Search Costs, Intermediation, and Trade: Experimental Evidence from Ugandan Agricultural Markets

Lauren Falcao Bergquist[1,*], Craig McIntosh[2], and Meredith Startz[3]

[1]Yale University and NBER
[2]University of California, San Diego
[3]Dartmouth College and NBER

August 2026

Abstract

We study the large-scale experimental rollout of a platform that reduced search and matching frictions in Ugandan agricultural markets by connecting buyers and sellers. Market integration improved substantially: trade increased and price gaps fell. Interpreting the experiment through a trade model, we estimate treatment effects accounting for equilibrium changes that impact control markets. The intervention reduced fixed trade costs by 20% and increased trade flows between treated markets by 7% and across all markets by 1%. Scale economies shaped engagement: few farmers used the platform, but equilibrium price convergence from improved arbitrage by larger traders passed through to farm revenue.

*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 and James Allen, Viyaleta Farysheuskaya, Russell Morton, Camilla Sarzo, and Jared Stolove for excellent research assistance. We thank Paul Nyende from AgriNet and the Kudu team, including Richard Ssekibuule, Nicole Immorlica, Kevin Leyton-Brown, Brendan Lucier, Neil Newman, and John Quinn, for their partnership. This trial was registered as AEARCTR-0000749, and is covered by UCSD IRB #140744. The data repository is openicpsr-251706 (see Bergquist et al., 2026). All errors are our own.

1 Introduction

Trade barriers in sub-Saharan Africa are among the highest in the world (Teravaninthorn and Raballand, 2009; Porteous, 2019). This impedes integration of agricultural surplus and deficit areas, putting a wedge between consumer and producer prices, with implications for food security, farmer income, and aggregate welfare (Barrett, 2008). While this is due in part to high transportation costs, spatial price variation is much larger than can be explained by transportation alone (Atkin and Donaldson, 2015). Evidence from the spread of mobile phones suggests that information frictions may be an important contributor – as phones reduced the cost of accessing information about distant places, market integration improved and price dispersion fell (Aker, 2010; Jensen, 2010; Allen, 2014). However, integration remains highly incomplete even in places with high mobile phone penetration (Moser et al., 2009). Furthermore, we know very little about *which* information matters – much of the literature has focused on prices, but interventions that disseminate price information have had limited impact (Fafchamps and Minten, 2012).

We investigate the role of a different information friction: search for transaction partners. Digital platforms aimed at facilitating matches are an increasingly popular tool in markets in low and middle income countries (LMICs) (Abate et al., 2023). We conduct a large-scale randomized control trial of a mobile marketplace in Uganda, which links potential buyers and sellers of agricultural goods through a simple SMS-based platform. The design of the platform and experiment allow us to isolate the causal effect of reducing buyer-seller search costs, distinct from other types of information or trade barriers. The at-scale randomization enables us to measure impacts on aggregate outcomes like trade flows and prices, accounting for equilibrium forces that affect both treated and control markets.

The mobile marketplace we study, called Kudu, acts as a clearinghouse. Those buying and selling agricultural commodities submit bids and asks, and the platform directly connects potential matches, who can then choose to transact or not off-platform. Kudu is free and promoted to both farmers and traders, with in-village support services to ensure ease of access. As part of the study, introduction was randomized across 110 Ugandan subcounties, covering a population of almost three million. We track outcomes over six harvest seasons, collecting high frequency data on crop prices in 236 markets, as well as multiple rounds of surveys with traders and farmers.

We begin with a simple treatment-control comparison, and find that there is more trade and price convergence among markets connected by Kudu. Trade increases on both extensive and intensive margins: treated market pairs are more likely to engage in any trade with each other, have a larger number of traders who trade between them, and see higher volumes traded. Prices are higher in treated markets relative to control markets in crop surplus areas, but lower in deficit areas. As a result, price gaps between treated markets are smaller, even though average price levels do not change. A separate treatment arm that provided price information rather than connections to new transaction partners had no impact, confirming that the channel for these impacts is likely buyer-seller search.

Randomization circumvents many of the inference problems that arise when studying potentially endogenous variation in trade costs, and randomizing at the subcounty level – rather than the individual, or even village level – enables us to see measurable changes in equilibrium outcomes like prices. However, a remaining challenge is that although control markets do not benefit directly from the reduction in search costs, they are affected by changes in these equilibrium outcomes through trade links. This means that simple treatment-control comparisons do not yield valid estimates of treatment effects, a challenge inherent to a trade setting. To correctly measure the impact of access to the platform, we therefore interpret the experimental variation through the lens of a trade model, estimating treatment effects in the presence of equilibrium forces that create spillovers between treatment and control units.

We build a model of domestic agricultural trade with fixed and variable trade costs between markets, and traders who are heterogeneous in both size and idiosyncratic match quality with specific transaction partners. We embed a search friction by requiring sellers to pay a fixed cost to find buyers in a particular market before selling there. This represents the costly search activities we observe in the survey data, such as time spent calling contacts or traveling to another market to look for buyers. Kudu potentially lowers this search cost by making it easier to connect with new partners. We estimate the direct effect of the platform on trade costs using experimental variation in bilateral treatment status between markets, holding market-level factors (such as prices) constant. We find that Kudu drove a 20% reduction in fixed trade costs between treated markets.

To quantify the total impact of Kudu, we estimate additional parameters of the model – such as the elasticity of substitution across sellers and the distribution of traders' operating costs – that govern the size of indirect effects operating through equilibrium forces. To do this, we make use of random variation in both own treatment status and in exposure to treatment in other markets. These estimates allow us to simulate a scenario in which no market has access to Kudu; in the presence of spillovers, this is the counterfactual needed to estimate treatment effects. With this in hand, we can identify the impact of Kudu on treated routes *and* on control routes. We find that Kudu increased trade flows on treated routes by 7% and in the aggregate across all markets (including negative spillovers on control routes) by 1%. While still substantial, this is only about a third of the treatment effect on aggregate trade flows implied by a simple treatment-control comparison that does not account for equilibrium effects.

Kudu's impact demonstrates that there are still large information frictions impeding agricultural trade in LMICs, and that these can be meaningfully alleviated by a low-cost clearinghouse technology that makes it easier to match with buyers. Many practitioners hope that platforms like Kudu will increase farmers' revenue, perhaps by enabling them to sell directly to buyers in higher priced markets. While we find that many farmers do benefit, the mechanism is more nuanced.

First, farmers very rarely engage in cross-market trade, either with or without Kudu. Instead, almost all activity on the platform and the resulting improvement in arbitrage is driven by professional traders. Our structural estimation sheds light on one possible reason why: the size of fixed costs is such that all but the very largest farmers operate at too small a scale to make cross-market trade profitable. The intervention lowers these fixed costs, allowing smaller traders to enter more trade routes on the margin – which is precisely what we observe in the survey data – but the average level of scale economies remains too large for most farmers.[1]

Despite their lack of direct engagement, farmers are still affected by changes in equilibrium prices. However, the reduction in search costs leads to price convergence, rather than a change in average prices. This is good for farmers and bad for consumers in surplus markets, where initially low prices rise, and the opposite in deficit markets, where initially high prices fall. There are net gains from trade, such that overall welfare in the study area increases. We see empirically that farmers in relative surplus areas enjoy significant increases in crop revenue

Finally, we use our structural model to explore how the total impact of the platform would change as access is counterfactually rolled out to more markets, holding the direct technological impact on trade costs across treated markets constant. At low levels of saturation, aggregate gains are small, and trade diversion effects are high, with large changes in the volume of trade on treated routes for a given change in fixed trade costs. When the platform is rolled out to all markets, treatment effects on average trade volumes are relatively small, but the change in the allocation of trade yields welfare gains that are roughly four times as large.

This paper contributes to a literature on interventions to increase smallholder farmers' revenues in LMICs. Much of the existing work in this space has focused on supply side interventions to increase farmer productivity, such as agricultural extension (Beaman et al., 2023), credit (Karlan et al., 2014), and insurance (Carter et al., 2017). We focus on market access and the role of search costs. The few other studies in this space have mostly focused on price search, with several randomized evaluations of programs to disseminate price information mostly to farmers via mobile phones finding mixed, and overall limited, impacts (Camacho and Conover, 2011; Fafchamps and Minten, 2012; Nakasone, 2014; Mitra et al., 2018; Hildebrandt et al., 2020; Wiseman, 2023). In contrast, we study buyer-seller search, which has received much less attention, especially in the context of agricultural markets (with the exception of Goyal (2010)).

Most of the previously mentioned experiments with smallholder farmers are randomized at the individual or village level, while prices are often determined in equilibrium across a wider area, complicating inference (Svensson and Yanagizawa, 2009; Svensson and Yanagizawa-Drott, 2012). We randomize across much larger units and interpret the experimental variation through the lens of a model that accounts for indirect effects on control markets. In doing so, we contribute to a growing development literature on accounting for equilibrium forces within experiments (Muralidharan and

[1] While some role for scale economies is needed to account for all the empirical patterns we observe, we note that this is not inconsistent with the fact that farmers likely also face other market failures, such as credit constraints or uninsured risk. These may interact with fixed costs and respond to different types of policy interventions.

Niehaus, 2017; Fink et al., 2020; Egger et al., 2022; Bergquist et al., 2025; Cai and Szeidl, 2024). We also connect to a non-experimental literature on the influence of trade costs on agricultural market integration and outcomes for small farmers in spatial equilibrium (Gollin and Rogerson, 2014; Shamdasani, 2021; Allen and Atkin, 2022; Kebede, 2024). By randomizing the introduction of Kudu, we join a small set of papers that randomize a component of route-specific trade costs (e.g., Brooks and Donovan (2025)) or combine experimental variation with spatial models (e.g., Kreindler (2024)).

We also contribute to a broader trade literature on search frictions in LMICs. Following the early work on the spread of mobile phone access (Aker, 2010; Jensen, 2010; Allen, 2014), a recent literature has explored non-price search frictions. Vitali (2024) looks at how consumer search affects the locations of garment firms in Uganda, Cai et al. (2024) study search for input suppliers in Chinese writing brush supply chains, and Startz (2024) measures product search frictions facing Nigerian importers of differentiated goods. Dillon et al. (2024) and Rudder and Dillon (2024) study the experimental dissemination of "yellow pages" telephone directories, and find that they increase SMEs' sales but do not necessarily lead to new links with other firms. Bai et al. (2021) document the importance of congestion in search on an e-commerce platform.

Finally, we build on a nascent literature that investigates how scale economies shape market participation and outcomes in LMICs. Bassi et al. (2022) and Kaboski et al. (2024) respectively study Ugandan SMEs and entrepreneurs, highlighting the importance of fixed cost investments, which when divisible can shape inter-firm relationships and when indivisible can affect firm risk-taking. Foster and Rosenzweig (2022) explore the relationship between farm size and machine capacity in determining productivity in LMIC agriculture. Grant and Startz (2024) document the role that fixed costs of trade play in determining the length of distribution chains in Nigeria. We similarly find that economies of scale affect which agents participate directly in cross-market trading opportunities. However, we also show that equilibrium effects can still generate gains for smaller market participants.

The rest of the paper is structured as follows: Section 2 discusses the setting and study design. Section 3 presents reduced-form effects. Sections 4, 5, and 6 present the structure, estimation, and results of the model respectively. Section 7 discusses implications for engagement by traders and farmers, and for full scale up outside a research context. Section 8 concludes.

2 Setting and Study Design

2.1 Sample and Data

The study took place in Uganda across 110 subcounties – administrative units that typically each contain about 30 villages – that were selected by our implementing partners as promising for platform roll-out. Maize is the most commonly grown and consumed crop in Uganda, as well as the one most traded on the platform, and so we focus most of our analysis and discussion on it throughout the paper; however, we will also briefly discuss impacts on other crops, as in principle the platform was open to a wide range of agricultural goods. Our study areas are spread across regions of the country and include a mix of major surplus and minor deficit areas for maize production (Figure A.1).

The study ran for three years, starting in 2015 and concluding in 2018, and spans six major agricultural seasons. Figure A.2 presents a timeline for the project. During this period, we gathered three core types of data, illustrated in Figure A.3: high-frequency market data, trader survey data, and farmer survey data. All impact analysis is based on these three data sources, although we also analyze administrative data from the platform to understand take-up and usage.

We began with a census of all markets within our study subcounties that were permanent (i.e., not meeting only on specific days of the week) and featured buying and selling of maize (as opposed to wet markets where only fruits and vegetables are sold). The 236 markets identified in this census were further classified as "hubs" (major local commercial centers) or "spokes" (more remote local markets). We gathered information biweekly (every two weeks) by calling an informant in each market, typically a trader whose store was based there. The resulting high-frequency market data captures the buying and selling prices, availability, and average quality of four major food crops (maize, nambale beans, matooke bananas, and tomatoes). Reporting was consistent, with 88% of the attempted (market by survey round) observations successfully recorded.

Our second dataset comes from surveys with traders. We first conducted a census of traders based in study markets who bought and sold at least one study crop. For markets with fewer than 10 such traders, we surveyed all traders; for markets with more than this, we randomly sampled 10 traders. Traders were administered a baseline survey prior to treatment in 2015, a midline survey in 2016 after one year of treatment, and an endline in 2018 after three years of treatment. For the trader midline, we were able to survey 1,358 of the 1,457 baseline traders (93.2%). The trader endline originally located 1,248 traders (85.7%), after which we randomly sampled 20% of attritors (41 individuals) for intensive tracking. We successfully located 37 of these (92.7%), bringing the weighted tracking rate for the trader endline to 98.6%. The analysis is weighted to make it representative of all traders in study markets. We also conducted a census of traders present in the markets at endline to study entry into and exit from the trading industry as a whole.

Finally, we surveyed a sample of agricultural households. We first listed all villages located in study subcounties. For each market in our study, we selected the village containing the market (which is typically 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 sampled 8-9 farming households in each market village and another 4 in each rural village. We imposed two eligibility criteria: the household had to (i) be engaged in agriculture, and (ii) have sold some quantity of any of the four crops included in the study (maize, beans, bananas, and tomatoes) in the previous year. In practice, these criteria excluded few households, with over 90% qualifying for study inclusion. Study households completed a baseline survey in 2015 and an endline survey in 2018, covering agricultural production, marketing, and sales. The household endline originally located 2,744 of the 2,971 baseline respondents (92.4%). We then randomly sampled 17% or 39 households for intensive tracking, of which 31 were successfully tracked (79.5%), giving us a weighted household tracking rate of 98.7%. Farmer analysis is weighted to make it representative of all farming households in the sampled study villages.

2.2 Production, Marketing, and Prices

As is common in the region, the farmers in our study area are smallholders and do very limited marketing of their crops. They are small net producers on average, though our sample contains net consumers as well. Farmers grow an average of 1,252kg/year of maize across two harvest seasons. The typical farmer sells 67% of their harvest, making two sales per year (one per season). Farmers sell very close to home, with 66% selling exclusively at the farm-gate. The remainder sell in nearby markets, averaging just 2km away from their farm-gate. Farmers typically sell to a single trader a year, with whom they have transacted in the past.

In contrast, traders operate on a much larger scale and do a great deal of active marketing. Although traders often trade in multiple crops, 93% of traders report maize as their most traded crop, which accounts for 13x the volume of beans, the next most commonly reported crop. The median study trader buys and sells 25 tons of maize per year, about 30 times the size of the median farmer. About two-thirds of traders (62%) use 5-10 ton trucks to move their inventory; the remainder use motorcycles or bicycles. Traders typically operate from a home base in a market where their store is located, making many small purchases within their subcounty, and selling downstream in fewer, larger transactions. The median market in our study area is home to six traders. Churn is low, with the median market seeing one new and one exiting trader over the three year study period. At baseline, traders bought maize at an average of 12.7 cents/kg and sold at 16.4 cents/kg, a buy-sell margin of 29%. From baseline monthly revenues of $2,243, traders report an average monthly profit of $297. By comparison, average total monthly household expenditure in our farmer sample is $65.

Mobile phones are an important technology used by traders, who report calling a seller, buyer, or other contact ahead of time to get information in 54% of their transactions. More traditional search-by-visiting is also common, with the remainder of transactions occurring when a trader goes to a market without calling ahead for any prior information. At baseline, about 50% of traders also report using radio broadcasts as a market discovery tool, and 10% some kind of SMS service. These services typically only offer simple price alerts; none to our knowledge provides direct connections to buyers or sellers. Use of these services is lower among farmers, with only 7% of farmers reporting listening to radio broadcasts for market information and 2% receiving information via SMS.

Market prices for maize show strong variation both over space and over time in our study area (see Table A.1). Figure 1 presents the absolute value of the gap in prices between each pair of markets in our sample (black solid line), based on our market surveys. The average price gap for maize is about 110 Ugandan shillings (UGX) per kg, or roughly 15% of average price. Unsurprisingly, price gaps rise with the distance between the markets. A major driver of these price gaps is transportation costs, which are among the highest in the world (Teravaninthorn and Raballand, 2009) in this region. Transportation costs alone, however, cannot explain the full gap in prices. Figure 1 also plots transport cost as a function of distance (dotted line), as reported in our trader surveys.[2] We observe price gaps that are roughly 50% higher than transport costs, suggesting pervasive violations of the law of one price.

2.3 Kudu: a Market-Matching Platform

The introduction and rapid spread of mobile phones across sub-Saharan Africa offered the promise of dramatically reducing search costs. Indeed, the rollout of cell-phone towers in the early 2000s 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 designed specifically for agriculture. The first generation of these initiatives focused on the dissemination of price information to farmers via mobile phone (Camacho and Conover, 2011; Fafchamps and Minten, 2012; Nakasone, 2014; Mitra et al., 2018; Hildebrandt et al., 2020).

In contrast to these price information interventions, Kudu facilitates matches between individual market participants. The system, which was developed at Kampala's Makerere University in partnership with researchers at Microsoft Research, is designed to make it easier for buyers and sellers of agricultural commodities to connect. Users can post asks (sale offers) and bids (purchase offers) 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.

The system then matches particular buyers and sellers as proposed trading partners. There are two processes by which buyers and sellers can be matched. The first, which accounted for 33% of successful matches, is an algorithm that clears the market each day, attempting to maximize what it calculates to be the global "gains" from matches. To do so, Kudu defines a scoring function for all possible matches formed by the asks and bids in the system for the same crop available at the same time. The score for a match of ask a and bid b is constructed as follows: $score_{ab} =$

[2] Traders reported transport costs and vehicle size used along each of their five most commonly traveled routes. 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. We use responses for the 93% of respondents whose main crop is maize.

$(p_b - p_a) * min\{q_a, q_b\} - d_{ab} * \lambda$, where p_a and p_b are the ask and bid prices respectively, q_a and q_b the ask and bid quantities, d_{ab} a measure of the km distance between the two parties, and λ an estimate of fuel costs per km, depending on road type.[3]

The second method, which accounted for 67% of successful matches, is a hand-matching process conducted by platform employees. Inherent to the fact that this was done manually, we do not have a precise formula to describe how employees constructed these matches; however, in Table C.1 we use Kudu administrative data, combined with our survey data, to shed more light. We see similar fundamentals considered by the algorithm — price gaps, quantities, and distance — are also predictive of manual matches. Further, Kudu employees may have been able to partially infer (through family name, location, or direct communication) other characteristics that may have mattered for a good match. For example, speaking the same language may facilitate communication between trading parties (over 40 languages are spoken in Uganda), while sharing ethnicity may facilitate trust; we see both are predictive of manual matches.

Users are then contacted by SMS to inform them that they have been matched and to provide them with the contact information of the other party. Buyers and sellers can then directly connect and arrange for the sale, although frequently a Kudu employee will help facilitate this communication, reaching out by phone to both parties to gauge their interest in the deal and coordinate next steps. At this point, any other relevant information can be shared between the potential trading partners, and all details can be (re-)negotiated, including price. If interested, buyers and sellers meet in person to exchange the goods and make payment. Kudu has no mechanism for processing payments, nor for enforcing that the terms of the original bid or ask are honored if a transaction takes place. See Appendix C.3 for further details.

We therefore interpret Kudu as lowering the cost of learning about a potential transaction partner. The platform itself screens for whether the match aligns in terms of availability (i.e., being ready to buy or sell a particular crop in a particular location on a particular day), quantity, and a variety of other features. It also provides contact information for this potential match that can be used to discuss any other factors that the parties may find relevant. We discuss how we think of these factors conceptually through our model in Section 4.

2.4 Randomization

The introduction of Kudu was randomly assigned at the subcounty level. At the time of our study, the typical subcounty in Uganda had about 30,000 residents. This is therefore an "at-scale"

[3]In the next section of the paper, we show strong patterns of treatment effects by distance. One may wonder whether these are mechanically induced by the weight placed on distance by the Kudu algorithm. However, as we show in Table C.1, even *conditional on being matched*, matches are more likely to be successful at shorter distances, suggesting the weight placed on distance by the matching procedure was non-binding (if anything, this suggests that Kudu's algorithm did not put enough weight on distance, given that its goal is to suggest trades that are likely to succeed). For a greater discussion of the role of distance in Kudu matching procedures, see Appendix E. For more technical details on the Kudu platform, see Newman et al. (2018).

randomization, conducted with the goal of being able to observe treatment-control differences in outcomes such as prices and trade flows that are determined in equilibrium at the market level. We blocked the randomization on whether the subcounty contained a hub market (17%) or not (83%), and then stratified by a subcounty-level price index (mean of the z-scores of the prices of each of the four focus crops at the markets in each subcounty).

The goal of the intervention was to make Kudu accessible to all farmers and traders operating within treatment subcounties. This included not only the individuals in our trader and farming household samples, but also non-sampled individuals in sample villages, as well as those in non-sample villages. This was accomplished through intensive promotion of the platform in treatment subcounties, both on the ground and via regular text messages. Staff were sent to treatment markets to serve as the on-the-ground agents promoting the platform to local farmers and traders. In addition, we promoted Kudu via text messages sent every two weeks.[4] These messages included an advertisement and information on how to trade on the platform, either by registering directly or by contacting a local Kudu staffer.[5] While nothing prevented traders or farmers outside of treated subcounties from registering to use Kudu, in practice we find that this rarely happened (as we discuss in Section 2.5 on take-up).

Table B.3 examines the balance of the market survey for the core variables measured (price, number of traders, and quality rating), while Table B.4 uses the market survey data in dyadic form and examines the baseline balance of the experiment on price gaps across markets. The experiment is well balanced at the market level. For the trader and household analysis, balance is analyzed using the sample still present at endline and is weighted using the sampling weights so as to mirror the structure of the outcome analysis. Table B.5 analyzes the baseline attributes of traders across seventeen different attributes and finds no evidence of baseline imbalance. Table B.6 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 chance. Figure B.1 and Tables B.1-B.2 show that attrition is low and similar across treatment arms.

[4]This message was sent to all treatment traders, as well as a randomly selected three-quarters of treatment farmer households. This farmer-level randomization, which was conducted at the household level (blocked on subcounty), was set up in order to generate exogenous variation in direct use of Kudu. The goal was to separately estimate the direct impact of using Kudu vs. the indirect impact of living in an area in which others were using Kudu and therefore being exposed to equilibrium effects of others' use, but not directly using the platform oneself. In practice, this second randomization offers a negligible first stage, as farmer take-up is close to zero. See Section 7.1 for further discussion.

[5]This message also included price information for maize, beans, matooke, and tomatoes in users' local market and main downstream hub market, as well as five randomly sampled treatment markets that rotated in each two-week round. However, as we describe in Section 3.4, we also conducted a separate experiment in which we sent only price information, and see no impact. It therefore appears that price information was not a core driver of our treatment effects.

2.5 Platform Usage

Over the three years of the study, Kudu received 29,308 unique asks (offers to sell) and 31,177 unique bids (offers to buy). Bids and asks posted on the platform climbed throughout the first year to reach a steady maximum of about 400 tons in bids and 200 tons in asks active per day. Figure C.1 shows the spatial distribution of asks, indicating study markets across the country posting upwards of 1,000 asks each. Among the individuals posting asks to sell on Kudu, 45% were sampled study traders, 14% were agents of an agribusiness firm contracted to promote Kudu, and 6% were sampled study farmers. For those posting bids to buy, the corresponding percentages are 48% and 11% for sampled traders and agents. Less than 1% of buyers are sampled farmers. The remainder of bids and ask came mostly from within the treated study areas, but from individual users who were not sampled in our data collection.

Tables C.2 and C.3 present take-up for traders and farmers by treatment status. Take-up of Kudu among treated traders was high, with 80% posting to the platform at least once and 22% successfully trading on the exchange. However, take-up was much lower among treated farmers, with only 26% of treated households ever posting to the platform and fewer than 2% successfully transacting on the platform. There was very little take-up of the platform in control: only 16% of control traders ever posted to the platform and only 3% traded on Kudu, while take-up among control farmers was even more negligible, with only 2% posting to the platform and close to 0% successfully transacting.

Therefore, almost all trade on Kudu was among treated traders, rather than among farmers selling their output directly. Tables C.4-C.5 present predictors of take-up for farmers and traders, respectively. We see that scale, in terms of quantity sold or traded, is an important predictor of adoption, a point to which we will return later.[6]

Overall, Kudu saw 7,300 tons of grain successfully transacted during the study period, worth a total value of 9.4b UGX (about $2.7 million USD). Figure C.2 shows the cumulative sales over the platform during the duration of the study. Although Kudu has the capacity to connect buyers and sellers of 76 types of crops and livestock, in practice, maize trade dominated the platform, accounting for 67% of asks and 5.6b UGX in trading volumes, swamping the value of all other crops put together (see Figure C.3). The remaining trade was spread across another nine crops, with the next most common being soybeans, beans, and rice, each of which account for much smaller volumes.

[6] For farmers, we also see evidence that the household's non-farm income is associated with a higher probability of trying Kudu, perhaps suggestive that access to capital increases adoption, though quantity remains a strongly significant predictor of take-up even when controlling for this. Interestingly, the other significant predictor of adoption is gender; female farmers are less likely to use Kudu. For traders, we see operating another business is associated with lower adoption, perhaps suggesting these traders have higher opportunity costs to trade. We return to this point on limited farmer adoption in Section 7.1 and explore it in great depth in Appendix K.

3 Reduced Form Impacts

We now turn to the reduced form effects of the platform on trade flows and market prices. We use the phrase "reduced form effects" to refer to the standard experimental measure of post-treatment differences in outcomes between treatment and control markets or trade routes (i.e. between pairs of markets). As we will address explicitly later in the paper, this may not capture the full treatment effect of the intervention, as prices and trade flows in control markets may be affected. However, this reduced form variation will be key to estimating our model and ultimately quantifying full treatment effects.

We focus primarily on maize, which accounted for the vast majority of platform activity. However, in Section 3.3, we explore impacts on other crops.

3.1 Route-level Effects

Figure 2 presents reduced form effects on several measures of trade between subcounties: whether any trade is occurring, the number of traders engaged in trade along the route, and the volume of trade flowing along the route. These outcomes are drawn from our trader survey data. Figure 3 presents reduced form effects on price gaps between markets, using data from our market-level price surveys. The top rows of Figures 2-3 present non-parametric local Fan regressions of each outcome on the distance between the subcounties, separately for treatment and control routes. Treatment routes are those between pairs of subcounties or markets in which both destination and origin are in the treatment group and can therefore possibly be connected via Kudu; control routes are all other routes. Distance is measured as the road distance of the shortest route connecting the two.

We begin by noting some important patterns observed on control routes. In the top left panel of Figure 2, we see that while the probability of any trade is high for short routes (i.e. those connecting subcounties that are close to one another), this diminishes rapidly with distance. The probability of any trade occurring is close to zero beyond 200km distance. Consistent with this, the number of traders serving a route (second column) also falls quickly with distance, as do total trade volumes (third column). Finally, price gaps are larger between markets located at further distances, as shown in Figure 3.

When we look at these outcomes along Kudu-treated routes, we observe a higher probability of any trade (first column of Figure 2) and a larger number of traders engaged in trade (second column). We also see larger volumes of trade flows (third column). These increases in trade flows come with reduced price gaps, as shown in Figure 3. Notably, these effects are concentrated among relatively short routes. Beyond about 200km, the introduction of Kudu does not appear to alter the pattern of trade.

To gauge the statistical significance of these differences, the bottom rows of Figures 2-3 present results from a randomization inference procedure, in which 1,000 placebo randomized treatment assignments are drawn. The "reduced form effect" – i.e., the difference between treatment and control line for the realized randomization in the top row – is shown in black. The gray lines show the same "effect" for each of the 1,000 placebo treatment assignments. Long and short-dashed lines indicate the 90th and 95th confidence intervals of these placebo "effects." Figures D.1 and D.2 present the resulting p-values from this randomization procedure. Consistent with the top row, which suggests treatment effects are concentrated among relatively nearby markets, we see statistically significant or marginally significant effects for routes at short distances.[7] Beyond 200km, we see that effects taper off to precise zeros.

We present the same results on trade flows and price gaps in regression form in Tables 1 and 2, respectively. We run the following specification:

$$Y_{dr} = \alpha + \beta_1 T_d + \beta_2 D_d + \varepsilon_{dr} \tag{1}$$

where Y_{dr} is the outcome of interest along route (market dyad) d in survey round r. Outcomes are regressed on T_d, a dummy for whether the route is treated and D_d, a measure of the shortest road distance between the pair of markets. Standard errors are clustered two-way by each subcounty (the unit of randomization).

The top panels (Panel A) of each table present results for the full sample of routes, while the middle (Panel B) and bottom (Panel C) panels break down results for short and long distance routes, respectively. Panels A of both tables suggest impacts across the full sample of routes that are consistent with those observed in Figures 2-3, with effects in the expected direction but which are only marginally or not significant. However, the middle panels document that effects are concentrated along short routes, with significantly higher probability of trade, a larger number of traders, and higher trade volumes (Panel B Table 1), as well as lower price gaps (Panel B Table 2), along routes less than 200km. In contrast, we see fairly precise zero impacts of the platform on trading outcomes (Panel C Table 1) and no significant effect on price gaps (Panel C Table 2) for far distance routes above 200km, as shown in the bottom panels. This pattern is confirmed in Table E.1, which presents a continuous interaction of distance and treatment. While many practitioners hope that platforms like Kudu will increase long distance trade (e.g., connecting farmers to urban markets), in reality, we find that the platform generates meaningful differences in trade flows and price gaps only among relatively nearby markets. This may be because the cost of trade is so high for far-apart markets that – even with Kudu – it continues to be unprofitable to trade between them. We return to the interpretation of these patterns in Section 5.1.

[7] Randomization inference p-values across nearby distances are well below 0.05 for any trade and number of traders, while they are as low as 0.15 for trade volumes and price gaps under this demanding specification, which requires significance at each distance. In Tables 1-2, which are better powered due to pooling the data, we see significant results for all outcomes at short distances.

3.2 Market-level Effects

We have shown that trade increased and price gaps fell between treated pairs of markets. We next turn to analyzing effects on "monadic" outcomes – those that are market-specific rather than "dyadic" or route-specific.

We investigate effects on market prices as well as outcomes for traders and farmers based in treated markets. Impacts are likely to be heterogeneous, and in particular, to differ between surplus markets (where we might expect treatment to generate increases in net outflows) and deficit markets (where we might expect the opposite) as trade increases and arbitrage improves. In the analysis that follows, we define the surplus status of each area based on the average quantity of maize sold per farmer (in tons), as measured in our farmer baseline survey. The average is just under 1 ton, with a standard deviation across markets of 0.86 tons.[8] We use the baseline survey measure, rather than the endline measure, to avoid treatment induced selection bias (though in practice, the two are strongly correlated, with a correlation coefficient of 0.7).

3.2.1 Price Impacts

Figure 4 presents reduced form effects on price levels in relative surplus versus relative deficit areas. First, we note in the left panel that, as expected, prices are higher in relative deficit areas and lower in relative surplus areas in the control group. However, we see a less steep relationship between prices and surplus-deficit status in the treatment group. Prices are relatively lower in treated deficit markets compared to control markets, and higher in treated surplus markets. The right panel presents this reduced form effect, along with the 90% and 95% bootstrapped confidence intervals. We see that prices are weakly lower in deficit areas and statistically significantly higher in surplus areas. These effects almost exactly offset each other for the median market.

Table 3 presents similar results in regression form. We see in Column 1 that the overall effect on average 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 4 shows that for the median trading center, the average price effect is roughly zero). Column 2 presents heterogeneity by baseline average marketed surplus, where we again see that prices are weakly lower in relative deficit areas and higher in relative surplus areas.

3.2.2 Trader Impacts

We saw in Section 3.1 that more traders are active along treated routes. Who are these additional traders? Table 4 documents that traders operating along treated routes are smaller on average than those operating on control routes, as measured by baseline profits, volume traded, and value

[8]The range spans just above 0 to 5.5 tons. The support is always positive, as the amount sold is bounded below by zero. This surplus measure is correlated with weather and soil conditions that determine the suitability of the climate for maize (see Table F.1).

traded (Columns 1, 3, and 5). Because this analysis combines pre-treatment measures of trader scale with post-treatment observation of whether they are actually trading on a route, it isolates *selection* effects over trader type, rather than treatment effects on trading volumes for a given trader. Columns 2, 4, and 6 show that, in particular, treatment allows a drop in the minimum size among traders active along a route. This is consistent with Kudu lowering the minimum scale needed to profitably operate on treated routes, and therefore drawing in smaller traders on those routes, a point to which we return in Section 4.

In Table 5, we turn to Kudu's effect on trader business outcomes. Columns 1-4 regress trader-level post-treatment outcomes on the trader's home market treatment status. We find that while treated traders' total quantity traded goes up modestly (Column 1), the buy-sell margin at which they trade goes down (Column 2). While these results are imprecisely measured, they are consistent with Kudu enabling traders to sell in additional markets and expand their trading volumes, but – as other traders are also exposed to Kudu, driving equilibrium price convergence – the reduction in price gaps between markets cuts into their buy-sell margins. Indeed, self-reported trader profits fall as a result of the intervention (Column 3).

One may wonder why traders use the platform, if it leads to a drop in their profits. The key distinction here is access to the platform at the market level versus individual take-up decisions. Individually, traders can only benefit from using the platform – it weakly lowers their costs of matching with buyers in distant markets. However, when a large number of traders are able to adopt the platform, the impact is to shift aggregate trade flows and market prices. The former, individual reduction in costs is the source of benefits for traders and drives adoption decisions. The latter, aggregate impact on market-level outcomes is a potential source of losses for traders, because on average price gaps between markets (which are the source of trader profits from arbitrage) fall. It is individually profitable for traders to adopt, regardless of what their competitors are doing.

These changes, however, are not sufficient to induce exit or entry into the trading industry as a whole (as opposed to serving particular routes). We see in Column 4 that there is no effect on the probability that treated traders are still in business post-treatment. We also re-conduct our trader census at endline, to measure whether treatment affects entry into the trading industry as a whole. In Column 5, run at the market-level, we see no effects on the number of new traders in business in treated markets.

In Appendix G, we explore several other dimensions of impacts on traders, on which we find null effects. For example, 64% of traders are themselves also engaged in farming and add some non-zero amount from their own harvest of maize to the amount they sell as part of their trading business. However, these volumes tend to be quite small, accounting for only 7% of the total volumes traded by the median trader who adds anything from his own harvest. In Column 1 of Table G.10, we test for treatment effects on this amount, but do not find any significant impacts. In Column 2, we explore treatment effects on trader storage behavior, testing whether Kudu affects the volumes of maize storage across 6-month seasons. We find no significant effect, perhaps unsurprising as traders do not store much; only 45% store anything across seasons, and those who do store only a small amount, about 15% of total volumes traded.

3.2.3 Farmer Impacts

We next turn to looking at outcomes for farming households. Consistent with the zero average effect on market prices documented in Section 3.2.1, we see in Table 6 that there is no statistically significant effect on average on total revenues, maize revenues, maize volumes sold, or prices received for farmers living in treated subcounties as a whole.[9] 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.

Although some practitioners might hope that Kudu would benefit farmers across the board, it seems more likely that, following impacts on price levels, outcomes should vary by the surplus-deficit status of the market. Recall that Kudu increased prices in surplus areas and reduced them in deficit areas. Figure 5 therefore displays treatment effects on farmer revenue by the relative surplus status of their home market. In the left panel we see that farmer revenue is higher in surplus areas in both treatment and control. However, the right panel shows that treatment significantly increases farmer revenues in surplus areas (and decreases revenues in deficit, although this effect is not statistically significant). This suggests that shifts in market prices can impact farming households, even though most did not directly use Kudu or engage in cross-market trade.

We also test whether Kudu affected production, either on average or differentially by surplus status. As shown in Col 1-2 of Table G.1, we find no significant impact of Kudu on harvest levels. Although point estimates go in the expected direction – positive in surplus areas and negative in deficit – they are far from significant. We also see no change in land, labor, or other intermediate inputs used (Tables G.2- G.3). Lastly, we see no shifts in production of other crops (Table G.6), nor in engaging in or income from non-agricultural activities (Table G.4). This suggests that the increase in trade we observe arises mainly from changes in the allocation of existing production, rather than through increases in productive specialization. Consistent with this, we observe an increase in the percent of production sold to the market in surplus areas, as shown in Col. 4 of Table G.1.

In principle, one might expect to see changes to production, given changes in prices. It is possible that there are small effects that we are simply not powered to detect, but also – consistent with the broader literature – that there is some stickiness that deters farmers from adjusting production in response to small price changes in a relatively short time frame. We explore these issues further in Appendix G, and see no clear evidence that production responses are correlated with farmer

[9] Note that farmer prices are available only for farmers who report sales, but we do not find a treatment effect on the probability of reporting a sale.

characteristics that might correspond to distortions. It is possible that with a bigger intervention or longer time-frame there might be more movement on these margins.

3.3 Effects on Other Crops

As described in Section 2.5, trade on Kudu was dominated by maize. However, there was more limited trade in nine other crops, with next most common being soybeans, beans, and rice. Of these three, we tracked off-platform trade flow for beans. In Appendix H, we replicate all of the maize analysis shown previously for beans.

Overall, we see treatment effects on trade flows, price gaps, and other outcomes for beans that are qualitatively similar to those for maize, but smaller and less precisely estimated, consistent with the smaller amount of beans trade. These results offer a reassuring validation of our findings for maize. They also imply that our main results, while an accurate description of the impact on maize trade and markets, may be a lower bound on the aggregate impact of the entire platform across all crops. However, because these other crops are traded in much smaller volumes and the magnitude of treatment effects appears to be much smaller, we believe our estimates focused on maize capture the majority of the quantitative impact of the platform overall.[10]

3.4 Unpacking Matching Frictions

Both the route- and market-level reduced form effects presented suggest that the intervention lowered bilateral trade costs and therefore improved arbitrage. In the following section, we will write down a model that is consistent with all the observed effects, in which Kudu reduced the cost of searching for potential transaction partners on treated routes by providing users with a quick, easy, and potentially profitable match.

The unique feature of Kudu is the reduced cost of forming new buyer-seller linkages. We also, however, provided randomized price information in a number of different ways during the three years of the study, and so have the ability to speak to the literature that has previously focused on the role of price information. Using these sub-experiments, we show that price information alone has very limited impact on market integration in our setting (perhaps because cell phones are already ubiquitous) and further positive evidence in favor of the role of buyer-seller matching frictions. These findings influence the way we model and interpret the treatment effects of Kudu.

We have two sub-experiments designed to test the role of price information alone. First, we randomly rolled out SMS-based price information alerts to a subset of control markets. We rolled in three control markets in each of the 12 market survey rounds between October 2016-March 2017

[10]Impacts on other crops could arise not just through direct use of the platform, but also through substitution across goods. Unfortunately, our experiment is not well-designed to formally identify cross-elasticities, as we lack independent variation across crops. However, we can explore the impact of treatment, which is predominantly in maize, on the production and consumption of other crops. We do so in Tables G.6-G.7, but see no adjustments on these margins.

(roughly the second half of the study period), and then, subsequent to the household and trader endline surveys, we rolled in an additional 36 control trading markets and observe a final four rounds of market surveys with this system in place. Because this roll-in did not include promotion of Kudu or any support in using the platform, its impact isolates the effect of price information alone. To analyze this price information-only experiment, we can then match the price gaps from the biweekly market data to the timing of the introduction of price information to the market. In Table I.2, we present the impact of this sub-experiment, comparing price gaps among markets rolled into the price-only intervention to pure controls, before and after the intervention using panel data with round fixed effects. Unlike the introduction of Kudu, which drove strong price convergence, here we find no evidence of a reduction in gaps in prices across markets that receive only price information (even for those at short distances). We also see no effects on trade flows at endline (see Table I.1, as well as Figures I.1-I.2).

A second sub-experiment shedding light on the role of price information comes from within-treatment variation. As described in Section 2, as part of the price information sent to treated traders and farmers in the main experiment, we randomly selected five treated markets in each round, and broadcasted information about prices in those markets to the entire treated network. Though more short-run, this intervention was quite powerful, as it sent the market's price information to thousands of traders and farmers simultaneously. However, this price-only intervention also had no effect on prices or dispersion in the featured markets (see Figure I.3).

These two sub-experiments suggest that price information did not drive the main treatment effects observed in the study. This is further corroborated by the fact that Kudu increased the flow of trade from net surplus to net deficit regions of the country, a form of heterogeneity that is relatively time-invariant and already well-understood by professional maize traders, the main users of Kudu. We also find no evidence that Kudu reduced high-frequency deviations from expected average prices, as one would expect if the type of search cost alleviated were price information (see Table I.3).

Instead, our evidence suggest that Kudu worked by reducing buyer-seller matching frictions, introducing users to specific other users who were available and willing to trade, and thereby reducing the cost of searching for potential transaction partners.[11] We now present a model with this logic at its core.

[11]In fact, many of the matches formed by Kudu were durable. We find that 43% of the traders who traded with their Kudu match report transacting again with that same individual beyond the initial transaction. These repeat transactions were substantial, accounting for 7 times the volume of the first deal. This is consistent with the fact that the total volume of trade generated by the intervention (based on trader surveys of trading patterns) is 4 times the volume observed in the administrative data from Kudu.

4 Model

The previous section showed that the intervention generated experimental variation in aggregate outcomes between treatment and control markets. However, these comparisons are not sufficient to correctly estimate the magnitude of treatment effects, because control markets are likely to have been affected indirectly via their trade connections. We can still make use of the experimental variation, but need to interpret it carefully. In this section, we describe a model that will inform how we separate direct effects on treated units from indirect equilibrium effects.

In order to capture the observed empirical patterns, the model must have two key features. First, it must allow for an extensive margin of trade between markets. At baseline, there are many market pairs that do not trade with one another at all, and we see that access to Kudu increases both the probability of any trade and the number of traders serving a given route. Second, there must be a role for costly matching between buyers and sellers, which is alleviated by Kudu.

We outline a model that has these two realistic features, while otherwise remaining as simple and comparable to standard trade models in the literature as possible. In particular, we make two types of simplifying assumptions in order to focus on forces that we are able to quantify using experimental variation. First, as we showed in Section 3.2.3, there are no significant treatment effects on production or non-farm income. Therefore, we model an endowment economy and an industry equilibrium in agriculture, rather than a full economy-wide general equilibrium. Second, we hold the number of traders fixed, because farmers very rarely use the platform or sell outside their local market, and the platform did not induce entry of new traders into the industry as a whole (as shown in Section 3.2.2). In both cases, our goal is to stay disciplined by the experimental results. While one could endogenize these additional margins of adjustment, we see little empirical evidence suggesting they matter for our quantification, and for that same reason, are unable to pin down relevant parameters using experimental variation. We provide additional discussion and evidence on these points in Appendix G.

4.1 Setup

There are J locations indexed by i and j. Each is populated by a continuum of households of measure Z_i, and a continuum of traders of measure N_i. We take N_i as exogenous, as we see no treatment effects on entry or exit of traders into the trading industry, in Table 5.

We also see no treatment effects on harvest levels, either on average or by surplus status (see Table G.1). For simplicity and to match this empirical reality, we therefore model an endowment economy. Households in i are endowed with quantity L_i of a homogeneous crop and D_i of an "outside" good whose price is normalized to one (i.e. a numeraire), which we can think of as non-farm income. Throughout, we denominate values in units of the numeraire, including trade costs and expenditure.

4.1.1 Households

Households maximize their utility, given by:

$$u(x_0, x_1, \omega) = x_0 + b * \mathbf{ln}(x_1 e^{\mu\varepsilon(\omega)})$$

subject to

$$y = x_0 + p(\omega)x_1$$

where x_0 is consumption of the outside good, x_1 is consumption of the crop, and b is a positive constant. Households have preferences over which trader they buy the crop from. This is captured by a match value $\varepsilon(\omega)$ with a particular trader indexed ω, which is drawn i.i.d from a Gumbel distribution with mean zero and unit variance. This represents non-price factors that buyers may care about, such as language, ethnicity, availability, personal characteristics, and location within the market (see Appendix C.1 for a discussion of which factors appear to matter empirically). Thus, while the crop itself is homogeneous, we allow for the possibility that sellers are imperfect substitutes from the perspective of buyers, and μ is a positive constant parameterizing the degree of substitutability. Households make a discrete choice over traders when buying.

In the budget constraint, y is household income, the price of the outside good is one, and $p(\omega)$ is the crop price charged by trader ω. We assume income is sufficient so that households consume a positive quantity of the outside good in equilibrium.

4.1.2 Traders

Traders buy the crop from households in their home market, and can resell anywhere, potentially in multiple locations.[12] They pay a variable cost of operation, and are heterogeneous in the level of this cost. We conceive of this operating cost broadly. It may reflect fundamental ability, but can also capture other underlying heterogeneity in a setting with many distortions, such as in the cost of working capital, access to labor, or opportunity costs. A trader indexed ω has operating cost $a(\omega)$, drawn from a Pareto distribution with shape parameter k and minimum a_L, so that the CDF is:

$$G(a) = 1 - \frac{a_L^k}{a^k} \text{ for } a \geq a_L$$

[12]This is informed by the empirical setting, in which traders purchase mostly within their subcounty, before selling onward to other markets. The majority, 60%, buy exclusively in their home subcounty, and for the rest, out-of-home market purchases are fairly small in volume, with an average of 76% of total purchases coming from within their home subcounty. Treatment does not change these stylized facts (see Table J.1). This may differ in other settings, in which case one would want to endogenize the choice of source location.

4.1.3 Trade and matching costs

Households sell their full crop endowment to traders based in their home location.[13] All agents are price-takers in this local wholesale market at a price p_i.

Traders can then choose to resell the crop in any destination locations by paying trade costs. Trading the crop between locations incurs a multiplicative variable cost, τ_{ij} and a fixed cost F_{ij}. The fixed cost may include information costs – as specified below – but may also include other, non-information related costs of market access that do not scale proportionally with quantity sold. The outside good is freely and costlessly traded.

Traders observe all prices and trade costs, but must also find specific buyers to sell to. We attribute part of the fixed cost of selling in a destination market, F_{ij}, to a search and matching cost, S_{ij}.[14] This represents the costs of the real search strategies traders report using, such as traveling to the market to find buyers or time spent making calls to find out who is buying on a given day. Once this cost has been paid, the trader is "available" to buyers in j, who can then freely observe their idiosyncratic match values and choose who to buy from. The measure of sellers available in equilibrium is Ω_j. Traders can match with multiple buyers in each destination, and in multiple destinations, and neither buyers nor sellers go idle for lack of a match.

The cost of a trader ω based in market i to buy and resell a quantity x in market j is therefore:

$$C\left(\omega, x_{ij}\right) = \tau_{ij} a\left(\omega\right) p_i x_{ij} + F_{ij}$$

4.2 Demand and supply

4.2.1 Demand

The form of utility implies that expenditure on the crop, e, will be constant ($e = b$) and does not depend on equilibrium income or match value with a particular seller. We can therefore simplify the household's optimization problem to one of choosing which trader to buy from, with indirect utility:

$$u\left(\omega\right) = y - e + e\left(\mathbf{ln} e - \mathbf{ln} p\left(\omega\right) + \mu\varepsilon\left(\omega\right)\right)$$

This discrete choice across sellers with logit match values gives rise to the same aggregate

[13]This is a simplifying assumption, but broadly consistent with the fact that most farmers both sell and buy maize, rather than storing the full consumption amount from their own harvest, as shown in Burke et al. (2019). This makes the household's farmer and consumer problems separable, without changing the intuition of any key results.

[14]We use this simple form of search costs both because our data is not well-suited to estimating parameters that would discipline more detailed mechanics (as we observe annual flows but not transaction-level outcomes, partners, or search behaviors), and because a fixed cost is a reasonable representation of common actual search costs, such as driving to a market to find out who is available to buy. This does not, in itself, assume what type of information is being accessed when the search cost is paid. Our interpretation that it is about partners rather than prices is driven not by model structure, but by the empirical evidence from Section 3.4.

demands as a standard CES utility function (Anderson et al., 1992) with elasticity of substitution $\sigma = \frac{\mu+1}{\mu}$. Therefore, at the market level, these preferences imply that aggregate demand for the crop in market i can be summarized by a CES price index P_i (defined in the following section) and constant aggregate crop expenditure, E_i.

4.2.2 Trader optimization

Facing these costs and demand, traders decide which destination markets to sell in, and what prices to charge. In any destination market j, a trader ω whose home market is i will optimally set price:

$$p_{ij}(\omega) = \left(\frac{\sigma}{\sigma-1}\right)\tau_{ij} a(\omega) p_i$$

where $\sigma = \frac{\mu+1}{\mu}$ is the own-price elasticity of demand, and $\frac{\sigma}{\sigma-1}$ is a constant (multiplicative) markup over marginal cost. Note that this implies that a trader will sell at a single price within a given destination market, but that prices will vary across traders within a market, and for a given trader across markets. They will then face demand:

$$x_{ij}(\omega) = \frac{p_{ij}(\omega)^{-\sigma} E_j}{P_j^{1-\sigma}}$$

and earn profits:

$$\pi_{ij}(\omega) = \left(\frac{1}{\sigma}\right)\left(\frac{\left(\frac{\sigma}{\sigma-1}\right)\tau_{ij} a(\omega) p_i}{P_j}\right)^{1-\sigma} E_j - F_{ij}$$

where E_j is aggregate expenditure on the crop and P_j is the price index in j.

A trader will serve any market where they can earn non-negative profits. There is a cutoff operating cost a^*_{ij} such that any trader in i with costs less than or equal to this level, $a(\omega) \leq a^*_{ij}$, will serve destination j, and otherwise will not:

$$a^*_{ij} = \left(\frac{P_j}{\left(\frac{\sigma}{\sigma-1}\right)\tau_{ij} p_i}\right)\left(\frac{\sigma F_{ij}}{E_j}\right)^{\frac{1}{1-\sigma}} \tag{2}$$

Given these individual trader choices, the price index in destination j is:

$$P_j^{1-\sigma} = \sum_i N_i \int_{a_L}^{a^*_{ij}} p_{ij}(a)^{1-\sigma} dG(a)$$

Note that our setup implies that traders may serve multiple destination markets – always sourcing in their home origin market – and that these decisions are separable across destinations.[15]

4.3 Equilibrium

4.3.1 Bilateral trade flows

We can characterize equilibrium trade flows in terms of three margins: whether there is any trade on an ij route, the number of traders serving the route, and the value of trade on the route. These map to the trade outcomes we explored in the empirical analysis in Section 3, and will provide the estimating equations for our structural analysis in Section 5.1.

It follows from Equation 2 that there will be trade on a given route if the trader with the lowest operating cost (a_L) can earn non-negative profits:

$$\left(\frac{1}{\sigma}\right)\left(\frac{\left(\frac{\sigma}{\sigma-1}\right)\tau_{ij}p_i a_L}{P_j}\right)^{1-\sigma} E_j > F_{ij} \tag{3}$$

Note that this does not depend on any trader-specific factors – only equilibrium features of the sending and receiving markets, bilateral trade costs, and the minimum of the trader operating cost distribution.

Again using the individual trader cost threshold, a^*_{ij}, we can find the number of traders and the value of trade on a route in equilibrium by integrating over the set of traders who have low enough costs to serve the route profitably.

The number of traders who will serve a route ij is:

$$N_{ij} = \begin{cases} N_i \int_{a_L}^{a^*_{ij}} dG(a) & \text{for } a^*_{ij} \geq a_L \\ 0 & \text{otherwise} \end{cases}$$

The value of trade on a route ij is:

$$M_{ij} = \begin{cases} N_i \left(\frac{\left(\frac{\sigma}{\sigma-1}\right)\tau_{ij}p_i}{P_j}\right)^{1-\sigma} E_j \int_{a_L}^{a^*_{ij}} a(\omega)^{1-\sigma}\, dG(a) & \text{for } a^*_{ij} \geq a_L \\ 0 & \text{otherwise} \end{cases}$$

Making use of the CDF of the bounded Pareto distribution from Section 4.1.2, this implies that, for $a^*_{ij} \geq a_L$ such that there is any trade:

$$N_{ij} = N_i - N_i a_L^k P_j^{-k}\left(\left(\frac{\sigma}{\sigma-1}\right)\tau_{ij}p_i\right)^k (\sigma F_{ij})^{\frac{-k}{1-\sigma}} E_j^{\frac{k}{1-\sigma}} \tag{4}$$

[15] The assumptions underlying separability across destinations do rule out some forces, including selling in multiple destinations on one trip. Because we observe annual trade flows rather than individual trips or transactions, we lack data that would allow us to discipline this type of trip-chaining empirically.

$$M_{ij} \quad = \quad N_i \left(\frac{\left(\frac{\sigma}{\sigma-1}\right)\tau_{ij}p_i}{P_j} \right)^{1-\sigma} E_j \frac{ka_L^{1-\sigma}}{1-k-\sigma} \left(\left(\frac{P_j}{a_L\left(\frac{\sigma}{\sigma-1}\right)\tau_{ij}p_i} \right)^{1-k-\sigma} \left(\frac{\sigma F_{ij}}{E_j} \right)^{-k} - 1 \right) \tag{5}$$

4.3.2 Equilibrium and welfare

Trader profits are rebated to households in each trader's home market. Aggregate income in i is therefore:

$$Y_i = D_i + p_i L_i + \Pi_i$$

where Π_i is the sum of profits earned by all traders based in i.

Equilibrium consists of prices p_i and cost thresholds a_{ij}^* such that crop markets clear in every location. The crop endowment equals aggregate purchases by traders in every origin market i:

$$L_i = \sum_j X_{ij}$$

where X_{ij} is the quantity traded on each route. Expenditure on the crop equals aggregate supply by traders in each destination market j:

$$E_j = \sum_i M_{ij}$$

where M_{ij} is the value of trade on each route.

In our analysis of platform impact, we will consider welfare (inclusive of rebated trader profits) given by:

$$W_j = Y_j - E_j + E_j \mathbf{ln} \left(\frac{E_j}{P_j} \right)$$

Our assumptions on utility, combined with CES aggregate demand, imply that welfare changes are driven by changes in the price level, changes in variety as traders enter and exit routes, and changes in income. They shut down reallocation of expenditure across goods.

5 Estimation

In order to use the model to guide the estimation of treatment effects, we first make more specific assumptions about the form of trade costs and how they may depend on access to Kudu. We then derive equations that correspond to the main outcomes in our data, and use them to estimate the parameters of the model using the experimental variation in trade costs.

5.1 Mapping the experiment to the model

Kudu enables a seller to match with buyers in a distant market by simply posting on the platform, which may substitute for other costly search strategies, such as traveling and making calls. In that sense, Kudu makes it cheaper for trading partners to reveal their match value $\varepsilon(\omega)$. We therefore interpret access to Kudu as affecting the search and matching component of the fixed cost of trade, if agents in both the origin and destination markets can use the platform.

We specify fixed trade costs as: $F_{ij} \equiv \exp\left(\phi + \theta \mathbf{ln} d_{ij} + \beta \mathbf{1}_{ij}^{K} - v_{ij}\right)$ where $\mathbf{1}_{ij}^{K}$ is an indicator variable equal to one if both i and j are treated. Fixed costs may also depend on distance between markets d_{ij} (e.g. due to the time needed to drive a truck from one place to another), as parameterized by θ, and an idiosyncratic route specific component $v_{ij} \sim N\left(0, \sigma_v^2\right)$.

Our primary parameter of interest is β, which captures the percent reduction in bilateral fixed costs driven by Kudu. Note that if the absolute value of the fixed costs of trade F_{ij} grows with distance (i.e., if $\theta > 0$, which is indeed what we will find empirically), this implies that Kudu induces a larger reduction in *absolute* fixed costs for markets that are further apart. Yet, this can still produce the reduced form pattern seen in the data, which is that effects on trade flows are concentrated among *nearby* markets. This is because the fixed costs of trade may be so large for markets far away from each other that – even with the greater absolute reduction in fixed costs induced by Kudu – the level of fixed costs still lies above potential variable profits for even the most efficient trader. This means that there may be no trade on the extensive margin along these routes, both at baseline and endline – and as a result, no treatment effects on trade flows.

Lastly, we specify variable trade costs as $\tau_{ij}^{\sigma-1} \equiv d_{ij}^{\gamma} e^{-u_{ij}}$, where γ parameterizes the dependence on distance between markets, and $u_{ij} \sim N\left(0, \sigma_u^2\right)$ is an idiosyncratic route-specific component.

5.2 Experimenting in Equilibrium

The advantage of our setting is that the experimental roll out of Kudu generates route-specific, exogenous variation in the search and matching component of fixed trade costs. Nonetheless, we face two types of challenges in using this experimental variation.

The first challenge is that, although Kudu only changes trade costs on treated routes, we expect this to lead to equilibrium changes that affect control markets and routes as well. To see why, recall that prices in each destination market depend on the number of traders from all other markets that sell in that destination. This is determined by route-specific cutoff operating costs, a_{ij}^{*}, which in turn depend on not only route-specific trade costs, but also on prices in source markets. This means that experimentally-driven changes in treated source market prices will change trade on routes linking treatment and control markets (even though trade costs don't change), and therefore change control market prices.

These equilibrium impacts on prices imply a Stable Unit Treatment Value Assumption (SUTVA) violation – we cannot obtain a correct estimate of the treatment effects of Kudu simply by comparing treated and control markets or routes. Importantly, this is not the type of spillover that could be avoided with tighter control of compliance with treatment assignment, or by randomizing over larger units. Rather, it is a challenge inherent to experimentation with trade costs, which by nature typically involves linkages between treated and control units. To address this, we combine two strategies: we estimate model-implied equations with market fixed effects, and we make use of experimental variation in exposure to treated markets.

The second challenge for our estimation is the importance of the extensive margin of trade in our setting. There are many zeros in the subcounty trading matrix, and large differences in the number of traders serving different routes even conditional on there being any trade. This means that the log-linear estimating equations typically derived from trade models may suffer from both a selection bias in the set of routes with non-zero trade, and a bias due to heterogeneity in the composition of traders serving each route. As Kudu reduces fixed costs along treated routes, traders with higher cost draws $a(\omega)$ – who are smaller and have lower variable profits – will find it profitable to trade on those routes, consistent with the findings presented in Table 4.[16] We adjust our estimating equations, below, to account for these potential selection biases using a method similar to Helpman et al. (2008).

5.3 Estimating Equations

Following the model and specification of trade costs, we derive three estimating equations. Each corresponds to an outcome observed in Figure 2 and therefore maps to the reduced form variation driven by the experiment. More detailed derivation steps between the model equations from Section 4 and the estimating equations below are shown in Appendix N.

5.3.1 Any trade

Equation 6 is derived from Equation 3 combined with our assumptions about the specification of trade costs (and in particular, that idiosyncratic route-level trade costs are normally distributed). It describes the probability that there is any trade observed between market i and market j:

$$\mathbf{Pr}\left(T_{ij}=1\right)=\Phi\left(\hat{\zeta}_0+\hat{\zeta}_j+\frac{(1-\sigma)}{\sigma_\eta^2}\ln p_i-\frac{(\gamma+\theta)}{\sigma_\eta^2}\ln d_{ij}-\frac{\beta}{\sigma_\eta^2}\mathbf{1}_{ij}^K\right) \tag{6}$$

The likelihood of trade is higher when the wholesale price in the sending market is lower (mediated by how price sensitive final consumers are, via σ), when markets are closer together in distance (mediated by the elasticities of fixed and variable trade costs with respect to distance, θ and γ), and when both markets have access to Kudu ($\mathbf{1}_{ij}=1$) (mediated by the treatment effect of Kudu on fixed costs, β). The combined variance of the idiosyncratic components of trade costs, σ_η^2, determines how many routes are on the margin of profitability, such that treatment induces

[16] See Appendix C.2 for further discussion of selection by traders into new routes and into adopting Kudu.

them to begin trading. Finally, $\hat{\zeta}_j$ is a destination market fixed effect, which subsumes market level factors Equation 3, including the price index and total expenditure.

5.3.2 Number of traders

Equation 7 is derived from Equation 4 combined with our assumptions about trade costs, and describes the number of traders serving a route. To deal with the potential for selection bias described above, we redefine the outcome to consider the fraction of traders based in market i that do *not* serve market j. This rearrangement enables us to estimate the equation using a standard log-linear form, and avoids the need to account for selection on the extensive margin of there being an any trade between markets, since there are no zeros in the outcome variable.

$$\ln\left(1-\frac{N_{ij}}{N_i}\right)=\varphi_0+\varphi_j+k\ln p_i-\frac{k\left(\gamma+\theta\right)}{1-\sigma}\ln d_{ij}-\frac{k\beta}{1-\sigma}\mathbf{1}_{ij}^{K}+\varphi_{ij} \tag{7}$$

This states that the share of traders in market i that do not sell in market j is lower when the source market is cheaper ($\ln\, p_i$), the distance is shorter, and when the route is treated. The shape of the trader operating cost distribution, κ, governs how many new traders are pulled across the threshold of profitability when trade costs fall. Again, φ_j is a destination market fixed effect subsuming the market size and price index.

5.3.3 Value of trade

Finally, Equation 8 is derived from Equation 5, and describes the value of trade per trader flowing from market i to market j.

$$\ln\left(\frac{M_{ij}}{N_i}\right)=\psi_0+\psi_j+(1-\sigma)\ln p_i-\gamma\ln d_{ij}+B\frac{\phi\left(\hat{z}_{ij}^{*}\right)}{\Phi\left(\hat{z}_{ij}^{*}\right)}+\ln\left\{exp\left[\delta\left(\hat{z}_{ij}^{*}+\hat{\bar{\eta}}_{ij}^{*}\right)\right]-1\right\} \tag{8}$$

The last two terms address selection into this equation due to treatment-induced extensive margin effects via a Heckman-style correction following Helpman et al. (2008)). The second-to-last term accounts for the fact that as treatment reduces trade costs, the set of routes with any trade changes, while the last term accounts for the fact that the composition of traders serving those routes changes. The selection term is a function of $\hat{z}_{ij}^{*}=\Phi^{-1}\left(\hat{\rho}_{ij}\right)$, where $\hat{\rho}_{ij}$ is the predicted value from Equation 12, and B, which is a coefficient to be estimated. The composition term depends on $\delta\equiv\frac{\sigma_\eta(k-\sigma+1)}{\sigma-1}$, $\hat{z}_{ij}$, and the Mills ratio $\hat{\bar{\eta}}\equiv\frac{\phi(\hat{z}_{ij}^{*})}{\Phi(\hat{z}_{ij}^{*})}$.

The remainder of the equation states that, once we have controlled for those extensive margins, the intensive margin of trade only depends on variable costs, which are themselves a function of sender price and distance, and destination market fixed effects captured by ψ_j.

5.4 Identification

Our goal is to use Equations 6-8 to obtain unbiased estimates of the model parameters. In doing so, we face the two challenges previously mentioned: first, the SUTVA violation implied by the equilibrium effects of the intervention; and second, the biases created by selection on the route- and trader-level extensive margins of trade. The latter is addressed through rearranging the outcome variable in Equation 7 and the inclusion of additional terms in Equation 8, as detailed in the previous section. The former requires further discussion, which we turn to now.

To estimate parameters relevant to bilateral flows while controlling for equilibrium effects of the intervention, we include market-level fixed effects, following the standard "structural gravity" approach in the trade literature. Comparing estimating Equations 6-8 to their model counterparts Equations 3-5, note that outcomes determined in equilibrium at the destination market level – in particular, the price index P_j – are now subsumed in destination market fixed effects. Otherwise, these would be in the error term and correlated with both trade outcomes and with the route-level treatment status and would bias our main coefficients of interest. With the inclusion of these fixed effects, our specification uses within-market, route-level variation in treatment to separate the direct effect of Kudu on bilateral trade costs from the indirect effect on market-level prices.

If our only goal were to measure unbiased impacts of route-level treatment on route-level outcomes, we could include sending market as well as destination market fixed effects, which would also subsume the origin price, p_i. However, we would like to separately identify the price elasticity of demand, σ, so that we can not only estimate direct effects of treatment, but also quantify the indirect effects in counterfactual scenarios.[17] However, sending market prices may be correlated with unobserved idiosyncratic components of trade costs. Therefore, we need a source of exogenous variation in price.

Fortunately, the experiment offers a second source of random variation in addition to the route-level treatment status: subcounty-level exposure to other treated subcounties. Formally, we define:

$$\lambda_i \equiv \sum_j \mathbb{1}_{\mathbb{j}}^{\mathbb{T}} \frac{X_j}{d_{ij}}$$

where λ_i is the exposure of subcounty i, defined as the sum of exposure to all other treated subcounties j ($\mathbb{1}_{\mathbb{j}}^{\mathbb{T}} = 1$), weighted by subcounty j's market size (as proxied by j's total bilateral trade flows with all other markets, X_j) and by the inverse of distance between i and j (d_{ij}). As noted by Miguel and Kremer (2004), and formalized by Borusyak and Hull (2023), even if treatment is random, *exposure* is only random conditional on economic geography and other features of the baseline economic environment. Therefore, we follow Borusyak and Hull (2023) and run 1,000 placebo randomization draws. For each placebo draw, we construct the exposure measure for each

[17]While there is no particular reason to believe that the estimates of γ and θ will be biased in our framework, we are also less concerned about these, as distance is policy invariant.

subcounty. We then demean our realized exposure measure λ_i by the average of the exposure measure under the 1,000 placebo randomization draws, $\bar{\lambda_i}$, such that $\tilde{\lambda_i} \equiv \lambda_i - \bar{\lambda_i}$. The resulting measure $\tilde{\lambda_i}$ captures variation in exposure that arises purely from the realized randomization draw, which is exogenous.

In Table 7, we estimate the impact of exposure to treatment in other subcounties on own price using:

$$lnp_i = \alpha_1 + \alpha_2 \tilde{\lambda_i} + \varepsilon_i$$

Column 1 shows a significant impact on own price of randomized exposure to treatment in other subcounties.[18] In Column 2, we construct this randomized variation in exposure separately based on whether the treated subcounties to which one is exposed are themselves in surplus or deficit at baseline. We see that (random) exposure to treated surplus subcounties reduces own price, as demand is diverted to treated surplus subcounties and away from one's own subcounty. Conversely, (random) exposure to treated deficit subcounties increases own price, as supply is diverted to treated deficit subcounties, albeit this latter effect is not significant. In Column 3, we interact own treatment status with these randomized exposure measures. We see the diversion effects are most pronounced among controls, as evidenced by the significant coefficients on exposure to treated surplus subcounties and exposure to treated deficit subcounties.

These results are useful in two ways. First, they are evidence of the existence of spillovers, providing empirical validation for our theoretical concern about SUTVA violations. Second, they offer exogenous variation in price that we can use in our main structural estimation. We use the specification from Column 3 in Table 7 as an instrument for sending market price. The exclusion restriction is that subcounty i's exposure to other treated subcounties (demeaned by its expected exposure) is uncorrelated with the error terms in Equation 6-8, meaning it affects the trade flows between i and j only through the price in market i, conditional on the other controls included. Critically, these controls include a fixed effect for market j, which absorbs any potential confounding effects such as market j's realized exposure. Because we use the same instrument for three outcomes in a simultaneous equation system, the exclusion restriction for each equation requires that the other equations in the system are specified correctly, and in particular, that the corrections terms in Equation 8 correctly account for selection on the other two margins.

In sum, intuitively, a route's own treatment status identifies the direct treatment effect on trade costs (β), while randomized exposure identifies indirect effects that operate through impacts

[18]Standard errors are clustered by subcounty, the level of randomization. However, Borusyak and Hull (2023) suggest that this may be insufficiently conservative, as variation in exposure may be correlated at a geographic level greater than that of randomization. We therefore follow Borusyak and Hull (2023) and present randomization inference p-values in the notes of Table 7. Although the coefficient in Column 1 (for which theory has no predicted sign) is no longer significant under randomization inference, results from Columns 2 and 3 (for which theory does predict effects) remain statistically significant under this more rigorous standard.

on market-level changes in prices.

5.5 Results and Parameter Estimates

We estimate the model parameters jointly via generalized method of moments (GMM). The moment conditions are standard, specifying that the residuals from Equations 6-8 are orthogonal to the exogenous variables and the price instruments.[19] We calculate GMM standard errors with two-way clustering on sender and receiver sub-counties.

Table 8 presents results. The main parameter of interest is β, which describes the direct treatment effect on trade costs. We estimate a β of -0.20, suggesting that the direct impact of Kudu is to reduce the size of fixed costs by one-fifth, on average, across treated markets. This implies that Kudu drove sizable reductions in fixed trade costs, but was far from eliminating scale economies in cross-market trade.

6 Aggregate impact of Kudu

6.1 Model Simulation

To understand the impact of Kudu in equilibrium, we simulate the full model. The simulation uses the parameter estimates from the previous section,[20] data from our study area at baseline (measuring production, expenditure, and the number of traders in each market, plus the distance between each pair of markets), and random draws of the idiosyncratic parts of variable and fixed trade costs. The final component is an assignment of treatment status to each market, which we vary across simulations as described below to capture actual and counterfactual scenarios. These pieces tell us what trade flows will be for all pairs of markets as a function of prices. We then iterate to find a fixed point for prices in all markets in order to solve for the equilibrium.

This set up allows us to see what trade flows and prices look like in equilibrium under a variety of scenarios about access to Kudu. The first we consider is an "actual Kudu" scenario, in which the treatment status of each market is the same as in the real RCT. Using this simulation, we construct a model-based version of our reduced form estimates in Table 1, by comparing treated

[19] Parameters that appear in multiple estimating equations have more than one corresponding moment condition. Thus, we are over-identified. To increase computational efficiency, we use a nested-loop approach. In the outer loop, we search over values for the structural parameters (β, σ, γ, θ, κ, σ_η^2, B). Conditional on parameter values, we then solve for fixed effects that minimize the mean residuals for each destination market. To address the incidental parameters problem arising from the large number of fixed effects, we estimate our final fixed effects as a weighted average between the initial estimates for each destination market and the average across all markets (i.e. Bayesian shrinkage). The optimal weighting between the two is calculated by repeating the procedure five times across different subsamples of the data and choosing a weight that maximizes the out-of-sample fit on the excluded subsamples. We find the optimal weighting is 95% on the individual estimate and 5% on the mean estimated fixed effect.

[20] We also calibrate three parameters to approximately match the average level of trade outcomes: the minimum operation cost a_L, the intercept of fixed costs, ϕ, and the relative variance of the fixed versus variable parts of the idiosyncratic components of trade costs, σ_u^2 versus σ_v^2.

to control routes. Table 9 presents results. While the model over-estimates the overall volume of trade, we see a close alignment between the model and data-based reduced form treatment effects, both overall and by distance. We present additional model validation results in Appendix J.

To correctly estimate treatment effects, we make use of a second scenario – critically, one that we cannot observe in the data – in which *none* of the markets in the simulation has access to Kudu. In the following section, we describe how we use this counterfactual to calculate the true treatment effects of Kudu, accounting properly for equilibrium forces and impacts on control markets.

6.2 Treatment effects on trade flows

In the "no Kudu" scenario, everything is the same as in the "actual Kudu" scenario, except that none of the markets is assigned to treatment, and so no route has the reduction in fixed costs induced by Kudu. We solve for trade flows and prices under this scenario as we did with the first.

We can then estimate treatment effects using the ideal comparison: comparing each route in a world with Kudu to itself in a counterfactual world without Kudu. We conduct this analysis by pooling the simulated "actual Kudu" and "no Kudu" outcomes into a single dataset, and then running a regression similar in spirit to Equation 1, but replacing "Treatment" with an indicator for the "actual Kudu" scenario. Notably, we can make this comparison both for routes that are treated in the "actual Kudu" scenario and for those that are control. This allows us to calculate the impact on treated routes "(IoT)" and the impact on control routes "(IoC)."

Table 10 presents these impacts on our three main outcomes: whether there is any trade along a route on the extensive margin (Column 1), the number of traders operating along a route (Column 2), and the volume of trade flowing across the route (Column 3). Panel A presents the impact on treated routes (IoT). We see treatment effects are positive and substantial. Treated routes see a 13% higher probability of any trade, an 17% increase in the number of traders operating along the route, and a 7% increase in the volume of trade flows, relative to their mean. Yet, these impacts are smaller than the estimates one would get by taking the simple difference between the treatment and control routes, which we refer to as the ITT (Intention to Treat), shown again in the table notes for ease of comparison. Impacts on the extensive margin of trade fall from 14% of its mean to 13%, impacts on the number of traders fall from 20% of its mean to 17%, and impacts on trade volumes from 10% of its mean to 7%.

These overestimates of impacts on treated routes are driven by spillovers to control routes. For the standard reduced form analysis of the experiment to be valid, one typically assumes that impacts on control units are zero. However, we see in Panel B that the impacts on control routes (IoC) are negative, albeit small. While the impact on the probability of any trade along control routes is small and not significant (the point estimate suggests a 0.5% reduction), we see significant negative effects on the number of traders and volume of trade flows along control routes. Control routes experience a roughly 1% drop in both number of traders and in trade volumes relative to their means, suggesting trade diversion away from control routes.

Panel C presents average impacts across all routes (IoA), which is simply the weighted average of the positive impact on treated routes (IoT) and negative impact on control routes (IoC).[21] Across all routes, Kudu increases the probability of trade by 3%, the number of traders by 4%, and the volume of trade flows by 1% relative to the mean. These estimates suggest that Kudu produced meaningful increases in aggregate trade flows by reducing the fixed cost of search and matching.

However, these effects are smaller than would be suggested by the "naive" IoA, presented in the table notes of Panel C, which assumes the impact on treated routes is the ITT and the IoC is zero. Tables notes also present the magnitude of this mismeasurement, comparing the "naive" IoA to the true IoA, as a percent of the true IoA. We see that naive average impacts are 20% larger than actual impacts on the probability of any trade, 38% larger than actual impacts on the number of traders, and 155% larger than the actual impact on trade volumes.

Though small per route, the negative spillovers on control units add up in the aggregate, and make up a substantial part of the bias in the IoA. We return to this point in Section 7.2.

6.3 Price, Welfare, and Distributional Impacts

Correctly estimating treatment effects on trade costs and therefore on bilateral trade is an important first step in understanding the role of search costs. However, the objects of interest to policymakers are not trade or trade costs per se, but rather how those translate into market-level outcomes relevant to welfare. Our structural approach allows us to calculate treatment effects on these outcomes in equilibrium.

Figure 6 presents treatment effects on prices, producer surplus, consumer surplus, and total welfare in treated markets.

We turn first to the effects on prices (presented in gray dots), which vary as a function of the surplus status of the market. In deficit areas (to the left in the figure), Kudu reduces prices on average, while in surplus areas (to the right in the figure), Kudu increase prices on average. Although this heterogeneity in price effects by surplus status is not targeted directly in the estimation, the model nonetheless captures well these patterns present in the reduced form evidence.

As a result of the reductions in prices that Kudu induces in deficit areas, the platform improves consumer surplus (presented in the long-dashed line), but lowers producer surplus (presented in the short-dashed line) in these areas. Conversely, by raising prices in surplus areas, Kudu improves producer surplus, but lowers consumer surplus in these areas.

By reducing trade costs, Kudu therefore generates classic winners and losers from trade, with net producers in surplus areas and net consumers in deficit areas gaining, and net consumers in

[21] Recall treated routes are those that are connected by Kudu. As Kudu was randomly introduced in half of the study subcounties, only one-quarter (0.5*0.5) of the origin-destination pairs are connected by Kudu. Therefore, IoA = 0.25*IoT + 0.75*IoC.

surplus areas and net producers in deficit areas losing. However, because surplus areas contain more net producers and deficit areas more net consumers, average welfare improves in both locations. On net, we estimate that aggregate welfare improves by 0.026% in treated markets (and 0.015% across all markets, including control). Although these aggregate welfare effects are small – both because equilibrium effects on prices are small and because maize makes up only a fraction of total consumption – they are meaningful for a lightweight and low-cost intervention like Kudu.

7 Policy Implications

7.1 Implications of scale economies for take-up

Platforms like Kudu are often motivated by policymakers as a way to "cut out the middleman" and improve farmer welfare by connecting the rural poor directly to high-priced urban markets. We find, however, that usage of the platform was very limited among farmers, suggesting that a rethink of this logic is in order. In Section 4, we showed that if Kudu reduces the search and matching component of fixed trade costs, we should expect to see smaller agents able to serve more routes profitably. This is precisely what we see among traders in Section 3.2.2. So, why are farmers not similarly being enabled to engage in more trade directly? Our model quantification sheds light on one possible answer to this question.

Figure 7 shows the distribution of transaction sizes for farmers (black line) and traders (grey line) from the survey data. We see that traders are substantially larger than farmers on average, although the two distributions overlap. Using the model simulations from Section 6, we can calculate the minimum "threshold" size necessary to make trade along each route profitable, given the fixed costs – that is, the quantity sold in equilibrium by a trader with exactly the cutoff operating cost, a^*_{ij}. We take the average of these size thresholds across routes, and superimpose this average on the farmer and trader size distributions in Figure 7. The long dash shows the average threshold size in the absence of Kudu, while the short dash presents the – now lower – average threshold size with Kudu.

We note two implications of this figure. First, the decline in the average threshold driven by Kudu shifts the minimum size for transacting to the left. This moves it further into the trader size distribution, suggesting that Kudu should lead to greater entry among traders at the route level, which is exactly what we see in Table 10. Further, Figure 7 implies that these new entrants should be smaller than incumbent traders serving the route. This is indeed what we observe, as shown in Table 4.

The second notable feature of Figure 7 is that the Kudu-induced reduction in the average threshold size still lies far to the right of the farmer size distribution. Even with a 20% reduction in fixed costs, almost all farmers are still simply too small to make direct engagement in cross-market trade profitable on most routes. Indeed, we find empirically that the 2% of farmers who succeed in selling through Kudu are not only large but also tend to engage in very short distance trade for which the threshold is lowest. This suggests that, despite the commonly cited motivation of using technology to allow smallholders to directly connect to markets, scale economies make this a tall order. Fixed costs would need to fall much further to make this possible. Importantly, search for buyers is only one component of these fixed costs – they may also be driven by transportation or other factors – and so there may not be any information-related intervention that reduces scale economies sufficiently to engage farmers directly in cross-market trade.

In sum, one explanation for the lack of farmer take-up is that most farmers are simply too small – in terms of the quantity they grow and sell themselves – to have any rationale to adopt Kudu, even given a substantial reduction in the minimum threshold size required for cross-market trade. Consistent with this, only 2% of farmers trade on Kudu and they tend to be the largest farmers (see Section 2.5). Of course, farmers could increase their "scale" by buying additional quantity from other farmers to resell; this would be equivalent to entering into the trading industry, which we also do not observe happening in response to treatment, as shown in Table 5. In Table G.9, we also show that Kudu does not affect the number of traders to which farmers sold in the 12 months prior to endline, the number of *new* traders to which farmers sold, or whether farmers made any sales at market (rather than farmgate). Overall, Kudu does not appear to alter farmers' traditional marketing channels.

There may be additional market frictions that either limit farmers' own production or deter them from trading: for instance, farmers may be heterogeneous in their access to credit or their risk tolerance. We explore these in greater detail in Appendix K. Some of these frictions are isomorphic to high trading costs, $a(\omega)$, in our model. Moreover, while frictions may exist in this setting, our model quantification shows they are not necessary to explain the lack of farmer take-up.

However, given the substantial equilibrium effects of the platform, this does not mean that farmers are not impacted by Kudu. As traders take up Kudu and cross-market arbitrage improves, shifting market prices can impact farming households. Indeed, we see in both the reduced form results and our model quantification that prices rise in surplus areas, and we showed in Figure 5 that this translated to increases in farmer revenue. Although this does not operate through the mechanism that policymakers may have had in mind – that mobile marketplaces like Kudu will enable the smallest, poorest farmers to directly access a wider market and bypass intermediaries – it appears that farmers can benefit from the passed through effects of arbitrage by those intermediaries.[22]

[22] It is worth noting that some pass-through of market price shifts to the farmgate will occur under any model of competition between traders, ranging from perfect competition to monopsony. Of course, the *degree* of pass-through – a key feature governing the *magnitude* of farmer gains from a primarily trader-used intervention – will be mediated by the degree of competition (and the shape of farmer supply). This is something we aim to explore in future work.

7.2 Impacts under scaled implementation

Network interventions like Kudu would typically be implemented universally; markets function best when deepest and there may be a natural monopoly dimension to this type of trading platform. Hence the ideal policy comparison is the impact at 100% saturation relative to the pure control at 0% saturation, neither of which has any empirical counterfactual, as there are no control units in the former and no treated units in the latter. One can therefore only use model-based inference to examine these questions.

We conclude by examining this question in more depth. Several caveats must be made here, in light of the fact that we have shut down some potentially important margins of adjustment in our model on which we saw no empirical movement within our experiment, but which may emerge at scale, including production responses and entry into the trading industry as a whole. However, the goal of this exercise is primarily to illustrate how standard experimental inference can mislead, as a function of scale.

The standard experimental estimate at any saturation s is just the difference between the treatment and control at that saturation, which we refer to as ITT_s. In situations in which SUTVA holds, then $IoT_s = ITT_s$, $IoC_s = 0 \; \forall \; s$, the IoA_s at any saturation is simply $s * ITT_s$, and the average projected impact under universal treatment is the observed ITT. In this context, scaling simply means "treating more units". With equilibrium effects, the evaluation of program scaling is more complex. Now, the $IoC_s \neq 0$, the $IoT_s = ITT_s + IoC_s$, and the $IoA_s = s * IoT_s + (1-s) * IoC_s$, or $s * ITT_s + IoC_s$. This latter formulation is informative because it shows that the spillovers to the control lead to incorrect experimental inference about everyone in the study (the IoC_s is not multiplied times s). The "naive" IoA_s is the one we would calculate if maintaining SUTVA, namely $s * ITT_s$, which will not equal the correct IoA_s in the presence of spillovers. With impacts on both treatment and control units in equilibrium, the IoA_s becomes the obvious basis for evaluation of the total effect of the program.

Figure 8 shows what happens to our three main outcomes – any trade, number of traders, and trade volumes – as we shift the intensity of treatment in counterfactual scenarios from 0% to 100% of markets. The panels in the top row present the per-route impact on treated routes (IoT) in the red short-dashed line, the per-route impact on control routes (IoC) in the long-dashed blue line, and average impact across all routes (IoA) in the solid purple line. We see that the gains per treated route decline with the fraction of routes that are treated, while the (absolute) magnitude of the negative spillovers to each control route increases. This is because the IoT is driven by the direct reduction in trade costs and counteracted by the indirect equilibrium effect on prices, while the control group has no reduction in trade costs and is only affected by the price change. As saturation increases, the size of the change in trade costs enjoyed by each treated route stays the same, while the equilibrium effect on prices increases. The impact across all routes (IoA) increases with treatment intensity, reflecting a composition shift in which more routes are treated and therefore enjoy the Kudu-induced reduction in trade costs. With universal treatment, we see that Kudu induces an 11% increase (over the mean) in the probability of trade, a 14% increase in the number of traders, and a 4% increase in trade volumes (see Table L.1).

The second row of Figure 8 compares the naive experimental estimate of the IoA_s (dashed purple line) with the true IoA_s (solid purple line). Since the negative spillovers to the control (IoC_s) increase with saturation, the error in the naive IoA likewise increases with the fraction treated. While this bias always arises from the IoC_s, higher saturation mechanically increases the extent to which the total bias is driven by mis-estimation of the treated outcome versus the control, as illustrated by the red and blue shading, respectively.

In the third row of Figure 8, we directly represent the extent of mismeasurement in the two key estimands, the IoT_s and the IoA_s. The dotted red line shows mismeasurement in the naive ITT as a percentage of the true IoT, while the solid purple line shows the mismeasurement in naive IoA as a percent of the true IoA. IoT mismeasurement is low at low saturation, and generally slopes upward. This validates the informal reasoning that often motivates arguments to disregard the potential for bias due to SUTVA violations in small-scale RCTs: if the prevalence of treated units in the population is low, then spillovers will be small per-control unit that is designed to serve as a counterfactual for treated units. Because the per-unit IoC is negligible, so is the bias in the IoT, which is what we see in Figure 8. If the goal is to estimate an accurate IoT for an intervention that will not be scaled up to higher saturation levels, then this reasoning is sound.

Importantly, however, this logic does not necessarily follow for estimating aggregate impacts, even at a low saturation level. While spillovers on control units are small and diffuse at low treatment saturation, by definition there are many control units. The purple solid line shows that adding these up can lead to mismeasurement of the total impact of the intervention that is sometimes larger at low saturation. Even when the IoT is quite accurate and there is no intention to extrapolate to a program of a different scale or saturation level, evaluations of low saturation experiments can lead to very misleading conclusions about aggregate impact *even at the scale studied*. In contrast, in our setting, higher saturation experiments yield very biased IoT, but, in fact, more accurate evaluations of total impact (IoA), at least for the number of traders and trade volumes.

Finally, the bottom row of Figure 8 demonstrates the scaling bias directly. For each saturation, we conduct the naive exercise of extrapolating the IoA under universal implementation from the experimental comparison at that saturation, which is simply the ITT_s. This can then be compared to the true model-based IoA_{100}. The quantity represented on the Y axis is the percent over- or under-estimation in the naive estimate of universal implementation that we would derive from an experiment run at each saturation, namely $100 * (\frac{ITT_s}{IoA_{100}} - 1)$. The naive scaled estimates are generally too large, but the degree of bias in the naive estimates falls as the saturation goes towards 100%. Nonetheless, these biases are very large relative to any others we find; with trade volumes for example even the smallest bias suggests a 50% overestimation of treatment effects at universal implementation based on an experiment run at 90% saturation. This bias will be a product both of two factors: the fact that the true IoT changes with s and so experiments run at smaller scale lack external validity for treatment effects observable only at high saturations, as well as the bias to internal validity in the ITT_s arising from spillovers to the control at any saturation. The fact that the overall bias in the scaled estimate decreases at high saturations implies that the former effect dominates: experiments run at close to universal treatment do a better job of capturing the high-saturation treatment effect, and this effect dominates the larger spillover bias present in high-saturation estimates of the ITT.

8 Conclusion

This study shows that search and matching frictions continue to inhibit trade in African agricultural markets. The introduction of a trading platform to facilitate connections between buyers and sellers resulted in greater trade flows on both the extensive and intensive margins. These increases in trade reduced price dispersion, increasing prices in surplus areas and decreasing prices in deficit areas. Results are consistent with the platform having reduced the fixed cost of trade by 20%. This led to price changes that yielded benefits for net producers in surplus areas, net consumers in deficit areas, and increased overall welfare in the study area.

Accounting for equilibrium effects is key to correct estimation of the impacts of trade cost interventions, even those randomized "at-scale." Control markets and routes are affected by trade and price effects in equilibrium, and therefore no longer serve as valid counterfactuals for the estimation of treatment effects. However, experimental variation interpreted through the lens of a model can provide estimates of accurate treatment effects, accounting for equilibrium forces. In our context, aggregate impacts on trade flows are only a third of those suggested by naive "reduced form" comparisons of treatment and control routes. Though the details of the model must of course be adapted to each setting, the approach of estimating an equilibrium model using experimental results – and then estimating counterfactuals that allow for the recovery of treatment effects – is one that can be applied to a variety of contexts, even beyond interventions that are narrowly focused on trade cost reductions.

We also see future uses for such an approach that go beyond estimating the corrected treatment effect of the intervention run in the experiment. For example, with more structure on the cost-side, such an approach could also be used for counterfactual policymaking questions such as: for a fixed budget, which markets are optimal to target, given those that are on the margin of trade? What is the optimal geographic spread of markets to include, given the trade-off of including markets with greater heterogeneity in surplus status but which are also located at greater distances from each other?

Finally, our results shed light on scale economies in trade. We find the platform was used almost exclusively by intermediaries. These scale economies may make interventions targeted at engaging farmers directly in cross-market trade unlikely to succeed, on the margin. However, farmers and consumers can benefit indirectly from market price changes due to trader adoption of such platforms. We see evidence that revenues increase significantly for farmers in surplus areas. This suggests that future market access-based interventions need to think carefully about incidence, and how the most effective programs and policies might target groups that are not necessarily the main intended beneficiaries.

References

Abate, Gashaw T, Kibrom A Abay, Jordan Chamberlin, Yumna Kassim, David J Spielman, and Martin Paul Jr Tabe-Ojong, "Digital tools and agricultural market transformation in Africa: Why are they not at scale yet, and what will it take to get there?," *Food Policy*, 2023, *116*, 102439.

Abodi, Maurine Adhiambo, Gideon Aiko Obare, and Isaac Maina Kariuki, "Supply and demand responsiveness to maize price changes in Kenya: An application of error correction autoregressive distributed lag approach," *Cogent Food & Agriculture*, 2021, *7* (1).

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.

_ **and David Atkin**, "Volatility and the Gains from Trade," *Econometrica*, 2022, *90* (5).

Anderson, Simon P., Andre De Palma, and Jacques-Francois Thisse, *Discrete Choice Theory of Product Differentiation*, MIT Press, 1992.

Atkin, David and Dave Donaldson, "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs," *NBER Working Paper 21439*, 2015.

Bai, Jie, Maggie Chen, Jin Liu, Xiaosheng Mu, and Daniel Yi Xu, "Search and information frictions on global e-commerce platforms: Evidence from Aliexpress," *NBER Working Paper 28100*, 2021.

Barrett, Christopher B, "Smallholder market participation: Concepts and evidence from eastern and southern Africa," *Food policy*, 2008, *33* (4), 299–317.

Bassi, Vittorio, Raffaele Muoio, Tommaso Porzio, Ritwika Sen, and Esau Tugume, "Achieving Scale Collectively," *Econometrica*, 2022, *90* (6), 2937–2978.

Beaman, Lori, Dean Karlan, Bram Thuysbeart, and Christopher Udry, "Selection into credit markets: evidence from agriculture in Mali," *Econometrica*, 2023, *91* (5).

Bergquist, Lauren Falcao, Benjamin Faber, Thibault Fally, Matthias Hoelzlein, Edward Miguel, and Andres Rodriguez-Clare, "Scaling Up Agricultural Policy Interventions: Theory and Evidence from Uganda," *Working Paper*, 2025.

_ , Craig McIntosh, and Meredith Startz, "Data and Code for: Search Costs, Intermediation, and Trade: Experimental Evidence from Ugandan Agricultural Markets," Inter-university Consortium for Political and Social Research [distributor] 2026. `https://doi.org/10.3886/E251706V1`.

Borusyak, Kirill and Peter Hull, "Non-Random Exposure to Exogenous Shocks," *Econometrica*, 2023.

Brooks, Wyatt and Kevin Donovan, "Market Access and Economic Development: Experimental Evidence from New Bridges in Rwanda," *Working Paper*, 2025.

Burke, Marshall, Lauren Falcao Bergquist, and Edward Miguel, "Sell Low and Buy High: Arbitrage and Local Price Effects in Kenyan Markets," *Quarterly Journal of Economics*, 2019, pp. 785–842.

Cai, Jing and Adam Szeidl, "Indirect Effects of Access to Finance," *American Economic Review*, 2024, *114* (8), 2308–51.

_ , Wei Lin, and Adam Szeidl, "Firm-to-Firm Referrals," *Working Paper*, 2024.

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.

Carter, Michael, Alain de Janvry, Elisabeth Sadoulet, and Alexandros Sarris, "Index insurance for developing country agriculture: a reassessment," *Annual Review of Resource Economics*, 2017, *9.*

Dillon, Brian, Jenny Aker, and Joshua Blumenstock, "Yellow Pages: Information, Connections, and Firm Performanace," *Working Paper*, 2024.

Egger, Dennis, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker, "General Equilibrium Effects of Cash Transfers: Experimental Evidence from Kenya," *Econometrica*, 2022, *90* (6), 2603–2643.

Fafchamps, Marcel and Bart Minten, "Impact of SMS-based agricultural information on Indian farmers," *The World Bank Economic Review*, 2012, *26* (3), 383–414.

Fink, Gunther, B. Kelsey Jack, and Felix Masiye, "Seasonal liquidity, rural labor markets and agricultural production," *American Economic Review*, 2020, *100* (11), 3351–3392.

Foster, Andrew and Mark Rosenzweig, "Are There Too Many Farms in the World? Labor Market Transaction Costs, Machine Capacities, and Optimal Farm Size," *Journal of Political Economy*, 2022, *130* (3).

Gollin, Douglas and Richard Rogerson, "Productivity, transport costs and subsistence agriculture," *Journal of Development Economics*, 2014, *107*, 38–48.

Goyal, Aparajita, "Information, direct access to farmers, and rural market performance in central India," *American Economic Journal: Applied Economics*, 2010, *2* (3), 22–45.

Grant, Matthew and Meredith Startz, "Cutting Out the Middleman: The Structure of Chains of Intermediation," *Working Paper*, 2024.

Helpman, Elhanan, Marc Melitz, and Yona Rubinstein, "Estimating Trade Flows: Trading Partners and Trading Volumes," *The Quarterly Journal of Economics*, 2008, *123* (2), 441–487.

Hildebrandt, Nicole, Yaw Nyarko, Giorgia Romagnoli, 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.

Kaboski, Joseph, Molly Lipscomb, Virgiliu Midrigan, and Carolyn Pelnik, "How Important are Indivisibilities for Development? Experimental Evidence from Uganda," *Working Paper*, 2024.

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

Kebede, Hundanol A., "Gains from market integration: Welfare effects of new rural roads in Ethiopia," *Journal of Development Economics*, 2024, *168.*

Kreindler, Gabriel, "Peak-Hour Road Congestion Pricing: Experimental Evidence and Equilibrium Implications," *Econometrica*, 2024, *92* (4).

Miguel, Edward and Michael Kremer, "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities," *Econometrica*, 2004.

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.

Moser, Christine, Christopher Barrett, and Bart Minten, "Spatial Integration at Multiple Scales: Rice Markets in Madagascar," *Agricultural Economics*, 2009, *40* (3), 281–94.

Muralidharan, Karthik and Paul Niehaus, "Experimentation at Scale," *Journal of Economic Perspectives*, 2017, *31* (4), 103–24.

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.

Porteous, Obie, "High Trade Costs and Their Consequences: An Estimated Dynamic Model of African Agricultural Storage and Trade.," *American Economic Journal: Applied Economics*, 2019, *11* (4), 327–366.

Rudder, Jess and Brian Dillon, "Search Cost and Relational Contracting," *Working Paper*, 2024.

Shamdasani, Yogita, "Rural road infrastructure and agricultural production: Evidence from India," *Journal of Development Economics*, 2021, *152.*

Shoko, RR, A Belete, and P Chaminuka, "Estimating the supply response of maize in South Africa: A Nerlovian partial adjustment model approach," *Agrekon*, 2016, *55* (3), 237–253.

Startz, Meredith, "The Value of Face-to-Face: Search and Contracting Problems in Nigerian Trade," *Working Paper*, 2024.

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.

Vitali, Anna, "Consumer Search and Firm Location: Theory and Evidence from the Garment Sector in Uganda," *Working Paper*, 2024.

Wiseman, Eleanor, "Border Trade and Information Frictions: Evidence from Informal Traders in Kenya," 2023.

Figures

Figure 1: **Price gaps and transport costs.** The y-axis presents the absolute difference in prices across each route 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 estimated transport costs. To generate this prediction, we asked surveyed traders to report the costs of traveling along each of their five most commonly traveled 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 function of the km traveled. These transport costs represent an upper bound on the price dispersion that should be observed if transport costs are the only trade costs. The gray area represents excess price dispersion, the portion of price dispersion that cannot be explained by transportation costs.

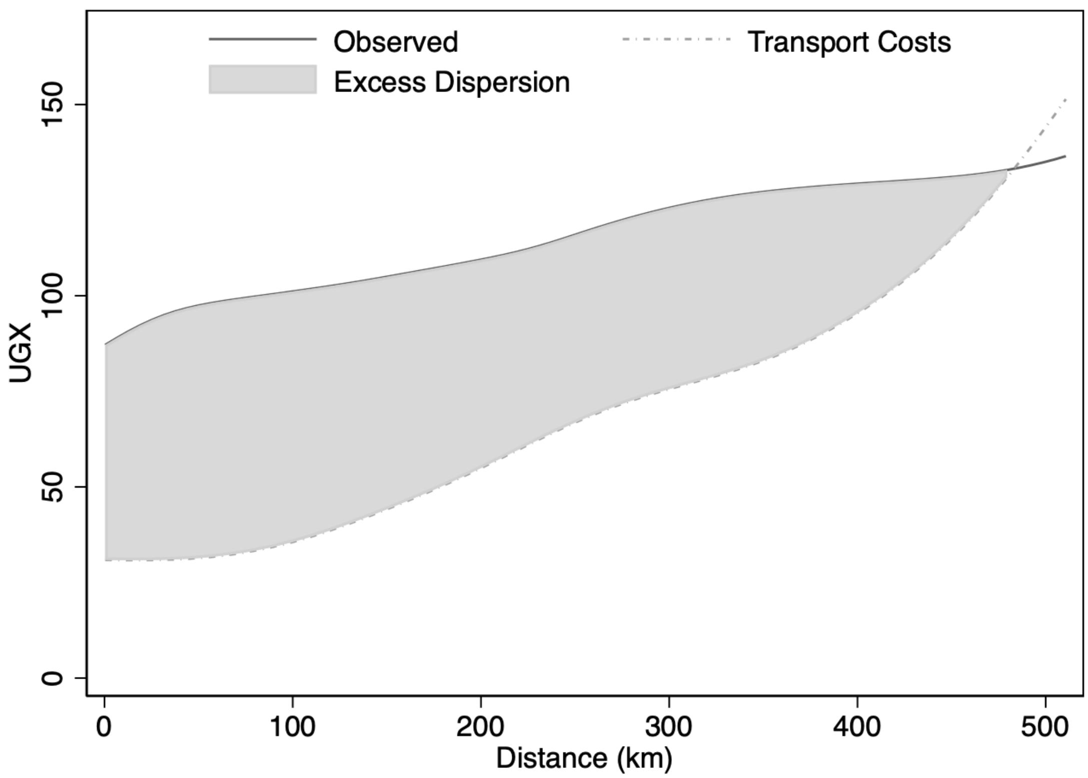

Figure 2: **Reduced form effects on trade flows.** The top lefthand figure presents a non-parametric local Fan regression of the probability of trade along any treated route (dotted line), in which both markets are treated and therefore are connected by Kudu, vs. control routes (solid line), as a function of distance. The bottom row presents results from a randomization inference procedure, in which 1,000 placebo randomized treatment assignments are drawn. The "reduced form effect," i.e., the difference between treatment and control line in the lefthand figure, from the realized randomization, is shown in black. The grey lines show the same effect for each of the 1,000 placebo treatment assignments. Long and short-dashed lines indicate the 90th and 95th confidence intervals of these placebo "reduced form effects." Subsequent columns present similar results for the number of traders trading along a route (second column) and trade volumes along the route (third column). See Figure D.1 for randomization p-values.

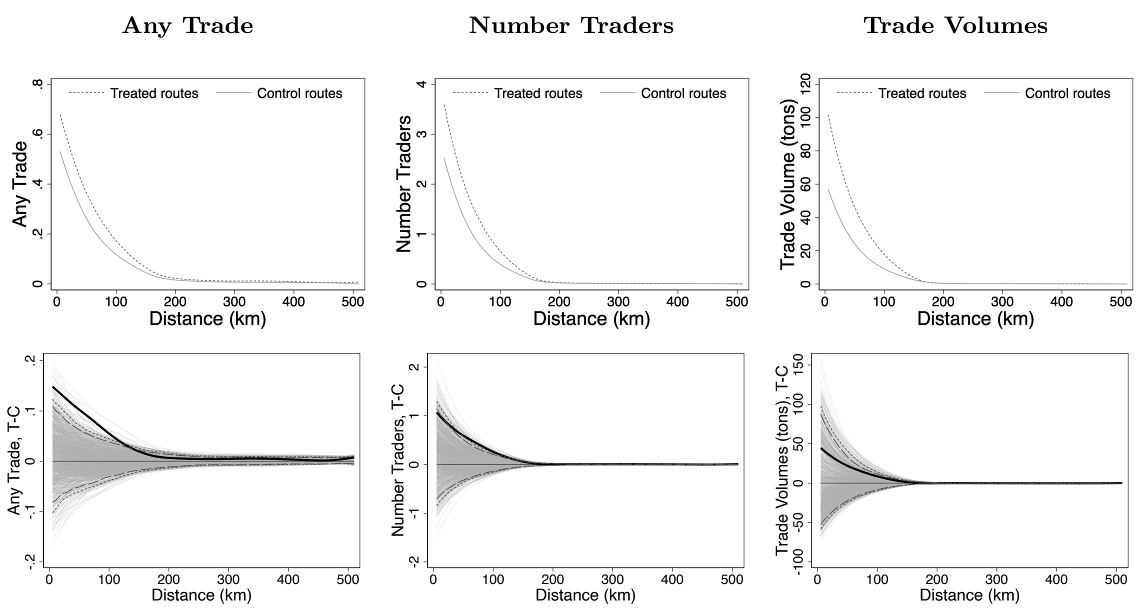

Figure 3: **Reduced form effects on price gaps.** The top panel presents a non-parametric local Fan regression of the price gaps across markets along any treated route (dotted line), in which both markets are treated and therefore are connected by Kudu, vs. control routes (solid line), as a function of distance. The bottom row presents results from a randomization inference procedure, in which 1,000 placebo randomized treatment assignments are drawn. The "reduced form effect," i.e. the difference between treatment and control line in the lefthand figure, from the realized randomization is shown in black. The grey lines show the same effect for each of the 1,000 placebo treatment assignments. Long and short-dashed lines indicate the 90th and 95th confidence intervals of these placebo "reduced form effects." See Figure D.2 for randomization p-values.

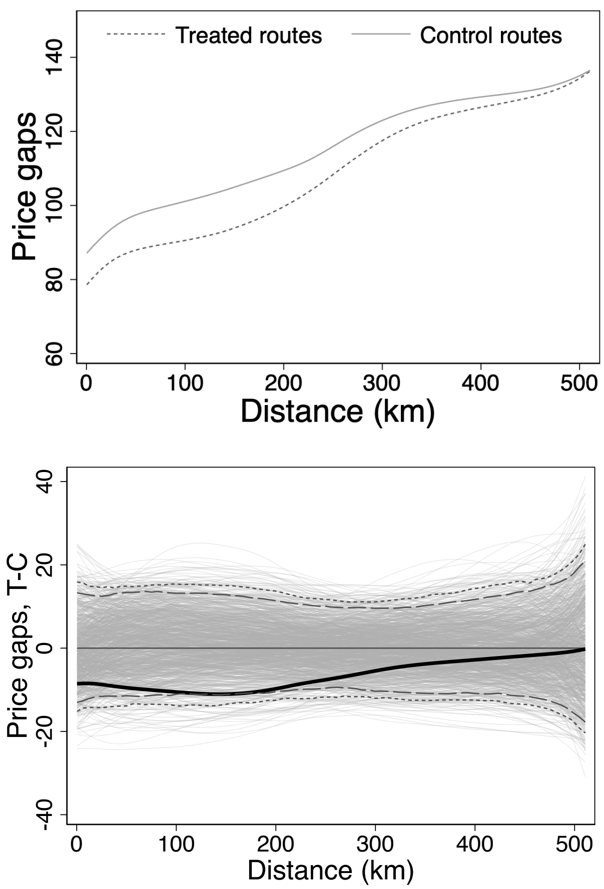

Figure 4: **Market price effects in surplus vs. deficit areas.** The left panel shows the level of market survey prices in treatment vs. control markets, with respect to the average market surplus per farmer, as measured in tons at baseline. The right panel shows the difference between the two (the treatment effect), along with the 90% and 95% confidence intervals from a bootstrap estimation and the density of observations.

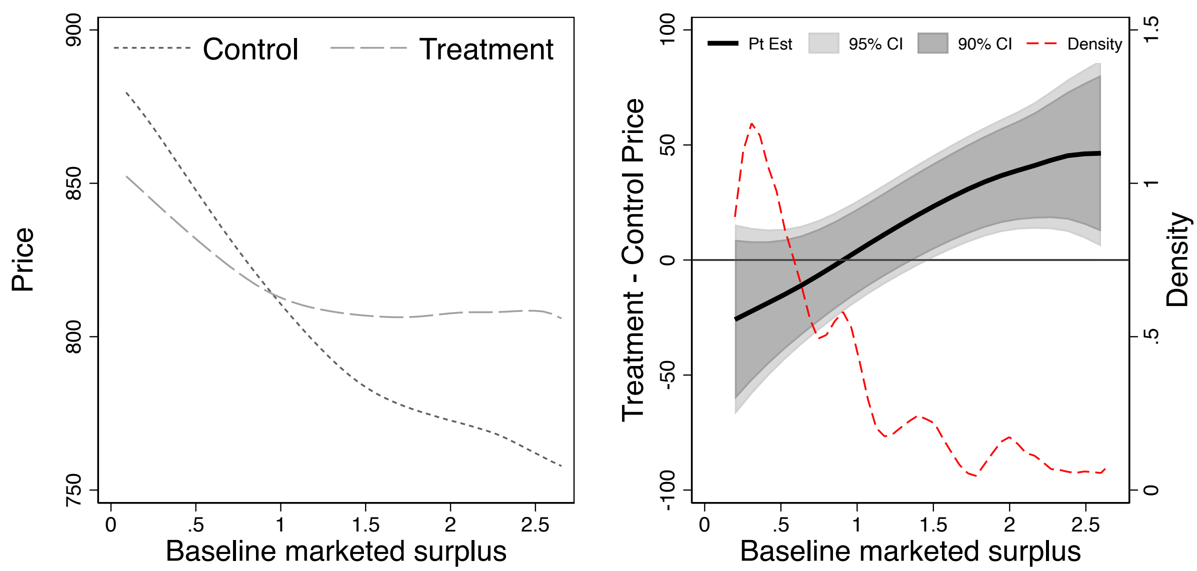

Figure 5: **Farmer revenue effects in surplus vs. deficit areas.** The left panel shows the level of farmer revenue from the household survey in treatment vs. control markets, with respect to the average market surplus per farmer, as measured in tons at baseline. The right panel shows the difference between the two (the treatment effect), along with the 90% and 95% confidence intervals from a bootstrap estimation and the density of observations.

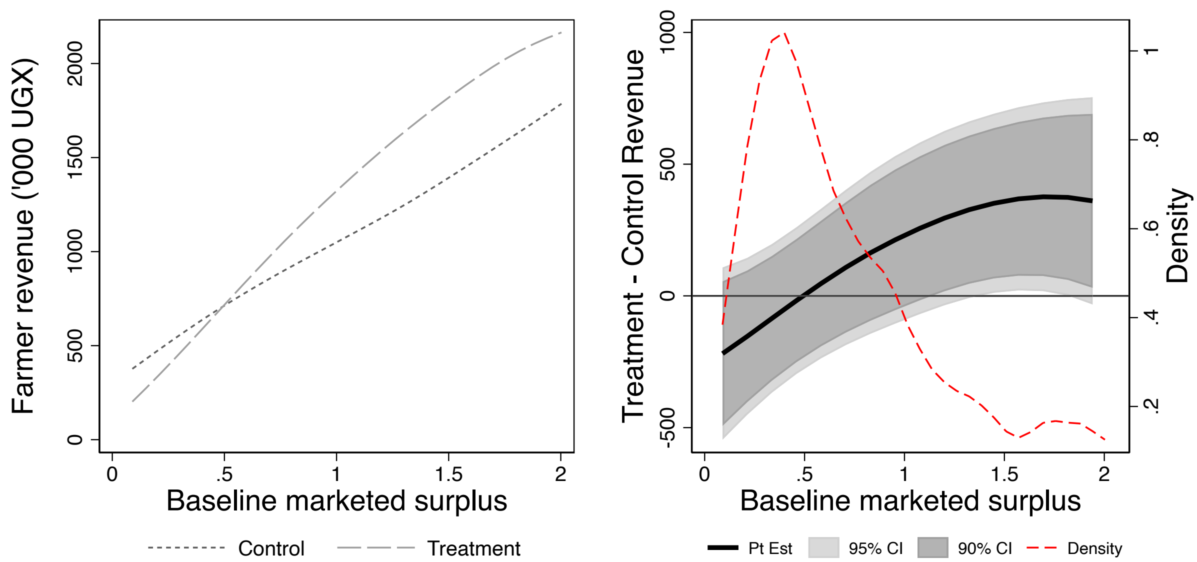

Figure 6: **Welfare effects.** Figure plots the total model-derived impacts of the intervention on producer surplus (short-dash line), consumer surplus (long-dash line), overall welfare (solid line) and prices (dots), as a function of standardized marketed surplus.

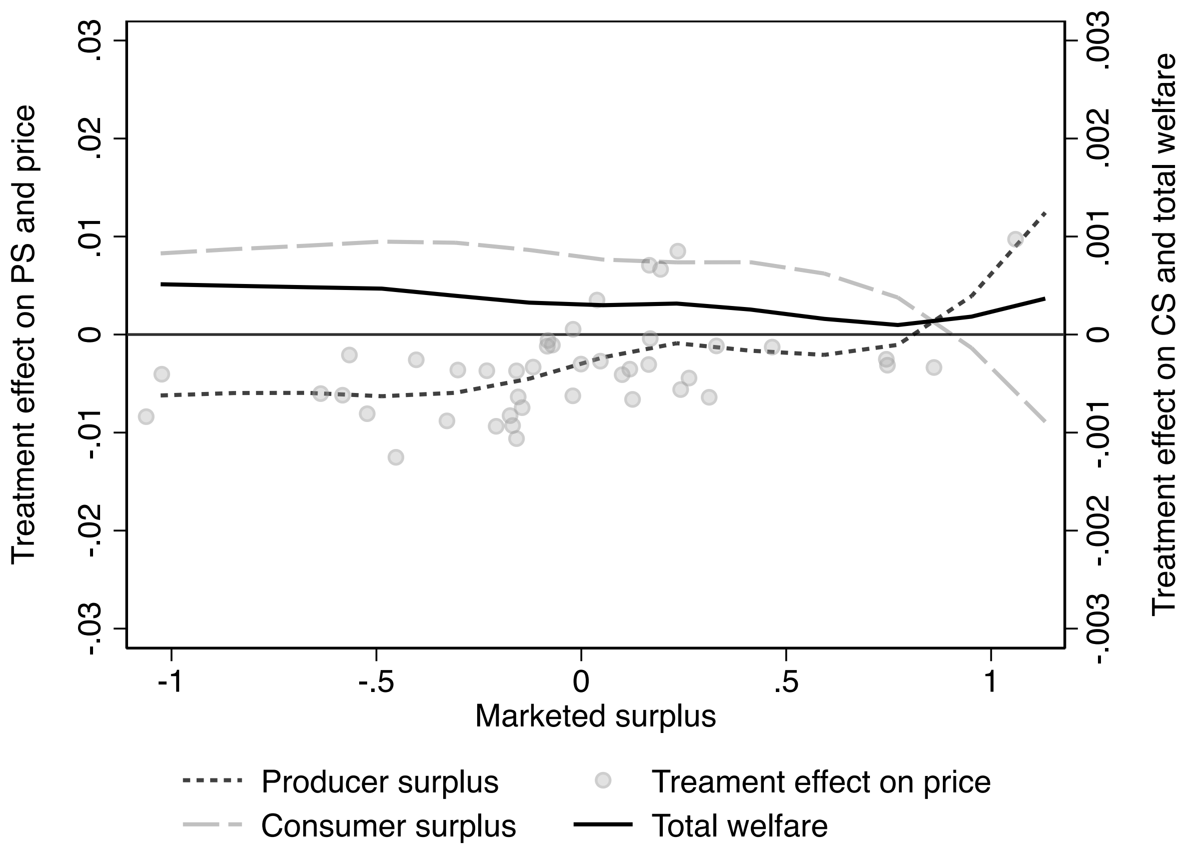

Figure 7: **Threshold size.** The densities represented in this picture are the log transaction size for farmers (left distribution) and traders (right distribution). The vertical lines represent the average across all routes of the model-derived threshold size for trade in the absence of the intervention (dashed vertical line) and in the presence of the intervention (dotted vertical line). This threshold size is the quantity sold in equilibrium by a trader with the highest operating cost, a_{ij}^*, that can still earn positive total profits on a particular route given the fixed cost.

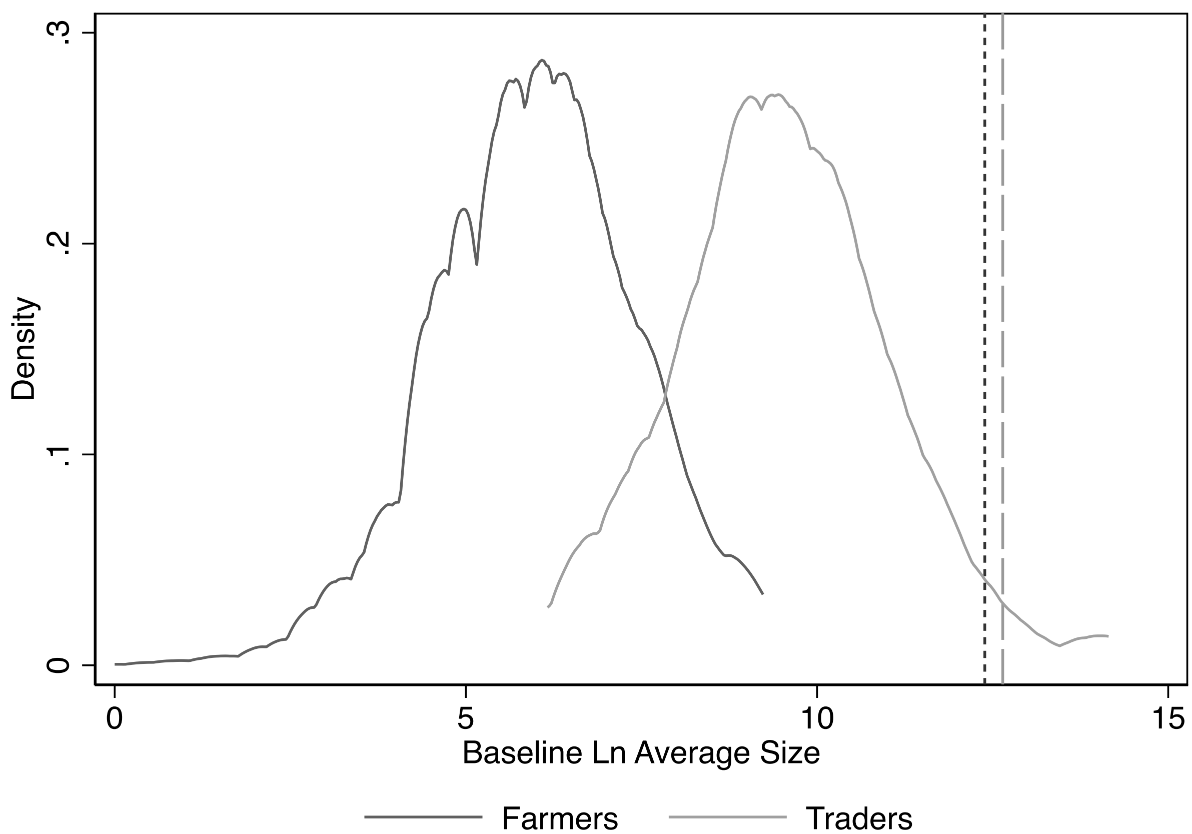

Figure 8: **Impacts by saturation rate.** The panels in the top row present the per-route impact on treated routes (IoT) in the red short-dashed line, the per-route impact on control routes (IoC) in the long-dashed blue line, and average impact across all routes (IoA) in the solid purple line, under counterfactuals representing various intensities of treatment, ranging from 0% to 100% (in steps of 10%). Outcomes are any trade along the route (left column), the number of traders trading along the route (middle column), and trade volumes along the route (right column). The second row of Figure 8 compares the naive experimental estimate of the IoAs (dashed purple line) with the true IoAs (solid purple line) under each counterfactual saturation rate. In red we show the IoA bias due to miscalculation of the impact on all treated routes and in blue the IoA bias due to miscalculation of the impact on all control routes. The third row presents the percent mismeasurement in the IoT in dotted red ($100 * \frac{naiveIoT - actualIoT}{actualIoT}$ and the percent mismeasurement in the IoA in solid purple ($100 * \frac{naiveIoA - actualIoA}{actualIoA}$. The bottom row calculates the naive estimate of the IoA at universal implementation that one would wrongly infer under the assumption that the ITT at a given saturation is the IoA at 100% universal implementation. The is presented as a percentage of the true model-based IoA at 100% saturation.

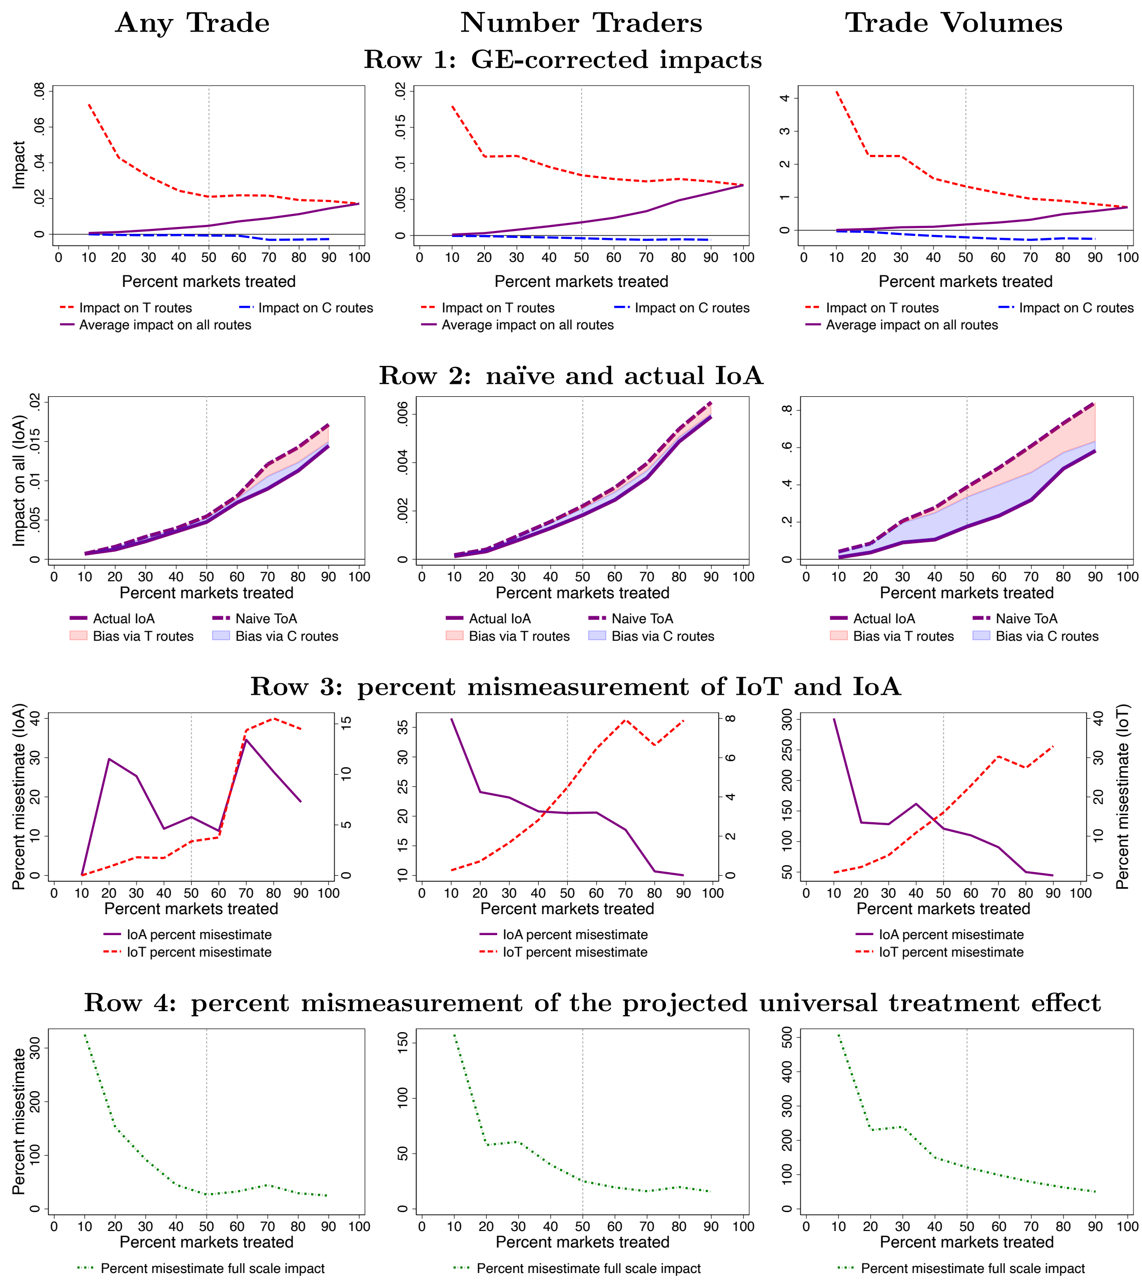

Tables

Table 1: **Reduced form effects on route-level trade flow outcomes.** Reduced form results on any trade along the route (first column), the number of traders trading along the route (second column), and trade volumes along the route (third column). Outcome is measured at a route level and regressed on a dummy for treatment (= 1 if both the markets are treated and therefore could possibly be connected by Kudu), distance between the pair of markets, and round fixed effects. Standard errors are clustered two-way by each subcounty (the unit of randomization). Panel A presents the results from the full sample, Panel B for nearby markets (routes less than 200km) and Panel C for far markets (routes more than 200km). "Mean DV" provides the mean of the dependent variable.

	Any Trade	Number Traders	Volume (tons)
		Panel A: All routes	
Treated route	0.02	0.08	2.86
	(0.01)	(0.05)	(1.76)
Dist (100km)	-0.05***	-0.21***	-5.10***
	(0.00)	(0.03)	(1.37)
Observations	11236	11236	11236
Mean DV	0.05	0.19	4.71
		Panel B: Routes less than 200km	
Treated route	0.06**	0.31**	10.79*
	(0.03)	(0.14)	(5.57)
Dist (100km)	-0.27***	-1.26***	-31.50***
	(0.02)	(0.16)	(9.14)
Observations	3429	3429	3429
Mean DV	0.16	0.61	15.00
		Panel C: Routes greater than 200km	
Treated route	0.00	0.00	-0.06
	(0.00)	(0.00)	(0.10)
Dist (100km)	-0.002*	-0.003*	-0.064
	(0.001)	(0.001)	(0.040)
Observations	7807	7807	7807
Mean DV	0.01	0.01	0.18

Table 2: **Reduced form effects on route-level price gaps.** Reduced form results on price gaps across markets. Outcome is regressed on a dummy for treatment (= 1 if both the markets are treated and therefore could possibly be connected by Kudu), the distance between the pair of markets, and round fixed effects. Standard errors are clustered two-way by each subcounty (the unit of randomization). Panel A presents the results from the full sample, Panel B for nearby markets (routes less than 200km) and Panel C for far markets (routes more than 200km). "Mean DV" provides the mean of the dependent variable.

	Price Gaps
Panel A: All routes	
Treated route	-6.36
	(5.09)
Dist (100km)	10.04^{***}
	(1.20)
Observations	1443093
Mean DV	116.30
Panel B: Routes less than 200km	
Treated route	-10.30^{*}
	(5.31)
Dist (100km)	7.63^{*}
	(4.21)
Observations	456507
Mean DV	97.43
Panel C: Routes greater than 200km	
Treated route	-4.29
	(5.62)
Dist (100km)	8.46^{***}
	(2.13)
Observations	986586
Mean DV	125.03

Table 3: **Reduced form price effects in surplus and deficit areas.** In Column 1, market-level data on price regressed on a dummy for treatment and round fixed effects. Column 2 interacts treatment with a dummy for being in a baseline marketed surplus subcounty (defined as average market surplus per farmer over 1 ton). Standard errors clustered by subcounty, the unit of randomization. "Mean DV" provides the mean of the dependent variable.

	(1)	(2)
Treatment	-4.801	-11.58
	(11.56)	(13.74)
Surplus		-48.60***
		(9.868)
Treat*Surplus		29.39*
		(15.10)
Observations	15211	15211
Mean DV	831.6	831.6
R2	0.844	0.847

Table 4: **New entrants on average smaller.** This table examines the baseline scale of business activity for traders who are active on a route post-treatment. It therefore studies selection into the trading route induced by the intervention. Columns 1, 3, and 5 show the average size across all traders active on the route post-treatment, while Columns 2, 4, and 6 show the minimum size. Importantly, size is measured by pre-treatment proxies at baseline. This table therefore shows how treatment induces into treated routes new types of traders (specifically, smaller traders). In Columns 1-2, size is measured by baseline profits of the trading enterprise. In Columns 3-4, it is measured by total tons traded (across all routes, at baseline) and in Columns 5-6 by total value traded (across all routes, at baseline). In all columns, standard errors are clustered two-way for each subcounty. "Mean DV" provides the mean of the dependent variable.

	Profits (mil UGX)		Tons traded		Value traded (mil UGX)	
	Mean	Min	Mean	Min	Mean	Min
Treated route	-1.89**	-1.99**	-74.55	-97.74	-38.07	-47.62
	(0.95)	(0.80)	(83.61)	(79.64)	(35.13)	(33.15)
Dist (100km)	0.98	1.73***	183.15	201.13*	81.76	89.98*
	(0.61)	(0.65)	(117.74)	(119.30)	(51.54)	(52.35)
Observations	774	774	774	774	749	749
Mean DV	7.41	5.35	256.59	190.94	126.81	92.65
P-value	0.05	0.01	0.37	0.22	0.28	0.15

Table 5: **Effects on trader business outcomes.** The first three columns present reduced form trader-level results: impacts on maize trading volumes, margins between maize buying and selling prices, and self-reported total profits, regressing each outcome on treatment, its baseline value, and controls selected from baseline covariates set using double lasso, as pre-specified. The fourth column runs an almost identical specification (omitting baseline value controls) for whether the trader is still in business (survey attritors are assumed out of business and given a zero) Finally, the fifth column uses market-level data (thus the drop in the number of observations) on the number of new traders in business (who have a store in the market) as measured in the endline trader census. In all columns, standard errors are clustered at the subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Maize Trading Volumes (tons)	Maize Buy-Sell Margins (UGX)	Total Self-Reported Profit (UGX)	Still in Business	Number New Traders in Business
Treatment	94.88	-10.96	-1010.77*	-0.02	-0.01
	(93.27)	(9.28)	(554.97)	(0.02)	(0.14)
Observations	2618	2268	2592	2898	236
Mean of DV	227.98	135.01	7279.21	0.85	1.05
R squared	0.03	0.02	0.12	0.01	0.00

Table 6: **Effects on farmer revenues, volumes sold, and prices.** Columns present reduced form impacts on total revenues (in thousands of UGX), maize revenues (in thousands of UGX), quantity of maize sold (in kgs), and price per kg of maize sold (in UGX). Outcomes are regressed on treatment, its baseline value, and controls selected from baseline covariates set using double lasso, as pre-specified. Standard errors are clustered at the subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Revenues Total ('000)	Revenues Maize ('000)	Qnt Sold Maize	Price Maize
Treat	99.2	72.0	61.3	18.0
	(90.9)	(68.5)	(118.2)	(14.0)
Observations	2775	2775	2769	1959
Mean DV	1019	672	1040	631
R2	0.32	0.31	0.33	0.02
Controls	Yes	Yes	Yes	Yes

Table 7: **Impact of experimental exposure.** Exposure is defined as $\lambda_i \equiv \sum_j \mathbb{1}_{\mathbb{j}}^{\mathbb{T}} \frac{X_j}{d_{ij}}$, where λ_i is the exposure of subcounty i, defined as the sum of exposure to all other treated subcounties j ($\mathbb{1}_{\mathbb{j}}^{\mathbb{T}} = 1$), weighted by subcounty j's market size (as proxied by j's total bilateral trade flows with all other markets, X_j) and by the inverse of distance between i and j (d_{ij}). We follow Borusyak and Hull (2021) and run 1,000 placebo randomization draws. For each placebo draw, we construct the exposure measure for each subcounty. We then demean our realized exposure measure λ_i by the average of the exposure measure under the 1,000 placebo randomization draws, $\bar{\lambda}_i$, such that $\tilde{\lambda}_i \equiv \lambda_i - \bar{\lambda}_i$. The resulting measure $\tilde{\lambda}_i$ therefore captures variation in exposure that arises purely from the realized randomization draw, which is exogenous. Column 1 regresses ln price on this exposure measure. Column 2 constructs this same procedure but separately for exposure to surplus and deficit market. Column 3 includes treatment interaction terms, allowing the impact of exposure to vary by own treatment status. The table presents standard errors clustered by subcounty, the unit of randomization. Table notes present p-values from randomization inference.

	Ln Price	Ln Price	Ln Price
Exposure to treated markets	-0.17***		
	(0.04)		
Exposure to treated surplus markets		-0.45**	-0.66***
		(0.17)	(0.25)
Exposure to treated deficit markets		0.16	0.69**
		(0.14)	(0.31)
Own market treated x Exposure to treated surplus markets			0.31
			(0.34)
Own market treated x Exposure to treated deficit markets			-0.73**
			(0.34)
Own market treated			-0.00
			(0.02)
Observations	22684	22684	22684
RI p-value: Exposure to treated markets	0.58		
RI p-value: Exposure to treated surplus markets		0.06	0.10
RI p-value: Exposure to treated deficit markets		0.67	0.08
RI p-value: Own market treated*Exposure to treated surplus markets			0.33
RI p-value: Own market treated*Exposure to treated deficit markets			0.08
RI p-value: Own market treated			0.91

Table 8: **Parameter estimates.** Parameter estimates from a joint estimation of Equations 5-7 using generalized method of moments (GMM). We calculate GMM standard errors (in parentheses) using two- way clustering on sender and receiver sub-counties.

Parameter	Estimate
β	-0.20 (0.13)
σ	4.44 (0.42)
γ	-0.67 (0.97)
θ	0.80 (1.53)
κ	0.04 (0.004)
σ_η^2	1.09 (0.19)
B	1.13 (1.08)

Table 9: **Model fit.** We use the simulated model of outcomes under an "actual Kudu" scenario, in which the treatment status of each market is the same as in the real RCT. We then construct a model-based version of our reduced form estimates in Table 1, regressing each outcome on a dummy for treatment (= 1 if both the markets are treated and therefore could possibly be connected by Kudu) and distance between the pair of markets. Standard errors are clustered two-way by each subcounty, the unit of randomization. Outcomes are: any trade along the route (first column), the number of traders trading along the route (second column), and trade volumes along the route (third column). Panel A presents the results from the full sample, Panel B for nearby markets (routes less than 200km) and Panel C for far markets (routes more than 200km) "Mean DV" provides the mean of the dependent variable.

	Any Trade	Number Traders	Volume (tons)
	Panel A: All routes		
Treated route	0.02	0.01	1.87
	(0.02)	(0.01)	(4.01)
Dist (100km)	-0.10***	-0.04***	-15.87***
	(0.01)	(0.01)	(2.63)
Observations	5671	5671	5671
Mean DV	0.16	0.05	19.22
	Panel B: Routes less than 200km		
Treated route	0.06**	0.04*	14.13
	(0.03)	(0.02)	(11.86)
Dist (100km)	-0.43***	-0.24***	-85.48***
	(0.02)	(0.03)	(13.85)
Observations	1731	1731	1731
Mean DV	0.37	0.13	50.04
	Panel C: Routes greater than 200km		
Treated route	0.01	-0.00	-2.18
	(0.02)	(0.00)	(1.83)
Dist (100km)	-0.030***	-0.005***	-3.238***
	(0.007)	(0.001)	(1.006)
Observations	3940	3940	3940
Mean DV	0.07	0.01	5.68

Table 10: **GE-corrected impacts.** To construct the GE-corrected impacts, we include two model-derived observations for each route; the "no Kudu" scenario and the "actual Kudu" scenario. We then run a regression similar in spirit to Tables 1 and 9, however, instead of regressing outcomes on a dummy for treatment, we now regress them on a dummy for being in the "actual Kudu" scenario (compared to being in the "no Kudu" scenario). This allows us to calculate the impact on treated routes "(IoT)," shown in Panel A (which restricts the regression to treated routes) and the impact on control routes "(IoC)," shown in Panel B (which restricts the regression to control routes). Panel C runs the same specification with the full sample of routes to calculate the average impact on all routes "(IoA)." As in Table 1, we control for distance between the pair of markets and cluster standard errors two-ways by each subcounty, the unit of randomization. Due to computational burden, we do not bootstrap the parameter estimates themselves, meaning standard errors reflect only uncertainty arising from the specific sample of markets and the draws of the idiosyncratic components of trade costs in the simulation, but not the error in the estimates of the parameters from the GMM process described in Section 5. Table notes present the "naive" impacts. The "IoT Naive" is the reduced form treatment effect from regressing outcomes on treatment under the "actual Kudu" scenario (this is the same as the estimates shown in Table 9). The "IoC Naive" is simply zero, as naive estimates assume control routes are unaffected by the intervention. Finally, the "IoA Naive" is the weighted average of the two. "Mean DV" provides the mean of the dependent variable.

	Any Trade	Number Traders	Volume (tons)
Panel A: Impact on treated routes (IoT)			
Kudu	0.0209***	0.0083***	1.3309***
	(0.0041)	(0.0009)	(0.2244)
Dist (100km)	-0.11***	-0.05***	-18.72***
	(0.01)	(0.01)	(3.86)
Observations	2970	2970	2970
IoT Naive (ITT)	0.0225	0.0100	1.8745
Panel B: Impact on control routes (IoC)			
Kudu	-0.0007	-0.0004***	-0.2118***
	(0.0005)	(0.0001)	(0.0525)
Dist (100km)	-0.10***	-0.04***	-14.83***
	(0.01)	(0.01)	(2.78)
Observations	8372	8372	8372
IoC Naive	0	0	0
Panel C: Average impact on all routes (IoA)			
Kudu	0.0049***	0.0019***	0.19***
	(0.0014)	(0.00039)	(0.062)
Dist (100km)	-0.10***	-0.04***	-15.82***
	(0.01)	(0.01)	(2.63)
Observations	11342	11342	11342
IoA Naive	0.0059	0.0026	0.49
Naive mismeasurement	0.20	0.38	1.55
Mean DV	0.16	0.05	19.03

Online Appendix

Appendix A Study setting and design

Figure A.1: **Maps of the Study Area.** The left-hand panel is USAID's FEWS-Net map of Surplus Maize Areas of Uganda, and the right-hand panel shows the 110 study subcounties.

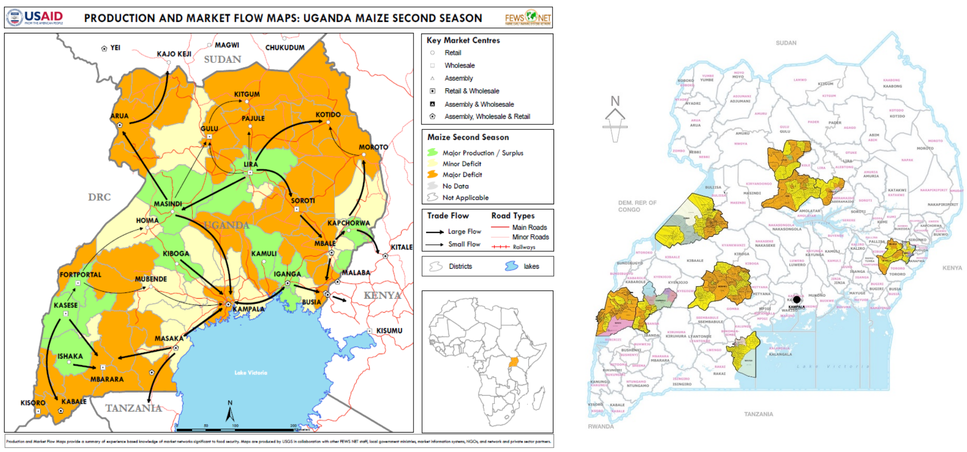

Figure A.2: **Study Timeline.** The study ran for three years (2015-2018), spanning six agricultural seasons (in orange). We collected high-frequency market data throughout the study period at the 236 markets in our sample. We surveyed traders at baseline in 2015, midline in 2016, and endline in 2018. We surveyed our farmer household sample at baseline in 2015 and at endline in 2018.

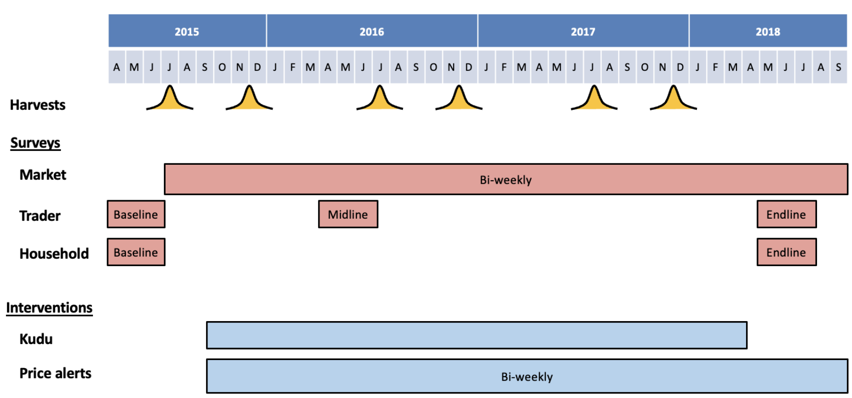

Figure A.3: **Sample and randomization design.** This figure presents the sample and randomization design.

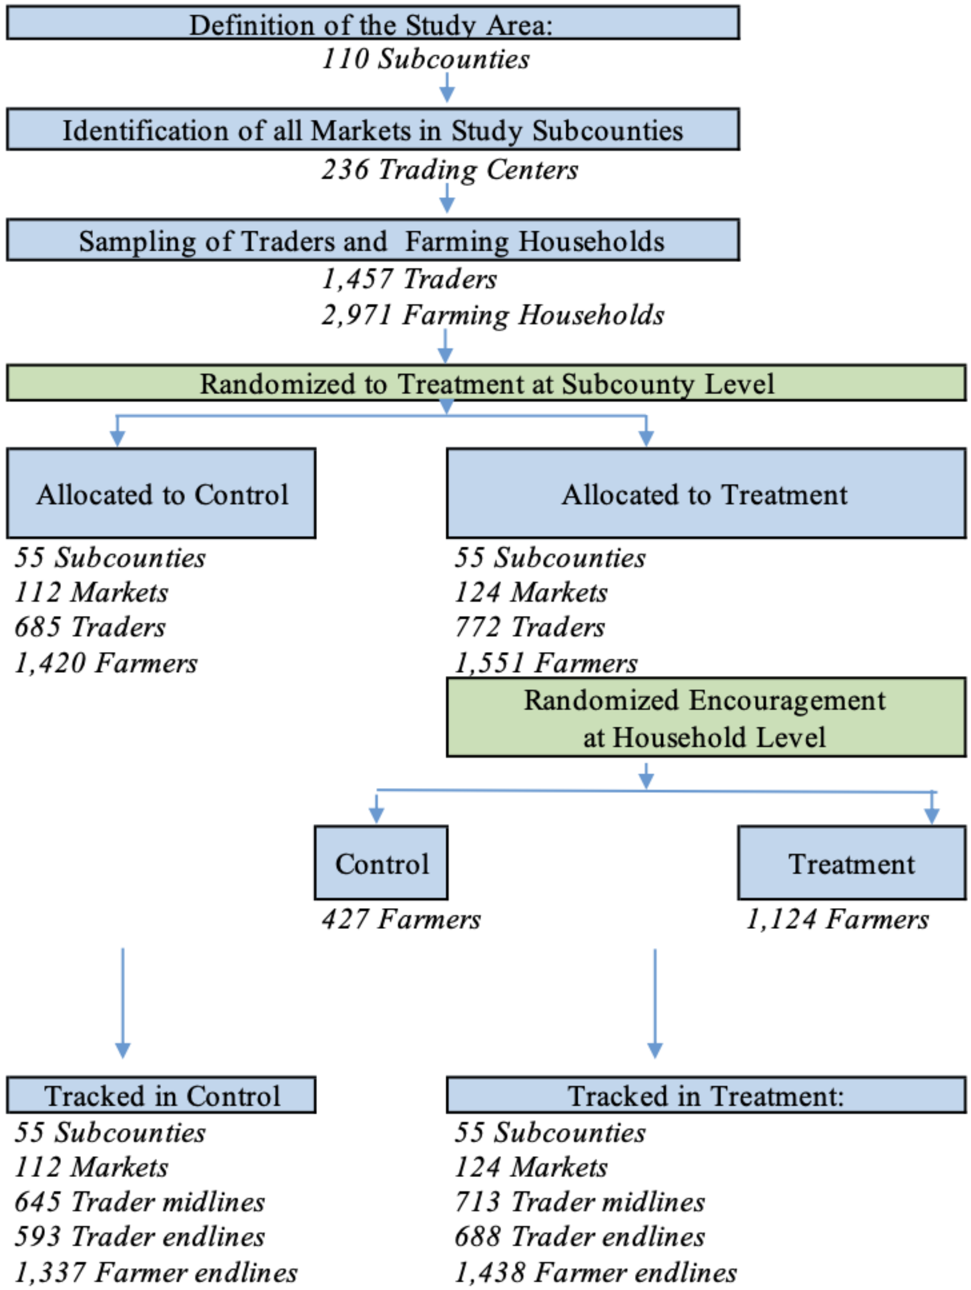

Table A.1: **Analysis of Variance in Market Prices.** Table reports the R-squared from a different dummy variable fixed effects regression of prices in the market survey. The first column uses market fixed effects (and so measures cross-sectional variation in prices), the second column month-of-year fixed effects (and so measures the degree of typical seasonality), the third column includes fixed effects for each two-week round of the survey (and so measures the extent of overall time-series variation), and the fourth column includes both market and round fixed effects.

	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

Appendix B Attrition and balance

Appendix B.1 Attrition

Figure B.1 and Tables B.1 and B.2 compare tracking rates in treatment control across the three data types. Overall, weighted attrition rates are very low and the overall unweighted attrition rate from the combination of standard and intensive tracking is similar across treatment arms for all data types. Among all the tests that we conduct, only the intensive tracking rate in the trader survey appears differential, but this arises from finding 14/14 control traders versus 24/27 treatment traders who were intensively tracked, which has relatively little influence on study-level effects.

Figure B.1: **Attrition from the Market Survey.** The figure shows the Cumulative Density Functions (CDFs) of the fraction of intended surveys completed for each market, separately for treatment and control. A Kolmogorov-Smirnoff test fails to reject equality of the two distributions.

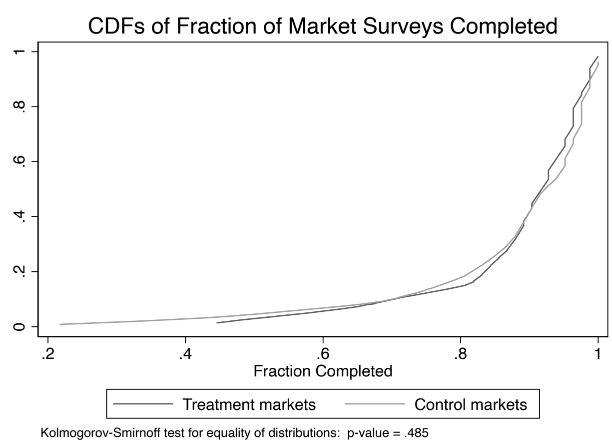

Table B.1: **Trader Survey Attrition.** Analysis uses the full baseline sample of traders with a measure of subsequent attrition as the dependent variable. The first row examines midline attrition, the second row endline attrition prior to intensive tracking, the third row attrition during intensive tracking, and the final row overall endline attrition. Standard errors are clustered by subcounty.

	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 B.2: **Household Survey Attrition.** Analysis uses the full baseline sample of households with a measure of subsequent attrition as the dependent variable. The first row examines attrition prior to intensive tracking, the second row attrition during intensive tracking, and the final row overall endline attrition. Standard errors are clustered by subcounty.

	Treat	Control	Obs	T-C	
				diff	*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

Appendix B.2 Balance

Table B.3: **Market Survey Balance.** Analysis uses the market-level average of outcomes from the two pre-treatment market survey waves to examine balance of the market survey. Standard errors are clustered by subcounty.

	Treat	Control	Obs	T-C	
				diff	*p-val*
Maize 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 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
Bananas price	12,201.52	12,632.30	207	-430.79	0.67
Number of bananas traders	5.70	6.11	233	-0.41	0.74
Bananas quality	1.71	1.69	207	0.02	0.78
Tomato price	164.38	163.97	231	0.41	0.96
Number of tomato traders	9.87	10.53	233	-0.66	0.63
Tomato quality	1.64	1.63	231	0.01	0.82

Table B.4: **Balance in route-level absolute price gaps.** Analysis examines balance in absolute price gaps at the route-level across the two pre-treatment rounds of the market survey. Standard errors are two-way clustered by each subcounty, the unit of assignment. "Mean DV" provides the mean of the dependent variable.

	Maize	Beans	Bananas	Tomatoes
Both treated	10.57	22.08	115.7	1.911
	(7.717)	(31.59)	(354.5)	(2.753)
Mean DV	131.5	559.4	6362.1	72.37
N	26213	21126	20192	26144

Table B.5: **Trader Survey Balance.** Analysis conducted in baseline trader data but using only the attrited endline sample, with weights reflecting survey sampling and intensive tracking as used in the main analysis. 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. Standard errors for the p-values are clustered by subcounty, the unit of assignment.

	Treat	Control	Obs	T-C	
				diff	*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	0.46
# 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 B.6: **Household Survey Balance.** Analysis conducted in baseline household data but using the attrited endline sample, with weights reflecting survey sampling and intensive tracking as in the main analysis. 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. Standard errors for the p-values are clustered by subcounty, the unit of assignment.

	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	0.10
Percent of time sold at market	0.29	0.28	2,775	0.01	0.84
Sell in market	0.36	0.33	2,775	0.03	0.59
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	0.09
Input exp (annual, UGX)	275,318.30	304,037.63	2,775	-28,719.33	0.46
Revenue, total (annual UGX)	637,169.99	555,801.61	2,775	81,368.38	0.43

Appendix C Kudu details

Appendix C.1 Matching procedures

There are two methods of matching through the Kudu platform. The first, which accounted for 33% of the matches made, is an algorithm that clears the market each day. To do so, it defines a scoring function for all possible matches (meaning asks, offers to sell, and bids, offers to buy, in the system for the same crop available at the same time). The score for a match of ask a and bid b is:

$$score_{ab} = (p_b - p_a) * min\{q_a, q_b\} - d_{ab} * \lambda$$

where p_a and p_b are the ask and bid prices respectively, q_a and q_b the ask and bid quantities, d_{ab} a measure of the km distance between the two parties, and λ an estimate of fuel costs per km.

The second approach, which accounted for 67% of the matches, is a hand-matching process conducted by employees who could view a dashboard of the business on the platform and attempt to match trades manually. Inherent to the fact that this was done manually, we do not have a precise formula to describe how employees constructed these matches; however, we can use the Kudu administrative data, combined with our survey data, to look at observables that predict these matches. This analysis also helps better understand which non-price features matter.

In Columns 1-2 of Table C.1 below, we study the effect of the features considered by the algorithm – price gaps, quantities, and distance between all possible pairs of asks and bids for the same crop listed at the same time on Kudu – on the probability of being matched by the algorithm (Col. 1) and the probability of being matched manually (Col. 2). We see matches are more likely when price gaps are higher, quantities are higher, and distances are shorter.

In Col. 4, we see that, even among those that are matched, higher quantities and shorter distances still predict success of the deal. This suggests that the patterns observed in which treatment effects are most pronounced along shorter trade routes are not merely an artifact of Kudu's matching procedure: even conditional on the platform's matching strategy, distance is an important factor determining where the search cost reduction offered by Kudu generates new trade.

Columns 3 and 5 add additional predictors, which were not available to the Kudu algorithm, nor formally to the Kudu employees in charge of manual matches, but which may have been intuitable to them through family name, location, or direct communication; moreover, these are factors that are likely observable to the potential trading parties themselves once they are connected by Kudu. They include proxies for communication and trust, including whether the trading parties match in ethnicity and language.[23] We see in Column 3 that manual matches are more likely when potential trading parties share a language or when they share ethnicity; this hints that Kudu employees may have considered communication or trust factors important for an optimal match. Column 5 suggests that, beyond this, language and ethnicity did not play a role in the success of the deal, conditional on the match. Therefore, location, availability, communication, and trust may all be important factors that matter to potential transaction partners. From the perspective of our model, these features are captured by the match value $\varepsilon(\omega)$ that a buyer draws when matching with a seller.

[23] We note two caveats to this analysis. First, this analysis suffers from a loss in power, due both to sampling in our surveys (we did not run the full household or trader surveys with every user of Kudu, so only have these demographic characteristics for at best a subset of users) and to the sensitivity of these particular questions, which resulted in some non-response (enumerators were instructed to skip the ethnicity question if respondents were uncomfortable). Moreover, language and ethnicity may be weak proxies for the communication, trust, or the myriad other intangible factors that may be shared between trading partners, especially once they can directly communicate.

Table C.1: **Predictors of matches and completed deals** Columns 1-2 regress whether a possible (ask, bid) pair is matched on the three factors considered by the algorithm: prices, quantities, and distance. In Col. 1 the outcomes variable is an algorithm-made matches, in Col. 2 the outcome is a manual match. Col. 3 adds additional features that may have been inferable by Kudu employees in the manual matches, including a binary dummy for whether the buyer and seller share a language and whether they share ethnicity, as measured in our household and trader surveys. Column 4 restricts attention those pairs that are matched, regressing whether the deal is completed on the three factors considered by the algorithm: prices, quantities, and distance. Column 5 add dummies for whether the pair share a language or ethnicity. "Mean DV" provides the mean of the dependent variable.

	(1) Match Algorithm	(2) Match Manual	(3) Match Manual	(4) Completed All	(5) Completed All
Price gap (1,000 UGX)	0.00017***	0.00009***	0.00015***	-0.01350***	-0.00414
	(0.00001)	(0.00001)	(0.00003)	(0.00306)	(0.00552)
Min quantity (100 tons)	0.00297***	0.00415***	0.00477***	0.18254***	0.03068
	(0.00014)	(0.00020)	(0.00061)	(0.03467)	(0.05937)
Distance (100 km)	-0.00026***	-0.00065***	-0.00056***	-0.03978***	-0.02247**
	(0.00000)	(0.00001)	(0.00002)	(0.00403)	(0.00868)
Language match			0.00059***		-0.00058
			(0.00005)		(0.01413)
Ethnic match			0.00018**		-0.00828
			(0.00006)		(0.02490)
Observations	11936526	11936526	1335108	13742	1749
Mean of DV	0.0004	0.0007	0.0008	0.0844	0.0349
Sample	All possible matches	All possible matches	All possible matches	Among those matched	Among those matched

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Appendix C.2 Take-up

Over the three years of the study, Kudu received 29,308 unique asks (offers to sell) and 31,177 unique bids (offers to buy). Bids and asks posted on the platform each day climbed through the first year to reach a steady maximum of about 400 tons of per day in bids and 200 tons in asks. Overall, Kudu saw 7,300 tons of grain successfully transacted during the study period, worth a total value of 9.4b UGX (about $2.7 million USD). Figure C.2 shows the cumulative sales over the platform during the duration of the study.

Figure C.1: **Volume of Bids and Asks.** Figure represents the spatial density of bids and asks over the study period.

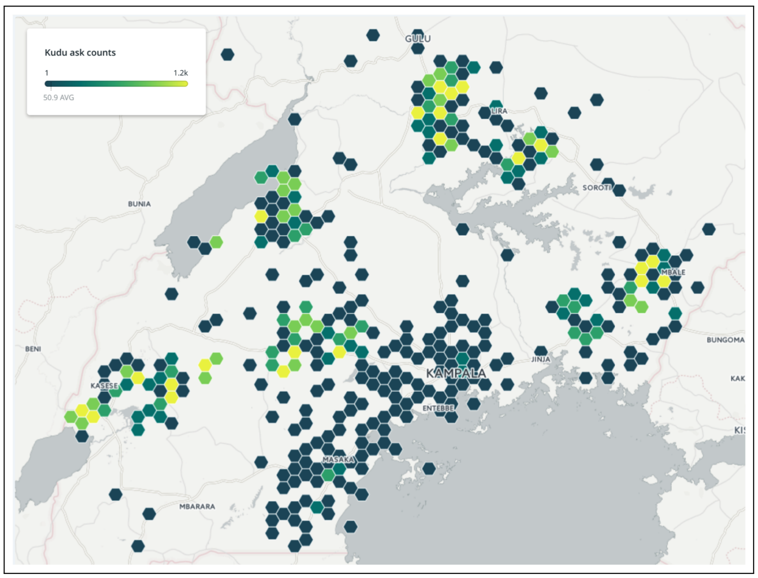

Figure C.2: **Cumulative transaction volumes on Kudu.** Cumulative transaction volumes of maize (in tons) on the Kudu platform over the study period.

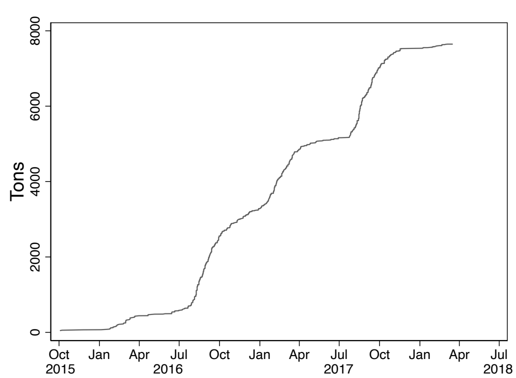

Figure C.3: **Value of Sales in Kudu by Crop**

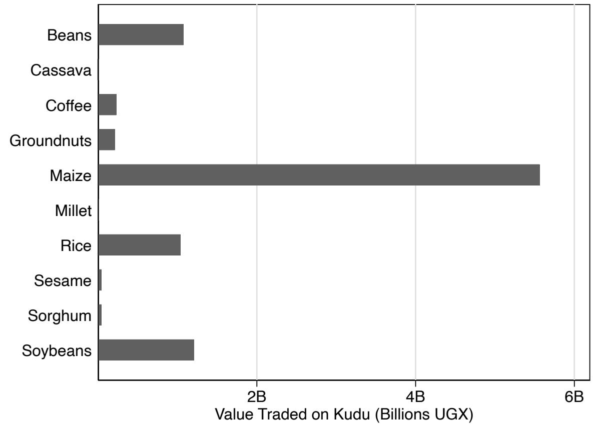

Although Kudu has the capacity to connect buyers and sellers of 76 types of crops and livestock, in practice, maize trade dominated the platform, accounting for 67% of asks and 5.6b UGX in trading volumes, swamping the value of all other crops put together (see Figure C.3). The remaining trade was spread across another nine crops, with the next most common being soybeans, beans, and rice, each of which account for much smaller volumes. Soybeans accounted for 1.2b UGX in trading volumes, beans for 1.1b UGX, and rice for 1b UGX.

Table C.2: **Take-up among traders by treatment status.**

	Treat	Control	Obs	T-C *diff*	T-C *p-val*
Post to Kudu	0.80	0.16	1,457	0.65	0.00
Complete deal on Kudu	0.22	0.03	1,457	0.19	0.00

Table C.3: **Take-up among farmers by treatment status.**

	Treat	Control	Obs	T-C *diff*	T-C *p-val*
Post to Kudu	0.26	0.02	2,775	0.24	0.00
Complete deal on Kudu	0.02	0.00	2,775	0.02	0.00

Table C.4: **Predictors of take-up among farmers.** Predictors of take-up among treated farmers, estimated via probit (with standard errors clustered at subcounty level). The outcome in Columns 1-4 is ever posting an ask or bid to Kudu, while the outcome in Columns 5-8 is successfully completing a transaction on Kudu. Regressors are all measured at baseline. In Columns 1 and 5 they are: quantity of maize sold (IHS), age of respondent, years of education completed, and a dummy for being female. Columns 2 and 6 adds dummies for whether the main activity is farming and whether the household operates a non-farm business, as proxies for opportunity costs. Columns 3 and 7 adds a standard survey measure of risk aversion as a proxy for ability to bear risk. Columns 4 and 8 adds household's total non-farm income from wages, non-farm businesses, gifts and remittances, government support, and other sources (IHS), as a proxy for credit constraints. "Mean DV" provides the mean of the dependent variable.

	Ever used Kudu				Completed deal on Kudu			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Quantity sold	0.035**	0.034**	0.035**	0.032**	0.032	0.032	0.031	0.028
	(0.014)	(0.014)	(0.014)	(0.015)	(0.041)	(0.041)	(0.041)	(0.042)
Age	-0.004	-0.003	-0.004	-0.003	-0.000	-0.001	-0.000	0.000
	(0.003)	(0.003)	(0.003)	(0.003)	(0.006)	(0.006)	(0.006)	(0.007)
Education	0.008	0.008	0.008	0.007	-0.015	-0.015	-0.015	-0.018
	(0.008)	(0.009)	(0.008)	(0.008)	(0.022)	(0.020)	(0.022)	(0.024)
Female	-0.303***	-0.305***	-0.303***	-0.296***	-0.216	-0.215	-0.226*	-0.267**
	(0.072)	(0.072)	(0.072)	(0.071)	(0.140)	(0.154)	(0.136)	(0.127)
Main activity farming		0.019				-0.069		
		(0.091)				(0.222)		
Operates other business		0.083				-0.126		
		(0.059)				(0.172)		
Risk averse			-0.003				0.057	
			(0.015)				(0.045)	
Non-farm income				0.012**				-0.013
				(0.005)				(0.011)
Observations	1435	1435	1435	1408	1435	1435	1435	1408
Mean DV	0.26	0.26	0.26	0.26	0.02	0.02	0.02	0.02

Table C.5: **Predictors of take-up among traders.** Predictors of take-up among traders, estimated via probit (with standard errors clustered at subcounty level). The outcome in Columns 1-5 is ever posting an ask or bid to Kudu, while the outcome in Columns 6-10 is successfully completing a transaction on Kudu. Regressors are all measured at baseline. In Columns 1 and 5 they are: quantity of maize traded (IHS), age of respondent, years of education completed, and a dummy for being female. Columns 2 and 6 adds hours spent per year in trading business (IHS) and a dummy whether the trader operates a non-farm business, as proxies for opportunity costs. Columns 3 and 7 adds a standard survey measure of risk aversion as a proxy for ability to bear risk. Columns 4 and 8 adds trader's total loan amount (IHS), as a proxy for credit constraints. Columns 5 and 10 adds a dummy for whether the trader owns a vehicle (truck, car, or motorcycle). "Mean DV" provides the mean of the dependent variable.

	Ever used Kudu					Completed deal on Kudu				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Quantity sold	0.128***	0.128***	0.127***	0.125***	0.130***	0.067**	0.071**	0.067**	0.066**	0.067**
	(0.032)	(0.032)	(0.032)	(0.032)	(0.032)	(0.034)	(0.033)	(0.034)	(0.034)	(0.034)
Age	-0.004	-0.004	-0.004	-0.004	-0.003	0.001	0.000	0.001	0.001	0.001
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Education	0.003	0.002	0.002	0.002	0.003	-0.016	-0.014	-0.016	-0.016	-0.016
	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)	(0.012)	(0.013)	(0.012)	(0.012)	(0.012)
Female	-0.009	-0.028	-0.010	0.005	0.028	-0.080	-0.021	-0.070	-0.076	-0.092
	(0.214)	(0.220)	(0.215)	(0.211)	(0.224)	(0.265)	(0.268)	(0.268)	(0.266)	(0.258)
Time trading business		-0.020					0.037			
		(0.068)					(0.075)			
Operates other business		0.071					-0.196*			
		(0.127)					(0.115)			
Risk averse			-0.034					-0.030		
			(0.022)					(0.025)		
Loan amount				0.008					0.003	
				(0.008)					(0.007)	
Owns vehicle					0.180					-0.050
					(0.117)					(0.122)
Observations	716	716	716	716	716	716	716	716	716	716
Mean DV	0.80	0.80	0.80	0.80	0.80	0.22	0.22	0.22	0.22	0.22

Tables C.4 and C.5 present predictors of Kudu take-up among farmers and traders respectively. We see size, as measured by quantity sold, is a strong predictor of adoption of Kudu. This is consistent with the model of scale economies presented in Section 4. Note that this is somewhat distinct from the question of the types of traders Kudu induces into treated routes. When considering route-level trade, we can think of traders as "always takers" (those who serve the route even without Kudu), "compliers" (those who are induced to start serving a treated route by the availability of Kudu), and "never takers" (those who do not serve the route in spite of the availability of Kudu). In our model, "always takers" will be those with the lowest cost draws $a(\omega)$, who are the largest traders and have the highest variable profits. They are able to earn profits on the route with or without Kudu. "Compliers" are those who could not earn positive total profits when fixed costs were higher, but can once Kudu lowers those costs. They sell smaller quantities and have lower variable profits relative to the always-takers. This means that traders active along treated routes will on average be smaller compared to traders active along control routes. Indeed, this is what we find empirically in Table 4. "Never takers" – those who cannot earn profits on a particular route, even with the reduced fixed costs brought about by Kudu – are the smallest traders with the highest cost draws. Route level compliance is distinct from individual-level take-up of the platform; a trader might use the platform in their business as a whole, but still decide to serve some routes and not others. However, the individual take-up patterns we see in Table C.5 are broadly consistent with the expectation that always takers and compliers will be traders who operate on a larger scale than never takers.

Appendix C.3 Other Kudu details

Kudu instructs users to post their reservation prices. The (non-binding) price recommended by the platform in the event of a match is the seller's price, so this is incentive compatible for the buyer. Perhaps because of this, seller ask prices were often higher than buyer bid and market prices. Qualitative interviews suggest that sellers often post strategic offers, much as they would in more traditional, in-person negotiations. Buyers' bid prices, on the other hand, track market prices closely (see Figure C.4). In practice, Kudu accommodates this sort of negotiating behavior by allowing for a small negative overlap of prices, enabling those that are close but not quite overlapping to match. Kudu also does not enforce the price recommendation, and in practice price can be renegotiated flexibly by the buyer and seller.

Figure C.4: **Maize Prices in Kudu vs. Market Survey.** Figure presents a comparison of average ask and bid prices on Kudu versus market prices.

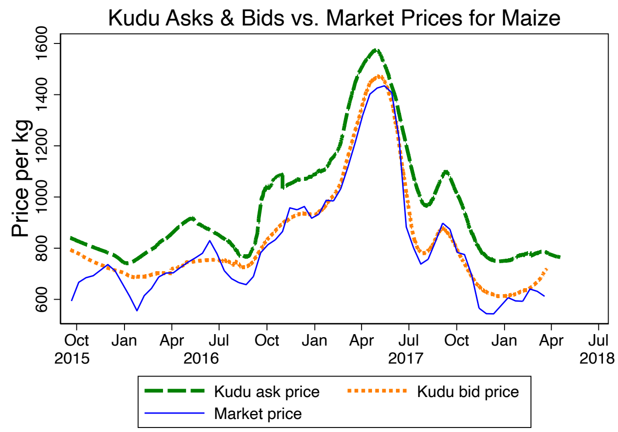

Appendix D Randomization inference

Figure D.1: **P-values from randomization inference for the reduced form effects on trade flows.** The figure shows the p-values resulting from a randomization inference procedure, in which 1,000 placebo randomized treatment assignments are draw. P-values from a two-sided test are presented in solid red and p-values from a one-sided test are presented in dash black. For the one-sided test, the null hypothesis being rejected is that the outcome is greater in control than in treatment. The left panel shows results for the probability of trade along any treated route, the middle panel for the number of traders trading along a route, and the right panel for trade volumes along the route.

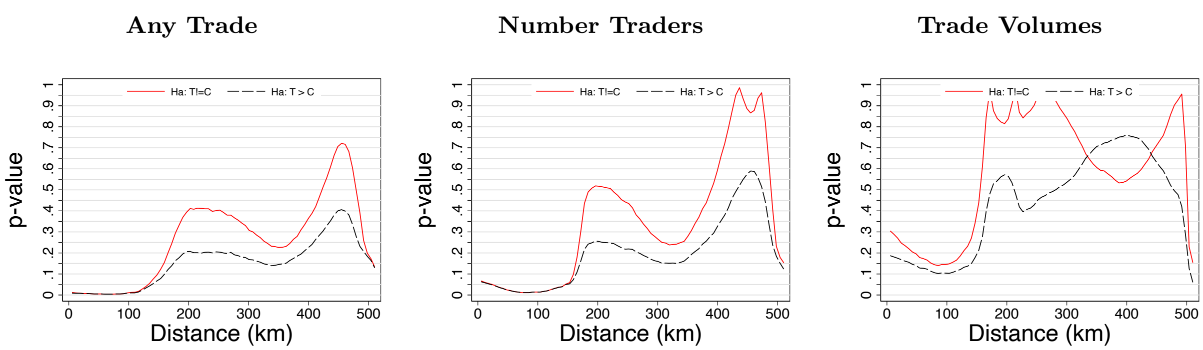

Figure D.2: **P-values from randomization inference for the reduced form effects on price gaps.** The figure shows the p-values resulting from a randomization inference procedure, in which 1,000 placebo randomized treatment assignments are draw. P-values from a two-sided test are presented in solid red and p-values from a one-sided test are presented in dashed black. For the one-sided test, the null hypothesis being rejected is that the outcome is greater in treatment than in control.

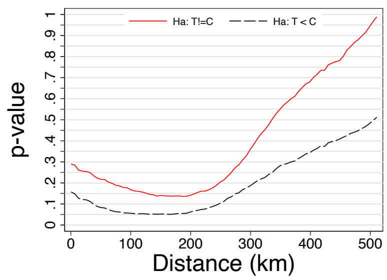

Appendix E Kudu effects by distance

Table E.1: **Reduced form effects on route-level trade flow and price gap outcomes, interacted with distance.** Reduced form results on any trade along the route (Col 1), the number of traders trading along the route (Col 2), trade volumes along the route (Col 3), and price gaps across the two markets connected by the route (Col 4). Outcome is measured at a route level and regressed on a dummy for treatment (= 1 if both the markets are treated and therefore could possibly be connected by Kudu), distance between the pair of markets, the interaction of treatment and distance, and round fixed effects. Standard errors are clustered two-way by each subcounty (the unit of randomization). "Mean DV" provides the mean of the dependent variable.

	Any Trade	Number Traders	Volume (tons)	Price Gaps
Treated route	0.07**	0.37*	13.24*	-10.80*
	(0.03)	(0.19)	(6.94)	(5.62)
Dist (100km)	-0.05***	-0.18***	-4.16***	9.63***
	(0.00)	(0.02)	(1.13)	(1.17)
Treated route*Dist (100km)	-0.02**	-0.10*	-3.72*	0.16
	(0.01)	(0.05)	(1.94)	(0.19)
Observations	11236	11236	11236	1443093
Mean DV	0.05	0.19	4.71	116.30

Appendix F Surplus definition and predictors

Surplus is calculated, for each market, as the average tons of maize sold per farmer, as measured in our baseline survey. The mean is 0.91, with a standard deviation of 0.86 and a range that spans just above 0 to 5.5 tons. The support is always positive, as the amount sold is bounded below by zero. Because this measure comes from our baseline survey, it is time-invariant. We avoid using endline surplus, as this could risk treatment-induced bias (though, in practice, baseline and endline surplus is strongly correlated, with a correlation coefficient of 0.7).

This surplus measure is correlated with weather and soil conditions. In Table F.1 below, we regress our surplus measure on three sources of gridded data: (i) 5-year average temperature data from the Climate Research Unit Time Series (CRU TU) dataset; (ii) 5-year average rainfall data from Climate Hazards Group InfraRed Precipitation with Station data CHIRPS; and (iii) The Food and Agricultural Organization's Global Agro-Ecological Zones (GAEZ) suitability index for maize. We see in Column 1 that overall higher temperature is associated with lower surplus, though Column 2 shows this relationship is quadratic, as one would expect given the mid-range optimal temperatures for production. Columns 3-4 also show a quadratic relationship with rainfall. Column 5 shows a positive correlation between surplus status and the FAO GAEZ maize suitability index, although the relationship is not statistically significant.

Table F.1: **Predictors of surplus** Surplus is regressed on gridded data on temperature (Col. 1-2), rainfall (Col. 3-4) and soil suitability (Col. 5) Temperature data is from the Climate Research Unit Time Series dataset, which provides gridded monthly average temperature data at a spatial resolution of 0.5 degrees; 5-year mean temperature values are calculated over the 2010–2014 period. The rainfall data is from CHIRPS, a global 5 km resolution gridded data set that combines satellite imagery with ground-based station observations to provide monthly precipitation estimates; these are averaged over the period 2010–2014 to construct 5-year average rainfall. The soil suitability is from the Food and Agricultural Organization's Global Agro-Ecological Zones (GAEZ) suitability index for maize, available under the Suitability and Attainable Yield theme (Theme 4); we use value cultivated with low input management and rainfed crops, averaged over the period 2001–2020 within each 10km grid.

	(1)	(2)	(3)	(4)	(5)
Temperature	-135.005^{***}	3079.892^{***}			
	(47.037)	(689.635)			
Temperature^2		-72.495^{***}			
		(15.964)			
Rainfall			-2.387^{***}	5.806	
			(0.297)	(4.879)	
Rainfall^2				-0.003^{*}	
				(0.002)	
Suitability Index					0.043
					(0.075)
Observations	234	234	234	234	234

Appendix G Other margins of adjustment

Appendix G.1 Empirical results

In this appendix, we explore treatment effects on agricultural inputs and outputs, as well as on engagement in and income from non-agricultural activities. We do not see significant movement on any of these margins, despite extensive data collection to measure such responses. We also investigate and find no evidence of cross-crop substitution in production and consumption. Finally, we see no changes in whether or where farmers market their crops, but we do see that they increase the fraction of their output that is sold in areas where market prices rise in response to the intervention.

In Table G.1, we first turn to production responses, either on average or differentially by surplus status. One might imagine that as farmers receive higher prices in surplus areas, this could drive increases in production, for example, through increased input usage. The converse might be true in deficit areas. However, as shown in Col 1-2 of Table G.1, we find no significant impact of Kudu on harvest levels. Although point estimates go in the expected direction – positive in surplus areas and negative in deficit – they are far from significant. We also see no change in land, labor, or other intermediate inputs used (Tables G.2- G.3). Lastly, we see no shifts in production in or out of other crops (Appendix Table G.6), nor in and out of non-agricultural activities (Appendix Table G.4). This suggests that the increase in trade we observe arises mainly from changes in the allocation of existing production, rather than through increases in productive specialization. Consistent with this, we observe an increase in the percent of production sold to the market in surplus areas, as shown in Col. 4 of Appendix Table G.1.

In principle, one might expect to see changes on these margins. It is an interesting question whether the limited production response is due to a small price change or a small supply elasticity, with the latter perhaps shaped by frictions preventing farmers from adjusting production. It is certainly possible that there are small effects that we are simply not powered to detect. This would be consistent with the small price changes induced by Kudu: point estimates suggest that prices increased by only 2.3% in surplus areas and dropped by only 1.4% in deficit areas. Indeed, our production estimates are noisy enough that we cannot rule out elasticities consistent with those in the literature, for example an elasticity of 0.43 for maize found by Abodi et al. (2021) in Kenya or 0.36 found for maize in South Africa found by Shoko et al. (2016); in fact, our point estimates are larger, suggesting at face value elasticity estimates of 1.2 in surplus areas and 3.5 in deficit areas, though they are statistically indistinguishable from zero.

Of course, it is also possible – consistent with both the broader literature and the lack of impact on production observed in our setting – that there is some stickiness that deters farmers from adjusting production in response to small price changes in a relatively short time frame. Because our study was not designed to induce experimental variation in any farmer production distortions, we are unable to speak to this definitively. However, we do in Table G.5 test whether production responses are greater among farmers who appear to be less constrained: those who have higher assets, larger landholdings, and hire labor. We see no differential treatment effects.

Ultimately, given the noise in our estimates, as well as the fact that our experiment was not designed to produce variation in production distortions, we are unable to attribute the lack of production response in our setting. It is possible that with a bigger intervention or longer time-frame there might be more movement on these margins. This is an interesting area for future research.

Table G.1: **Impacts on farmer production and share of production sold.** Analysis uses endline farmer surveys to measure impacts on production and share of production sold. Column 1 presents treatment effects on maize harvests (in kg), while column 2 presents an interaction of treatment and baseline surplus status of the area. Column 3 presents treatment effects on share sold, while column 4 presents an interaction of treatment and baseline surplus status of the area. Standard error are clustered by subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Production (kg harvested)		Share of production sold	
	(1)	(2)	(3)	(4)
Treat	-15.54	-45.93	0.03	0.01
	(131.67)	(97.16)	(0.03)	(0.04)
Surplus		577.61**		0.10**
		(250.74)		(0.04)
Treat*Surplus		108.86		0.10**
		(329.57)		(0.05)
Observations	2768	2768	2560	2560
Mean of DV	1380.20	1380.20	0.53	0.53
R squared	0.34	0.36	0.15	0.18

We also test for and fail to find evidence of substitution across crops, though we caution that our experiment is not well-designed to look at this directly, as we lack independent variation across crops that would be needed to formally identify cross-elasticities. However, as treatment effects were greatest in maize, we can look for direct evidence of treatment effects *away* from other crops.

For production, in our farmer surveys, we measure production of three crops (besides maize): beans, tomatoes, and matooke. As discussed above, Table G.6 presents the impact of treatment on production of these crops, overall and by surplus area. We find no evidence of substitution in production (consistent with the fact that we find no main production response on maize).

We similarly find no evidence of substitution effects on the consumption side in Table G.7, though we caution that here there are some data limitations (we lack direct data on consumption, but instead infer it from survey modules that were designed to capture trade and production). The lack of significant shifts in production or consumption of other crops is perhaps not surprising in light of the small experimentally-generated price effects for maize. Of course, we cannot rule out that a larger intervention or one implemented at scale might generate stronger adjustments on these margins.

Finally, we see no adjustment on other margins, such as farmers' marketing channels (Table G.9) or trader behavior such as the amount they contribute from their own harvest to their trading business or the amount they store (Table G.10).

Table G.2: **Impacts on land and material inputs used in production.** Analysis uses endline farmer surveys to measure impacts on land and intermediary inputs. Odd columns present overall treatment effects, while the columns present an interaction of treatment and baseline surplus status of the area. Columns 1-2 present effects on acres cultivated, columns 3-4 on rental prices per acre in thousands of UGX (as estimated by farmers in response to a question asking "Imagine that your household owns all the pieces of land it cultivated/farmed in the past 12 months and that it were to rent out this land to someone else for one year, how much would your household charge?"), and columns 5-6 on the total amount spent on all intermediate inputs (fertilizer, seeds, etc.) over the last 12 months. Standard error are clustered by subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Acres Cultivated		Land Rental		Intermediate Inputs	
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	-0.04	-0.06	-8.29	-11.47	-1.58	-9.09
	(0.18)	(0.20)	(15.52)	(16.68)	(15.42)	(16.83)
Surplus		-0.30		53.77***		-62.17***
		(0.29)		(19.71)		(21.89)
Treat*Surplus		0.09		12.98		28.24
		(0.36)		(25.78)		(33.89)
Observations	2771	2771	2708	2708	2775	2775
Mean of DV	4.91	4.91	240.10	240.10	142.65	142.65
R squared	0.23	0.23	0.04	0.06	0.09	0.09

Table G.3: **Impacts on labor inputs used in production.** Analysis uses endline farmer surveys to measure impacts on labor inputs. Odd columns present overall treatment effects, while the columns present an interaction of treatment and baseline surplus status of the area. Columns 1-2 present effects on the person-months worked on the farm by household labor in the past 12 months, columns 3-4 on the number of hired workers who worked on the farm in the past 12 months, columns 5-6 on the total amount paid to all hired labor in the past 12 months. Standard error are clustered by subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Family Labor		Workers Hired		Labor Expenditure	
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	-0.01	0.03	-1.72	-2.52	0.25	-6.57
	(0.29)	(0.33)	(1.28)	(1.57)	(32.91)	(26.86)
Surplus		-1.05*		-3.21*		69.20
		(0.56)		(1.62)		(55.80)
Treat*Surplus		-0.17		3.03		25.61
		(0.66)		(1.94)		(87.33)
Observations	2770	2770	2770	2770	2761	2761
Mean of DV	3.39	3.39	11.51	11.51	363.47	363.47
R squared	0.04	0.05	0.17	0.18	0.23	0.23

Table G.4: **Impacts on farmer non-farm enterprises.** Analysis uses endline farmer surveys to measure impacts on non-farm enterprise activity. Columns 1-2 presents treatment effects on operating any enterprise and Columns 3-4 present treatment effects on annual enterprise income (in thousands of UGX). Columns 1 and 3 present overall treatment effects, while Columns 2 and 4 present an interaction of treatment and baseline surplus status of the area. Standard error are clustered by subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Any Enterprise		Enterprise Income	
	(1)	(2)	(3)	(4)
Treat	-0.03	-0.05	134.75	61.26
	(0.04)	(0.05)	(118.15)	(126.61)
Surplus		-0.13***		-155.68
		(0.04)		(124.94)
Treat*Surplus		0.09		281.43
		(0.06)		(261.62)
Observations	2774	2774	2726	2726
Mean of DV	0.49	0.49	812.27	812.27
R squared	0.13	0.13	0.10	0.10

Table G.5: **Role of Constraints in Driving Producer Output Response**. Baseline productive attributes of farming households are interacted with the treatment to test if there are differential impacts of the treatment on production. The first column examines heterogeneity by wealth, using the IHS of the baseline asset index. The second column examines heterogeneity by landholding, using winsorized baseline farm area in acres. The third column interacts using a dummy for whether the household was using any hired labor at baseline. "Mean DV" provides the mean of the dependent variable.

	(1) Harvest (kg)	(2) Harvest (kg)	(3) Harvest (kg)
Treat	-331.34	-70.72	-13.78
	(1378.00)	(145.62)	(117.62)
Treat x Assets	23.12		
	(97.03)		
Assets	217.83***		
	(71.01)		
Treat x Land		12.78	
		(21.58)	
Land		34.75***	
		(12.94)	
Treat x Labor			16.23
			(189.50)
Labor			573.01***
			(143.11)
Observations	2768	2522	2768
Mean of DV	1380.20	1380.20	1380.20
R squared	0.27	0.30	0.27

Table G.6: **Impacts on production of other crops.** Analysis uses endline farmer surveys to measure impacts on production of other crops: beans, tomatoes, and matooke (kg harvested). Odd columns present overall treatment effects, while even columns present the interaction of treatment and baseline surplus status of the area. Standard error are clustered by subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Beans		Tomatoes		Matooke	
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	-4.47	-2.66	0.87	1.14	0.04	0.03
	(26.06)	(27.66)	(4.76)	(5.96)	(0.06)	(0.06)
Surplus		92.30**		-10.99*		0.06
		(44.46)		(5.62)		(0.08)
Treat*Surplus		-6.20		-0.92		0.03
		(59.46)		(9.32)		(0.11)
Observations	2773	2773	2763	2763	2761	2761
Mean of DV	181.67	181.67	16.61	16.61	0.28	0.28
R squared	0.22	0.23	0.04	0.04	0.03	0.03

Table G.7: **Impacts on consumption of other crops.** Analysis uses endline farmer surveys to measure impacts on consumption of other crops: beans, tomatoes, and matooke (in kg). Odd columns present overall treatment effects, while even columns present the interaction of treatment and baseline surplus status of the area. Standard error are clustered by subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Beans		Tomatoes		Matooke	
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	2.59	7.20	1.11	1.25	0.04	0.03
	(12.53)	(14.15)	(1.06)	(1.33)	(0.05)	(0.05)
Surplus		46.20**		-2.32***		0.04
		(19.04)		(0.85)		(0.07)
Treat*Surplus		-17.46		-0.52		0.06
		(25.47)		(1.80)		(0.09)
Observations	2775	2775	2775	2775	2775	2775
Mean of DV	73.99	73.99	2.30	2.30	0.21	0.21
R squared	0.10	0.11	0.00	0.01	0.00	0.01

Table G.8: **Impacts on other crop prices.** Analysis uses market surveys to measure impacts on prices of other crops: beans, tomatoes, and matooke (kg harvested). Odd columns present overall treatment effects, while even columns present the interaction of treatment and baseline surplus status of the area. Standard error are clustered by subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Beans		Tomatoes		Matooke	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-26.60	-26.68	0.106	-6.931	313.4	474.3
	(51.57)	(55.11)	(7.160)	(8.213)	(830.6)	(1038.0)
Surplus		-218.4***		-2.786		-1482.0*
		(60.20)		(7.524)		(820.8)
Treat*Surplus		49.04		24.39*		-139.0
		(85.19)		(12.50)		(1210.2)
Observations	10685	10685	15710	15710	11758	11758
Mean DV	2146.8	2146.8	180.3	180.3	14943.1	14943.1
R2	0.331	0.363	0.0588	0.0699	0.0385	0.0537

Table G.9: **Impacts on marketing.** Columns 1, 3, and 5 present treatment effects on the number of traders to which the farmer sold in the last 12 months, the number of *new* traders to which the farmer sold in the last 12 months, and whether the farmer sold anything at market in the last 12 months. Columns 2, 4, and 6 present an interaction of treatment and baseline surplus status of the area. Standard error are clustered by subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Number Traders		Number New Traders		Any Sales at Mkt	
Treat	-0.02	-0.02	0.02	0.05	0.00	0.01
	(0.08)	(0.09)	(0.03)	(0.03)	(0.03)	(0.03)
Surplus		0.39***		0.16***		-0.05
		(0.10)		(0.05)		(0.03)
Treat*Surplus		-0.01		-0.09		-0.04
		(0.13)		(0.09)		(0.05)
Observations	2768	2768	2768	2768	2774	2774
Mean of DV	1.05	1.05	0.22	0.22	0.26	0.26
R squared	0.05	0.07	0.01	0.01	0.04	0.05

Standard errors in parentheses

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

Table G.10: **Effects on amount added from traders' own harvest and amount stored.** The reduced form impacts on the amount added to trading business from own harvest (in tons) and the amount in storage between seasons (in tons), regressing each outcome on treatment and controls selected from baseline covariates set using double lasso. Outcomes are only available in endline survey. In all columns, standard errors are clustered at the subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Own Harvest Added (tons)	Storage (tons)
Treat	-0.69	-2.51
	(0.78)	(2.48)
Observations	1194	1194
Mean of DV	3.92	9.17
R squared	0.02	0.05

Appendix G.2 Implications for modeling assumptions

Given the lack of empirical evidence of movement along these margins, we do not incorporate them into our model. We could, for example, endogenize production, but if we use our estimates of treatment effects to discipline this quantitatively, we will end up with an answer equivalent to the endowment economy since we do not observe significant treatment effects. Another option would be to calibrate these elasticities from the literature, but this would be a departure from the key feature of this paper, which is to use well-identified experimental variation to discipline the parameters driving equilibrium impacts.

Because we wish to retain a tight connection between the experimental variation and estimation of key model parameters, we therefore choose only to include margins on which we see movement empirically, which leads us to employ an industry equilibrium model. We see this as key to credibly using model-based estimates to correct for spillovers in the context of RCTs.

However, we must caveat that our estimates for impacts at scale in Section 7.2 continue to shut down these margins. While we do not see movement on these margins in our experiment, such adjustments may of course emerge at scale. As we note in the text, the goal of the exercise is not (just) to offer a (qualified) estimate of effects at scale, but moreover, to illustrate how standard experimental inference can mislead, as a function of scale.

Appendix H Beans impacts

As described in Section 2.5, trade on Kudu was dominated by maize. However, as seen in Figure C.3, there was more limited trade in nine other crops, with next most common being soybeans, beans, and rice. Of these three, we tracked off-platform trade flow for beans. In this appendix, we replicate all of the maize analysis shown previously for beans.

We see in Table H.1 significant, albeit small, increases in the probability of any beans trade and the number of beans traders active along treated routes, with effects concentrated along shorter routes, as for maize. Effects on trading volumes are also positive, but are small and not statistically significant. Table H.2 suggests that treatment may have reduced bean price gaps across markets, though these effects are imprecisely estimated. Table H.3 explores effects on bean price levels as a function of (beans) surplus status. Point estimates suggest that, as for maize, treatment may lower bean prices in deficit areas (which tend to otherwise see high prices) and increase bean prices in surplus areas (which tend to otherwise see low prices), but these effects are far from statistically significant. Table H.4 hints at similar trader-level reduced form effects, with Cols. 1-2 showing point estimates that are positive for trading volumes and negative for buy-sell margins, but effects are again not significant. We see no effect on entry or exit of beans traders in Cols. 3-4. Similarly, farmer level reduced form effects in Table H.5 are positive but not statistically significant. As with maize, we see no effects on production of beans, either overall or by surplus status, in Table H.6.

Overall, we see treatment effects on trade flows, price gaps, and other outcomes for beans that are qualitatively similar to those for maize, but smaller and less precisely estimated, consistent with the smaller amount of beans trade. These results offer a reassuring validation of our findings for maize. They also imply that our main results, while an accurate description of the impact on maize trade and markets, may be a lower bound on the aggregate impact of the entire platform across all crops. However, because these other crops are traded in much smaller volumes and the magnitude of treatment effects appears to be much smaller, we believe our estimates focused on maize capture the majority of the quantitative impact of the platform overall.

Table H.1: **Reduced form effects on route-level beans trade flow outcomes.** Reduced form results on any beans trade along the route (first column), the number of beans traders trading along the route (second column), and beans trade volumes along the route (third column). Outcome is measured at a route level and regressed on a dummy for treatment (= 1 if both the markets are treated and therefore could possibly be connected by Kudu), distance between the pair of markets, and round fixed effects. Standard errors are clustered two-way by each subcounty (the unit of randomization). Panel A presents the results from the full sample, Panel B for nearby markets (routes less than 200km) and Panel C for far markets (routes more than 200km). "Mean DV" provides the mean of the dependent variable.

	Any Trade	Number Traders	Volume (tons)
	Panel A: All routes		
Treated route	0.01**	0.01**	0.01
	(0.00)	(0.00)	(0.06)
Dist (100km)	-0.004***	-0.006***	-0.318
	(0.002)	(0.002)	(0.293)
Observations	11236	11236	11236
Mean DV	0.00	0.01	0.32
	Panel B: Routes less than 200km		
Treated route	0.02**	0.02**	0.12
	(0.01)	(0.01)	(0.16)
Dist (100km)	-0.03***	-0.04***	-1.52
	(0.01)	(0.01)	(1.37)
Observations	3429	3429	3429
Mean DV	0.01	0.02	1.00
	Panel C: Routes greater than 200km		
Treated route	0.00	-0.00	-0.02
	(0.00)	(0.00)	(0.02)
Dist (100km)	-0.00	-0.00	-0.02
	(0.00)	(0.00)	(0.02)
Observations	7807	7807	7807
Mean DV	0.00	0.00	0.02

Table H.2: **Reduced form effects on route-level bean price gaps.** Reduced form results on bean price gaps across markets. Outcome is regressed on a dummy for treatment (= 1 if both the markets are treated and therefore could possibly be connected by Kudu), the distance between the pair of markets, and round fixed effects. Standard errors are clustered two-way by each subcounty (the unit of randomization). Panel A presents the results from the full sample, Panel B for nearby markets (routes less than 200km) and Panel C for far markets (routes more than 200km). "Mean DV" provides the mean of the dependent variable.

	Price Gaps
Panel A: All routes	
Treated route	-12.21
	(10.42)
Dist (100km)	39.64***
	(6.76)
Observations	723649
Mean DV	443.29
Panel B: Routes less than 200km	
Treated route	-5.43
	(13.22)
Dist (100km)	50.30***
	(11.38)
Observations	228792
Mean DV	350.71
Panel C: Routes greater than 200km	
Treated route	-14.53
	(13.48)
Dist (100km)	5.09
	(9.74)
Observations	494857
Mean DV	486.10

Table H.3: **Reduced form beans price effects in (beans) surplus and deficit areas.** In Column 1, market-level data on beans price regressed on a dummy for treatment and round fixed effects. Column 2 interacts treatment with a dummy for being in a baseline marketed surplus subcounty for beans. Standard errors clustered by subcounty, the unit of randomization.

	(1)	(2)
Treatment	-26.60	-30.51
	(51.57)	(54.02)
Beans Surplus		-258.4***
		(55.26)
Treat*Beans Surplus		39.82
		(81.46)
Observations	10685	10685
Mean DV	2146.8	2146.8
R2	0.331	0.386

Table H.4: **Effects on trader business outcomes for beans.** The first two columns present reduced form trader-level results: impacts on beans trading volumes and margins between beans buying and selling prices, regressing each outcome on treatment, its baseline value, and controls selected from baseline covariates set using double lasso, as pre-specified. The third column runs an almost identical specification (omitting baseline value controls) for whether the trader is still in the beans business (survey attritors are assumed out of business and given a zero) Finally, the fourth column uses market-level data (thus the drop in the number of observations) on the number of new traders in business (who have a store in the market) as measured in the endline trader census. In all columns, standard errors are clustered at the subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Beans Trading Volumes (tons)	Beans Buy-Sell Margins (UGX)	Still in Beans Business	Number New Beans Traders in Business
Treatment	0.98	-61.65	0.01	-0.00
	(2.33)	(135.94)	(0.02)	(0.06)
Observations	2463	178	2898	236
Mean of DV	4.43	-224.92	0.14	0.15
R squared	0.00	0.02	0.01	0.00

Table H.5: **Effects on farmer beans revenues, volumes sold, and prices.** Columns present reduced form impacts on total revenues (in thousands of UGX), beans revenues (in thousands of UGX), quantity of beans sold (in kgs), and price per kg of beans sold (in UGX). Outcomes are regressed on treatment, its baseline value, and controls selected from baseline covariates set using double lasso, as pre-specified. Standard errors are clustered at the subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Revenues Total ('000)	Revenues Beans ('000)	Qnt Sold Beans	Price Beans
Treat	99.2	3.2	4.5	32.2
	(90.9)	(28.7)	(16.1)	(61.3)
Observations	2775	2771	2770	684
Mean DV	1019	170	106	1567
R2	0.32	0.15	0.20	0.02
Controls	Yes	Yes	Yes	Yes

Table H.6: **Impacts on farmers' beans production.** Analysis uses endline farmer surveys to measure impacts on production. Column one presents treatment effects on beans harvests (in kg), while column 2 presents an interaction of treatment and baseline (beans) surplus status of the area. Standard error are clustered by subcounty, the level of randomization.

	Harvests Beans KG	Harvests Beans KG
Treat	-4.474	15.08
	(26.06)	(19.54)
Beans Surplus		174.9***
		(37.66)
Treat*Beans Surplus		-68.67
		(49.63)
Observations	2773	2773
Mean of DV	181.7	181.7
R squared	0.224	0.251

Appendix I Price information-only interventions

Table I.1: **Impact of price information alone on trade flows.** Reduced form results on any trade along the route (first column), the number of traders trading along the route (second column), and trade volumes along the route (third column). Table uses only control routes from the main intervention and analyzes the roll-in of the SMS-based price information-only intervention. Panel A presents the results from the full sample, Panel B for nearby markets (routes less than 200km) and Panel C for far markets (routes more than 200km). "Mean DV" provides the mean of the dependent variable.

	Any Trade	Number Traders	Volume (tons)
	Panel A: All routes		
Treated route - price info	0.01	0.20	-1.31
	(0.03)	(0.24)	(3.62)
Dist (100km)	-0.07***	-0.30***	-5.07**
	(0.01)	(0.06)	(2.22)
Observations	1378	1378	1378
Mean DV	0.08	0.30	4.93
	Panel B: Routes less than 200km		
Treated route - price info	-0.01	0.44	-6.43
	(0.06)	(0.61)	(12.20)
Dist (100km)	-0.41***	-1.90***	-33.00**
	(0.05)	(0.37)	(16.01)
Observations	423	423	423
Mean DV	0.23	0.95	15.44
	Panel C: Routes greater than 200km		
Treated route - price info	-0.00	-0.00	-0.32
	(0.01)	(0.01)	(0.24)
Dist (100km)	-0.00	-0.00	-0.12
	(0.00)	(0.00)	(0.10)
Observations	955	955	955
Mean DV	0.01	0.01	0.27

Table I.2: **Impact of price information alone on price gaps.** This analysis uses route-level data in panel format with fixed effects for market survey wave, and analyzes the roll-in of the SMS price information-only intervention to control markets during the second half of the study. The sample for this analysis consists only of routes in which neither market was treated in the original experiment, and 'Both Treated' means that both markets on the route were included in the roll-in treatment sample. Standard errors are two-way clustered at the subcounty level. "Mean DV" provides the mean of the dependent variable.

	Price Gaps
Panel A: All routes	
Treated route - price info	8.63
	(7.46)
Distance (100km)	10.27***
	(1.97)
Observations	140667
Mean DV	135.38
Panel B: Routes less than 200km	
Treated route - price info	14.03
	(9.21)
Distance (100km)	13.89
	(12.00)
Observations	43429
Mean DV	116.82
Panel C: Routes greater than 200km	
Treated route - price info	6.12
	(6.75)
Distance (100km)	8.33*
	(4.65)
Observations	97238
Mean DV	143.67

Figure I.1: **Reduced form effects of price information alone on trade flows.** The top lefthand figure presents a non-parametric local Fan regression of the probability of trade along any treated route (dotted line), in which both markets are treated only by price information in the staggered roll-in of the SMS-based price information, vs. control routes (solid line), as a function of distance. The bottom row presents results from a randomization inference procedure, in which 1,000 placebo randomized treatment assignments are drawn. The "reduced form effect," i.e., the difference between treatment and control line in the lefthand figure, from the realized randomization, is shown in black. The grey lines show the same effect for each of the 1,000 placebo treatment assignments. Long and short-dashed lines indicate the 90th and 95th confidence intervals of these placebo "reduced form effects." Subsequent columns present similar results for the number of traders trading along a route (second column) and trade volumes along the route (third column).

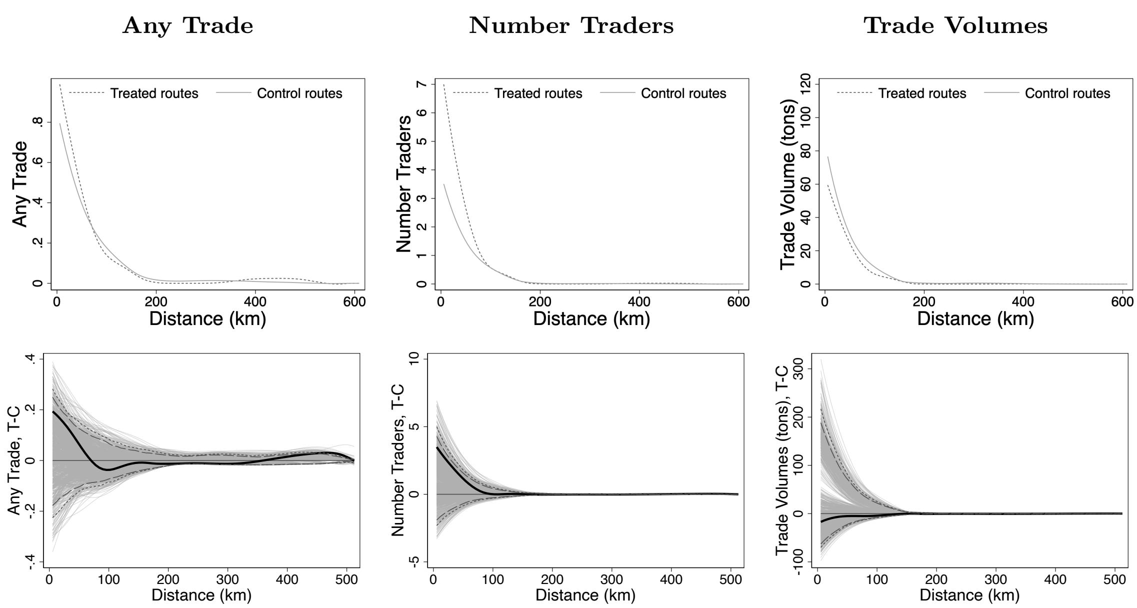

Figure I.2: **Reduced form effects of price information alone on price gaps** The top panel presents a non-parametric local Fan regression of the price gaps across markets along any treated route (dotted line), in which both markets are treated by the staggered roll-in of the SMS-based price information, vs. control routes (solid line), as a function of distance. The bottom row presents results from a randomization inference procedure, in which 1,000 placebo randomized treatment assignments are drawn. The "reduced form effect," i.e. the difference between treatment and control line in the lefthand figure, from the realized randomization is shown in black. The grey lines show the same effect for each of the 1,000 placebo treatment assignments. Long and short-dashed lines indicate the 90th and 95th confidence intervals of these placebo "reduced form effects."

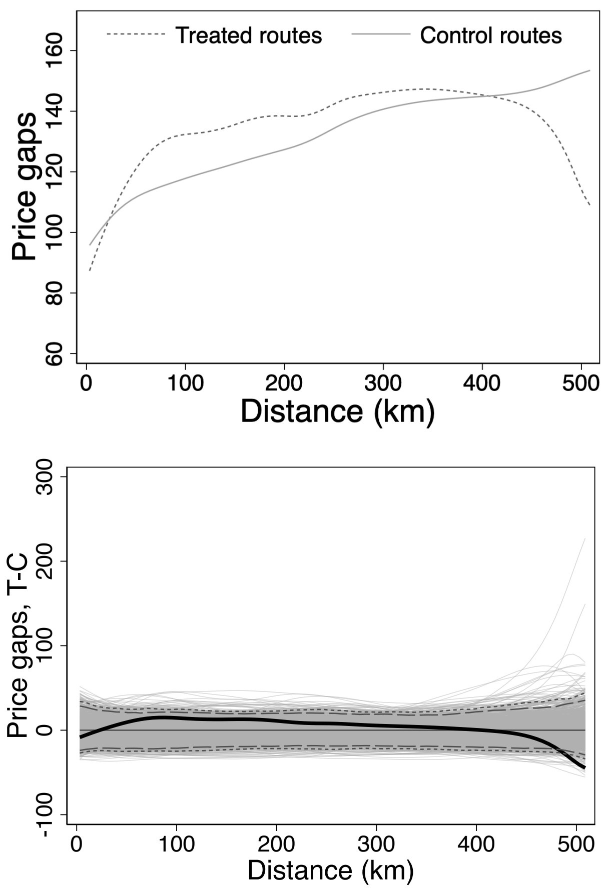

Figure I.3: **Impact of the Random Blast**. This figure presents an event study of the impact of the "Random Blast," a second sub-experiment to study the role of price information alone. In the "Random Blast," we randomly selected five treated markets in each round and broadcasted information about prices in those markets to the entire treated network. Though more short-run, this intervention was quite powerful, as it sent the market's price information to thousands of traders and farmers simultaneously. However, this price-only intervention also had no effect on price levels or gaps in the featured markets. The top panel presents monthly leads and lags of the treatment on price levels. The bottom panel examines impacts on price gaps between the Random Blast market and all other treatment markets that received information on that market's price.

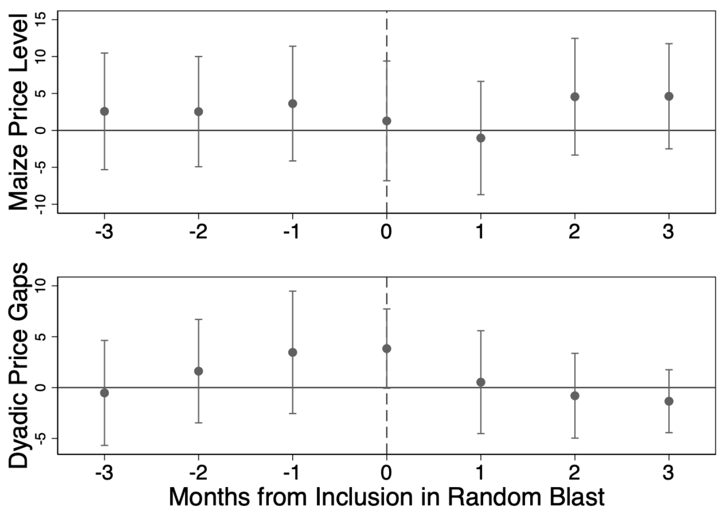

Table I.3: **Heterogeneity by unexpected price deviation.** This table examines whether Kudu induce faster convergence of prices when unexpected shocks hit. To calculate an unexpected route-level price deviation, we first compute the average price for each district x month-of-year. Then, for each round of the market survey, we calculate the deviation in that round and district from the typical price at that location and time of year. Moving into the route-level data structure, we then calculate the absolute value of the difference in these deviations across the dyadic pair ("Dyadic shock dispersion"). This then represents the component of the price gap between two districts that would not have been anticipated given the typical seasonal differences and hence provides a measure of the informational innovation present in the price discovery from the treatment. We then interact these price gaps with route-level treatment status to examine whether being treated causes faster price convergence for those treated routes where the price information is revealing larger than expected price gaps. We look at this interaction starting with the contemporaneous shocks, and then lag the shocks by one, two, and three periods to allow for the fact that arbitrage will take time to lower these price gaps. "Mean DV" provides the mean of the dependent variable.

	Contemporaneous	First Lag	Second Lag	Third Lag
Both Treated x Shock	-0.0158	-0.0222	0.00969	0.00659
	(0.0452)	(0.0419)	(0.0308)	(0.0317)
Dyadic shock dispersion	0.459***	0.211***	0.0805***	0.0756***
	(0.0400)	(0.0338)	(0.0240)	(0.0278)
Both treated	-6.192	-5.344	-6.692	-6.424
	(4.910)	(4.532)	(4.559)	(4.420)
Distance (100km)	5.840***	8.135***	9.414***	9.528***
	(1.155)	(1.154)	(1.164)	(1.178)
Mean DV	116.3	116.7	117.2	117.6
N	1443093	1418308	1392372	1368072

Appendix J Additional details on model assumptions and validation

Appendix J.1 Empirical basis for model assumptions

Table J.1: **Purchases from home market.** The model assumes purchases are made from traders' home market (subcounty). In Column 1, the outcome variable is a dummy for whether the trader buys exclusively from his home subcounty, while in Column 2, the outcome is the percent of total volumes bought from within trader's home subcounty. We see the mean in the table notes are high, both for the extensive margin in Column 1 and for the intensive margin in Column 2. Further, we see no effect of regressing each outcome on treatment (and controls selected from baseline covariates selected using double lasso). Standard errors are clustered at the subcounty, the level of randomization. "Mean DV" provides the mean of the dependent variable.

	Buy only inside home subcounty	Percent bought in home subcounty
Treat	0.03	0.02
	(0.04)	(0.04)
Observations	2139	2139
Mean of DV	0.60	0.76
R squared	0.17	0.05

Appendix J.2 Model Validation

In this appendix, we present additional model validation results. We find that the model does a good job of matching the extent of price dispersion across markets, as well as the relevant cross-sectional price patterns. It does less well at matching the average level of the maize price, which is lower in the simulation than in the real data; however, the size of the average inter-market price gap relative to the price level is similar (20% in our model versus 15% in the data).

In Figure J.1 below, we show the relationship between inter-market price gaps and distance in the lefthand panel and the relationship between market surplus status and price in the righthand panel, both using the results from the model simulation. We see that the relationships to both distance and surplus closely map those in the real data, as shown in Figures 1 and 4 in the paper. As shown in Figure 6, we also correctly capture the (non-targeted) relationship between surplus status and the direction of treatment effects on prices.

We can also replicate the reduced form exposure regressions within our model results as a validation test. Results are presented in Table J.2 below. Reassuringly, we see patterns that align with those in the data presented in Table 7 in the paper. Column 1 shows a significant impact on own price of randomized exposure to treatment in other subcounties. In Column 2, we see that (random) exposure to treated surplus subcounties reduces own price, as demand is diverted to treated surplus subcounties and away from one's own subcounty. Conversely, (random) exposure to treated deficit subcounties increases own price, as supply is diverted to treated deficit subcounties, albeit this latter effect is not significant. In Column 3, we interact own treatment status with these randomized exposure measures. We see the diversion effects are most pronounced among controls, as evidenced by the significant coefficients on exposure to treated surplus subcounties and exposure to treated deficit subcounties. Although the magnitudes differ across the model and data results, they reassuringly match direction and significance for almost all coefficients in the table.

Figure J.1: **Model-predicted price gaps by distance and price levels by surplus.** Left panel shows price gaps and right panel price levels, as a function of distance.

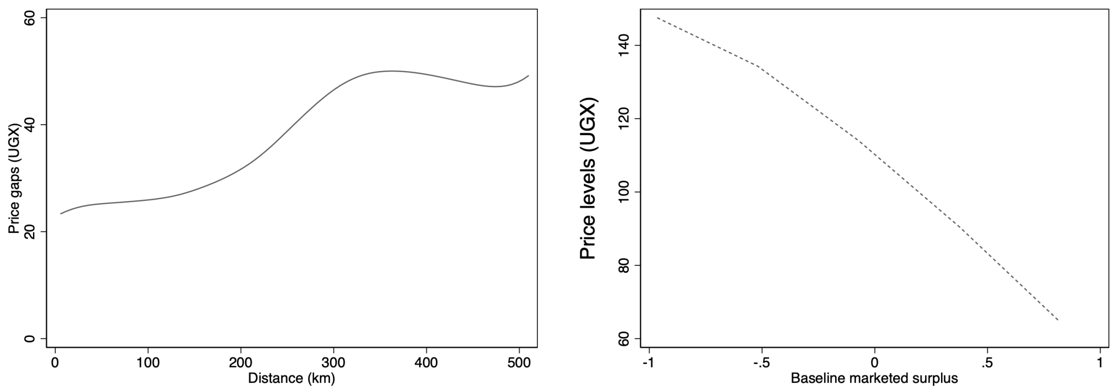

Table J.2: **Impact of experimental exposure, replication in model.** Results of Table 7 in paper replicated in the model.

	Ln Price	Ln Price	Ln Price
Exposure to treated markets	-1.45***		
	(0.27)		
Exposure to treated surplus markets		-3.24***	-4.12***
		(0.98)	(0.76)
Exposure to treated deficit markets		0.65	2.28***
		(0.75)	(0.70)
Own market treated x Exposure to treated surplus markets			1.59
			(1.86)
Own market treated x Exposure to treated deficit markets			-2.28*
			(1.22)
Own market treated			-0.08
			(0.08)
Observations	107	107	107

Appendix K Role of Distortions in Scale Economies

As discussed in Section 7.1, one explanation for the lack of farmer take-up is that most farmers are simply too small to have any motivation to adopt Kudu, because they will not be able to earn profits selling in other markets, even with the Kudu-induced reduction in the minimum threshold size required for cross-market trade. One may wonder how the existence of other distortions fits, if at all, into our model of scale economies.

Our model captures scale economies in the presence of heterogeneity in variable cost of operating a trading business, $a(\omega)$. While in the model farmers and traders are distinct agents, it is relatively straightforward to assume instead that, ex ante, there is a single pool of agents who farm and who have some potential productivity in trading, $a(\omega)$. Based on their draw of $a(\omega)$, agents decide whether to trade, in which case we call them "traders," or not, in which case we call them farmers. "Farmers" are those who do not end up trading in equilibrium because their cost draws make them too high cost / small scale to profitably operate on any cross-market route.

The key is that these heterogeneous variable costs $a(\omega)$ will be compared to the fixed cost of trade, F, to determine whether the agent engages in cross-market trade. There are many possible contributors to these heterogeneous variable costs. Of course, one factor may be that some agents simply have higher ability when it comes to trading. However, these heterogeneous $a(\omega)$ can also be micro-founded with other cost heterogeneity or even market distortions. The key question is whether these can be well represented in reduced form as higher variable costs.

To give a few examples: if farmers are not engaging in cross-market trade because they have higher opportunity costs of their time or capital, that would be appropriately modeled as a micro-foundation for having higher operating costs. Lacking the necessary capital to trade can also factor into $a(\omega)$, depends on how this is modeled. If there are credit market frictions that lead some people to pay a higher effective interest rate on working capital than others – for example, if farmers would have to go to a loan-shark who would charge extremely high interest rates in order to rent a truck to move grain across markets – then this would be well-modeled as farmers having a higher draw of variable operating costs to trade, through higher cost of capital. However, a more nuanced model in which farmers face an absolute cap on the amount of credit they can access would fit less well within our current framework. An analogy can be made extending this argument to frictions in labor markets: if someone has access to cheaper or more productive labor, this also falls neatly into the current framework, as long as it can be modeled as facing higher implicit wages. Risk fits less well under our current framework. An extremely reduced form representation of this, in which more risk averse farmers require higher risk premia to trade would fit within our set-up. However, a richer treatment of risk aversion would require modeling uncertainty, which is currently not featured in our model.

Given that some explanations fit well under our model of scale and others less so, it may be interesting to explore which explanations are correlated with cross-market trade (keeping in mind, again, that we do not mean to claim that scale economies are the *only* friction possibly at play). Though our study was not designed to speak directly to these issues, but we have some (very imperfect) proxies for these dimensions of heterogeneity in the data. Ideally, we would consider cross-market trade as the outcome, since that is what the model predicts responds to scale economies. However, so few farmers sell beyond their nearest market that we lack empirical variation for that analysis. (As we note in the main text, two-thirds sell at their farm gate, and the remainder at a market averaging 2km away. While we do not ask specifically whether these markets are outside the farmer's home subcounty, our best estimate given the distance is that 2.5%

of farmers sell directly to locations outside their subcounty.) Instead, we look at heterogeneity in attempts to use the platform.

In Table C.4 in Appendix C.2, we showed the predictors of posting to Kudu (which 26% of farmers do) and completing a deal on Kudu (which 2% of farmers do). In Column 1, we see that farmer size, as measured by log quantity sold, is a strong predictor of posting to Kudu. Column 5 suggests a noisily estimated relationship of similar magnitude between size and completing a deal on Kudu. In the remainder of the columns, we add proxies for some of the potential differences listed above. Column 2 adds proxies for opportunity costs, including a dummy for whether the head of the household's main activity is farming and whether the household owns another business. We see neither is a significant predictor of use of Kudu, nor does their addition change the point estimate on farmer size. Column 3 adds a survey measure for risk aversion, in which respondents choose between a lottery and a guaranteed amount of decreasing size. Again, we see no relationship with use of Kudu, nor change in the coefficient on size. Unfortunately, we lack a good measure of credit access, but as a rough proxy, in Column 4 we add the household's total non-farm income from wages, non-farm businesses, gifts and remittances, government support, and other sources. Here we see a significant relationship between non-farm income and use of Kudu, though it does little to reduce the relationship between size and use of Kudu. Columns 6-8 add the same proxies for completing a deal on Kudu. We find similar results in magnitude, but because only 2% of farmers sell on Kudu, the analysis in these columns is underpowered.

How does this conception of the heterogeneous costs potentially including distortions affect our policy implications? On one hand, if variable costs are high due to distortions, then this implies that lowering these distortions could generate entry into cross-market trade. However, as Figure 7 shows, the scale of most farmers is far from the threshold needed to enter into cross-market trade, suggesting that dramatic shifts in underlying variable trading costs would be needed to see substantial movement on this margin. Thus while distortions may be present, understanding them in the context of scale economies has important policy implications.

Appendix L 100% Scaled Treatment Effects

Table L.1: **100% Scaled Treatment Effects.** Treatment on All (ToA) at 100% treatment saturation. "Mean DV" provides the mean of the dependent variable.

	Any Trade	Number Traders	Volume (tons)
Kudu	0.017***	0.0070***	0.70***
	(0.0019)	(0.00066)	(0.086)
Dist (100km)	-0.10***	-0.04***	-16.01***
	(0.01)	(0.01)	(2.65)
Observations	11342	11342	11342
Mean DV	0.16	0.05	19.03

Appendix M Additional pre-committed analyses

For transparency, we describe here any additional analyses referred to in our pre-analysis plan, written in 2015, that we did not present above. We initially partnered with a private sector agribusiness firm to employ and train 210 Commission Agents (CAs) to promote Kudu. However, CAs were not reliable promoters, perhaps because they were recruited from among local traders and therefore lacked the incentive to promote the platform to others, although they were regular users of Kudu for their own trading operations. Ultimately, the study team hired staff to do this promotion, which was much more successful in generating broad awareness of Kudu. We had intended to conduct an experiment to test credit constraints among traders by offering loans to a randomly selected subset of CAs. 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 also 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 at baseline supposed.

Appendix N Derivation of estimating equations

We specify variable trade costs as $\tau_{ij}^{\sigma-1} \equiv d_{ij}^{\gamma} e^{-u_{ij}}$, where d_{ij} is the distance between markets i and j. Fixed trade costs are: $F_{ij} \equiv \exp\left(\phi + \theta \mathbf{ln} d_{ij} + \beta_1 \mathbf{1}_{ij}^K + \beta_2 \mathbf{ln} d_{ij} \mathbf{1}_{ij}^K - v_{ij}\right)$ where $u_{ij} \sim N\left(0, \sigma_u^2\right)$ and $v_{ij} \sim N\left(0, \sigma_v^2\right)$, and $\mathbf{1}_{ij}^K$ is an indicator variable equal to one if both i and j are treated.

Using these assumptions about the form of trade costs, we can derive three estimating equations, as follows.

Any trade

Taking logs of Equation 3, we get:

$$\ln\frac{1}{\sigma} + (1-\sigma)\left(\ln\frac{\sigma}{\sigma-1} + \ln\tau_{ij} + \ln p_i - \ln P_j\right) + \ln E_j + (1-\sigma)\ln a_L - \ln F_{ij} > 0 \tag{9}$$

and using the parameterization of τ_{ij} and F_{ij}, we get:

$$\ln\frac{1}{\sigma} + (1-\sigma)\ln\frac{\sigma}{\sigma-1} + (1-\sigma)\ln a_L + (1-\sigma)\ln p_i + (\sigma-1)\ln P_j + \ln E_j \\ + (\gamma-\theta)\ln d_{ij} + \mathrm{u_{ij}} - \phi - \beta_1 \mathbf{1}_{ij}^K - \beta_2 \mathbf{ln} d_{ij} \mathbf{1}_{ij}^K + v_{ij} \tag{10}$$

Finally, collecting terms and combining the destination-specific components into a fixed effect, we get:

$$\zeta_0 + \zeta_j + \frac{(1-\sigma)}{\sigma_\eta^2}\ln p_i - (\gamma+\theta)\ln d_{ij} - \beta_1 \mathbf{1}_{ij}^K - \beta_2 \mathbf{ln} d_{ij} \mathbf{1}_{ij}^K + \eta_{ij} > 0 \tag{11}$$

where ζ_j is a destination market fixed effect, and $\eta_{ij} = v_{ij} + u_{ij}$ so that $\eta_{ij} \sim N\left(0, \sigma_\eta^2\right)$ where $\sigma_\eta^2 = \sigma_v^2 + \sigma_u^2$

Dividing through by σ_η^2 and making use of the normal distribution of η_{ij}, we get an equation for the probability of there being any trade from i to j:

$$\rho_{ij} = \Phi\left(\hat{\zeta}_0 + \hat{\zeta}_j + \frac{(1-\sigma)}{\sigma_\eta^2}\text{ln}p_i - \frac{(\gamma+\theta)}{\sigma_\eta^2}\text{ln}d_{ij} - \frac{\beta_1}{\sigma_\eta^2}\mathbf{1}_{ij}^K - \frac{\beta_2}{\sigma_\eta^2}\mathbf{ln}d_{ij}\mathbf{1}_{ij}^K\right) \tag{12}$$

where hatted variables are equal to their counterparts in 11, divided by σ_η^2, e.g. $\hat{\zeta}_0 = \frac{\zeta_0}{\sigma_\eta^2}$.

This yields our first moment condition:

$$\mathbf{Pr}\left(T_{ij}=1\right) - \Phi\left(\hat{\zeta}_0 + \hat{\zeta}_j + \frac{(1-\sigma)}{\sigma_\eta^2}\text{ln}p_i - \frac{(\gamma+\theta)}{\sigma_\eta^2}\text{ln}d_{ij} - \frac{\beta_1}{\sigma_\eta^2}\mathbf{1}_{ij}^K - \frac{\beta_2}{\sigma_\eta^2}\mathbf{ln}d_{ij}\mathbf{1}_{ij}^K\right) = 0$$

Number of traders

Estimation of a log-linear equation describing the number of traders serving a route will be biased by selection on the any trade margin – the set of routes with non-zero traders have unobserved features that are systematically different from those with zero traders. To avoid this problem, we rearrange Equation 5 to consider the fraction of traders based in market i that do not serve market j, and take logs:

$$\text{ln}\left(1-\frac{N_{ij}}{N_i}\right) = k\left(\text{ln}\frac{\sigma}{\sigma-1} + \text{ln}a_L - \frac{1}{1-\sigma}\text{ln}\sigma\right) + k\text{ln}p_i - k\text{ln}P_j + \frac{k}{1-\sigma}\text{ln}E_j + k\text{ln}\tau_{ij} - \frac{k}{1-\sigma}\text{ln}F_{ij}$$

Again using the parameterization of trade costs from the previous section, we get:

$$\text{ln}\left(1-\frac{N_{ij}}{N_i}\right) = k\left(\text{ln}\frac{\sigma}{\sigma-1} - \text{ln}a_L - \frac{1}{1-\sigma}\text{ln}\sigma\right) + k\text{ln}p_i - k\text{ln}P_j + \frac{k}{1-\sigma}\text{ln}E_j + \frac{k}{1-\sigma}\left(\gamma\text{ln}d_{ij} + u_{ij}\right) - \frac{k}{1-\sigma}\left(\phi + \theta\mathbf{ln}d_{ij} + \beta_1\mathbf{1}_{ij}^K + \beta_2\mathbf{ln}d_{ij}\mathbf{1}_{ij}^K - v_{ij}\right) \tag{13}$$

And collecting terms and combining the destination-specific components into a fixed effect, we get:

$$\text{ln}\left(1-\frac{N_{ij}}{N_i}\right) = \varphi_0 + \varphi_j + k\text{ln}p_i - \frac{k\left(\gamma+\theta\right)}{1-\sigma}\text{ln}d_{ij} - \frac{k\beta_1}{1-\sigma}\mathbf{1}_{ij}^K - \frac{k\beta_2}{1-\sigma}\mathbf{ln}d_{ij}\mathbf{1}_{ij}^K + \varphi_{ij}$$

where $\varphi_{ij} = \frac{k}{1-\sigma}\left(v_{ij} + u_{ij}\right)$ is a mean zero normally distributed error term.

Value of trade on route

Both the any trade margin and the composition of heterogeneous traders serving a route are important for determining the value of trade on a route, and we need to account for both factors when taking the model to the data.

Starting from Equation 5, taking logs, and making use of the parameterization of τ_{ij}, we get:

$$\text{ln}M_{ij} = (1-\sigma)\text{ln}\frac{\sigma}{\sigma-1} + (1-\sigma)\text{ln}p_i + \text{ln}N_i + (\sigma-1)\text{ln}P_j + \text{ln}E_j + (1-\sigma)\text{ln}\tau_{ij} + \text{ln}V_{ij}$$

where $V_{ij} = AW_{ij}$ and $W_{ij} = \max\left\{\left(\frac{a_{ij}^*}{a_L}\right)^{1-k-\sigma} - 1, 0\right\}$ and $A = \frac{ka_L^{1-\sigma}}{1-k-\sigma}$ is a constant.

Now using the parameterization of τ_{ij} from the previous section, we get:

$$\ln M_{ij} = (1-\sigma)\ln\frac{\sigma}{\sigma-1} + (1-\sigma)\ln p_i + \ln N_i + (\sigma-1)\ln P_j + \ln E_j + \gamma \ln d_{ij} + u_{ij} + \ln V_{ij} \quad (14)$$

Finally, collecting terms and combining the destination-specific components into a fixed effect, we get:

$$\ln M_{ij} = \psi_0 + \psi_j + (1-\sigma)\, p_i - \gamma \ln d_{ij} + u_{ij} + \ln V_{ij} \quad (15)$$

Note that $\ln M_{ij}$ is only defined when $T_{ij} = 1$. Conditioning on $T_{ij} = 1$ means that both $\ln V_{ij}$ and u_{ij} are correlated with $\ln d_{ij}$, and so if they are in the unobserved error term, estimates of γ will be biased. Therefore, we need estimates of these variables to include on the right side of the estimating equation. We can do this through a Heckman type correction for selection on both the any trade and trader composition margins, following the approach in Helpman et al. (2008):

$\ln\left\{exp\left[\delta\left(\hat{z}_{ij}^* + \hat{\bar{\eta}}_{ij}^*\right)\right] - 1\right\}$ is a consistent estimate for $\mathbb{E}\left[\ln W_{ij}|., T_{ij} = 1\right]$, where where $\delta \equiv \frac{\sigma_\eta(k-\sigma+1)}{\sigma-1}$, $\hat{z}_{ij}^* = \Phi^{-1}\left(\hat{\rho}_{ij}\right)$ and $\hat{\rho}_{ij}$ is the predicted value from Equation 12, and δ is a coefficient to be estimated.

$B\frac{\phi\left(\hat{z}_{ij}^*\right)}{\Phi\left(\hat{z}_{ij}^*\right)}$ is a consistent estimate for $\mathbb{E}\left[\ln u_{ij}|., T_{ij} = 1\right]$, where $\hat{z}_{ij}^*$ is defined as above and $B = \frac{\mathrm{corr}(u_{ij}, \eta_{ij})}{\left(\frac{\sigma_u}{\sigma_\eta}\right)}$ is a coefficient to be estimated.