positive spillover effects we see.²². A second possibility is transshipment. Consider two traded markets, Markets A and B. A trader based in Market B may, as a result of treatment, start buying maize in Market A. If that trader usually sells from Market B to Market C, a control market, he may as a result of treatment ship more to Market C, because he has been able to buy more from Market A. Therefore, trade between Markets B and C will go up, even though only one market is treated. A third option is a shopping trip model. Consider again two treated markets, Market A and B. A trader in Market B may start buying in Market A. If control Market C is on the way to treated Market A, he may decide to buy from Market C as well. Again, this would generate an increase in trade between Market B and C, even though only one market is treated.²³ We currently cannot distinguish between the second two explanations, but in future work we plan to bring in data on baseline trading networks and the

 $^{^{21}}$ There are alternatively in which spillovers could be negative: for example, if the intervention spurred a diversion of trade from control markets into treated markets. This is a possibility that we tackle explicitly in Section 6.3. However, the fact that we see an *increase* in trade between one-treated-market dyads is prima facie evidence against this concern. If there were diversion, we would expect a decrease in one-treated-market dyad trade, as a given treated market shifts from trading with a given control market into trading with treated markets. We further explore the impact of treatment on traders' trade volumes with markets outside of our sample. If diversion is at play, we should see negative effects here. Appendix Table A.10 presents these results. If anything, we see that trade volumes with out of sample markets *increase* for treated traders.

 $^{^{22}}$ Even at the upper end, assuming take-up rates in control communities that are at 20% of what we see in treated communities, this cannot explain the full magnitude of the one-treated dyad effects, which are typically about 50% of that seen in both-treated dyads

 $^{^{23}}$ A related possibility is that due to spatial correlation in prices, sending information about Market A may also provide some information about nearby Market C, such that a trader from Market B decides to buy in both Markets A and C.

geography of markets to distinguish these stories. Section 6.3 uses heterogeneity in the intensity of baseline trading networks to study spillover effects among traders.

4 Trader Effects

We now turn to our trader surveys to unpack how the platform affected intermediaries. We have already seen in the previous section that the platform encouraged greater intermediary activity in treated subcounties and markets. Here, we explore in greater detail trader take-up of the platform and effects on their businesses.

Table 4 presents trader take-up results. We see that, by the endline survey, 91% of treated traders report having heard of Kudu, while only 32% of control traders have heard of the platform. Therefore, while Kudu was not restricted to be operational only in treated areas, we do see a significant and large difference in awareness of the platform generated by our encouragement design. We also observe a 42 percentage point increase in the likelihood of receiving any price information via SMS (from any source).

However, in terms of knowledge of prices, we do not see a substantial treatment effect. We ask traders to report their best guess of the current market price in their local market, their hub market, and their superhub market, which we then compare the the actual price as measured by our market surveys. We call the absolute value of the gap the "error" in price knowledge. Although traders' knowledge of nearby local and hub markets is slightly better than their knowledge of superhub prices, we see no differences between treatment and control traders in knowledge for any market type. This may be because knowledge in our control is already quite high, as demonstrated by the relatively small error size.

We do, however, see strong treatment effects in terms of self-reported impacts on negotiations, both with farmers from whom traders buy and with buyers to whom they sell. Treated traders are more likely to report that they are aware of farmers and buyers receiving price information via SMS. They are also more likely to state that this information changed how they negotiated with their trading partner.

Finally, in terms of Kudu take-up, 80% of treated traders used Kudu (meaning they placed an

ask or a bid), while 22% successfully completed a deal on the platform. In comparison, only 12% of control traders tried Kudu, and only 3% successfully completed a transaction.

Treat	Control	Obs	Т	-C
			$di\!f\!f$	p- val
0.91	0.32	$1,\!281$	0.59	0.00
0.86	0.44	$1,\!281$	0.42	0.00
90	82	$1,\!277$	8	0.39
83	90	$1,\!270$	-7	0.47
117	125	$1,\!248$	-9	0.96
0.26	0.08	$1,\!281$	0.18	0.00
0.42	0.13	$1,\!281$	0.28	0.00
0.16	0.05	$1,\!281$	0.10	0.00
0.22	0.09	$1,\!281$	0.13	0.00
0.80	0.16	$1,\!457$	0.65	0.00
0.22	0.03	$1,\!457$	0.19	0.00
	Treat 0.91 0.86 90 83 117 0.26 0.42 0.16 0.22 0.80 0.22	Treat Control 0.91 0.32 0.86 0.44 90 82 83 90 117 125 0.26 0.08 0.42 0.13 0.16 0.05 0.22 0.09 0.80 0.16 0.22 0.03	$\begin{array}{c cccc} {\rm Treat} & {\rm Control} & {\rm Obs} \\ \hline 0.91 & 0.32 & 1,281 \\ 0.86 & 0.44 & 1,281 \\ 90 & 82 & 1,277 \\ 83 & 90 & 1,270 \\ 117 & 125 & 1,248 \\ 0.26 & 0.08 & 1,281 \\ 0.42 & 0.13 & 1,281 \\ 0.42 & 0.13 & 1,281 \\ 0.16 & 0.05 & 1,281 \\ 0.22 & 0.09 & 1,281 \\ 0.80 & 0.16 & 1,457 \\ 0.22 & 0.03 & 1,457 \\ \end{array}$	$\begin{array}{cccccccccccccccccccccccccccccccccccc$

Table 4: Trader take-up

Table 5 presents effects on trader profits, volumes traded, markups, and prices.²⁴ The main specification pools the post-treatment survey rounds and runs:

$$Y_{ir} = \alpha + \beta Treat_i + \gamma Y_{i0} + \delta_r + X_i + \varepsilon_{dr}$$
⁽²⁾

In which Y_{it} is outcome Y for trader *i* in round *r* (either midline or endline), $Treat_i$ is a dummy for being a treated trader, Y_{i0} is the baseline level of the outcome variable, δ_r is a dummy for survey round, and X_i is a vector of controls.²⁵ Treatment effects are given by the coefficient β .

By reducing search costs and encouraging traders to enter into new markets, mobile marketplaces like Kudu have often been promoted as fostering greater competition among intermediaries.

 $^{^{24}}$ Since 92% of sample is comprised of maize traders, for variables for which we must specify the crop – i.e. volumes, markups, and prices – we present result for maize.

²⁵Our pre-analysis plan specified that we would include baseline controls that are most predictive of the outcome. We do this by identifying controls to include via a double lasso procedure, set to predict endline profits, our main trader outcome. Those covariates considered were: gender, age, education, length of time in business, number of subcounties in which purchase, number of subcounties in which sell, profits, net revenues, annual costs, annual revenue, monthly costs, and markups, quantities purchased and sold, prices at which purchased and sold, revenue, net revenue, and cost per kg for maize and beans, all as measured at baseline. Those selected by the lasso procedures and therefore included in X_i are: baseline profits, baseline annual costs, and baseline monthly costs.

Indeed, we do see that traders located in treated subcounties see a significant reduction in profits, by about 14% of their average value (Column 1). This appears to come mainly from a reduction in trader markups (the difference between the price at which a specific trader buys and sells); point estimates suggest a reduction of about 8%, though this effect is measured with imprecision and is not significant (Column 3). Volumes traded appear to increase, perhaps sizably, though again, this point estimate is not significant (Column 2).

Columns 4-5 present effects on the price at which traders sell maize, while Columns 6-7 present treatment effects on the price at which traders purchase maize. Similar to results presented in Table 3, we see no significant effects on the level of prices (Columns 4 and 6). However, looking at heterogeneity by relative deficit and surplus areas (as proxied by baseline marketed surplus), we see that in relative deficit areas (those with low baseline marketed surplus, where prices tend to be high), treatment results in trader sale prices that are significantly lower (Column 5). Conversely, in areas that are relative deficit (those with high baseline marketed surplus, where prices tend to be low), we observe that treated traders sell at higher prices. We see similar, albeit slightly muted, effects for the trader purchase price in Column 7. We will return later to discuss the relative magnitudes of the sale vs. purchase price treatment effects.

	Profits ('000)	Tons Traded	Markups	Sell	Sell Price		Price
Treat	-1025.1^{*}	133.2	-10.6	-10.4	-25.5^{*}	-3.4	-14.7
	(555.5)	(106.3)	(9.3)	(10.0)	(13.9)	(8.4)	(11.7)
Baseline marketed surplus					-18.0^{**}		0.6
					(7.0)		(6.4)
Treat*Baseline marketed surplus					15.0^{*}		10.5
					(8.9)		(6.8)
Observations	2592	2370	2268	2295	2295	2282	2282
Mean DV	7279	243	135	735	735	602	602
Mean Baseline Marketed Surplus					1.00		1.00
R2	0.12	0.04	0.06	0.06	0.07	0.02	0.03
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 5. Ffort	a on trad	lon profita	volumos	monlauna	and	nnioog
Table 5: Effect	s on trac	ter proms.	, volumes,	markups	, and	prices

Finally, we explore treatment effects based on baseline heterogeneity. Figure 4 presents treatment effects on profits, markups, and trade volumes based on their baseline levels (as measured in the baseline survey). The top panels plot Fan regression estimates of the outcome at endline on baseline levels separately by treatment and control, while the bottom panel presents the difference (i.e. the treatment effect), along with the 90% and 95% confidence intervals. Density in the baseline measure is presented in red.

While the negative effects on profits and positive effects on volumes traded appear fairly consistent across their baseline distribution, we do interestingly see that markups are higher among treated traders at the low end of the baseline markup distribution and lower among treated traders at the high end of the baseline markup distribution. These estimates therefore suggest that the introduction of the mobile marketplaces appears to lead to convergence in markups, helping low markup traders and harming high markup traders. This is consistent with the idea that trading in the absence of our intervention is skill- and human capital-intensive, in which case those with stronger trading networks and superior information can reliably reap greater profits. Our intervention, by reducing the cost of information to a symmetric low level appears to have removed a substantial portion of the heterogeneity in trader markups.



Figure 4: Heterogeneous effects on trader profits, markups, and volumes by baseline levels

5 Farmer Effects

We have seen thus far that the introduction of a mobile clearinghouse platform induces greater market integration and lowers intermediaries' profits. These results are often seen as stepping stones along a causal chain ending in the ultimate goal of improving the welfare of smallholder farmers. We turn now to effects of the platform on farmers.

First, we explore measures of awareness and take-up of the platform among farmers. Table 6 presents these results. We see that 52% percent of farmers in treated subcounties have heard of Kudu, compared to only 12% in control subcounties. Similarly, 55% of treated farmers have received price information via some form of an SMS-based platform, compared to only 16% of control farmers.

Next, we explore impacts on price knowledge. We first note that overall price knowledge is lower among farmers than among traders, with error rates that are 42-84% higher than observed among traders. This is consistent with the presence of information asymmetries between farmers and traders. Similar to those of traders, farmers' error rates grow with distance. In contrast to traders, however, we do find some suggestive evidence that information to farmers improves their knowledge of prices; errors rates are smaller among treated farmers than among control farmers, albeit not quite significantly so (with p-values ranging from 0.12 to 0.16). 23% of farmers in treated subcounties report using price information received via SMS when negotiating prices in the past year, compared to only 1% in control subcounties.

Turning to take-up of Kudu itself, we see that 26% of treated farmers have ever used Kudu (meaning placing an ask or bid), compared to 2% of control farmers. However, the success rate for study farmers managing to transact on Kudu is relatively low; only 2% of treated farmers completed a transaction on Kudu. Though this rate is significantly higher than the 0% observed in the control group, this low adoption of Kudu is notable, and suggests that farmers either see low benefits or high barriers to adoption of the platform. We will explore determinants of adoption below.

	Treat	Control	Obs	Т	-C
				diff	p- val
Heard of Kudu	0.52	0.12	2,707	0.40	0.00
Received SMS price info	0.55	0.16	2,774	0.39	0.00
Ever used Kudu	0.26	0.02	2,775	0.24	0.00
Completed deal on Kudu	0.02	0.00	2,775	0.02	0.00
Local market price (abs error)	108	117	2,761	-8	0.12
Hub market price (abs error)	122	125	$2,\!671$	-3	0.15
Superhub price (abs error)	214	230	$2,\!491$	-17	0.16
Used SMS price info when negotiating price	0.23	0.08	2,774	0.15	0.00

Table 6: Farmer take-up

We turn now to effects on farmer revenue, volumes sold, and prices. We run the following specification in Columns 1, 3, 5, and 7:

$$Y_i = \alpha + \beta Treat_i + \gamma Y_{i0} + X_i + \varepsilon_i \tag{3}$$

In which Y_i is outcome Y for farmer *i* at endline, $Treat_i$ is a dummy for being a treated farmer, Y_{i0} is the baseline level of the outcome variable, and X_i is a vector of controls.²⁶ Treatment effects are given by the coefficient β .

Table 7 presents results. We see no statistically significant effect on total revenues, maize revenues, maize volumes sold, or price received for maize sold. Point estimates are positive and, in some cases, quite large (for example, the point estimate on total revenues is 9.7% of average revenue), but estimates are imprecise.

Our pre-analysis plan specified that we would analyze heterogeneity in treatment effects by propensity to use Kudu. To do so, we first model the likelihood of adopting Kudu as a probit function of the controls used in Equation 3, estimated for the treatment group only. We then

²⁶Our pre-analysis plan specified that we would include baseline controls that are most predictive of the outcome. We do this by identifying controls to include via a double lasso procedure, set to predict endline total revenue, our main farmer outcome. Those covariates considered were: gender, age, high level of educational attainment, number of household members, revenues, quantity sold, land holdings size, quantity harvested, number of times told, any sales at the market, percent sold to market, distance to market, distance to Kampala, total value of all assets, expenditures in the last 30 days, food expenditure in the past 30 days, value of inputs used in the last year, all as measured at baseline. Those selected by the lasso procedures and therefore included in X_i are: baseline revenues, baseline quantity sold, and baseline value of inputs used in the last year.

predict likelihood of adoption for the entire sample, based on these controls (this prediction is called $PropensityScore_i$ below). Finally, we run the following specification:

$$Y_i = \alpha + \beta_1 Treat_i + \beta_2 Propensity Score_i + \beta_3 Treat_i * Propensity Score_i + \gamma Y_{i0} + \varepsilon_{dr}$$
(4)

Columns 2, 4, 6, and 8 of Table 7 present results. The coefficient on β_2 helps to characterize likely adopters. Who is likely to adopt Kudu depends on who faces the largest benefits to adoption, relative to the cost. One might predict that smaller, poorer farmers who currently receive low prices would have the most to gain from access to a new platform on which to sell their crops. However, one could also imagine that the platform would have a hard time finding a buyer for farmers who sell relatively small surpluses, and therefore that it may be the larger, wealthier farmers who are best positioned to use Kudu.

Consistent with Allen (2014), our evidence suggests the latter interpretation. We see that takeup is significantly higher among farmers with higher total revenues and higher prices received, as evidenced by the statistically significant correlation between the propensity score and these outcomes. Maize revenues and quantity sold are also positively correlated with propensity to adopt Kudu, though not significantly so. It seems, therefore, that larger, better off farmers are more likely to use Kudu.

Looking at the interaction term between treatment and the propensity score, we see the that positive revenue effects are concentrated among likely adopters, significantly so when we look specifically at maize revenues (Column 4). Likely adopters also see significantly larger maize quantities sold (Column 6). We see no significant difference in price for likely adopters, perhaps because prices are determined in general equilibrium in the village, and are therefore not distinct between adopters and non-adopters.

	Revenue	s, Total ('000)	Revenue	s, Maize ('000)	Qnt Sol	d, Maize	Price S	old, Maize
Treat	99.2	54.8	72.0	65.9	61.3	37.1	18.0	19.5
	(91.2)	(95.0)	(68.4)	(73.7)	(118.2)	(122.1)	(14.2)	(14.1)
Propensity		1131.0*		272.9		354.1		136.4^{*}
		(585.2)		(401.9)		(682.2)		(74.7)
Treat*Propensity		430.4		1235.7**		1578.6^{*}		-136.7
		(771.4)		(539.8)		(854.1)		(98.9)
Observations	2775	2745	2775	2745	2769	2739	1959	1941
Mean DV	1019	1019	672	672	1040	1040	631	631
R2	0.32	0.30	0.31	0.29	0.33	0.32	0.02	0.02
Controls	Yes	No	Yes	No	Yes	No	Yes	No

Table 7: Effects on farmer revenues, volumes sold, and prices

Finally, Figure 5 presents heterogeneity in treatment effects based on baseline levels of the outcome, as measure in the baseline survey. Similar to Figure 4, the top panels plot Fan regression estimates of the outcome at endline on baseline levels separately by treatment and control, while the bottom panel presents the difference, along with the 90% and 95% confidence intervals. Density in the baseline measure is again presented in red.

Consistent with the results from Table 7, we see null effects for the vast majority of farmers, who tend to lie to the left-hand side of the distribution in total revenues, maize revenues, and quantifies sold. This is unsurprising given the low adoption rates among these farmers. However, for the minority of farmers who lie to the right-hand side of the distribution – who are more likely to adopt Kudu – we see significant increases in total revenues and maize revenues specifically. Though not significant, point estimates on quantity sold are also large and positive for farmers who already sold large volumes at baseline. Perhaps surprisingly, we see a negative and significant effect on prices for those farmers in the far-right side of the distribution of baseline prices. This may reflect some of the broader convergence in prices observed upon introduction of this marketplace. Consistent with this, the top right-most panel suggests that prices are higher for those with initially low prices (albeit not significantly so) and lower for those with initially higher prices.



Figure 5: Heterogeneous effects on farmer revenues, volumes sold, and prices by baseline levels

In sum, we see disappointingly little in the way of benefits from use of the platform for the average farmer. In particular, we find little evidence to support the hypothesis that mobile marketplaces like Kudu will enable the smallest, poorest farmers to access a wider market and earn higher revenues. Though we do see some evidence of benefits for those who adopt Kudu, these tend to be the larger farmers who are better off at baseline.

6 Unpacking the Intervention

6.1 Isolating the effect of information

We performed a number of sub-experiments to isolate components of the overall intervention. First, we randomized at baseline whether study farmer households in the treatment would receive the SMS Blast (three quarters) or not (one quarter). This comparison varies the extent to which farming households were receiving biweekly price information and promotions of Kudu to their phones, allowing us to understand the additional effect of direct connection to the mobile component of the project. The control group for this sub-experiment are farmers in treatment villages who may already be experiencing general equilibrium changes from prices as well as indirect information effects from the treated farmers around them.

Second, to exploit the power from the panel nature of the market survey, we then further randomly rolled tranches of control markets into the SMS Blast treatment. In each of the 12 market survey rounds between October 21, 2016 and March 24, 2017 we rolled in three control markets to the SMS Blast, treating all study traders and farmers with the Blast. Then, subsequent to the household and trader endline surveys, we rolled in an additional 36 control trading markets to the SMS Blast and so observe four final rounds of market surveys with this system in place.

We begin by investigating the nature of the farmer-level SMS Blast experiment by looking at the extent to which farmer-level uptake of the components of the experiment differs between the SMS Blast treatment and the within-village controls. Table A.11 shows that the treatment-village control group has uptake rates that are in general about three-quarters of the way between the pure control group and the SMS Blast treatment (for example, the fraction of pure control farmers using Kudu was 2%, in the SMS Blast control it was 21%, and in the SMS Blast treatment it was 28%). This suggests that there was fairly widespread dissemination of information about the intervention in treatment villages. Given this, it is perhaps not surprising that when we examine the first row of Table A.12 there are few strong differences in outcomes between the SMS Blast treatment and control. While the point estimates generally suggest a positive effect of being treated, only for the variable 'know better' (indicating that farmers reported they were able to use price information to drive a better bargain) do we see significant differences.

The third row of this table provides a test of the roll-in to control villages, including a dummy for those farmers that as of the endline survey had already started to receive the Blast. Here, we see much smaller point estimates and no evidence that simply including control farmers in the Blast had any influence on outcomes at endline. This impression is confirmed by lookiing at Tables A.13 (using monadic market survey data to study the roll-in) and Table A.14 (using dyadic data to look at the effect of the roll-in on price dispersion). To exploit the panel nature of this treatment, we analyze these impacts using market and round-level data with two-way fixed effects and standard errors clustered at the subcounty level. In all cases effects on market-level outcomes appear very limited. The general lack of significance is confirmed visually in Figure B.11.²⁷ Overall, then, we conform with the broader literature in finding no large improvements stemming from information-only market price interventions, whether these are implemented individually within villages otherwise generally treated, or rolled in over time to untreated markets.

6.2 Transport Risk

We have a number of windows into the effects of the transport guarantees, because they were randomized both at the individual and at the contract (match) level. First, we can look at the cross-buyer experiment, asking whether those buyers who were permanently assigned to receive the Basic or Comprehensive guarantee transact more business on the platform. Table A.15 shows that they do not. Next, within the original control group who were assigned no permanent guarantee, we can examine the effect of the random fraction of bids they post assigned to each guarantee group on the overall amount of business conducted. Table A.16 shows that increasing the fraction guaranteed does not increase business transacted. Finally, at the bid level we can ask whether having a specific bid assigned to a guaranteed increases the chance of doing business, both among the original experimental control group as well as among other buyers entering the platform subsequent to the experiment. Table A.17 examines the bid-level data and shows that for the original control, having a specific bid guaranteed increases the number of successful deals, the amount transacted, and the value transacted.

Unfortunately, this pattern of results appears most consistent with a general lack of significance of the guarantees at generating new business, with the control group having come to understand the system well enough to game it (meaning that they re-posted bids until they were randomly assigned insurance, and then transacted only on the covered bid). While we certainly do not take these results as suggesting that transport and contractual risk are unimportant in anonymous technological marketplaces, it does not appear that these guarantees, backed by AgriNet, were effective at removing them.

²⁷This figure plots average maize buying prices across time, breaking the roll-in into five aggregated tranches that enter the treatment from the pooled control counterfactual.

6.3 Spillover Effects

We have provided a number of pieces of evidence suggesting that this intervention may have generated meaningful spillovers onto the control group. Most basically, the summary statistics on adoption and usage show that there was some uptake of the intervention, and an increase in receipt of SMS price information, even in control villages. Second, the evidence of convergence at the market level indicates that prices move in general equilibrium. If the spatial boundaries of these impacts do not satisfy SUTVA across sub-counties, spillovers to the control will exist. Finally, the "one treated" result in the dyadic data is a direct indication of the fact that trade flows increase between dyads in which only one of the markets is treated, indicating that trade patterns are affected even in markets that are not themselves treated.

What does the presence of these positive spillovers mean for the accuracy of our overall treatment effects? They suggest that our market integration effects and farmer treatment effects may, if anything, be an underestimate on the true gains for the platform. For traders, for whom the platform seems to generate competitive pressures that squeeze profits, the implications of spillovers may be more complicated. To examine the effects of spillovers on traders, we follow Hildebrant et al. (2020) in using baseline data to map the trading linkages between study clusters, and then exploiting the incidental randomized variation in the intensity of treatment within these trading networks to measure spillovers. The average subcounty is connected to 1.8 other subcounties by baseline buying networks, and the average number of treated connections is .9. Because we expect these indirect effects from treatment in the trading network to differ depending on whether a market is itself treated, we use an interaction term to separate the Spillovers on the Treated (ST) from the Spillover on the Non-Treat (SNT). The resulting specification then uses a dummy for own treatment status T_{ij} for subcounty j and observation i, the observational number of other subcounties a market's subcounty was trading with at baseline η_j , and the (then conditionally randomized) number of these subcounties that were treated η_j^T , and estimates the following ANCOVA regression:

$$Y_{ijr} = \alpha_j + \beta_1 T_{ij} + \beta_2 \eta_j + \beta_3 \eta_J^T + \beta_4 (T_{ij} * \eta_j) + \beta_5 (T_{ij} * \eta_j^T) + \rho Y_{ij0} + \epsilon_{ij}.$$
 (5)

Given stratification-block FE α_j , β_1 measures the effect of the treatment in a village with no baseline trading connections (within our study sample), β_2 gives the (observational) slope coefficient on number of total trading partners, β_3 gives the spillover effect of treatment in the trading network for control markets (the SNT), β_4 gives the treatment heterogeneity by number of total trading partners (when none of the trading partners is treated), and β_5 gives the additional spillover effect of having treated trading partners for markets that are themselves treated (that is, ST-SNT). Given the complexity of estimating standard errors in this data structure, we use Randomization Inference to randomly reassign the subcounty-level experiment 1000 times, re-calculating the network treatment intensity and the regression above for each iteration.

In Appendix Table A.18, we perform a balance test by pointing this same specification (minus the ANCOVA term) at the baseline data. The results confirm a generally well-balanced experiment on the network intensities. Table A.19 shows the results of looking at the first year of implementation. These short term effects indicate quite strong and heterogeneous spillovers. The overall effect of treatment is negative (although the intercept treatment effect is estimated with low precision in this specification), and exposure to the treatment for control markets is also bad for traders. The SNT in the third row of this table shows that as control traders become more exposed to the treatment – perhaps facing greater competition in their buying markets, which are treated – they pay higher buying prices, achieve lower markups, and as a result see lower profits despite their quantities traded not changing much. Relative to this, the differential ST - SNT in the first row shows that treated traders are able to offset these spillover-related losses, perhaps due to greater access themselves to other markets when both are plugged into the Kudu network. We see a significant negative coefficient on buying price and a significant positive coefficient on markups that offset the spillover-related erosion of markups observed among the control group. Moreover, we see large and significant effects on both purchases and sales, suggesting that treated traders benefit in terms of sales volumes from exposure to a greater number of treated markets. These effects results in a differential ST - SNT profit effect that is positive, significant, and of a similar magnitude (though oppositely signed) to that of the SNT effect. This suggests the for treated traders, the negative effect of being more exposed to treatment is offset by the positive effect of being able to access other markets that are within the Kudu network, resulting in limited total spillover effects (ST) for this group. Given that the net treatment effect on trader profits is negative and exposure to the treatment is mostly negative for the controls, this suggests that the negative Intention to Treat (ITT) estimates on traders presented throughout the rest of the paper are if anything biased towards zero by the presence of spillovers.

Looking over the longer term, we can pool across midline and endline as is done in the rest of the paper for the trader and farmer outcomes (see Appendix Table A.20). We find that, over the longer term, the negative SNT effects are moderated (and even become marginally positive for net revenue). The differential spillover on treated traders remains significant for the volume of purchases and the signs remain positive for revenues and profits, but are no longer significant. We therefore find that spillover effects for traders are strongest immediately after the introduction of the platform; over the longer run, most of the spatial inequalities created may have been arbitraged away.

Overall, then, our study suggests that we experienced spillover effects that are in the direction of treatment effects (including in Table 1), and hence we may be underestimating the gains from the platform in terms of increases in trade flows, price convergence, and impacts on traders. This is an inherent challenge for any randomized experiment designed to increase trade flows in an environment in which markets are not in perfect autarky. Short of randomizing at the country level, avoiding spillovers due to spatial and trade-network based connections may be impossible.²⁸ What the randomization at scale buys us is the ability to detect *any* differences in GE outcomes – trade flows, prices, etc.– which may be impossible to detect with an experiment randomized at the individual or village level. However, it does not necessarily provide the full treatment effect. In future work, we therefore intend to combine these results with a structural model, as well as with data on the geography and trading networks observe in our survey data, to estimate the full treatment effects on trade flows and market integration.

²⁸Even if one could randomize at country-level, most countries are not in autarky.

6.4 Pre-committed analyses not presented

Our pre-analysis plan, written in 2015, refers to a number of forms analysis that we do not present. For transparency, we describe them briefly here. First, we had intended to conduct an experiment to test credit constraints among traders by offering loans to a randomly selected subset of Commission Agents. We conducted a pilot for this experiment in the first season, issuing 62 short-term working capital loans to a group randomly selected from 124 CAs who expressed a desire for credit. In the end, the repayment rate on these loans was poor (78%) and our partner decided not to move this experiment to the intended scale, so we do not analyze it. Our PAP specifies a set of hypotheses about convergence between spokes and hubs, and the differential effect of treatment for spokes in which the hub is and is not treated. In the end we were only able to map 84% of our spokes to hubs, and the analysis conducted within this reduced sample is typically inconclusive, suggesting that the trading networks may be more complex than our simple hub-and-spoke mapping supposed. So while we emphasize deviations from the superhub in the text, we do not present analysis relative to hubs.

7 The Case for Investing in Trading Platforms

The operating cost for running the platform during the three years of the project was \$927,190. Making up these costs were program administration, including compensation for managers at IPA and AgriNet, along with the deal coordinators and the program staff in the field, was \$560,112. Targeting, including call center operations and all village-level promotion activities, cost \$168,105. Participant training of CAs and AN supervisors was \$39,784. Program material costs, including airtime costs and the money required to run the guarantee system, were \$53,648. Monitoring costs, primarily the eight staff members who supervised transactions on the ground and implemented the guarantees, were \$46,757. Kudu's costs, not borne by the project, consisted of salary for the lead programmer and manager of the platform, short-code fees, and radio ads, and totaled \$58,784.

7.1 The Business Case

Our platform has three separable components, and we consider the business case for each of them in turn. First of these is Kudu. The core issue for the standalone Kudu model is that, due to limited use of mobile money in rural Uganda, the platform does not have a mechanism to collect commissions on transactions.²⁹ Hence, it appears that the most logical model to make Kudu sustainable would be a user fee model where individuals pay to post bids and asks on the platform. Given a total number of bids and asks of approximately 54,000 and costs of \$58,000, this fee would need to be approximately a dollar per use. While this is a tiny amount of money relative to the sums transacted in agricultural deals, it is likely that such a fee would sharply curtail use of the system by farmers and lead to paucity of asks. Further, the usage numbers recorded in the study reflect the influence of the (much more expensive) call center and on-the-ground staff. An alternative business different model would be for Kudu to sell its up-to-the-minute price information. However, to generate reliable and sufficient data, it would have to operate at a massive scale, which presents a chicken-and-egg problem in terms of how to build up to a platform with sufficient scale to make this kind of market information service profitable. Hence, while Kudu represents a substantial potential boon to welfare from market participation, monetizing this benefit is not straightforward.

A second component is the SMS Blast system. The costs of collecting the market price data and sending out the SMS Blast was \$5,857 per month, although as a part of the study we were collecting data on many smaller spoke markets that likely would not make sense from a profit perspective for a commercial system, which may be better off focusing on only larger markets. Our baseline survey asks a question about WTP for market information from traders; the mean stated WTP for an SMS service providing information on spoke, hub, and superhub markets was \$0.42 per month, indicating that our market information system could have broken even with 14,000 users. Had it been optimized to operate in fewer and larger markets, that threshold would fall. So, while our results do not indicate that price-only systems have large benefits for market participants, this business model may be the easiest to construct.

²⁹In fact, the standalone Kudu platform does not even have a mechanism to track if actual transactions occur. Enumerators from the research team conducted follow-up on all matches to supplement Kudu's administrative data and track deal success for the duration of the project.

Finally we have the most costly component of the study platform, which is the AgriNet call center, network of CAs, deal coordinators, and monitoring agents to track transactions on the ground. While this hands-on approach appears to be a necessary part of launching an online trading platform, it is costly and raises the core question of how it can be paid for, given that the core value proposition of the platform to traders and farmers is a *lack* of intermediation costs on the platform. Given that a) the number of highly profitable trades on Kudu that AgriNet was able to intermediate directly was small, and b) substantial expense is required to put the logistics in place to be able to collect commissions on brokered trades, the project was fundamentally unable to develop a model through which brokerage fees could cover the costs of operating the system. A subscription model would be available either to Kudu or to a market price information system, but intermediation costs seem inherently to be linked to commissions on trade. Therefore, we conclude that this type of intermediary platform is not straightforward to make viable as a commercial enterprise at the scale observed in this study.

7.2 The Welfare Case

Our 1,457 sampled study traders were representative of a broader population of 1,752 traders in study districts, meaning that we capture within the study 83% of the people on whom the harm of decreased trading margins fell. Trader profits fell by an average of \$292 per year, or almost \$900 over the three years of the study. Therefore study traders lost a total of \$1.3 million in profits, and the broader sample of which they are representative lost a total of \$1.53 million. Combined with the direct cost of running the platform, we therefore estimate the social cost of the platform to be \$2.42 million dollars. The fundamental welfare question is therefore whether the total benefits to farmers exceed this amount.

The extrapolation of the total farmer benefits from our study sample requires careful consideration. Imprecision issues aside, it is easy to calculate the aggregate the estimated benefit of the intervention to farmers in our study sample. However, because we see evidence that intervention moved general equilibrium outcomes, like total trade volumes and prices, we must consider the effect of the intervention on the broader population of farmers, including those in our study catchment area but who were not sampled in our household surveys.

How can we best estimate the impact of the intervention on this population? First, we focus on treated households that did not receive the Blast, as the Blast was only targeted to a subset of individuals in our study and was not available to the broader population. Second, we estimate effects separately for those in the "Near" village (or "LC1"), who are representative of a smaller population of households in the more urban village containing the TC, and for those in the "Far" village (or "LC1"), who are representative of a much larger population of more rural households in the surrounding subcounty.³⁰

To estimate these ingredients, we present in Table A.22 the core farmer impacts broken out by main treatment status, SMS Blast treatment status, and "Near" vs. "Far" LC1 status, with dummies for each of these three categories and full interactions between them. We can then use the coefficients from Table A.22 to calculate the total revenue effect in each of the four relevant strata.³¹ For the two strata treated by the Blast (near and far LC1s) the study sample represents the population experiencing this effect. For the near stratum not receiving the Blast, the study sample of 1,280 should be representative of the 16,297 households in the same LC1s from which they are sampled. For the far stratum not receiving the Blast, the study sample of 567 should be representative of the much larger sample of 919,697 households in all 'far' parishes (including those containing no study participants).

We start by restricting our benefit calculation to the benefit of farmers in our study sample only. For these farmers, we calculate an aggregate benefit of \$124,000, far less than the costs. However, applying the per-household benefits to the populations for which they should be representative, the outcome in the "Far" Blast control dominates the welfare calculation and drives our estimate of total benefits to farmers to \$34 million dollars, thirteen times as large as the total social cost (details of these calculations are provided in Table A.1. Another way of putting these results is that the net welfare benefit of the platform would be positive if the larger population of farmers in

³⁰This assumes that there were no spillovers from the Blast. If there were spillovers, then welfare effects for non-Blast households in villages in which some farmers go the Blast would not represent the welfare gains to non-Blast households in villages in which no farmers received the Blast.

³¹Because we are undertaking an exercise in the spirit of estimating small benefits spread out over large numbers of households, we ignore significance in this table and use only the point estimates.

non-study "Far" LC1s received a revenue benefit of only \$.79 per year, an effect that is only 6% of the benefit we observed in our study sample that should be representative of this population.

8 Conclusion

This study shows that search costs continue to inhibit trade in African grain markets, and that mobile-based agricultural trading platforms can reduce these frictions and facilitate greater trade. As a result of this increased commerce, price dispersion across markets decreases. The effects of the platform are most pronounced over smaller distances where the pecuniary barriers to trade are smaller. Nearby markets see a reduction in price dispersion of 8% and 15% as one and both markets are treated respectively.

This study also demonstrates that search costs play a part in generating rents that sustain intermediaries. As such, it provides unusually direct experimental evidence that market power can be generated by incomplete information. Traders visit more locations and conduct more trade in markets in which a novel trading platform is introduced to reduce search costs. Yet, because they face reduced margins, traders end up with 14% lower profits and more homogenous trading margins.

The overall picture that emerges from our study is one of concentrated harm from improved arbitrage falling on a defined set of intermediaries (a large fraction of whom are in our sample) and a dispersed benefit accruing to a very large number of farmers (only a small fraction of whom we sample). In terms of statistical power, the concentrated harms are easier to see than the more diffuse benefits. Despite the very large geographical scale of our study, we still face the typical challenge facing all market-level interventions: that of observing and treating a sufficient number of market clusters to detect small benefits arising from general equilibrium effects with precision. With only 55 subcounties assigned to treatment and 55 to control, our ability to reject the null for farmer effects is fairly weak. That said, the qualitative magnitudes of the impacts we uncover are large; farmer revenues rise by 6.5% even in our least-treated group of households (those who reside in rural LC1s and do not receive price and trading information directly).

Unfortunately, the trading platform proved difficult to monetize as a commercial prospect. The

spread between bid and ask prices was small, with ask prices being above bid prices on average for the whole three years of operation. While a subscription model might be viable for either an SMS-based price information system or the online trading platform, the margins on the platform were not sufficiently large or easily arbitraged to permit our private sector partner to make money off the system. This does not mean, however, that the intervention did not generate large welfare consequences. We estimate revenue benefits to farmers that are almost forty times as large as the operating cost of the system and loses to traders. This suggests that platforms increasing the efficiency of core grain trading markets have a public good dimension, and building them at public expense can be an efficient way of improving rural welfare.

References

- Aker, Jenny C, "Information from markets near and far: Mobile phones and agricultural markets in Niger," *American Economic Journal: Applied Economics*, 2010, 2 (3), 46–59.
- and Isaac M Mbiti, "Mobile phones and economic development in Africa," Journal of economic Perspectives, 2010, 24 (3), 207–32.
- Allen, Treb, "Information Frictions in Trade," Econometrica, 2014, 82 (6), 2041–2083.
- Antras, Pol and Arnaud Costinot, "Intermediated trade," The Quarterly Journal of Economics, 2011, 126 (3), 1319–1374.
- Barrett, Christopher B, "Smallholder market participation: Concepts and evidence from eastern and southern Africa," *Food policy*, 2008, *33* (4), 299–317.
- Bergquist, Lauren Falcao and Michael Dinerstein, "Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya," *American Economic Review*, 2020, 110 (12), 3705–3747.
- Camacho, Adriana and Emily Conover, "The Impact of Receiving Price and Climate Information in the Agricultural Sector," *IDB Working Paper Series, No. IDB-WP-220*, 2011.
- Fafchamps, Marcel and Bart Minten, "Impact of SMS-based agricultural information on Indian farmers," The World Bank Economic Review, 2012, 26 (3), 383–414.
- Goyal, Aparajita, "Information, direct access to farmers, and rural market performance in central India," American Economic Journal: Applied Economics, 2010, 2 (3), 22–45.
- Hildebrant, Nicole, Yaw Nyarko, Giorgia Romangnoli, and Emilia Soldani, "Price Information, Inter-Village Networks, and "Bargaining Spillovers": Experimental Evidence from Ghana," *Working Paper*, 2020.
- Jensen, Robert T, "Information, efficiency, and welfare in agricultural markets," Agricultural Economics, 2010, 41, 203–216.
- Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria, "Asymmetric information and middleman margins: An experiment with Indian potato farmers," *Review of Economics and Statistics*, 2018, 100 (1), 1–13.
- Nakasone, Eduardo, "The Role of Price Information in Agricultural Markets: Evidence from Rural Peru," *Working Paper*, 2014.
- Newman, Neil, Lauren Falcao Bergquist, Nicole Immorlica, Kevin Leyton-Brown, Brendan Lucier, Craig McIntosh, John Quinn, and Richard Ssekibuule, "Designing and evolving an electronic agricultural marketplace in Uganda," in "Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies" 2018, pp. 1–11.
- Rashid, Shahidur and Nicholas Minot, "Are staple food markets in Africa efficient? Spatial price analyses and beyond," Technical Report 2010.

- Svensson, Jakob and David Yanagizawa, "Getting prices right: the impact of the market information service in Uganda," *Journal of the European Economic Association*, 2009, 7 (2-3), 435–445.
- and David Yanagizawa-Drott, "Estimating Impact in Partial vs. General Equilibrium: A Cautionary Tale from a Natural Experiment in Uganda," *Working Paper*, 2012.
- Teravaninthorn, Supee and Gael Raballand, "Transport Prices and Costs in Africa," World Bank, 2009.

Appendix A Tables

	Trading Center	Month of Year	Survey Round	TC and Round
Maize	0.04	0.18	0.84	0.87
Beans	0.26	0.15	0.34	0.55
Matooke	0.55	0.01	0.06	0.60
Tomato	0.30	0.04	0.09	0.39

Notes: Each coefficient in this table reports the R-squared from a different dummy variable fixed effects regression of prices in the panel market survey. The first column uses Trading Center FE (and so measures cross-sectional

variation in prices), the second column month-of-year fixed effects (and so measures the degree of typical seasonality), the third column includes fixed effects for each round of the market survey (and so measures the extent of pure time-series market-level price variation), and the fourth column includes both TC and round fixed effects.

	Price	Traders	Seasonal variance	Idiosyncratic variance
TC is Hub Market	57.71^{***} (16.60)	$\begin{array}{c} 4.958^{***} \\ (0.803) \end{array}$	-3276.9 (2260.5)	-19.52^{**} (9.573)
Distance to Superhub	-0.627^{***} (0.0811)	$\begin{array}{c} -0.00828^{**} \\ (0.00392) \end{array}$	-3.918 (11.04)	-0.122^{***} (0.0468)
Number of Lending Institutions	-3.663^{**} (1.617)	$0.101 \\ (0.0783)$	235.6 (220.2)	-0.588 (0.933)
Grain Storage Capacity in TC	0.0605^{*} (0.0338)	$\begin{array}{c} 0.00477^{***} \\ (0.00163) \end{array}$	-6.152 (4.599)	$0.000695 \\ (0.0195)$
TC Accessible by Lorry all Year	16.56 (23.96)	$\begin{array}{c} 0.0949 \\ (1.160) \end{array}$	1108.9 (3263.6)	2.733 (13.82)
Mean DV N	824.1 214	$\begin{array}{c} 4.935\\214\end{array}$	$\begin{array}{c} 10941.1\\ 214 \end{array}$	139.8 214

Table A.2: Descriptive Analysis of Maize Markets.

Notes: The outcome in the first column is the average maize selling price in each TC. The second column gives the average number of maize traders. The third column calculates month-of-year FE for each TC and then takes the variance of these FE, so is a measure of the extent of regular seasonal variation in prices. The fourth column regresses prices on TC and survey round FE, and then squares the residual, so is a measure of the degree of idiosyncratic variation around the TC and round average price.

	Treat	Control	Obs	T	-C
				diff	p- val
Baseline trader completes midline	0.92	0.94	$1,\!457$	-0.02	0.35
Tracked in original endline exercise	0.87	0.85	$1,\!457$	0.02	0.38
Found in Intensive Tracking	0.89	1.00	41	-0.11	0.08
Baseline trader completes endline	0.89	0.87	$1,\!457$	0.03	0.18

Table A.3: Trader Survey Attrition.

Notes: Analysis uses the full baseline sample of traders to study different definitions of attrition. Standard errors are clustered by subcounty.

	Treat	Control	Obs	T	-C
				$di\!f\!f$	p- val
Tracked in standard exercise	0.92	0.93	2,971	-0.01	0.44
Tracked in intensive tracking	0.74	0.85	39	-0.11	0.35
Successfully tracked	0.93	0.94	$2,\!971$	-0.01	0.31

Table A.4: Household Survey Attrition.

Notes: Analysis uses the full baseline sample of farming households to study different definitions of attrition. Standard errors are clustered by subcounty.

	Treat	Control	Obs	Т-(C
				diff	p- val
Maize buying price	506.06	494.02	232	12.04	0.37
Maize selling price	625.92	628.50	232	-2.58	0.86
Number of maize traders	8.82	8.84	233	-0.03	0.98
Maize quality	1.58	1.61	232	-0.04	0.57
Beans buying price	1,566.81	$1,\!573.55$	214	-6.74	0.92
Beans selling price	$1,\!884.94$	$1,\!994.98$	214	-110.04	0.16
Number of beans traders	5.14	5.01	233	0.13	0.88
Beans quality	1.45	1.50	214	-0.06	0.51

Table A.5: Market Survey Balance

Notes: Analysis uses the cross-sectional average of outcomes from the two pre-treatment market survey waves to examine balance of the market survey. Standard errors are clustered by subcounty.

	Maize	Beans	Bananas	Tomatoes
One treated	$\begin{array}{c} 0.0566 \ (0.0564) \end{array}$	$\begin{array}{c} 0.0349 \\ (0.0704) \end{array}$	$\begin{array}{c} 0.0571 \\ (0.0895) \end{array}$	0.00314 (0.0817)
Both treated	$0.0905 \\ (0.103)$	$0.0104 \\ (0.115)$	$0.103 \\ (0.158)$	$0.0567 \\ (0.145)$
Mean DV N	$4.592 \\ 26218$	$5.978 \\ 21129$	$8.265 \\ 20196$	$3.478 \\ 26149$

Table A.6: Market Survey Balance in Price Dispersion

Notes: Analysis uses dyadic averages in the two pre-treatment rounds of the market survey to examine balance in price dispersion. Standard errors are clustered by subcounty.

	Treat	Control	Obs	T-C	
				$di\!f\!f$	p- val
Female	0.07	0.06	$1,\!281$	0.01	0.64
Age	37.16	37.39	$1,\!281$	-0.23	0.76
Education	7.68	7.32	$1,\!281$	0.36	0.24
Age of business	10.86	10.92	$1,\!178$	-0.07	0.92
# of subcounties in which bought	1.15	1.12	$1,\!281$	0.03	0.44
# of subcounties in which sold	1.27	1.31	$1,\!281$	-0.03	0.65
Net revenue, mz & bn	$21,\!946,\!001.68$	$28,\!474,\!012.42$	$1,\!275$	-6,528,010.74	0.54
Business costs per month	$6,\!290,\!868.45$	$6,\!050,\!540.21$	$1,\!281$	$240,\!328.24$	0.80
Annual Revenue	$47,\!550,\!250.45$	$45,\!657,\!411.81$	$1,\!278$	$1,\!892,\!838.64$	0.82
Annual Costs	$43,\!068,\!736.38$	40,790,579.76	$1,\!281$	$2,\!278,\!156.62$	0.72
Volume buy (kgs), mz	$112,\!323.01$	$100,\!580.90$	$1,\!281$	11,742.10	0.63
Volume buy (kgs), bn	$6,\!174.67$	4,936.33	$1,\!281$	$1,\!238.34$	0.49
Volume sold (kgs), mz	$157,\!676.55$	$161,\!821.69$	$1,\!281$	-4,145.14	0.94
Volume sold (kgs), bn	$6,\!667.08$	$5,\!906.31$	$1,\!281$	760.77	0.71
Trade maize	0.92	0.94	$1,\!281$	-0.02	0.42
Trade beans	0.28	0.25	$1,\!281$	0.03	0.54
Annual profits	$5,\!617,\!367.84$	5,717,231.92	$1,\!274$	-99,864.08	0.92

Table A.7: Trader Survey Balance

Notes: analysis conducted using the endline sample, with weights reflecting survey sampling and intensive tracking. Standard errors are clustered by subcounty, the unit of assignment. The first two columns give the means in the control and treatment group respectively. The third column gives the total number of observations across the two groups. The last two columns give differences in means and the corresponding p-value.

	Treat	Control	Obs	T-C	
				diff	p- val
Number HH members	6.17	6.19	2,775	-0.03	0.86
Female	0.37	0.39	2,775	-0.01	0.66
Age	41.81	41.81	2,774	0.00	1.00
Highest grade completed	7.60	7.08	2,775	0.51	0.07
Food expenditure (month)	$93,\!295.98$	$79,\!337.38$	2,743	$13,\!958.61$	0.04
Land size (acre)	5.65	5.88	2,529	-0.23	0.60
Qtny sold, total (annual, kg)	$1,\!133.87$	$1,\!056.52$	2,775	77.35	0.70
Qtny harvest, total (annual, kg)	1,862.26	1,751.21	2,775	111.05	0.68
Number times sell	3.12	2.75	2,775	0.37	0.11
Percent of time sold at market	0.29	0.28	2,775	0.01	0.85
Sell in market	0.36	0.33	2,775	0.03	0.59
Distance to market	2.02	2.21	$2,\!420$	-0.20	0.63
Distance to Kampala	175.00	172.15	$2,\!437$	2.85	0.80
Assets (UGX)	$2,\!508,\!859.59$	$2,\!297,\!790.91$	2,775	$211,\!068.68$	0.62
Total exp (monthly, UGX)	$219,\!099.68$	$191,\!827.79$	2,775	$27,\!271.89$	0.08
Input exp (annual, UGX)	$275,\!318.30$	$304,\!037.63$	2,775	-28,719.33	0.47
Revenue, total (annual UGX)	$637,\!169.99$	$555,\!801.61$	2,775	$81,\!368.38$	0.43

Table A.8: Household Survey Balance

Notes: analysis conducted using the endline sample, with weights reflecting survey sampling and intensive tracking. Standard errors are clustered by subcounty, the unit of assignment. The first two columns give the means in the control and treatment group respectively. The third column gives the total number of observations across the two groups. The last two columns give differences in means and the corresponding p-value

	Buy Price	Sell Price	Sell-Buy Margin	No. of Traders	Quality
Treatment	-1.781	-1.326	0.478	0.158	-0.0271
	(11.65)	(13.45)	(5.664)	(0.481)	(0.0483)
Mean DV	721.9	820.6	98.71	5.013	1.642
Ν	236	236	236	236	236

Table A.9: Impact on Levels in Maize Markets.

Notes: Analysis uses the cross-sectional average of all post-treatment waves of the market survey to study the impact of the intervention on market-level outcomes. Standard errors are clustered at the subcounty level.

Table A 10.	Impact on	Out-of-sample	Volumes	Purchased	and Sold
1abic 11.10.	impact on	Out-of-sample	Volumes	i urchascu	and bolu.

	(1)	(2)
	Out-of-sample buy	Out-of-sample sell
Treat	3996*	3543
	(2102)	(5187)
Observations	2914	2914
Mean of DV	10634	32780
R squared	0.03	0.06
Controls	Yes	Yes

Notes: The outcome variable is volumes purchased (Column 1) and sold (Column 2) in out-of-sample markets by traders, regressed on trader treatment status. Standard errors are clustered at the subcounty level.

	Blast	No Blast	Obs	Ĥ	Ŋ
				diff	p-val
Heard of Kudu	0.56	0.41	1,370	0.15	0.00
Received SMS price info	0.58	0.46	$1,\!437$	0.13	0.05
Ever used Kudu	0.28	0.21	$1,\!438$	0.07	0.32
Completed deal on Kudu	0.02	0.02	1,438	-0.00	0.95
Error Spoke Price	104	119	1,428	-15	0.01
Error Hub Price	119	131	1,378	-11	0.14
Error Superhub Price	209	227	1,274	-19	0.33
Used SMS price info when negotiating price	0.25	0.18	1,437	0.07	0.05

Table A.11: Impact of SMS Blast on Farmer Uptake. SMS Treatment vs control within treatment markets.

Notes: Table uses all farming households in treatment villages, and compares those who receive the SMS Blast directly to those who were excluded from the blast.

Mz Rev in Treatment 51948.7 (68368.0 (68368.0 26761.1 (91485.6 in Control -8540.5 (82726.7 (82726.7 671870.8	 Qty Sold 90.15 90.15 (96.23) -2.131 (136.9) 61.24 (189.3) (189.3) 	Price 15.04 (14.03) 8.573 8.573 (15.73) 9.351 9.351 (24.59) 631.0	Know price 0.0329 (0.0424) 0.0576 (0.0582) 0.00691 (0.0674) 2.711	Know better 0.0787* (0.0468) -0.0411 (0.0575) -0.00214 (0.0646) 2.556	Notes: Table includes both treatment and
2775	2769	1959	2774	2774	

Table A.12: Impact of SMS Blast. Farmer Outcomes, SMS Blast and SMS rollin separately estimated.

Blast. The dummy 'Information in Control' indicates that a household was in a control cluster that had been rolled into the SMS Blast as of the subcounty level.

	(1) Buy Price	(2) Sell Price	(3) Sell Buy Margin	(4) No of Traders	(5) Quality
Roll in Treatment	$1.843 \\ (7.196)$	$3.364 \\ (9.234)$	$1.355 \\ (4.298)$	-0.127 (0.446)	$0.00524 \\ (0.0364)$
Mean DV N	$719.4 \\ 6928$	$\begin{array}{c} 819.7\\ 6936\end{array}$	$100.3 \\ 6928$	$4.861 \\ 7445$	$1.657 \\ 6936$

Table A.13: Maize: Impact of the SMS Blast Roll-in.

Standard errors in parentheses

* p < 0.10, ** p < 0.05, *** p < 0.01

Notes: This analysis uses monadic market data in panel format with fixed effects for market and market survey wave, and analyzes the roll-in of the SMS Blast to control markets during the second half of the study. Standard errors are clustered at the subcounty level.

	(1) Maize	(2) Beans	
One treated by Roll-in	0.0288 (0.0296)	-0.0715 (0.0525)	
Both treated by Roll-in	$0.0262 \\ (0.0676)$	-0.0823 (0.0917)	Notes: This analysis uses dyadic market d
Dist (10km)	$0.00238 \\ (0.00451)$	-0.00363 (0.00812)	
Mean DV N	$4.601 \\ 330828$	$6.024 \\ 165528$	

Table A.14: Impacts of SMS Roll-in on Buying Price, 2-way FE.

Standard errors in parentheses

* p < 0.10, ** p < 0.05, *** p < 0.01

in panel format with fixed effects for market and market survey wave, and analyzes the roll-in of the SMS Blast to control markets during the second half of the study. Standard errors are two-way clustered at the subcounty level.

	(1)	(2)	(3)	(4)	(5)
	Num Bids	Num Matches	Num Trans	Qnt Trans	Value Trans
Basic	-0.375	-0.122	0.019	100.096	101479.290
	(1.357)	(0.760)	(0.061)	(424.860)	(448148.426)
Comprehensive	0.863	0.158	0.012	-63.482	-2.255e+05
	(1.877)	(1.055)	(0.089)	(543.738)	(551463.574)
Observations	1095	1095	1095	1095	1095
Mean of DV	10.05	5.42	0.20	1246.51	1362961.46
R squared	0.00	0.00	0.00	0.00	0.00

Table A.15: Impact of the Buyer-level Transport Guarantees.

Standard errors in parentheses

* p < 0.10, ** p < 0.05, *** p < 0.01

Notes: table presents an analysis of the experiment that assigned the initial group of buyers on the Kudu platform to receive either no transport guarantees (control), a Basic guarantee (against quantity shortfalls), or a Comprehensive guarantee (also against quality and price shortfalls). Analysis is conducted at the buyer level, treatment dummies indicate buyer-level randomized assignment. The first three columns give the number of bids posted, the number of bids matched in Kudu, and the number of transactions conducted. Columns 4 and 5 give the total quantity and the total value transacted on Kudu during the study.

	(1)	(2)	(3)	(4)	(5)
	Num Bids	Num Matches	Num Trans	Qnt Trans	Value Trans
Fraction Basic	3.357	4.108	0.004	189.198	66685.904
	(5.375)	(3.106)	(0.167)	(1113.505)	(1.092e+06)
Fraction Comp	1.424	5.375	0.022	-304.297	-8.012e+05
	(7.232)	(4.842)	(0.201)	(1227.848)	(1.250e+06)
Observations	273	273	273	273	273
Mean of DV	10.05	5.42	0.20	1246.51	1362961.46
R squared	0.00	0.01	0.00	0.00	0.00

Table A.16:	Impact of	f the Buyer-level	Transport	Guarantees amor	g Original	Controls.
					0 - 0 -	

Standard errors in parentheses

* p < 0.10, ** p < 0.05, *** p < 0.01

Notes: analysis is conducted only within the original control group for the buyer-level transaction guarantee experiment. RHS variables are the fraction of bids that were assigned to the Basic or the Comprehensive guarantee for each control buyer. The first three columns give the number of bids posted, the number of bids matched in Kudu, and the number of transactions conducted. Columns 4 and 5 give the total quantity and the total value transacted on Kudu during the study.

		Original Co	ntrol		Subsequent B	uyers
	(1)	(2)	(3)	(4)	(5)	(6)
	Deals	Quantity	Value	Deals	Quantity	Value
Basic	0.025**	104.305^{**}	94446.186**	-0.011	-47.781	-41587.874
	(0.012)	(40.930)	(37389.973)	(0.014)	(40.394)	(36034.253)
Comprehensive	0.042^{**}	135.361**	111227.273**	-0.028	-54.511	-47387.490
	(0.019)	(61.461)	(53619.444)	(0.018)	(56.674)	(50488.238)
Observations	1468	1468	1468	3092	3092	3092
Mean of DV	0.03	82.09	76818.18	0.15	367.99	327168.36
R squared	0.01	0.01	0.01	0.00	0.00	0.00

Table A.17: Impact of the Transaction-level Transport Guarantee.

Standard errors in parentheses

* p < 0.10, ** p < 0.05, *** p < 0.01

Notes: analysis is conducted at the level of the match within Kudu, which is the step at which the transport guarantees were assigned to those not originally given the treatment from the buyer-level experiment. Columns 1-3

use the original control group from the buyer-level experiment, and Columns 4-6 use all other buyers who subsequently entered the platform. Controls are dummies for whether the match was assigned to receive the Basic or Comprehensive guarantee, with those receiving no guarantee as the remaining category. The outcomes give the number of successful transactions, as well as the quantity and value of successful transactions on Kudu.

	Ann Profit	b_rev_ann_th_win	Markup	Purchases	Sales	Buy Price	Sell Price
Own T x Treated Links	846.24	22290.97^{***}	-4.96	-2287.37	13153.82	5.63	0.39
	(0.101)	(0.00)	(0.477)	(0.853)	(0.788)	(0.352)	(0.950)
Own T x Number Links	-332.88	-5003.84	0.44	3482.98	-3717.52	-3.62	-2.71
	(0.289)	(0.195)	(0.936)	(0.686)	(0.926)	(0.466)	(0.531)
Treated Links	237.84	-3402.02	10.21^{*}	5775.27	-48309.19	1.99	12.81^{***}
	(0.571)	(0.261)	(0.053)	(0.509)	(0.267)	(0.693)	(0.002)
Number Links	63.83	3927.58	-5.36	-6300.79	31295.05	-0.34	-6.38*
	(0.797)	(0.121)	(0.205)	(0.318)	(0.374)	(0.938)	(0.051)
Own Treatment	149.55	-6803.38	10.29	16273.67	19258.13	-5.76	2.93
	(0.862)	(0.453)	(0.295)	(0.576)	(0.722)	(0.531)	(0.845)
Observations	1449	1453	1341	1457	1457	1343	1347
Mean of DV	5846.70	49659.73	114.62	105072.08	144885.24	432.76	548.90
R squared	0.02	0.08	0.01	0.00	0.00	0.01	0.02
Notes: Table analyzes baselin	ie balance of the n	nain outcomes using the	specification for	or studying spillo	ver effects. Num	<u>aber of Links is t</u>	he number of
markets with which a TC trad	led at baseline (n_i)), and Number of Treated	d Links is the	number of these	links treated (n)	T^{T}). The third row	v of the table

Outcomes.
Balance,
Spillovers
Baseline
Table A.18:

gives the Spillovers on the Non-Treated (SNT) and the top row the difference between the Spillovers on the Treated (ST) and the SNT.

	Annual Profit	Net Revenue	Markup	Purchases	Sales	Buy Price	Sales Price
Own T x Treated Links	1,014.09*	$68, 185.15^{**}$	15.92^{***}	$215,220.17^*$	$156,052.16^{**}$	-13.18*	2.19
	(0.10)	(0.01)	(0.00)	(0.01)	(0.01)	(0.07)	(0.69)
	[0.07]	[0.04]	[0.01]	[0.05]	[0.05]	[0.07]	[0.66]
Own T x Links	-447.95	-6,719.91	-12.78***	-28,867.53	-13,111.15	10.77^{***}	-2.70
	(0.44)	(0.47)	(0.00)	(0.34)	(0.56)	(0.03)	(0.45)
	[0.14]	[0.99]	[0.00]	[66.0]	[0.99]	[0.00]	[0.25]
Treated Links	-761.10^{**}	4,072.72	-14.59^{***}	9,333.44	9,210.77	13.19^{***}	-1.00
	(0.17)	(0.20)	(0.00)	(0.36)	(0.36)	(0.02)	(0.81)
	[0.02]	[0.37]	[0.00]	[0.36]	[0.34]	[0.00]	[0.34]
Number Links	552.58^{**}	-10,264.10	8.96^{***}	-20,746.81	-20,352.84	-9.16^{***}	0.15
	(0.31)	(0.02)	(0.00)	(0.14)	(70.0)	(0.03)	(0.96)
	[0.03]	[0.37]	[0.00]	[0.43]	[0.35]	[0.00]	[0.76]
Own Treatment	-1,320.49	-16,096.17	-5.52	47,602.32	8,411.00	-16.63	-16.02
	(0.07)	(0.62)	(0.62)	(0.68)	(0.91)	(0.19)	(0.13)
	[0.14]	[0.72]	[0.66]	[0.71]	[0.91]	[0.22]	[0.15]
Observations	1,344	1,327	1,193	1,355	1,355	1,199	1,209
Mean of DV	7,116.18	54,927.69	166.83	172,695.56	160,641.03	593.04	758.33
R squared	0.08	0.03	0.02	0.04	0.04	0.02	0.06
		;	,				

Table A.19: Short-term Trader Spillovers.

refer to maize only. Number of Links is the number of markets with which a TC traded at baseline (η_j) , and Number of Treated Links is the number of these links treated (η_i^T) . The third row of the table gives the Spillovers on the Non-Treated (SNT) and the top row the difference between the Spillovers Notes: Table analyzes pooled Trader outcomes in the midline survey. Columns 1 and 2 are total profits and revenues ('000), and the remaining columns on the Treated (ST) and the SNT. The table presents the p-values calculated from a standard regression clustered at the subcounty level in the parentheses, as well as the p-values derived from Randomization Inference in brackets. All outcomes winsorized.

	Annual Profit	Net Revenue	Markup	Purchases	Sales	Buy Price	Sales Price
Own T x Treated Links	524.89	69.94	2.13	$26,518.05^{*}$	23, 230.55	10.66	13.59
	(0.32)	(0.98)	(0.77)	(0.13)	(0.19)	(0.20)	(0.01)
	[0.31]	[0.91]	[0.69]	[0.09]	[0.13]	[0.19]	[0.12]
Own T x Links	15.04	2,675.19	-4.60	1,345.94	4,162.57	-6.81^{*}	-11.61^{***}
	(0.97)	(0.10)	(0.44)	(0.88)	(0.66)	(0.32)	(0.01)
	[0.96]	[0.26]	[0.34]	[1.00]	[0.99]	[0.08]	[0.01]
Treated Links	-46.95	$1,610.34^{*}$	6.47	-4,816.02	-1,859.15	-2.37	-5.61
	(0.90)	(0.25)	(0.35)	(0.36)	(0.74)	(0.75)	(0.18)
	[0.88]	[0.05]	[0.15]	[0.97]	[0.57]	[0.94]	[0.18]
Number Links	64.74	$-1,082.75^{**}$	-3.49	3,799.28	1,403.12	2.18	5.76^{**}
	(0.86)	(0.29)	(0.53)	(0.47)	(0.80)	(0.71)	(0.07)
	[0.58]	[0.02]	[0.41]	[0.80]	[0.45]	[0.49]	[0.03]
Own Treatment	-1,482.01	-5,744.85	2.11	-20,247.12	-26,011.12	2.09	2.84
	(0.06)	(0.17)	(0.86)	(0.29)	(0.21)	(0.84)	(0.82)
	[0.10]	[0.33]	[0.92]	[0.43]	[0.38]	[0.87]	[0.91]
Observations	2,592	2,179	1,993	2,453	2,453	2,056	2,016
Mean of DV	7,279.21	12,583.56	129.24	70,743.79	76,518.06	587.48	730.44
R squared	0.08	0.01	0.03	0.05	0.04	0.01	0.02

Table A.20: Pooled Trader Spillovers.

of Treated Links is the number of these links treated (η_{j}^{T}) . The third row of the table gives the Spillovers on the Non-Treated (SNT) and the top row the Notes: Table analyzes pooled Trader outcomes across the midline and endline as in the rest of the paper. Columns 1 and 2 are total profits and revenues (000) and the remaining columns refer to maize only. Number of Links is the number of markets with which a TC traded at baseline (η_j) , and Number difference between the Spillovers on the Treated (ST) and the SNT. The table presents the p-values calculated from a standard regression clustered at the subcounty level in the parentheses, as well as the p-values derived from Randomization Inference in brackets. All outcomes winsorized.

	Mz Sell pr	Mz Buy pr	Mz Quality	Mz Traders	Bn Traders
Own T x Treated Links	7.045	5.541	-0.009	1.268^{***}	0.806^{***}
	(0.334)	(0.440)	(0.784)	(0.001)	(0.000)
	[0.336]	[0.460]	[0.890]	[0.006]	[0.002]
Own T x Links	-4.176	0.933	-0.003	-0.240	-0.224
	(0.444)	(0.871)	(0.896)	(0.286)	(0.091)
	[0.202]	[0.722]	[0.928]	[0.852]	[0.472]
Treated Links	4.505^{**}	7.130^{**}	0.004	-0.147	-0.053
	(0.513)	(0.272)	(0.896)	(0.338)	(0.561)
	[0.036]	[0.024]	[0.874]	[0.814]	[0.334]
Number Links	-2.959^{*}	-5.053^{***}	-0.008	0.048	0.001
	(0.567)	(0.338)	(0.684)	(0.694)	(0.988)
	[0.064]	[0.008]	[0.732]	[0.692]	[0.510]
Own Treatment	0.496	-6.260	-0.019	-0.385	0.037
	(0.972)	(0.606)	(0.716)	(0.489)	(0.945)
	[0.964]	[0.572]	[0.672]	[0.526]	[0.974]
Observations	708	708	708	708	708
Mean of DV	830.56	730.19	1.65	5.10	2.77
R squared	0.00	0.01	0.01	0.10	0.05

Table A.21: Pooled Market Survey Spillovers across all 3 years.

Notes: Table analyzes pooled Market Survey outcomes, collapsed to annual averages as in the trader spillover estimation. Number of Links is the number of markets with which a TC traded at baseline (η_j) , and Number of Treated Links is the number of these links treated (η_j^T) . The third row of the table gives the Spillovers on the Non-Treated (SNT) and the top row the difference between the Spillovers on the Treated (ST) and the SNT. The table presents the p-values calculated from a standard regression clustered at the subcounty level in the parentheses, as well as the p-values derived from Randomization Inference in brackets.

	(1)	(2)	(3)
	Revenue	Qty Sold	Price
Treat	12.29	76.08	-0.00612
	(26.66)	(159.4)	(0.00680)
Blast Treatment	7.697	0.795	0.0115**
	(24.47)	(137.9)	(0.00512)
Treat x Near LC1	-9.117	-170.3	0.0144^{*}
	(30.65)	(165.3)	(0.00836)
Blast x Near LC1	14.18	166.4	-0.0128*
	(30.99)	(173.4)	(0.00736)
Near LC1	40.90***	210.1**	0.00380
	(14.11)	(87.01)	(0.00598)
Observations	2775	2769	1959
Mean of DV	192.0	1040.1	0.180
R squared	0.283	0.309	0.0227
Controls	Yes	Yes	Yes

Table A.22: Farmer impacts for Costing.

Standard errors in parentheses

* p < 0.10, ** p < 0.05, *** p < 0.01

Notes: This table uses farmer endline data to provide the treatment effects needed for the welfare analysis. 'Treat' is a dummy for treatment at the subcounty level, 'Blast treatment' indicates the household received the SMS Blast directly, and 'Near LC1' is a dummy indicating that the

household was located in the village adjacent to the market, rather than the randomly sampled Far LC1s that were more rural. Standard errors are clustered at the subcounty level.

Figure A.1: Welfare Calculations

Project Costs:							-\$887,464.00	-\$887,464.00
Traders:								
	Profit				# traders in	# traders in	Total prof effect in	Total profit effect
	effect		# years		sample	рор	sample	in population
	-292			3	1457	1752	-\$1,276,332.00	-\$1,534,752.00
Total Costs:								-\$2,422,216.00
Farmers:								
	Revenue					# households		
Treatment Cell:	Coefficient	Revenue total	4 yoars		# households	in population	Total Rev effect in	Total rev effect in
ffedunient cen.	from	effect	# years		in sample	of sampled	sample	population
	regression					parishes		
Far Blast Control	12.29	\$12.29		3	567	919,697	\$20,905.29	\$33,909,246.39
Far Blast	7.697	\$19.99		3	360	360	\$21,585.96	\$21,585.96
Near Blast Control	-9.117	\$3.17		3	1,280	16,297	\$12,184.32	\$155,131.14
Near Blast	14.8	\$30.26		3	764	764	\$69,362.80	\$69,362.80
						Farmer Total:	\$124,038.37	\$34,155,326.29
						ı.	In Sample	In Population

Overall:

ample In Population -\$2,039,757.63 \$31,733,110.29

Appendix B Figures



Figure B.1: Maps of the Study Area

Notes: The left-hand panel is USAID's FEWS-Net map of Surplus Maize Areas of Uganda, and the right-hand panel shows the 11 study districts.



Figure B.2: Map of Hub and Spoke structure of trading centers

depicts each study spoke TC connected by a straight line to its respective hub market.

A M J J A S O N D J F M A M J J A S O N D J F M A M J J A S O N D J F M A M J J A S O N D J F M A M J J A S Harvests **Surveys** Market Bi-weekly Baseline Midline Endline Trader Baseline Endline Household **Interventions** Kudu SMS Blast Bi-weekly

Figure B.3: Study Timeline



Figure B.4: CONSORT Diagram of Study Recruitment and Attrition



Figure B.5: Attrition from the Market Survey

Notes: The figure shows the Cumulative Density Functions (CDFs) of the fraction of intended surveys (83) that were completed for each TC, separating out the treatment and control TCs. The KS test fails to reject that the two distributions are the same.

Figure B.6: Amount of Grain newly Posted to Kudu per Day



Figure B.7: Volume of Bids and Asks





Figure B.8: Maize Prices in Kudu vs. Market Survey

Figure B.9: Distribution of Ask and Bid Prices, by Season









Figure B.11: Impact of the Roll-in of Price Information

Notes: This

figure shows the average maize selling price at the TC/market survey level. All TCs not yet receiving the Mobile Price Information service at a moment in time are group together into the Control, and then each tranche is broken out as a separate average once it is rolled in to the information service.