

Designing and Evolving an Electronic Agricultural Marketplace in Uganda

Neil Newman
University of British Columbia
Vancouver, BC, Canada

Lauren Falcao Bergquist
University of Chicago
Chicago, IL, USA

Nicole Immerlica
Microsoft Research
Cambridge, MA, USA

Kevin Leyton-Brown
University of British Columbia
Vancouver, BC, Canada

Brendan Lucier
Microsoft Research
Cambridge, MA, USA

Craig McIntosh
University of California San Diego
San Diego, CA, USA

John Quinn
Makerere University
Kampala, Uganda

Richard Ssekibuule
Makerere University
Kampala, Uganda

ACM Reference Format:

Neil Newman, Lauren Falcao Bergquist, Nicole Immerlica, Kevin Leyton-Brown, Brendan Lucier, Craig McIntosh, John Quinn, and Richard Ssekibuule. 2018. Designing and Evolving an Electronic Agricultural Marketplace in Uganda. In *COMPASS '18: ACM SIGCAS Conference on Computing and Sustainable Societies (COMPASS), June 20–22, 2018, Menlo Park and San Jose, CA, USA*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3209811.3209862>

1 INTRODUCTION

A significant challenge facing rural development is inefficiency in agricultural markets. One major 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, as 80% or more of the population in many African countries work in agriculture [8]), but also intra-seasonal and cross-locational price fluctuations that distort the market and reduce incentives for investing in productivity-enhancing inputs. Prior work [22] 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 via 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, 12]. 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. If farmers were more informed about the conditions of the market and better empowered to reach out to buyers beyond their immediate social network, they would have a stronger position from which to negotiate. Unfortunately, there is a massive hurdle to setting up an electronic marketplace in this setting: our potential userbase consists 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 (discussed in more detail in Section 3.2)—is high. 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. Even though five years have passed since our initial pilot study in 2013, feature phones remain pervasive throughout Uganda today.

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 [14] or [5] 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 to limit communication due to airtime costs, cultural resistance to adopting electronic markets, and high travel costs.

We introduced an electronic market platform for agricultural trade, branded in Uganda as *Kudu*, that is designed to address these challenges. In brief, our system operates as follows. Farmers and traders use their mobile devices to place *bids* (requests to buy) and *asks* (requests to sell) into a centralized nationwide database. Kudu identifies profitable trades, which are then proposed to the corresponding participants. Users' trust in the system is enhanced by the availability of in-village support services, provided by Agrinet, a private-sector Ugandan intermediary; users are supported by a call center. Our platform also gathers price data and broadcasts it back to farmers and traders using SMS, drawing from a large set of national, regional,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

COMPASS '18, June 20–22, 2018, Menlo Park and San Jose, CA, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5816-3/18/06...\$15.00

<https://doi.org/10.1145/3209811.3209862>

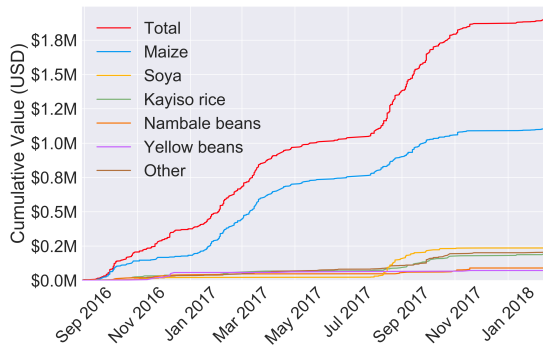


Figure 1: Cumulative value of verified transactions between September 2016 and February 2018.

and local markets and providing a uniquely tailored information set to each user.

Kudu was first piloted in 2013 [22]; after a brief hibernation, it rebooted in partnership with Agrinet and Innovations for Poverty Action in May 2015. Since then, the marketplace has been active for two and a half years and has attracted over 21,000 users through radio ads, village promotion meetings, and word of mouth. Users have 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.¹ Figure 1 shows the cumulative value of transactions on our platform, broken down by crop. Figure 2 shows the verified transactions on Kudu plotted geographically. Figure 3 illustrates the users active on our platform over time, separating existing and new users.

Over the past couple of years, we have been running Kudu as part of a multi-year randomized control trial to assess its role on farmer welfare. While anyone is free to register and use Kudu, we only advertise and offer in-village support in certain regions. We do not currently charge users anything to use Kudu; instead, a combination of grants and self-funding have covered Kudu’s expenses (which are dominated by the cost of human employees). We are exploring various monetization options, including charging commissions on proposed trades, but such mechanisms would be nontrivial to institute in our setting because nearly all trades are conducted using cash. Grant funding for the project finished in March 2018, but our local partners plan to continue operations and to transition the market towards for-profit operation.

1.1 Related Work

There have been attempts in the past to improve agricultural markets through price advisory systems. Examples include Esoko’s commodity index [4], Farmgain Africa [9], and Infotrade Market Information Services [11].² These services typically offer SMS subscriptions and radio based market information. However, most experimental evidence concludes that price advisory systems have been ineffective in

¹For comparison, Uganda’s agricultural sector is responsible for 24.5% of GDP, having a total value of about \$6.5 billion USD.

²See Section 1.2 of [21] for a survey of agricultural price information systems in Uganda.

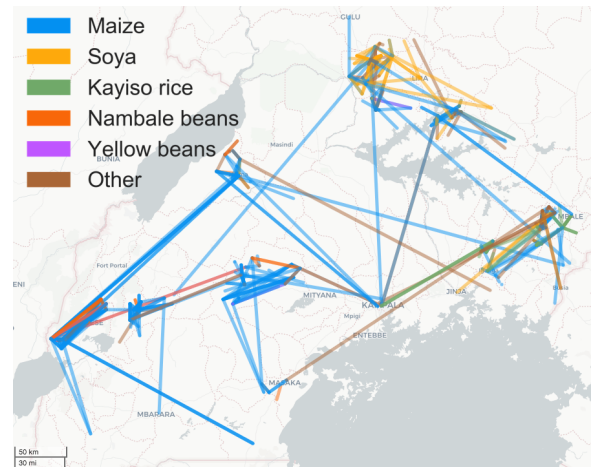


Figure 2: Geographical range of our verified transactions, with edges linking the reported parishes (small villages) of buyers and sellers respectively. About a third of verified transactions occurred within the same parish and are not visible on this map. The large clusters correspond to the 11 districts throughout Uganda in which Kudu is supported through in-village services. We note that Kudu has spread beyond these treatment districts.

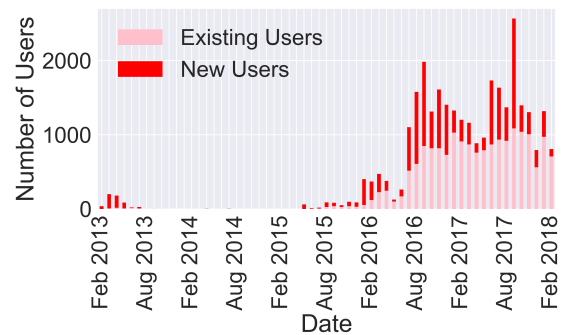


Figure 3: Active users over time. Each bar represents the number of unique active users in a one month interval. New users that month are highlighted. An active user is defined as a user that has used any of Kudu’s services that month.

improving farmer welfare [7]. 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 [16]. 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 is aimed at comprehensively addressing this set of barriers to market access. It goes beyond previous services, offering nuanced market information, direct market connections, and wraparound services needed to provide smallholder farmers truly improved market access.

One alternative system design—which we rejected very early on—would simply offer classified ads: a database of bids and asks that users must search manually. We rejected this idea for two main reasons. First, we do not believe that searching through listings can be made effective on feature phones. Second, such a design would not take advantage of the fact that the underlying market is in *commodities*: i.e., that many different sellers offer goods that are effectively perfect substitutes from buyers’ perspectives. In such settings, asking market participants to browse a list of potential trading partners is highly suboptimal. One agricultural smartphone application currently using the classified ad approach is AgroMarketDay [13]. Artificial intelligence has been recognized as playing an increasingly important role in market design [15], for example to reduce search frictions. As will describe later in Section 5, Kudu leverages AI to decide what matches to propose to users.

The remainder of this paper is organized around the information flow in our platform. We start by describing how users register for Kudu and Kudu’s bidding language in Section 2. We then go over the various interfaces to Kudu in Section 3. In Section 4 we discuss how matches flow through our system and reasons why they may fail. Next, we cover how we decide which bids and asks to pair in Section 5. We then describe additional support we offer to facilitate trade in Section 6. Finally, we discuss future plans for Kudu in Section 7 and conclusions in Section 8. The web interface for the implementation described in this paper can be accessed at <http://kudu.ug>.

2 BIDDING & REGISTRATION

In this section, we discuss how Kudu gathers bids and asks from users. When a user places a bid or ask, they tell Kudu what crop they want to sell, their requested buy (sell) prices, and desired (available) quantities. The units in which quantities are specified depends on the crop and reflects how they would usually be advertised. For example, bids for maize specify the desired weight in kilograms, whereas bids for potatoes specify the desired number of sacks. Kudu also optionally gathers “tags” that narrow down what a buyer is looking for or what a seller has to offer: for example, beans can be shelled or unshelled, and maize can be wet or cleaned. Our services are available in four languages: English, Luganda, Luo, and Runyakitara. Our marketplace currently supports 76 crop types.³ Crops differ in quality. This is problematic, because we want 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 are enforcement and inconsistency in users’ abilities to grade crops effectively. Instead, we solicit 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

³We refer to anything sold on Kudu as a crop, but a small fraction of our supported commodities are not plants, such as eggs, fish, and livestock.

quality level. Despite its inelegance, this system works well in the Ugandan cultural context where point-of-sale bargaining is already common, and has not been the source of significant pushback from users.

Another way we simplified bidding over our initial design was to remove “location filters” that specify that a buyer will only consider traveling within a specific geographic region. We dropped this feature because it was not well utilized; instead, we take travel costs into account when proposing matches.

When a new user attempts to place a bid or ask on our system, we “register” the user. One key fact we record about a user is location, which we currently store as a single fixed location in the center of the user’s *parish* (the smallest administrative unit in Uganda). There are about 5,000 parishes, each of which typically spans only a few kilometers. Parishes are grouped together into (nearly 1,000) *subcounties*, which form around 120 *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 is reasonable for about 85% of them, and hence decided that a more complex system would not be worth the cost.

To avoid hassling potential users with a complicated authentication system, we do not require users to set a username and password. Instead, users on Kudu are identified by their phone numbers. This does not work well when a user changes their phone number, or when multiple people with different devices want to share a single account, but it works well in the common case.

3 INTERFACES TO KUDU

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

3.1 SMS

Users can send a toll-free SMS to 8228 using any of the following templates:

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

If the user enters a crop name that does not exactly match one known in the system, we search for a close match and automatically correct it. The user receives a confirmation SMS with their bid or ask and can call us if anything is wrong.

If a message strays too far from the template, we are unable to automatically parse it. To date we have received more than 3,000 SMS messages that could not be parsed. We fail to parse messages for many reasons: some have nothing to do with agriculture (questions about voting and health), some are written in an unsupported language, some are missing required information (e.g., the SMS consists of the single word “buy” or “sell”), and some mirror the template (a user literally sends “buy crop quantity price”). Other common mistakes are including units or descriptive information (“SELL DRYED CASSAVA 5000KGS 1500 PER KG”), or treating the SMS like a

classified ad (“BUY GINGERS, LOCAL-40000SHS,FOREIGN-30000SHS,PER SACK OR 120KGS, ANY NO. OF SACKS OR KGS.CALL 256*****”). We assemble all of the messages that cannot be parsed and our staff correct these messages as they are able, phoning users when necessary.

Even when an SMS matches a template exactly, we may still fail to capture the user’s intent. If the crop name is misspelled, for example, our system may make the wrong correction. Users may also reverse the ordering of the positional quantity and price arguments, and both numbers can 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), so first-time SMS users receive a phone call from our call center to confirm this information. This and other trials have taught us that our SMS templates are 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 is inexpensive to run. However, given the presence of USSD, described below, we do not expect SMS to be a popular method for interacting with Kudu. In the past six months (September 2017–February 2018), only 0.17% of bids (13) and 1% asks (72) were delivered via SMS. We have therefore shifted to focusing more heavily on other interfaces.

3.2 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 4 for a sample Kudu interaction.

USSD solves many of the issues with SMS: a user can learn to navigate the interface independently; bids can be previewed before submission (allowing a user to confirm that the information is accurate); error messages can be reported in response to nonsensical inputs (e.g., 0 quantity). All of this is in principle possible with SMS but would be 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. USSD also has a prestige that SMS lacks, and in addition to being useful can act as a strong signal to users that a service is backed by a serious enterprise. For all of these reasons, USSD has also been used in other development projects [19].

Unfortunately, despite all the positives just discussed, USSD comes with its own set of issues. One key problem is that messages can be no longer than 182 characters. This is very restrictive in practice: selecting from a long list is difficult, such as when disambiguating parishes with similar names. Furthermore, sessions longer than 2 minutes time out, leaving the user to start from scratch. This can be a very frustrating experience for long menu sequences, or when a user has not prepared answers to all of the questions in advance. A first-time USSD user trying to buy or sell is prompted for additional information to register, further exacerbating the time limit

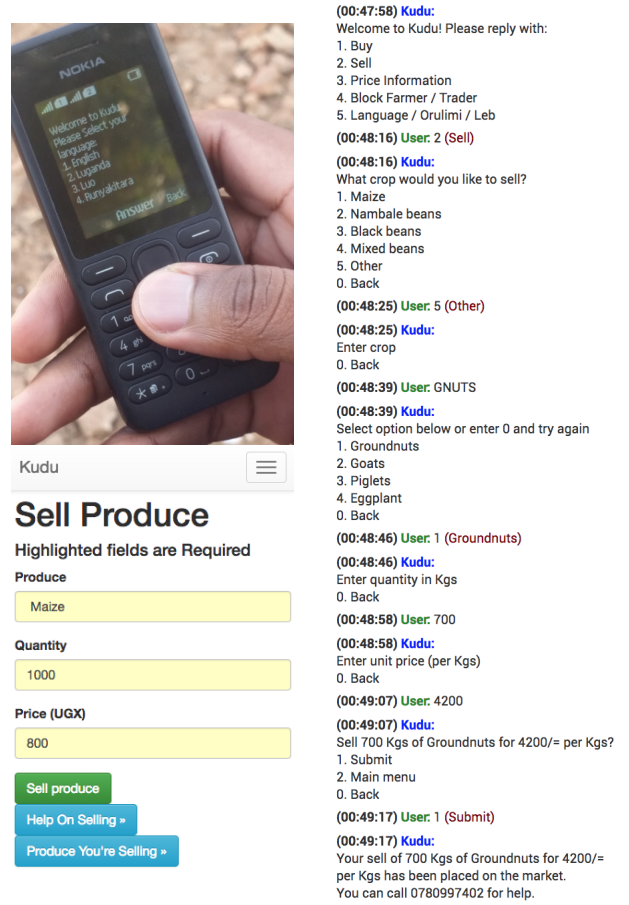


Figure 4: Top Left: Kudu’s USSD interface running on a feature phone. Right: Sample USSD interaction for selling groundnuts (peanuts). Bottom Left: A user placing an ask on Kudu’s web interface.

issue. In the end, we still have 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 has been to check price information, but it also produced 1.2% (74) of our bids and 5.4% (383) of our asks over the past six months. Our USSD service went offline on January 31, 2018 because our provider shut down all USSD services. We are currently investigating new providers.

3.3 Web

Since the pilot, we have provided a web interface to Kudu as shown in Figure 4. While we do not expect this option to be used by small-holder farmers, the web interface is important for discovery and may be appealing for more technologically sophisticated users. The marginal implementation cost of maintaining these web forms given the existence of the other interfaces is also very low. In the past six months, 0.6% (35) of all bids and 0.6% (42) of all asks were

placed via the web interface. Users can register online: in the past six months, around 70 users have done so. The website is the most straightforward interface for enriching the bidding language (e.g., users could select locations they would travel to on a map).

3.4 Call Center

The vast majority of our bids and asks in our current system do not come from the above three “self-serve” options. Instead, in the past six months 97.7% (6167) of bids and 91.6% (6471) of asks were solicited from our call center. Monday through Friday, agents call traders and farmers and ask them if they have anything to buy or sell. This has been very effective at thickening the market, and does not require users to be technologically sophisticated. However, it is very labor intensive — scaling linearly in the number of users on Kudu — so it is unlikely to be an economical approach in the long term.

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

4 FROM PROPOSALS TO PHYSICAL TRANSACTIONS

Once Kudu has gathered bids and asks, the next step is to make matches. We will discuss the matching process in more detail in the following section; this section will focus on what happens after we decide 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. We now employ *deal coordinators* to shepherd matches along and introduce a human element into the system. By the time phone numbers are exchanged, the deal coordinators have already spoken to the seller and buyer individually and can vouch for one to the other, incorporating past experiences when available.

When a match is proposed, a deal coordinator first calls the seller. If the seller is interested in the match, they next call the buyer. Pending buyer interest, phone numbers are exchanged. The deal coordinator checks in with both parties to see if an agreement was reached. Finally, they follow up after the date the deal was supposed to go through and record what transpired. Of course, either side can pull out at any step in this process.

An additional job for deal coordinators is to look at asks that they are 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.

4.1 Why Proposals Fail

Only a small fraction of the trades that we propose actually occur: e.g., over the past six months, only 7% of the asks that we proposed matches for led to deals. While we should certainly expect some failures, given that these are trades that Kudu identifies as “profitable”, it is worth asking why the fraction of failures is so large. The first, and by far most common, reason is that by the time we propose a match, at least one of the parties is no longer in a position to trade. We will discuss this in Section 4.1.1. The second is that Kudu’s assessment

of what trades are profitable does not necessarily correspond to users’ own assessments. Sometimes such mismatches are because of price or quantity issues, but most commonly the problems arise because of travel costs; we discuss these issues further in Section 4.1.2.

4.1.1 *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 notifies us that they would like to sell, we need to move quickly; a common failure mode of our proposals is that the seller has already sold their crop by the time we send 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 jumps to 16%. 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 hear otherwise from the user, we assume that a bid or ask is no longer valid.

Even if we do propose a match in time, and negotiations are successful, transacting outside of the system is still a concern. When a buyer travels to meet their matched farmer, on the way they may encounter another farmer selling exactly what they want. If this occurs, they may not feel obligated to continue onwards to transact with the intended recipient, and instead take the closer trade and turn around. The only defense we have against this behavior right now is that the buyer might not want to jeopardize their reputation with the deal coordinators.

4.1.2 *Geographic Constraints.* After discounting trades that fail because one party has already sold, the next most common failures have to do with one party not wanting to travel. In order to accurately estimate trade profitability, it is important to develop a model of transport costs. We have coordinates for the location of every Parish. One might guess that transportation costs are linear in Euclidean distance, but this turns out to be highly inaccurate because of bodies of water and road quality issues. It is also tempting to use Google Maps to estimate travel times, but unfortunately it does not know about many of the smaller roads that connect parishes. Instead, we used the road network data available for Uganda on OpenStreetMap [18] to model the road geography and generate approximate travel times between Parishes. Given an estimated cost per kg per hour of transport, we can then roughly estimate the transportation cost for any given trade proposal. Figure 6 shows that most of our successful trades involved short distances (but that a long tail represents significant travel).

There are other reasons to refuse a trade proposal based on its location other than worry that transport costs will swamp potential gains from trade. We asked our users to give reasons for which they might avoid traveling to an unfamiliar area based on a Kudu recommendation. The most popular answers (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

Matches - Call Seller

Show 5 entries

Match	Follow Up	Match Date	Produce	Seller Name	Seller Number	Seller Comments	Buyer Name	Buyer Number	Buyer Comments	Ask Quantity	Bid Quantity	Ask Price	Bid Price	Matched By	Interested?	Next Follow-Up
15417	Add	02/11/2018 11:57 a.m.	Kanyebwa beans	Mbabazi Aswankwire	256555555555	—	Jorum Godfrey	256555555555	—	50	3,000	1,500	1,600	Deal Coordinator #1	<input type="checkbox"/> Yes <input type="checkbox"/> No	02/13/2018 10 a.m. confirm crop
17027	Add	02/15/2018 12:02 p.m.	Kayiso rice	Ashraf Bendicto	256555555555	—	Mujuni Onesmus	256555555555	—	2,000	10,000	2,500	2,300	Deal Coordinator #2	<input type="checkbox"/> Yes <input type="checkbox"/> No	
17572	Add	02/07/2018 11:44 a.m.	Millet	Okello Straji	256555555555	—	Omara Costanziya	256555555555	—	200	1,000	1,000	1,700	Deal Coordinator #1	<input type="checkbox"/> Yes <input type="checkbox"/> No	
17575	Add	02/13/2018 11:45 a.m.	Millet	Olwi Budala	256555555555	—	Nabwonya Oyo	256555555555	—	200	1,000	1,000	1,800	Deal Coordinator #1 (Kudu AI)	<input type="checkbox"/> Yes <input type="checkbox"/> No	02/15/2018 3 p.m. Call back confirm crop
17576	Add	02/13/2018 11:46 a.m.	Millet	Mesulamu Davidson	256555555555	—	Glacia Kinemata	256555555555	—	200	1,000	1,000	1,700	Deal Coordinator #3 (Kudu AI)	<input type="checkbox"/> Yes <input type="checkbox"/> No	

Showing 1 to 5 of 20 entries

Previous 1 2 3 4 Next

Figure 5: A dashboard recording match proposals where the seller has yet to be called. A deal coordinator can schedule reminders to call the seller again if they do not pick up, confirm that they are interested, or cancel the match. Similar dashboards exist for matches in other states of completion.



Figure 6: An empirical cumulative distribution function for successful and unsuccessful trades by estimated travel time. Most of our successful trades involve short distances.

sentiments like the very entrepreneurial “If a trade is profitable, nothing can stop me.”) Needless to say, Kudu does not model all of these concerns when judging that a match is “profitable”.

5 MATCHING

A significant technological challenge in the Kudu ecosystem is choosing which bids should match with which asks. We would like to do so in a way that maximizes the value that Kudu brings to the marketplace.

To quantify the value of a trade, we define a scoring function that maps 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. The scoring function should capture the 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

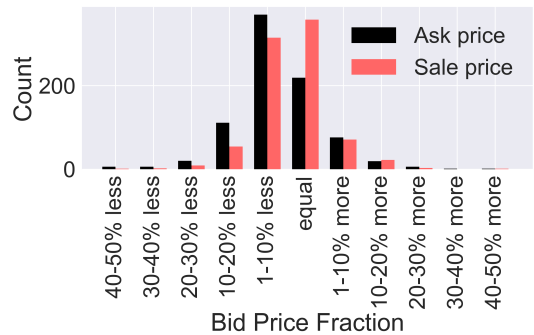


Figure 7: A histogram of bid prices as a fraction of ask prices and final sale prices for verified transactions. If users were truthful, no trades would ever occur with a bid price lower than the ask price; clearly this is not the case. However, most of the mass is distributed around the center, suggesting that the user prices still convey useful signaling information (most of the time). The buyer usually pays slightly less than their bid price in the final sale.

during deployment that, because users bid and bargain strategically, this scoring function can assign negative values to trades that actually go through in practice (and therefore must actually generate positive gains). This strategic behavior is evident in Figure 7, which shows that users sometimes trade despite notionally negative gains. Our current implementation therefore assumes that bid and ask prices are merely signals rather than binding constraints. We use these declared prices to fit probability distributions that model each user’s true price, and then use these to calculate potential gains from trade. That is, we

sample from both price distributions, reject samples where the bid price is lower than the ask price, and compute the expected gains.

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

5.1 Manual Matching

When the number of bids and asks is low, an effective method to clear the market is via *manual matching*. Deal coordinators manually look at the database of bids and asks and decide which parties to match. Kudu currently employs five deal coordinators, each specializing in particular treatment districts and supporting local languages.

Such a system leverages human intelligence and human relationships in the matching process. Deal coordinators develop intuition about which parties make good matches. This intuition can be based on features that are hard to quantify, e.g., the personalities of the trading participants. Furthermore, deal coordinators interact repeatedly with the same parties, developing valuable social capital and trust. A participant who has a personal relationship with a deal coordinator might be more willing to submit bids or asks to Kudu than one who only interacts with the system electronically. (An illustration: one coordinator told us that her conversations typically involve an initial discussion about the participant’s family before any discussion of the potential match.)

However, the downsides to manual matching are significant. First, it requires (expensive) human employees, the number of which must scale linearly with the number of participants on the platform. Second, it is unlikely that deal coordinators are effective at optimizing a global objective such as overall gains from trade. Discussions with the deal coordinators suggest they follow 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 is based on their estimate of how likely the buyer is to accept a trade at that given time. The selection of the seller is then entirely based on the value of the seller to that buyer. Importantly, this process ignores the value of the seller to other potential buyers. Finally, as the system grows, search frictions, like the size of the database and the limited sorting tools, make it difficult for deal coordinators to find the best matches, even according to their own metrics.

5.2 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 is that, as the automatic matching component improves, the participants will be able to transact the proposed matches without the intervention of the deal coordinators. Such a system would easily scale as the market grows, can optimize global objectives, and is not significantly hindered by search frictions. However, a fully automatic system sacrifices 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 [22]). 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 minimum competitive (i.e., Walrasian) prices for the matching market [10], making truthful bidding a dominant strategy for buyers but giving farmers incentives to manipulate their sale prices.⁴ In our idealized market, buyers and farmers would trade 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. This system again failed to leverage the human intelligence of the deal coordinators.

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 problem⁵ 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 4.1.1). 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 offers matches continuously. The system assigns a priority to each buyer, equal to the highest potential surplus from a trade involving that buyer. When a deal coordinator enters the system, she is presented with a list of buyers, sorted by priority. Our intention is that the deal coordinator will choose to work with the highest-priority buyer in this list.⁶ Once the deal coordinator selects a buyer, the algorithm selects five possible sellers for the buyer, in decreasing order of gains from trade. Deal coordinators can accept or reject any seller. Rejections come with reasons that help us improve the algorithm. If used as intended, this process mimics the greedy

⁴As demonstrated by a celebrated theorem due to Myerson and Satterthwaite [17], it is impossible to make truthful bidding a dominant strategy for both sides of the market. We focus on incentivizing buyers because they typically constitute the short side of our market and because they have access to a more robust array of outside options.

⁵Another 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.

⁶The deal coordinator can also search the system for a specific buyer; this will defeat the global optimization guarantees.

Ask	Ask Date	Seller Name	Seller District	Seller Subcounty	Seller Parish	Ask Quantity	Ask Price	Ask Matched Times	Seller Comments	Bulk
30081	Feb 02, 2018	Sam Kiara	Kasese	Maliba	Mubuku	4,000	600	2		<input type="checkbox"/> Yes <input checked="" type="checkbox"/> No
30408	Feb 09, 2018	Oker Kyotalimye	Kamwenge	Bwizi	Bwizi	10,000	600	1		<input type="checkbox"/> Yes <input checked="" type="checkbox"/> No
30410	Feb 09, 2018	George Mugole	Kamwenge	Nkoma	Nkoma	1,000	600	1		<input type="checkbox"/> Yes <input checked="" type="checkbox"/> No
30421	Feb 12, 2018	Alex Ibrahim	Mubende	Kasambya	Kigando	30,000	700	2		<input type="checkbox"/> Yes <input checked="" type="checkbox"/> No
30477	Feb 13, 2018	Chwula Peseptwa	Mubende	Kasambya	Lusiba	15,000	670	2		<input type="checkbox"/> Yes <input checked="" type="checkbox"/> No

Bid	Bid Date	Buyer Name	Buyer District	Buyer Subcounty	Buyer Parish	Bid Quantity	Bid Price	Bid Matched Times	Buyer Comments
27552	Feb 02, 2018	Ronah Kakunyule	Oyam	Ngai	Akuca	1,000	900	4	
27615	Feb 13, 2018	Tushar Rwejemara	Iganga	Nakigo	Wairama	10,000	750	4	
27707	Feb 11, 2018	Frank Apuka	Apac	Cegere	Cegere	3,000	800	4	
27734	Feb 07, 2018	Robin Moses	Kamwenge	Kahunge	Kiyagara	500	1,200	4	
27735	Feb 09, 2018	Okwio Emmanuel	Kasese	Kisinga	Kagando	25,000	800	4	

Figure 8: The manual matching interface on Kudu. Asks are shown on the left and bids on the right. A deal coordinator selects one from each column to create a match. This approach breaks down when the number of bids and asks grows large.

algorithm for maximum matching and hence captures a constant fraction of the gains from trade in a static system.⁷ Furthermore, it gives deal coordinators more flexibility, in the form of rejecting sellers or processing buyers in an order other than the recommended one. We can use this flexibility to improve our algorithm’s scoring function by leveraging the human intelligence of the deal coordinators.

6 FACILITATING TRADE

We have 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 has rolled out their agent model into the communities in which we are introducing Kudu. Agents promote Kudu by advertising the service to both farmers and local traders, typically via house-to-house visits and announcements via loudspeaker in markets. They then follow up this advertisement with a village-based meeting in which they provide information on Kudu services and training on how to use the system. Agents also distribute their phone numbers so that users can call them if they have questions about the service or need help registering an ask or bid later in the season.

In addition to promotion and training, AgriNet offers several additional services designed to address issues that can hinder transactions between buyers and sellers, even once they have 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 [20]. This requires both coordination and access to capital. Kudu has the capacity to note and electronically bulk lots of the similar crops available for sale in nearby locations; AgriNet agents are available to provide on-the-ground coordination of this bulking. In order to finance this bulking, AgriNet offers its agents access to Cash on Bag (COB) credit. Agents in turn may 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.⁸

⁷However, unlike the greedy algorithm, our system is dynamic. In upcoming work we discuss how the dynamism impacts this guarantee.

⁸In addition to this bulking procedure, we have future plans to implement automated bundling as part of the Kudu matching process. We discuss this in more detail in our conclusions.

Another challenge that may limit buyers’ willingness to trade on the platform is the risk inherent in directly trading with farmers in remote villages with whom they have not yet developed trust. Buyers must make up-front investments in transportation out to rural villages without guarantee that any agreements made in advance regarding quantity or quality of available crops will be carried out as promised once they arrive. Buyers may instead choose to trade only with trusted brokers or other traders with whom they have repeated interactions, resulting in a fractured chain of many short-distance, relationship-based exchanges [6]. 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, offers transport cost compensation for buyers who travel to a rural sale point and are disappointed.

As a further measure to address the risks involved with remote trading, local monitoring agents are used to certify the details of the transactions. After receiving a call from a deal coordinator about an agreed-upon deal, a monitoring agent visits the seller to check the quality and quantity and communicates his findings to the buyer. The agent is present at physical transaction and oversees exchange of money, providing regular updates to deal coordinators throughout the transaction process. These monitoring agents provide other services as well, such as recruiting and training new users through visiting local markets and village promotion meetings. Finally, the monitoring agents can help smooth out price negotiation by being physically present. For example, we have heard from deal coordinators that there can be a tension after both parties have exchanged phone numbers regarding who will call the other first, out of fear of looking desperate. The monitoring agent can help address this issue by mediating the negotiation.

6.1 Price Information

Soliciting bids and asks on Kudu is challenging due to technological and informational constraints. We have already discussed technological constraints, which are largely outside our scope to change: e.g., few farmers have smart phones, which forces us to solicit information through limited interfaces; many are illiterate, forcing us to rely on human intervention. The informational constraints are rooted in the users’ reluctance to participate in the system without

Kudu Match Proposals

More match proposals will appear as you accept and reject the current proposals

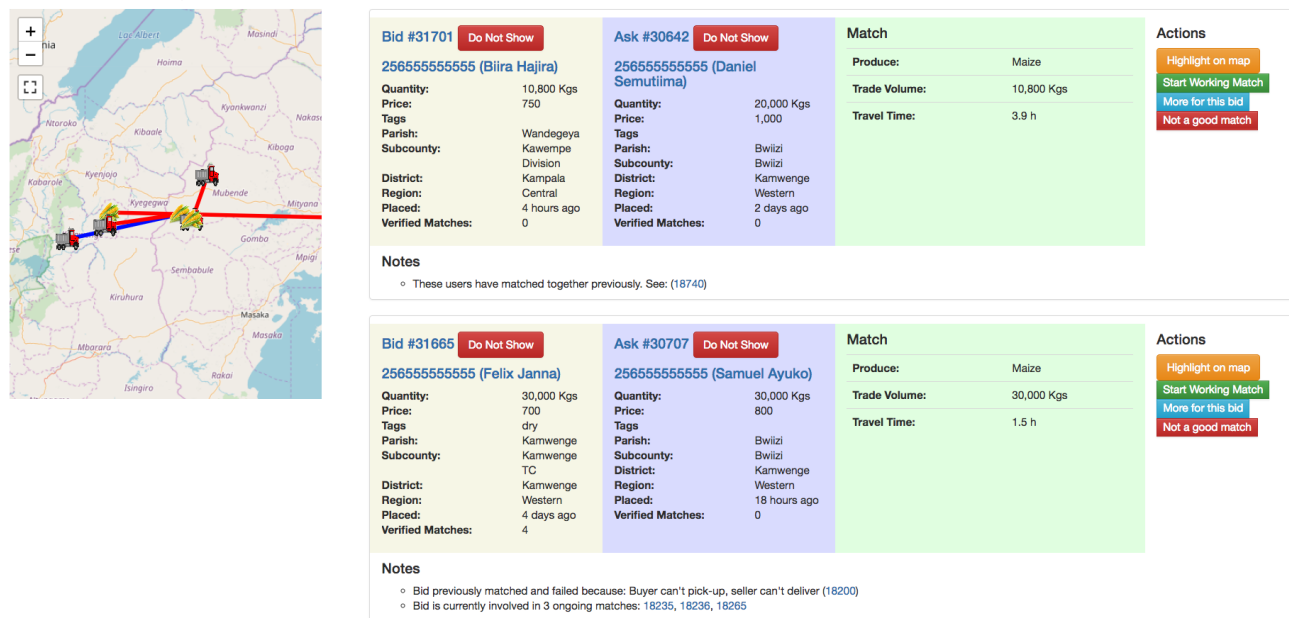


Figure 9: The Kudu AI interface. Deal coordinators are shown matches corresponding to the largest gains from trade. Matches are overlaid on a map of Uganda. Characteristics of the match and the involved users are also shown: for example, we highlight if the two users have matched together before or are engaged in any ongoing transactions. Deal coordinators can accept or reject proposals, refreshing the available choices.

knowledge of going market rates. To tackle this challenge, we provide price quotes. Ideally our price information would come from verified transactions that have occurred on Kudu, but our system is still too small to consistently have sufficient data in enough districts to be useful.

A next hope would be to use the bids and asks, which are more plentiful than verified transactions. However, we should expect the bids and asks we receive to be somewhat biased, since our platform exists in the context of outside options. Our users are only interested in using the platform if it can get them a better deal than they could otherwise find. This means that ask prices are usually inflated and bid prices are usually shaved. Broadcasting this data can be particularly problematic because it can create a feedback loop wherein farmers receive overly optimistic price information from Kudu and then price their crops accordingly.

One way to dampen this feedback loop is to use the recommended prices of our market algorithm to generate hypothetical transactions instead of broadcasting statistics of bids and asks directly. Of course, as these transactions might not actually occur, these prices are bound to be inaccurate as well.

An initial version of this 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 current implementation of this service uses biweekly survey data to determine market prices for select crops in the treatment districts and reports the 25-75

percentile of wholesale prices at markets. Collecting data in this way yields coarse information, is expensive, and results in stale quotes.

Ideally we would combine these systems. In the near future we plan to build a system that uses a combination of verified transactions and hypothetical transactions based on bid/ask data. The price quote will then be some weighted convex combination of these data points. The weights can be functions of the features of the transaction: whether it was verified or hypothetical, the time of the transaction, the time of the involved bid/ask, etc. The technical challenge is to tune these weights algorithmically.

7 FUTURE PLANS

Kudu remains under active development. When we began this journey, we thought the most important thing would be to settle on the right market design and to back the market with an effective clearing algorithm. We continue to consider these pieces important, and we remain focused on unresolved problems in the market design space: proposing high quality matches, making the system financially self-sustaining, and creating trust without expensive deal coordinators or in-village support services. On the other hand, we have come to appreciate the importance of additional issues that have little to do with market design, such as identifying a reliable USSD operator and balancing our local partner's competing interests against our own.

On the market design side, we are actively engaged in improving our core matchmaking services. In the short term, we are working to

improve models of transport time and to provide power users with a richer bidding language. In what follows, we outline some of the longer-term directions we are actively pursuing.

7.1 Machine Learning

Our system currently relies on a structural (i.e., hand-crafted) model that aims to capture participants' likely gains from trade under a proposed match, taking into account factors like profit margin and travel costs. However, this model is not highly effective: proposed matches correspond to a low probability of trade. We would like to use machine learning methods to predict whether a trade will be successful, augmenting our structural model. At the moment, our biggest problem is lack of data: we only have approximately a thousand positive examples of successful trades. A second issue is that the distribution of our match proposals is far from stationary. Third, it is challenging to reliably adapt learned models over time. Despite these challenges, we expect that machine learning will play an important role in Kudu in the future. We are currently investigating different features upon which a model might depend and learning how helpful each is to predictive performance. This has helped us to recognize the importance of bid and ask "freshness" and users' history of consummating previous trades on the system.

7.2 Automating Proposed Trades

We are also exploring ways to improve the process of implementing a trade. Right now we can reach users via phone calls and SMS messages. Calling a user is flexible and eliminates ambiguity about their intentions, but is expensive and requires reaching the user over the phone. SMS messages are cost effective and can be managed on the user's own time but are problematic when users are unresponsive, e.g., due to illiteracy, lack of battery power, or sharing a phone among multiple family members, and require significantly more trust in the electronic matching system. Moving forward, we hope to augment phone calls with an Interactive Voice Response (IVR) system to help illiterate users.

7.3 Bundling

An individual smallholder farmer typically has less produce available than a given buyer would like to purchase. Trades are not frictionless: hiring trucks is expensive and time consuming. This can dissuade buyers from trading with small farmers, even at favorable prices. We can help by bundling the asks of nearby farmers and presenting them as a package. Bundling poses difficulties because it embeds a complex combinatorial problem into the matching process, not to mention causing a combinatorial explosion in the number of possible trades. Moreover, bundled trades require a high degree of coordination and thus present a significant risk of unraveling. Kudu's manual match interface already allows deal coordinators to bundle asks, but this feature has never led to a successful trade and is currently disabled. Instead, we currently see bundling performed mostly by aggregators, who indeed account for a large volume of trade on the system. Going forward, we would like to increase Kudu's ability to bundle trades directly, since aggregators often offer farmers much worse prices than buyers are willing to pay.

7.4 Trust and Reputation

Our system includes no mechanism for ensuring that traders honor agreements they make. Buyers can renege on previously accepted deals or renegotiate at the time of transaction by threatening to leave (which is particularly problematic when crops are perishable). Sellers can attempt to renegotiate at the last minute, leveraging the fact that a buyer cannot easily walk away after paying for a truck rental and driving a long distance. Escrow is 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 is in place; and the buyer could tell Kudu to release the money when the goods are transferred. However, various practical hurdles make this idea more challenging in practice: traders could be exposed to mobile money phishing scams; kiosks to withdraw mobile money are not yet prevalent enough in rural areas; and (probably most importantly) fees for mobile money transactions are currently too high.

Another way of increasing users' trust in the system would be to integrate a more robust reputation system. Kudu already lets users blacklist anyone with whom they have had a poor experience, but we could go much further. Since we have a transaction history for every user on Kudu, we could reward completion of successful matches by increasing users' priorities, and conversely deprioritize users who do not follow through. Deal coordinators could also use such reputation information to convince users to trade with each other. One challenge is that it is expensive to verify successful trades. Another is that "whitewashing" is easy: Kudu identities are tied to phone numbers, and new SIM cards are inexpensive.

8 CONCLUSIONS

This paper has described Kudu, an electronic market for agricultural trade in Uganda. Traders can post bids and asks using a feature phone—via SMS, USSD and voice—and we also offer a web interface. Kudu then proposes matches, leveraging a combination of optimization algorithms, data-driven models, and human expertise. Our system is augmented by a rich variety of support services that help to facilitate trade. The system has been active for over two years, involving tens of thousands of users and yielding verified trades totaling almost \$2 million USD.

9 ACKNOWLEDGMENTS

None of the practical results described in this paper 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. 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 fruitful business conversations with Paul Nyende. We obtained map data from OpenStreetMap <https://www.openstreetmap.org>. This project received institutional support from Microsoft and UBC and financial support from NSERC via an E.W.R. Steacie Fellowship and a Discovery Grant.

REFERENCES

- [1] Jenny C Aker. 2010. Information from markets near and far: Mobile phones and agricultural markets in Niger. *American Economic Journal: Applied Economics* 2, 3 (2010), 46–59.
- [2] The World Bank. 2018. Mobile cellular subscriptions (per 100 people). <https://data.worldbank.org/indicator/IT.CEL.SETS.P2>. Accessed: 2018-03-02.
- [3] Lauren Falcao Bergquist. 2017. Pass-through, Competition, and Entry in Agricultural Markets: Experimental Evidence from Kenya. (2017).
- [4] Olayinka David-West. 2010. Esoko Networks: facilitating agriculture through technology. GIM Case Study No. B061.
- [5] Liran Einav, Chiara Farronato, and Jonathan Levin. 2016. Peer-to-peer markets. *Annual Review of Economics* 8 (2016), 615–635.
- [6] Marcel Fafchamps and Bart Minten. 1999. Relationships and traders in Madagascar. *The Journal of Development Studies* 35, 6 (1999), 1–35. <https://doi.org/10.1080/00220389908422600>
- [7] Marcel Fafchamps and Bart Minten. 2012. Impact of SMS-based agricultural information on Indian farmers. *The World Bank Economic Review* 26, 3 (2012), 383–414.
- [8] Farm Africa. 2018. Our work in Uganda. <https://www.farmafrika.org/uganda/uganda>. Accessed: 2018-03-02.
- [9] Farmgain Africa. 2018. Farmgain Africa. <http://farmgainafrica.org/>. Accessed: 2018-03-02.
- [10] Faruk Gul and Ennio Stacchetti. 1999. Walrasian Equilibrium with Gross Substitutes. *Journal of Economic Theory* 87, 1 (July 1999), 95–124.
- [11] Infotrade. 2018. Infotrade - Market Information Services. <http://www.infotradeuganda.com/>. Accessed: 2018-03-02.
- [12] Robert Jensen. 2007. The digital divide: Information (technology), market performance, and welfare in the South Indian fisheries sector. *The quarterly journal of economics* 122, 3 (2007), 879–924.
- [13] L Katusiime, I Omiat, et al. 2014. AgroMarketDay. In *RUFORUM Fourth Biennial Conference, Maputo, Mozambique, 19-25 July 2014*. RUFORUM, 371–375.
- [14] Scott Duke Kominers, Alexander Teytelboym, and Vincent P. Crawford. 2017. An invitation to market design. *Oxford Review of Economic Policy* 33, 4 (Nov. 2017), 541–571. <https://doi.org/10.1093/oxrep/grx063>
- [15] Paul R Milgrom and Steven Tadelis. 2018. *How Artificial Intelligence and Machine Learning Can Impact Market Design*. Technical Report. National Bureau of Economic Research.
- [16] Sandip Mitra, Dilip Mookherjee, Maximo Torero, and Sujata Visaria. 2017. Asymmetric information and middleman margins: An experiment with Indian potato farmers. *Review of Economics and Statistics* 0 (2017).
- [17] Roger B Myerson and Mark A Satterthwaite. 1983. Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29, 2 (1983), 265 – 281. [https://doi.org/10.1016/0022-0531\(83\)90048-0](https://doi.org/10.1016/0022-0531(83)90048-0)
- [18] OpenStreetMap contributors. 2017. Planet dump retrieved from <https://planet.osm.org>. <https://www.openstreetmap.org>.
- [19] Trevor Perrier, Brian DeRenzi, and Richard Anderson. 2015. USSD: The Third Universal App. In *Proceedings of the 2015 Annual Symposium on Computing for Development*. ACM, 13–21.
- [20] Nicholas J. Sitko and T.S. Jayne. 2014. 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 (2014), 56 – 67. <https://doi.org/10.1016/j.worlddev.2013.07.008>
- [21] Richard Ssekibuule. 2017. *Market Design for Agricultural Trade in Developing Countries*. Ph.D. Dissertation. Radboud University Nijmegen.
- [22] Richard Ssekibuule, John A Quinn, and Kevin Leyton-Brown. 2013. A mobile market for agricultural trade in Uganda. In *Proceedings of the 4th Annual Symposium on Computing for Development*. ACM, 9.