Kudu: An Electronic Agricultural Marketplace in Uganda

Neil Newman[1], Nicole Immorlica[2], Kevin Leyton-Brown[1], Brendan Lucier[2], John Quinn[3], and Richard Ssekibuule[3]

[1] University of British Columbia, Vancouver BC, Canada
[2] Microsoft Research, Cambridge MA, USA
[3] Makerere Univeresity, Kampala, Uganda

1 Introduction

This chapter describes our experiences designing an electronic market platform for agricultural trade, branded in Uganda as *Kudu* [24].

A significant challenge facing rural development is inefficiency in agricultural markets. One potential driver of such inefficiency is farmers lacking information about the national market for their crops and therefore selling in local markets at suboptimal prices. The result is not only lower prices for farmers (often a huge group, since 80% or more of the population in many African countries work in agriculture [12]), but also intra-seasonal and cross-locational price fluctuations that distort the market and reduce incentives for investing in productivity-enhancing inputs. Prior work [29] has demonstrated the existence of arbitrage opportunities both via buying and selling in different parts of the country as well as via paying for crop storage between seasons. Such inefficiencies are driven by information failures: market discovery occurs almost entirely through word-of-mouth interactions; buyers and sellers settle on prices through negotiation. Most gains from trade are captured by better-informed intermediaries [3]. Worse still, when both parties are insufficiently well informed, mutually beneficial trades simply may not occur [1,17]. In the long run, without accurate knowledge of nationwide agricultural demand, it is difficult for farmers to make good decisions about which crops to plant.

The internet has revolutionized many two-sided markets by making it easy for market participants to discover current conditions and to find each other. We were motivated by the idea that if farmers were both better informed about market conditions and better empowered to reach out to buyers beyond their immediate social network, they would achieve better market outcomes. Unfortunately, there was a massive hurdle to setting up an electronic marketplace in small-scale Ugandan agriculture: our potential user base consisted of smallholder farmers—farmers growing mainly for subsistence who occasionally have crops to sell—who have limited or no access to the web. However, penetration of feature phones—phones capable of sending and receiving voice calls and SMS messages, and running USSD applications—was nevertheless high. For example, the World Bank estimated that there were 55 mobile subscriptions per 100 people in Uganda in

2016 [2]. We therefore set out in 2011 to design an electronic marketplace in which a user could fully participate using only a feature phone.

In brief, our system operated as follows. Farmers and traders used their mobile devices to place *bids* (requests to buy) and *asks* (requests to sell) into a centralized nationwide database. Kudu identified profitable trades, which were then proposed to the corresponding participants. Users' trust in the system was enhanced by the availability of in-village support services, provided by Agrinet, a private-sector Ugandan intermediary; users were supported by a call center. Our platform also gathered price data and broadcasted it back to farmers and traders using SMS, drawing from a large set of national, regional, and local markets and providing a uniquely tailored information set to each user.

Kudu was first piloted in 2013 [29]; after a brief hibernation, it rebooted in partnership with Agrinet and Innovations for Poverty Action in May 2015. In this iteration, Kudu was part of a multi-year randomized control trial to assess its role on farmer welfare. While anyone was free to register and use Kudu, we only advertised and offered in-village support in certain parts of the country. We did not charge users anything to use Kudu; instead, a combination of grants and self-funding covered Kudu's expenses (which were dominated by the cost of human employees). The marketplace was then active for nearly three years and registered over 21,000 users through radio ads, village promotion meetings, and word of mouth. Users submitted nearly 30,000 asks and over 30,000 bids, resulting in more than 850 verified completed transactions involving over 5,000 tons of grain with a value of more than $1.9 million USD[4][5]. Grant funding for the project concluded in March 2018 and the market has been largely inactive since this time; efforts to commercialize operations were non-trivial to institute in our setting (see Section 6 for more details). Major takeaways from the randomized control trial were that Kudu had a measurable impact on reducing price dispersion between nearby markets. A back-of-the-envelope calculation by our collaborators [4] suggests that Kudu may have yielded net positive welfare benefits on the order of over $30M (after taking into account a projected reduction in trader profits of $1.53M and the $1.39M cost of running the intervention). While the error bars on this calculation are large and so the raw numbers should be taken with a grain of salt, the key observation is that so many farmers were affected by Kudu that the system would have yielded positive overall welfare effects as long as the average farmer benefited even slightly. We discuss these results in more detail in Section 5 .

[4] There are some (relatively small) discrepancies between the numbers reported above and those reported by our collaborators in [4]. Substantively there are two main causes: first, a subset of market activity from users belonging to Agrinet was not posted directly to Kudu and is therefore excluded from our metrics. Second, we differed in whether we recorded users who asked to place recurring "persistent" bids as having placed a single bid or as having placed multiple separate bids.

[5] For comparison, Uganda's agricultural sector is responsible for 24.5% of GDP, having a total value of about $6.5 billion USD.

2 Problem Statement

The problem motivating our work is that agricultural markets in developing countries such as Uganada have very high search costs, leading to inefficiencies. After factoring in transportation costs, there can be significant variation in commodity prices between locations in violation of the "law of one price". Using self-reported transportation costs from traders to construct a per-kg estimate of the cost of moving crops, our collaborators found differing prices between nearby markets that could not be explained by the cost of transportation alone [4]; we reached a similar conclusion in an earlier study [29]. Our hypothesis was that an electronic marketplace could empower smallholder farmers by connecting them to new trading partners and improving their awareness of prevailing prices. Our main design constraint was that the market had to be entirely operable using a feature phone.

There have been attempts in the past to improve agricultural markets through price advisory systems. Examples include Esoko's commodity index [7], Farmgain Africa [13], and Infotrade Market Information Services [16].[6] These services typically offer SMS subscriptions and radio based market information. However, experimental evidence that price advisory systems have been effective in improving farmer welfare is mixed [11,15,5,23,21]. These systems are typically based on manually gathered quotes that are sparse, geographically coarse, and biased by participants seeking to skew the reported statistics. Also, these systems often report only a single number (e.g., mean price) rather than distributional information, which can inform farmers about how to price their crop based on how urgently they need to sell. Moreover, evidence suggests that simply providing price information may be insufficient for farmers who do not have the means of actually accessing the better markets about which they may learn [20]. Smallholder farmers may need connections to specific buyers in these new markets or, in the likely event that they lack the ability to transport their crops themselves, they may even need those buyers to come to them. Kudu was aimed at comprehensively addressing this set of barriers to market access. It went beyond previous services, offering nuanced market information, direct market connections, and wraparound services needed to provide smallholder farmers truly improved market access.

A market designer needs to do more than just provide a means for people to interact with the market: they must encourage participation by making the market simple to use and its benefits obvious, ensure that strategic gaming does not undermine the market, and make certain that even as the market grows, finding a trading partner does not become overwhelming. Solutions to these challenges take different forms in different marketplaces: see [18] or [9] for surveys of how marketplaces tailor solutions according to their unique constraints. In addition to the aforementioned technological hurdles, unique challenges in our setting include technically unsophisticated and even illiterate users, the need

[6] See Section 1.2 of [28] for a survey of agricultural price information systems in Uganda.

to limit communication due to airtime costs, cultural resistance to adopting electronic markets, and high travel costs.

One alternative system design—which we rejected very early on—would simply have offered a database of bids and asks that users would have had to search manually. We rejected this idea because we believed that solely offering self-serve ads would not be enough to instill credibility and because we did not believe that searching through listings could be made effective on feature phones, especially when a prime consideration was location.

Artificial intelligence has been recognized as playing an increasingly important role in market design [19], for example to reduce search frictions. As we will describe later in Section 3.3, the key AI problem in Kudu was to decide what matches to propose to users and when to propose them. This involved both selecting the proper matching algorithm and accurately predicting whether proposed trades would be successful.

3 Method

In this section, we provide a detailed breakdown of how Kudu operated and some of the major challenges we faced. We begin with how Kudu gathered bids and asks (Section 3.1), including some reflections on pain points of various interfaces. Then, we discuss how matches flowed through our system and reasons why our proposed matches often failed to become deals (Section 3.2)). Next, we explain the evolution of our matching procedure in Section 3.3. Finally, we describe additional support we offered to facilitate trade (Section 3.4).

3.1 Gathering Bids and Asks

To place bids (asks) on Kudu, a user needed to tell us what crop they wanted to buy (sell), their requested buy (sell) prices, and desired (available) quantities[7]. Our services were available in four languages: English, Luganda, Luo, and Runyakitara. Our marketplace supported 76 crop types.[8] Crops differ in quality. This was problematic, because we wanted Kudu to be able to treat competing asks as interchangeable. After much reflection and user feedback, we did not adopt a quality grading system; two key hurdles were enforcement and inconsistency in users' abilities to grade crops effectively. Instead, we solicited bids and asks in terms of "fair average quality," inviting traders to negotiate a price adjustment at transaction time to deal with deviations from this quality level. Despite its inelegance, this system worked well in the Ugandan cultural context where point-of-sale bargaining is already common; this design choice did not lead to significant pushback from users.

[7] The units in which quantities were specified depended on the crop and reflects how they would usually be advertised. For example, bids for maize specified the desired weight in kilograms, whereas bids for potatoes specified the desired number of sacks.

[8] We refer to anything sold on Kudu as a crop, but a small fraction of our supported commodities were not plants, such as eggs, fish, and livestock.

When a new user placed a bid or ask on our system, we “registered” the user. One key fact we recorded about each user was their location, stored at the *parish* level (an administrative unit in Uganda made up of a small number of villages; Uganda has about 5,000 parishes). Parishes are grouped together into (nearly 1,000) *subcounties*, which form 136 *districts*, which in turn combine into 4 *regions*. Our assumption that people occupy fixed (and arbitrary) locations within a parish is obviously a coarse one; however, in a survey of our users, we determined that this assumption was reasonable for about 85% of them, and hence decided that a more complex system would not justify its cost.

Over the course of the project, we experimented with ways of enhancing the bidding language: for example, at one point we allowed “location filters" that specified that a buyer would only consider traveling within a specific geographic region. We dropped this feature because it was not well utilized; instead, we eventually accounted for travel costs when proposing matches (see Section 3.2).

To avoid hassling potential users with a complicated authentication system, we did not require users to set a username and password. Instead, users on Kudu were identified by their phone numbers. We worried about this breaking down when a user changed their phone number, or when multiple people with different devices wanted to share a single account, but it worked well in the common case.

There were four ways that users could interact with Kudu: sending an SMS, using our USSD application, visiting our website, or speaking to our call center. Each of these interfaces could be used to buy, sell, or request price information.

SMS

Users could send a toll-free SMS to our short code using any of the following templates:

buy [crop] [quantity] [unit price]
sell [crop] [quantity] [unit price]
price [crop]

If a user entered a crop name that did not exactly match one known in the system, we searched for a close match and automatically corrected it. The user received a confirmation SMS with their bid or ask and could call us if anything was wrong.

With the SMS interface, we were often unable to automatically parse messages because they strayed too far from the template. We received more than 3,000 SMS messages that could not be parsed. We failed to parse messages for many reasons: some had nothing to do with agriculture (e.g., questions about voting and health), some were written in an unsupported language, some were missing required information (e.g., the SMS consists of the single word “buy” or “sell”), and some mirrored the template (i.e., literal texts of “buy crop quantity price”). Other common mistakes were including units or descriptive information (“SELL DRYED CASSAVA 5000KGS 1500 PER KG”), or treating the SMS like a newspaper ad (“BUY GINGERS, LOCAL-400000SHS,FOREIGN-300000SHS,PER SACK OR

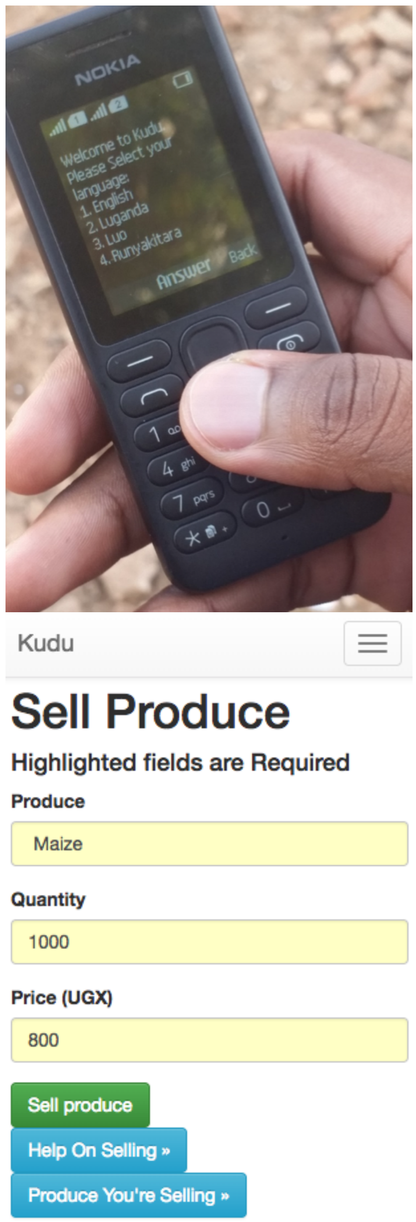

(00:47:58) Kudu:
Welcome to Kudu! Please reply with:
1. Buy
2. Sell
3. Price Information
4. Block Farmer / Trader
5. Language / Orulimi / Leb

(00:48:16) User: 2 (Sell)

(00:48:16) Kudu:
What crop would you like to sell?
1. Maize
2. Nambale beans
3. Black beans
4. Mixed beans
5. Other
0. Back

(00:48:25) User: 5 (Other)

(00:48:25) Kudu:
Enter crop
0. Back

(00:48:39) User: GNUTS

(00:48:39) Kudu:
Select option below or enter 0 and try again
1. Groundnuts
2. Goats
3. Piglets
4. Eggplant
0. Back

(00:48:46) User: 1 (Groundnuts)

(00:48:46) Kudu:
Enter quantity in Kgs
0. Back

(00:48:58) User: 700

(00:48:58) Kudu:
Enter unit price (per Kgs)
0. Back

(00:49:07) User: 4200

(00:49:07) Kudu:
Sell 700 Kgs of Groundnuts for 4200/= per Kgs?
1. Submit
2. Main menu
0. Back

(00:49:17) User: 1 (Submit)

(00:49:17) Kudu:
Your sell of 700 Kgs of Groundnuts for 4200/=
per Kgs has been placed on the market.
You can call 0780997402 for help.

Fig. 1: Top Left: Kudu's USSD interface running on a feature phone. Right: Sample USSD interaction for selling groundnuts (peanuts). Bottom Left: A user placing an ask on Kudu's web interface.

120KGS, ANY N0. OF SACKS OR KGS.CALL 256*********"). We assembled all of the messages that could not be parsed and our call center staff corrected these messages as they were able, phoning users when necessary.

Even when an SMS matched a template exactly, our system could still fail to capture the user's intent. If the crop name was misspelled, for example, our system could make the wrong correction. Users could also accidentally reverse the ordering of the positional quantity and price arguments, and both numbers could sometimes be in the same ranges making this difficult to identify (perhaps advocating for named arguments).

One of the main disadvantages of SMS is that it is not intuitive and requires training. An initial trial found that it was too difficult to register users via SMS (we requested that users send a "parish [parish]" message, but few did and it was hard to disambiguate between similar sounding parishes). Ultimately, first-time SMS users received a phone call from our call center to confirm this information. This and other trials have taught us that our SMS templates were not very flexible, limiting our ability to make changes to the bidding language over time.

We found that many users were able to grasp the SMS format after training, and the SMS system was inexpensive to run. However, our expectation was that it would be made obsolete by the presence of USSD, described below. In a six month study between September 2017–February 2018, SMS accounted for only 0.17% of bids (13) and 1% of asks (72).

USSD

Unstructured Supplementary Service Data (USSD) allows the user of a feature phone to open a real-time connection to an application and to engage in two-way data exchange, creating a responsive experience. A familiar example is an application for purchasing airtime. See Figure 1 for a sample Kudu interaction.

USSD solved many of the issues with SMS: a user could learn to navigate the interface independently; bids could be previewed before submission (allowing a user to confirm that the information was accurate); error messages could be reported in response to nonsensical inputs (e.g., 0 quantity). All of this was in principle possible with SMS but would have been unwieldy, requiring multiple back and forth messages. USSD has further advantages that are not implementable via SMS: e.g., one can implement a password login; sensitive messages are not stored on the device. We also found that having a USSD application was a sign of prestige, and in addition to the advantages we have described it acted as a strong signal to users that our service was backed by a serious enterprise. For all of these reasons, USSD has also been used in other development projects [26].

Unfortunately, despite all of the positives just discussed, USSD came with its own set of issues. One key problem was that it allowed messages no longer than 182 characters. This was very restrictive in practice: it made selecting from long lists difficult, such as when disambiguating parishes with similar names. Furthermore, and most importantly, sessions longer than 2 minutes timed out, leaving the user to start from scratch. This could lead to very frustrating experiences when menu sequences were long and when users had not prepared answers to all of the questions in advance. First-time USSD users trying to buy or sell was prompted for additional information to register, further exacerbating the time limit issue. In the end, we still had to dedicate call center employees to identifying incomplete USSD sessions and calling users back to place their bids for them.

We launched our toll-free USSD application in November 2015. Most USSD usage was to check price information, but in the six month study period mentioned above it also produced 1.2% (74) of our bids and 5.4% (383) of our asks. Our USSD service went offline on January 31, 2018 because our provider unexpectedly shut down all USSD services. We never revived the USSD application.

Web

Since our initial pilot, we have provided a web interface to Kudu as shown in Figure 1. While we did not expect this option to be used by smallholder farmers, the web interface was important for discovery and may have been appealing for more technologically sophisticated users. In the six month study period, 0.6% (35) of all bids and 0.6% (42) of all asks were placed via the web interface.

Call Center

Ultimately, as the observant reader will have noticed, the vast majority of our bids and asks in our current system did not come from any of our three "self-serve" options. Instead, in our six-month period of investigation, 97.7% (6167) of bids and 91.6% (6471) of asks were solicited by our call center. Monday through Friday, agents called traders and farmers and asked them if they had anything to buy or sell. This was very effective at thickening the market, and did not require users to be technologically sophisticated. However, it was very labor intensive—scaling linearly in the number of users on Kudu—so was not an economical approach in the long term.

Lastly, we note that a small fraction of bids and asks also came into our system by users calling in. We suspect that some users found this approach more convenient than dealing with our other interfaces.

The above described ways in which users contributed information into the system. We also sometimes needed to contact users at other times, for which we relied on phone calls and SMS messages. Calling a user is flexible and eliminates ambiguity about their intentions, but was expensive and required reaching the user over the phone. SMS messages were cost effective and could be managed on the user's own time but were problematic when users were unresponsive, e.g., due to illiteracy, lack of battery power, or sharing a phone among multiple family members, and required significantly more trust in the electronic matching system. We were strongly considering augmenting phone calls with an Interactive Voice Response (IVR) system to help illiterate users if the project had continued.

3.2 From Proposals to Physical Transactions

Once Kudu gathered bids and asks, the next step was to make matches. We will discuss the matching process in more detail in the following section; this section will focus on what happened after we decided to propose a match. During the 2013 pilot, when we matched two users, we sent them a text with each other's phone numbers, wished them a successful matching, and left them to their own devices. For some, this was enough to spark a transaction, but for many an automated text telling them to contact a random stranger did not instill enough confidence to lead to a trade. In the second wave, we employed *deal coordinators* to shepherd matches along and introduce a human element into the system. By the time phone numbers were exchanged, the deal coordinators had already spoken to the seller and buyer individually and could vouch for one to the other, incorporating past experiences when available.

When a match was proposed, a deal coordinator first called the seller. If the seller was interested in the match, they next called the buyer. Pending buyer interest, phone numbers were exchanged. The deal coordinator checked in with both parties to see if an agreement was reached. Finally, they followed up after the deal date and recorded what transpired. Of course, either side could pull out at any step in this process.

An additional job for deal coordinators was to look at asks that they were unable to match at the end of every day and give feedback about why they did not match (e.g., price or quantity too low) and allow users to change their prices.

Why Proposals Failed

We proposed trades that we genuinely thought would be profitable, yet only a small fraction of the trades that we proposed actually occured: e.g., in a detailed investigation between September 2017–February 2018, only 7% of the asks that we proposed matches for led to deals. While we certainly expected some failures, given that these were trades that Kudu identified as "profitable", it is worth asking why the fraction of failures was so large. The first, and by far most common, reason is that by the time we proposed a match, at least one of the parties was no longer in a position to trade. The second was that Kudu's assessment of what trades were profitable (based on the existence of a bid–ask spread) did not necessarily correspond to users' own assessments. Sometimes such mismatches were because of price or quantity issues, but most commonly the problems arose because of travel costs.

Transacting Outside The System Farmers are highly liquidity constrained. They tend to sell crops when they urgently need money, e.g., for school fees or medical bills. Cash kept on hand can be stolen or preyed upon by extended family and friends looking for hard-to-refuse loans. Therefore, when a user notified us that they wanted to sell, we needed to move quickly; a common failure mode of our proposals was that the seller had already sold their crop by the time we sent them a match. Conditioning on only those trades where the seller was still interested in transacting upon being contacted with a match, our success rate jumped to 16% (i.e., more than double). Understanding the need to sell quickly prompted us to institute expiry dates for bids and asks (7 and 3 days respectively). After this duration, unless we heard otherwise from the user, we assumed that a bid or ask was no longer valid.

Even if we did propose a match in time, and negotiations were successful, transacting outside of the system was still a concern. When a buyer traveled to meet their matched farmer, on the way they could encounter another farmer selling exactly what they wanted. If this occurred, they sometimes did not feel obligated to continue onwards to transact with the intended recipient, and instead took the closer trade. The only defense we had against this behavior was a reluctance among buyers to jeopardize their reputations with our deal coordinators.

Geographic Constraints After discounting trades that failed because one party had already sold, the next most common failures had to do with one party not wanting to travel. In order to accurately estimate trade profitability, it was important to develop a model of transport costs. We had coordinates for the location of every Parish. One might guess that transportation costs would have been roughly linear in Euclidean distance (we initially did), but this turned out to be highly inaccurate because of bodies of water, mountainous regions, and road quality issues. It was also tempting to use Google Maps to estimate travel times, but unfortunately that service did not know about many of the smaller roads that connect parishes. Instead, we used the road network data available for Uganda on OpenStreetMap [25] to model the road geography and generated approximate travel times between Parishes, using some assumptions about mean travel speeds on different types of roads. Given an estimated cost per kg per hour of transport, we could then roughly estimate the transportation cost for any given trade proposal. We estimated a travel time of less than one hour for more than 80% of our successful trades.

There are other reasons that some of our users refused trade proposals based on their locations beyond concern that transport costs would swamp potential gains from trade. Based on user surveys, the most common additional reasons (in descending order) were: bad roads and weather; risk of being robbed; uncertainty about the trustworthiness of business partners in the new area; not having any contacts and connections in the area; the reputation of the quality of crops in some areas (some areas are known for having poor quality crops); language barriers; too hard to determine if the journey would be profitable; worries about tax rates and local competition; war, insurgency, and epidemics. (Conversely, other responses contained sentiments like the very entrepreneurial "If a trade is profitable, nothing can stop me.") Needless to say, Kudu did not model all of these concerns when judging that a match was "profitable". Even if we had elaborated our profitability model to capture more of these concerns, a further complication is that agents' preferences about all of these issues were heterogeneous and difficult to elicit.

Trust and Reputation Our system contained no mechanism for ensuring that traders honored agreements they made. Buyers could renege on previously accepted deals or renegotiate at the time of transaction by threatening to leave (which was particularly problematic in the case of perishable crops). Sellers could also attempt to renegotiate at the last minute, leveraging the fact that a buyer could not easily walk away after paying for a truck rental and driving a long distance. Escrow was a natural solution to both problems: a buyer could deposit some fraction of the trade's price into an account managed by Kudu; the system could notify the seller that the money was in place; and the buyer could tell Kudu to release the money when the goods were transferred. One issue was that this would still not have eliminated risk on the buyer side if the goods were not as promised, since the farmer would not have had to put anything in escrow; Kudu would need to mediate disputes. Various other practical hurdles made this idea challenging to implement in practice: traders could be exposed to mobile money phishing scams; kiosks to withdraw mobile money were not yet prevalent enough in rural areas; and (probably most importantly) fees for mobile money transactions were far too high.

Another way of increasing users' trust in the system might have been to integrate a more robust reputation system. Kudu let users blacklist anyone with whom they had a poor experience, but we could have gone further. Since we had a transaction history for every user on Kudu, we could have rewarded completion of successful matches by increasing users' priorities, and conversely deprioritized users who did not follow through. Deal coordinatiors could also have used such reputation information to convince users to trade with each other. Challenges to a more sophisticated reputation system were that it was expensive to verify successful trades and that "whitewashing" was easy: Kudu identities were tied to phone numbers, and new SIM cards were inexpensive. Also, while traders transact frequently, farmer usage could be more sporadic (e.g., they often made only one big sale on Kudu per season), so it would have taken a long time to develop a reliable signal of farmer reliability.

3.3 Matching

A significant technological challenge in the Kudu ecosystem was choosing which bids should match with which asks. We wanted to do so in a way that maximized the value that Kudu brought to the marketplace.

To quantify the value of a trade, we defined a scoring function that mapped each possible trade to a real number based on the bid and ask quantities and prices, as well as other factors such as the distance between both parties and reputation considerations. Our aim was for the scoring function to capture gains from trade. As a first attempt we set it equal to the price differential times the quantity minus a linear function of the distance between the participants, representing the cost of travel due to fuel consumption. We noticed during deployment that, because users bid and bargained strategically, this scoring function sometimes assigned negative values to trades that actually went through in practice (and therefore must actually have generated positive gains from trade). This strategic behavior was evident in Figure 2, which shows that users sometimes traded despite notionally negative gains. We therefore refined our implementation to assume that bid and ask prices were merely signals rather than binding constraints. We used these declared prices to fit probability distributions that modeled each user's true price,

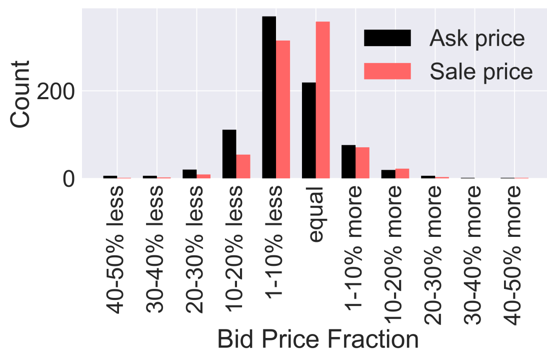

Fig. 2: A histogram of bid prices as a fraction of ask prices and final sale prices for verified transactions. If users were truthful, no trades would ever have occurred with a bid price lower than the ask price; clearly this was not the case. However, most of the mass is distributed around the center, suggesting that user prices still conveyed useful signaling information in most cases. The buyer usually payed slightly less than their bid price in the final sale.

and then used these to calculate potential gains from trade. That is, we sampled from both price distributions, rejected samples where the bid price was lower than the ask price, and computed the expected gains.

We employed two separate strategies for choosing matches. Our initial solution involved deal coordinators accessing our database to make *manual matches*. We then migrated to a *hybrid* of manual matches and *automatic matches*, proposed by an optimization algorithm.

Manual Matching

When the number of bids and asks was low, an effective method for clearing the market was via *manual matching*: having deal coordinators manually look at the database of bids and asks and decide which parties to match. Kudu employed five deal coordinators, each specializing in particular treatment districts and local languages.

Such a system leveraged human intelligence and human relationships in the matching process. Deal coordinators developed intuition about which parties made good matches. This intuition was based in part on features that were hard to quantify, e.g., the personalities of the trading participants. Furthermore, deal coordinators interacted repeatedly with the same parties, developing valuable social capital and trust. A participant who had a personal relationship with a deal coordinator might have been more willing to submit bids or asks to Kudu than one who only interacted with the system electronically. (An illustration: one

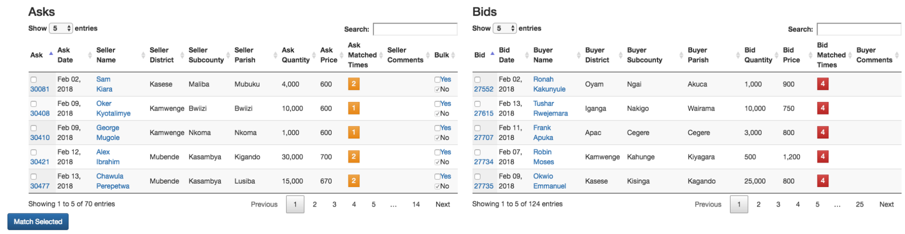

Fig. 3: The manual matching interface on Kudu. Asks are shown on the left and bids on the right. A deal coordinator selected one from each column to create a match. This approach began to break down when the number of bids and asks grew large.

coordinator told us that her conversations typically involved an initial discussion about the participant's family before any discussion of the potential match.)

However, the downsides to manual matching were significant. First, it required (expensive) human employees, the number of which needed to scale at least linearly with the number of participants on the platform. (Observe that the number of potential matches grows superlinearly, even after geographic and other constraints are taken into account.) Second, it was unlikely that deal coordinators were effective at optimizing a global objective such as overall gains from trade. Discussions with the deal coordinators suggested they followed a local greedy heuristic, selecting a single buyer and then searching for the best seller that might match with that buyer. The selection of the buyer was based on their estimate of how likely the buyer was to accept a trade at that given time. The selection of the seller was then entirely based on the value of the seller to that buyer. Importantly, this process ignored the value of the seller to other potential buyers. Finally, as the system grew, search frictions, like the size of the database and the limited sorting tools, made it difficult for deal coordinators to find the best matches, even according to their own metrics.

Automatic and Hybrid Matching

An *automatic matching* algorithm takes as input a set of bids and asks and algorithmically proposes trades. Deal coordinators follow up on the trades recommended by the algorithm. In a *hybrid matching* system, deal coordinators revert to manual matching once all automatic matches have been processed. Our hope was that, as our automatic matching component improved, the participants would have been be able to carry out proposed matches without intervention by the deal coordinators. Such a system would have scaled easily as the market grew, could have optimized global objectives, and would not have been significantly hindered by search frictions. However, a fully automatic system would have sacrificed the human intelligence and social capital of the deal coordinators.

Our initial 2013 pilot ran a heuristic algorithm that periodically went over all of the bids in the system in an arbitrary order and matched each bid with an unmatched ask with high score according to our scoring function (see [29]). This approach addressed the issue of facilitating search for deal coordinators. However, because it did not always intelligently choose the order in which bids were processed, it did not optimize gains from trade. Furthermore, it was not able to leverage deal coordinators' background knowledge, forcing them to concentrate on automatically selected matches.

In 2015, we introduced an improved match optimization algorithm, which ran three times a day. At run time, the algorithm simultaneously considered all bids and asks in the system and proposed a feasible set of trades that maximized the total gains from trade, according to our scoring function. (This amounted to running a maximum weight matching algorithm in a bipartite graph; the optimization could thus be performed efficiently.) Our solution also attempted to help the participants find a "fair" price. We set the recommended price of a transaction to the the minimum competitive (i.e., Walrasian) prices for the matching market [14], making truthful bidding a dominant strategy for buyers but giving farmers incentives to manipulate their sale prices.[9] In our idealized market, buyers and farmers would have traded at our recommended price. In reality, buyers and farmers typically negotiated prices outside the system, and the recommended price was not even communicated to participants when the deal coordinators found it unhelpful.

Our automated matching system was nowhere near as successful as the manual matching system in terms of producing deals. By comparing the workflow of automatic matching to manual matching, we identified a large problem[10] that our system faced: as we proposed trades only three times per day, many participants faced long wait times before matching. Additionally, many of these matches quickly unraveled: most commonly when a deal coordinator called the seller and found out that their ask was no longer valid (see Section 3.2). There was little the system could do for the matched buyer in such cases—even if they had a strong bid, the algorithm would likely already have matched the strongest other asks to other bids.

In November 2017, we shifted to an automatic system, called *Kudu AI*, that offered matches continuously. The system assigned a priority to each buyer, equal to the highest potential surplus from a trade involving that buyer. When a deal coordinator entered the system, she was presented with a list of buyers, sorted by priority. Our intention was that the deal coordinator would choose to

[9] As demonstrated by a celebrated theorem due to Myerson and Satterthwaite [22], it is impossible to make truthful bidding a dominant strategy for both sides of the market. We focused on incentivizing buyers because they typically constituted the short side of our market and because they had access to a more robust array of outside options.

[10] Another potential problem with this and the 2013 system was the interplay of manual and automatic matches: we worried during manual matching periods, the deal coordinators might cherry-pick the best matches, leaving less for the automatic matching to work with.

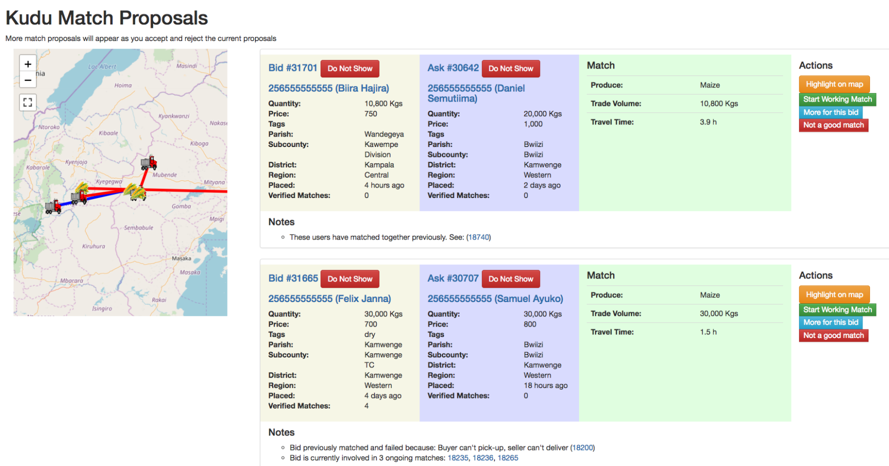

Fig. 4: The Kudu AI interface. Deal coordinators were shown matches corresponding to the largest gains from trade. Matches were overlaid on a map of Uganda. Characteristics of the match and the involved users were also shown: for example, we highlighted if the two users have matched together before or are engaged in any ongoing transactions. Deal coordinators could accept or reject proposals, refreshing the available choices.

work with the highest-priority buyer in this list[11]. Once the deal coordinator selected a buyer, the algorithm selected five possible sellers for the buyer, in decreasing order of gains from trade. Deal coordinators could accept or reject any seller. Rejections came with reasons that help us improve the algorithm. When used as intended, this process mimics the greedy algorithm for maximum matching and hence captures a constant fraction of the gains from trade in a static system. Furthermore, it gave deal coordinators more flexibility, in the form of rejecting sellers or processing buyers in an order other than the recommended one. We hoped to use this flexibility to improve our algorithm's scoring function by leveraging the human intelligence of the deal coordinators. Ultimately, the project went dark before we were able to thoroughly learn from this interface.

3.4 Facilitating Trade

We partnered with AgriNet, one of Uganda's largest private-sector brokerage companies, to promote Kudu with farmers and facilitate trades with on-the-ground services. As part of this collaboration, AgriNet rolled out their agent model into the communities in which we introduced Kudu. Agents promoted Kudu by advertising the service to both farmers and local traders, typically

[11] The deal coordinator could also search the system for a specific buyer; this would defeat the global optimization guarantees.

via house-to-house visits and announcements via loudspeaker in markets. They then followed up this advertisement with a village-based meeting in which they provided information on Kudu services and training on how to use the system. Agents also distributed their phone numbers so that users could call them with questions about the service or if they needed help registering an ask or bid later in the season.

In addition to promotion and training, AgriNet offered several additional services designed to address issues that could hinder transactions between buyers and sellers, even once they had found each other on the Kudu platform. First, because many farmers in Uganda operate at a small scale, surpluses are often quite diffuse, and aggregation is necessary to attract large national buyers [27]; this requires both coordination and access to capital. Kudu had the capacity to note and electronically bulk lots of the similar crops available for sale in nearby locations; AgriNet agents were available to provide on-the-ground coordination of this bulking. In order to finance this bulking, AgriNet offered its agents access to Cash on Bag (COB) credit. Agents in turn could use this credit to pay cash-constrained farmers for 50% of the value of their crop upon bulking with the agents and 50% upon sale to the buyer.[12]

Another challenge that may have limited buyers' willingness to trade on the platform was the risk inherent in directly trading with farmers in remote villages with whom they had not yet developed trust. Buyers needed to make up-front investments in transportation out to rural villages without any guarantee that agreements made in advance regarding quantity or quality of available crops would be carried out as promised once they arrived. Buyers sometimes instead choose to trade only with trusted brokers or other traders with whom they have had repeated interactions, resulting in a fractured chain of many short-distance, relationship-based exchanges [10]. To address these risks, AgriNet offered a "transaction guarantee" service. This service, designed to reduce the risk to buyers inherent in engaging in a more anonymous marketplace, offered transport cost compensation for buyers who traveled to a rural sale point and were disappointed.

As a further measure to address the risks involved with remote trading, local monitoring agents were used to certify the details of the transactions. After receiving a call from a deal coordinator about an agreed-upon deal, a monitoring agent visited the seller to check the quality and quantity and communicate his findings to the buyer. The agent was present at physical transaction and oversaw exchange of money, providing regular updates to deal coordinators throughout the transaction process. These monitoring agents provided other services as well, such as recruiting and training new users through visiting local markets and village promotion meetings. Finally, the monitoring agents helped smooth out price negotiation by being physically present. For example, we have heard from deal coordinators that there was sometimes a tension after both parties had exchanged phone numbers regarding who would call the other first, out of fear

[12] In addition to this bulking procedure, we had future plans to implement automated bundling as part of the Kudu matching process.

of looking desperate. The monitoring agent could help address this issue by mediating the negotiation.

3.5 Price Information

Soliciting bids and asks on Kudu was challenging due to technological and informational constraints. We have already discussed technological constraints, which were largely beyond our power to change: e.g., few farmers have smart phones, which forced us to solicit information through limited interfaces; many are illiterate, which forced us to rely on a call center. The key informational constraint was that users were reluctant to participate in the system without knowing current market rates for the crops they were interested in trading. To tackle this challenge, we provided price quotes. Ideally our price information would have come from verified transactions that had occurred on Kudu, but our system was too small to consistently have sufficient data in enough districts to be useful.

A next hope would have been to use the bids and asks, which were more plentiful than verified transactions. However, we expected the bids and asks we received to be somewhat biased, since our platform existed in the context of outside options. Our users were only interested in using the platform if it could get them a better deal than they could otherwise find. This meant that ask prices were usually inflated and bid prices were usually shaved. Broadcasting this data could have been particularly problematic because it might have created a harmful feedback loop in which farmers received overly optimistic price information from Kudu and then priced their crops accordingly.

As a result, an initial version of our price information system simply reported the median ask price, over the previous week, for a given crop (by default, nationally, but optionally scoped to a given location). Our final implementation of this service used biweekly survey data to determine market prices for select crops in the treatment districts and reported the 25–75 percentile of wholesale prices at markets. Collecting data in this way yielded coarse information, was expensive, and resulted in stale quotes. It was nevertheless highly popular among our users.

4 Resource Requirements

Kudu's software requirements were minimal—it ran on a virtual machine with 2 CPUs and 4GB of memory. It was built as a Django [8] web application and did not utilize any commercial software. Matching algorithms came from the iGraph [6] library. As described in Section 3.2, OpenStreetMap [25] was used for distance calculations in our transportation cost model. Costs to run the market (prior to any wraparound services) consisted of a salary for the platform manager, short-code fees, and some radio ads.

The operating costs for running the platform from May 2015 to March 2018 (during the randomized control trial) including all of the wraparound services were $927,190, with most of this going towards salaries for deal coordinators, call center operations, monitoring agents, and other program staff. More detailed accounting is available in [4].

5 Field Evaluation

Over the course of our operations, users submitted nearly 30,000 asks and over 30,000 bids, resulting in more than 850 verified completed transactions involving over 5,000 tons of grain with a value of more than $1.9 million USD.

Figure 6 shows the cumulative value of transactions on our platform, broken down by crop. Figure 7 shows all verified transactions on Kudu, plotted geographically. Figure 5 illustrates the users active on our platform over time, separating existing and new users.

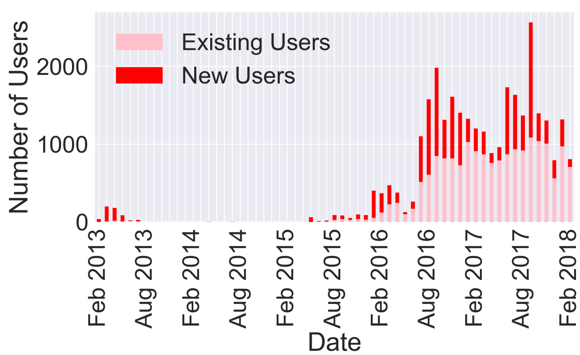

Fig. 5: Active users over time. Each bar represents the number of unique active users in a one month interval. New users that month are highlighted. An active user is defined as a user who used any of Kudu's services during the given month. In total, 11,861 unique users engaged with the platform.

Kudu was embedded within an extraordinarily large scale randomized control trial conducted between May 2015–March 2018, spanning six agricultural seasons [4]. The trial encompassed 11 districts, amounting to 110 subcounties and 236 trading centers. Half of the subcounties were reserved as a control, the other half were given the in village services described in the sections above. The study also recruited informants that gathered market prices in 236 markets every two weeks for four major crops (crops, maize, nambale beans, matooke bananas, tomatoes). These prices were blasted to a subset of the treated farmers and to all treated traders. Impacts on farmers and traders were gauged through extensive

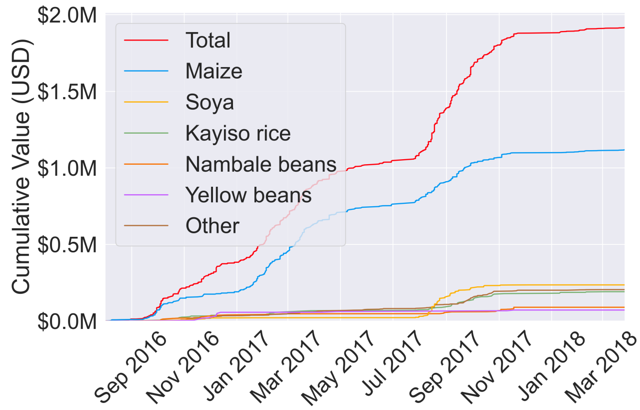

Fig. 6: Cumulative value of verified transactions between September 2016 and March 2018.

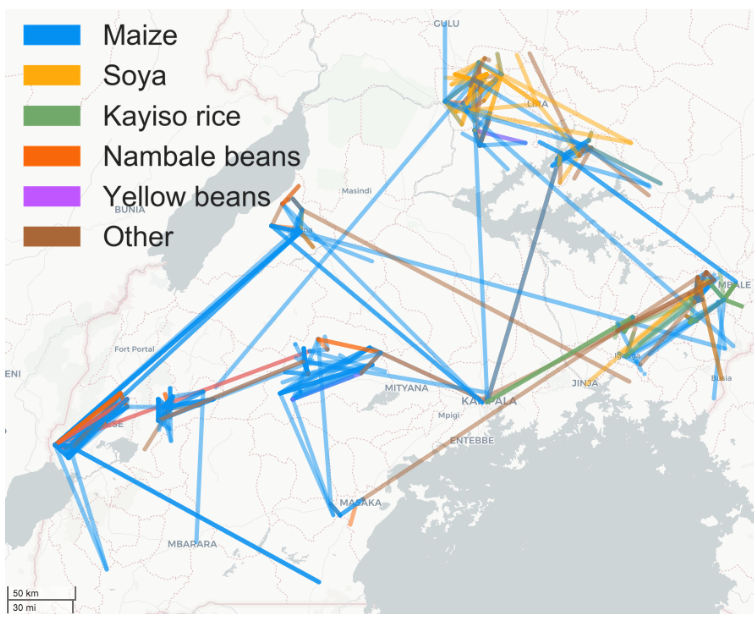

Fig. 7: Geographical range of our verified transactions, with edges linking the reported parishes (small villages) of buyers and sellers respectively. About a third of verified transactions occurred within the same parish and are not visible on this map. The large clusters correspond to the 11 districts throughout Uganda in which Kudu was supported through in-village services. We note that Kudu spread beyond these treatment districts.

surveys before, during, and after the intervention. Surveys tracked 1,457 traders[13] and 2,971 farming households). Farmers were stratified by their proximity to a marketplace.

We refer the reader to [4] for detailed results of the study, but share some insights here:

- By far, most activity on the platform from study participants came from traders, both on the ask and bid sides. Our platform was mainly used to arrange sales between traders, not to arrange farmgate sales with smallholder farmers as we had initially envisioned. The quantities that smallholder farmers had available to sell were typically too small to be of interest to their counterparties on the other side of the market.
- Trade flow between treated markets increased and nearby markets saw a reduction in price dispersion of 8% and 15% as one and both markets were treated respectively.
- Simply blasting SMS price information without providing the rest of the services was ineffective. Facilitating trade was essential for impact.

Perhaps the most interesting output of the study is a back-of-the-envelope welfare calculation for Kudu. It works as follows: Based on survey data, the intervention reduced trader profits, amounting to a net harm of $1.53M. When the costs of running the intervention are added, the social cost of the platform was estimated to be $2.92M. Revenue effects were then estimated for farmers by regression using the following features: whether or not they lived in a treated subcounty, lived near a marketplace, whether they were receiving SMS blasts[14], and interaction terms. The computation found a $124K welfare benefit to farmers in the study sample, at first glance much less than the costs. However, if those benefits are extended to non-surveyed households of which they should be representative in treated subcounties (which can be justified given evidence that the intervention moved trade volumes and prices), then the calculation is dominated by the 919,697 households that are "far" from a marketplace and did not receive SMS blasts. Such households were estimated to each receive a $12.29 benefit, leading to total benefits of $34M and a net benefit of over $30M.

While we caution that this is a back-of-the-envelope analysis using non-statistically significant values (the standard error on the $12.29 benefit is $26.66), the important observation is simply that there are so many farmers, even improving their situation a small amount ($0.79/year) would have positive overall welfare effects.

6 Redesigning the Market To Minimize Human Intervention

Funding for Kudu ended in March 2018. Without a call center or deal coordinators, market activity ceased. Sustaining Kudu required commercial viability, which

[13] This is a high degree of coverage: 83% of the traders meeting the survey criteria in the study's districts.

[14] Only possible when living in a treated subcounty.

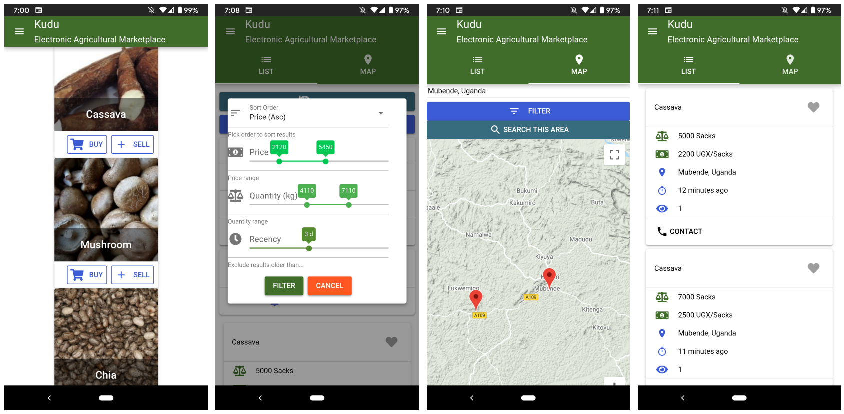

Fig. 8: Screenshots of the Kudu application running on an Android smartphone. Users select from a menu of crops and can filter on features of the listing such as price, quantity, recency, and geography. Users interested in a listing receive the owner's contact deals and manually initiate a trade.

proved to be a tricky puzzle: we had no practical way to implement commissions (we could not verify transactions without a physical agent present at point-of-sale, and most deals were made in cash so there was no money flowing through the platform from which we could have extracted a cut). Furthermore, we did not expect users to tolerate listing fees: [4] estimate the required break-even fee even for a bare-bones Kudu with no wraparound services at a dollar per listing, a value small relative to the scale of most transactions but certainly too high for farmers. A final business model we considered would be be to sell up-to-date price information derived from Kudu's database, but this would only have been possible if Kudu was already operating at a scale large enough to reliably generate data.

We reflected on how we might revive our marketplace with significantly lower operating costs, having concluded that commercialization was unlikely. Our biggest expenses were employee salaries for deal coordinators, the call center, and monitoring agents. These services gave our brand credibility and made interaction with our system possible for many users for whom it would otherwise have been too technologically complex. Since a fully automated system was impractical for our target users, our goal was not to eliminate it entirely but rather to minimize its size. Betting on smartphones slowly penetrating the market, and in acknowledgement of the heterogeneity and complexity of user preferences, we designed an application that allows users to browse and create listings themselves. As discussed in Section 2, while conceptually these services could be offered through a complex USSD application on a feature phone, practically we believe that a smartphone UI is required for users to effectively search through listings given their lack of substitutability due to factors like location. The application allows users to filter products by quantity, price, and recency. Instead of placing "bids", users specify search criteria and receive notifications when new matching asks are posted to the platform. We envisioned this application being used by traders or large aggregators; a tiny call center could support smallholder farmers who would place their listings over SMS. Screenshots of the application are provided in Figure 8.

We released the application on the Play Store (https://play.google.com/store/apps/details?id=com.kudu.market) and began a very small (hundreds of dollars) marketing campaigns via Facebook advertisements to get a small number of users on the platform. Our plan was to work through initial bugs and then expand the scope of the advertising campaign as well as use radio ads to recruit more users. Our initial Facebook ad campaign had limited impact: we were not able to create enough thickness for the market to be useful; we never moved on to radio ads. Without a local partner it was difficult to gain much insight into how the app was being used and how we might promote it. We met with several candidate organizations but never arrived at a partnership. Around March 2019, as the chaos of COVID-19 hit, we stopped actively working on the application. We have posted our source code publicly at https://github.com/newmanne/KuduApp; we hope that interested end users and/or researchers will build upon it in the future.

7 Lessons Learned

When we began this journey, we believed that it was most critical to settle on the right market design and to back the market with an effective clearing algorithm. We planned for market growth that would render it impossible to manually intervene in each transaction and often discussed the exciting AI challenges we would soon face: How machine learning on large datasets of successful and unsuccessful proposals would allow us to replace our hand-crafted model of gains from trade; how we would embed travelling salesman problems into our matching framework to bundle several small nearby asks; or how we should best implement reputation systems. We have since come to appreciate the importance of additional issues that have little to do with market design: identifying reliable mobile operators; balancing our local partner's competing interests against our own; tensions between wanting to act like an agile startup and ensuring the sanctity of a randomized control trial; and the difficulties of working on a problem that was, for most authors, located far afield.

The last point cannot be overstated: we often had little visibility into how the software was being used on the ground and in retrospect could probably have learned a lot more by more frequently checking in with our deal coordinators, who were much more knowledgeable regarding cultural norms (for example, an expectation that price negotiation continues right up to the point of trade) and local conditions. Our experiences with matching taught us about the effectiveness of human-in-the-loop designs, both in terms of deal coordinators' abilities to identify profitable trades but also in terms of their abilities to elicit preferences from traders. We learned that human capital is essential to running a marketplace in a location such as Uganda given the current state of technological integration.

We came to appreciate the extent to which electronic markets in the developed world depend upon the ubiquity of smartphones and credit cards as we struggled with their absence.

In retrospect, we were too slow to respond to the fact that traders were our primary market participants, instead clinging to the vision that smallholder farmers could successfully sell through our system. In general, their listings were simply too small for most traders to care about, and perhaps we should have spent fewer resources trying to popularize Kudu among this demographic. Furthermore, given that our main players were traders, we might have thought differently about the resources available to and level of sophistication of our users. One possibility is that such a trader-oriented system could still benefit smallholder farmers as they would be selling into more integrated markets, even if they did not actually use the system themselves.

Another takeaway from the project is that information is most useful when it is actionable: the randomized control trial showed that our centralized marketplace enabled significant welfare gains beyond price blasts alone. Information should also be timely: many of our earlier proposals failed because we did not appreciate that listings were highly time sensitive and required quick responses.

8 Conclusions

This chapter has described Kudu, an electronic market for agricultural trade in Uganda designed to combat inefficiencies in rural agricultural markets. Traders and farmers posted bids and asks using a feature phone. Kudu then proposed matches, leveraging a combination of optimization algorithms, data-driven models, and human expertise. Our system was augmented by a rich variety of support services that help to facilitate trade. The system was active for several years, involving tens of thousands of users and yielding verified trades totaling more than $1.9 million USD. Results from a multi-year randomized control trial demonstrated that Kudu and all of the wraparound services accompanying it was successful at reducing arbitrage opportunities between nearby markets. A rough welfare calculation suggests that Kudu achieved significant net welfare benefits.

Acknowledgements

We are profoundly appreciative of our collaborators Craig McIntosh and Lauren Bergquist for their deep involvement in Kudu, including but not limited to designing and analyzing the randomized control trial and all of the wraparound services described in this chapter. Without their partnership, Kudu would not have been able to achieve the scale it did and we would have been unable to evaluate the project as thoroughly. We also acknowledge their time and effort in providing useful feedback on this chapter.

None of the practical results described in this chapter would have been possible without the hard work of Kudu employees, especially our five deal coordinators who gave us a lot of useful feedback and domain knowledge on matching in this setting: Miriam Akera, Doreen Alupo, Brian Tibenderana, Rose Kadondi, and Nuuru Kabahuma. Stephanie Annijas and Laza Razafimbelo at Innovations for Poverty Action made critical contributions to maintaining daily market operations. We gratefully acknowledge the support of Agrinet and in particular its Director, Paul Nyende.

We are thankful for contributions from past research assistants Laura Carvajal, Arman Raina, and Peter West.

We obtained map data from OpenStreetMap https://www.openstreetmap.org. This project received institutional support from Microsoft and UBC and financial support from Google via a Faculty Award and NSERC via an E.W.R. Steacie Fellowship and a Discovery Grant.

References

1. Aker, J.C.: Information from markets near and far: Mobile phones and agricultural markets in Niger. American Economic Journal: Applied Economics **2**(3), 46–59 (2010)
2. Bank, T.W.: Mobile cellular subscriptions (per 100 people). https://data.worldbank.org/indicator/IT.CEL.SETS.P2 (2018), accessed: 2018-03-02
3. Bergquist, L.F.: Pass-through, competition, and entry in agricultural markets: Experimental evidence from Kenya (2017)
4. Bergquist, L.F., McIntosh, C., Startz, M.: Search costs, intermediation, and trade: Experimental evidence from Ugandan agricultural markets. CEGA Working Paper Series No. WPS-173 (2021). https://doi.org/https://doi.org/10.26085/C3759K
5. Camacho, A., Conover, E.: The impact of receiving price and climate information in the agricultural sector. IDB Working Paper Series, No. IDB-WP-220 (2011)
6. Csardi, G., Nepusz, T.: The igraph software package for complex network research. InterJournal **Complex Systems**, 1695 (2006), https://igraph.org
7. David-West, O.: Esoko networks: facilitating agriculture through technology. GIM case study no. b061 (2010)
8. Django Software Foundation: Django, https://djangoproject.com
9. Einav, L., Farronato, C., Levin, J.: Peer-to-peer markets. Annual Review of Economics **8**, 615–635 (2016)
10. Fafchamps, M., Minten, B.: Relationships and traders in Madagascar. The Journal of Development Studies **35**(6), 1–35 (1999). https://doi.org/10.1080/00220389908422600
11. Fafchamps, M., Minten, B.: Impact of sms-based agricultural information on Indian farmers. The World Bank Economic Review **26**(3), 383–414 (2012)
12. Farm Africa: Our work in Uganda. https://www.farmafrica.org/uganda/uganda (2018), accessed: 2018-03-02
13. Farmgain Africa: Farmgain Africa. http://farmgainafrica.org/ (2018), accessed: 2018-03-02
14. Gul, F., Stacchetti, E.: Walrasian Equilibrium with Gross Substitutes. Journal of Economic Theory **87**(1), 95–124 (July 1999)
15. Hildebrandt, N., Nyarko, Y., Romagnoli, G., Soldani, E.: Price information, inter-village networks, and "bargaining spillovers": Experimental evidence from Ghana. Working Paper (2020)
16. Infotrade: Infotrade - market information services. http://www.infotradeuganda.com/ (2018), accessed: 2018-03-02
17. Jensen, R.: The digital provide: Information (technology), market performance, and welfare in the South Indian fisheries sector. The quarterly journal of economics **122**(3), 879–924 (2007)
18. Kominers, S.D., Teytelboym, A., Crawford, V.P.: An invitation to market design. Oxford Review of Economic Policy **33**(4), 541–571 (Nov 2017). https://doi.org/10.1093/oxrep/grx063, https://academic.oup.com/oxrep/article/33/4/541/4587951
19. Milgrom, P.R., Tadelis, S.: How artificial intelligence and machine learning can impact market design. Tech. rep., National Bureau of Economic Research (2018)
20. Mitra, S., Mookherjee, D., Torero, M., Visaria, S.: Asymmetric information and middleman margins: An experiment with Indian potato farmers. Review of Economics and Statistics (0) (2017)
21. Mitra, S., Mookherjee, D., Torero, M., Visaria, S.: Asymmetric information and middleman margins: An experiment with Indian potato farmers. Review of Economics and Statistics **100**(1), 1–13 (2018)

22. Myerson, R.B., Satterthwaite, M.A.: Efficient mechanisms for bilateral trading. Journal of Economic Theory **29**(2), 265 – 281 (1983). https://doi.org/https://doi.org/10.1016/0022-0531(83)90048-0, http://www.sciencedirect.com/science/article/pii/0022053183900480
23. Nakasone, E.A.: The role of price information in agricultural markets: Evidence from rural Peru. Working Paper (2014)
24. Newman, N., Bergquist, L.F., Immorlica, N., Leyton-Brown, K., Lucier, B., McIntosh, C., Quinn, J., Ssekibuule, R.: Designing and evolving an electronic agricultural marketplace in uganda. In: Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies. pp. 1–11 (2018)
25. OpenStreetMap contributors: Planet dump retrieved from https://planet.osm.org . https://www.openstreetmap.org (2017)
26. Perrier, T., DeRenzi, B., Anderson, R.: Ussd: The third universal app. In: Proceedings of the 2015 Annual Symposium on Computing for Development. pp. 13–21. ACM (2015)
27. Sitko, N.J., Jayne, T.: Exploitative briefcase businessmen, parasites, and other myths and legends: Assembly traders and the performance of maize markets in eastern and southern Africa. World Development **54**, 56 – 67 (2014). https://doi.org/https://doi.org/10.1016/j.worlddev.2013.07.008, http://www.sciencedirect.com/science/article/pii/S0305750X13001770
28. Ssekibuule, R.: Market Design for Agricultural Trade in Developing Countries. Ph.D. thesis, Radboud University Nijmegen (2017)
29. Ssekibuule, R., Quinn, J.A., Leyton-Brown, K.: A mobile market for agricultural trade in Uganda. In: Proceedings of the 4th Annual Symposium on Computing for Development. p. 9. ACM (2013)