

# Developing World I: Kudu

## A Mobile Market for Agricultural Trade in Uganda

Professor Kevin Leyton-Brown  
University of British Columbia  
Department of Computer Science

11/27/17

*These notes are based off a presentation by Professor Kevin Leyton-Brown, for the section on the developing world in the Mechanism Design for Social Good Reading Group. The notes are taken by members of the reading group with some figures and texts taken from the accompanying paper and project website [1]. Questions and comments from reading group members during the presentation are labeled as such. Please contact the reading group organizers with any questions or comments.*

### Introduction

There are large market inefficiencies in agricultural markets in developing nations. Kudu began with research looking at how to use mobile technology to make agricultural markets more effective in Uganda. There were large price discrepancies and arbitrage opportunities, which indicated problems with the efficiency of the market. Farmers faced problems in finding buyers for their produce, being threatened with spoilage of their goods when no buyers could be found and often having little negotiating power. Traders also faced uncertainty in being able to locate produce, relying on word-of-mouth networks. Mobile price advisory services provided some help, particularly to farmers, but seemed to have problems with accuracy and timeliness. Conventional auction or listings services were not accessible to the majority of farmers and traders in Uganda who have basic phones with SMS functionality only.

Kudu is a collaboration between many organizations and individuals including: Richard Ssekibuule (Makerere University), John Quinn (Makerere University), Neil Newman (UBC), Nicole Immorlica (Microsoft Research), Brendan Lucier (Microsoft Research), Craig McIntosh (UCSD), Lauren Bergquist (University of Chicago), AgriNet Uganda Ltd and ITA Uganda. One of the themes of this talk will be that we began with a naive and optimistic vision of agricultural trade mediated electronically. But, in the end, making the stuff happen in practice has brought in many people into the loop.

## The Circumstances

The insight behind Kudu was that rather than providing a listing of items offered which buyers could bid on (known as a single auction), a different type of system called a double auction might be more appropriate to users with basic phones. In this type of auction, buyers and sellers separately communicate their requirements and the prices they are willing to trade at. The system then matches compatible buyers and sellers. Therefore, our users only have to send a single text message, and the system takes price, location and other factors into account to automatically find the best matches.

Google provided support for development in the form of a Research Award in 2011, and the system went live in 2012 as a free service. We began by holding meetings with groups of farmers in Bukomansimbi district and with traders in Kampala. By continuing such meetings in other areas, and combined with radio broadcasts and other forms of publicity, we have had an enthusiastic response from users and rapid user growth, and are matching increasingly large numbers of buyers and sellers across Uganda daily.

Leyton-Brown was on sabbatical for a semester in Uganda. His goal while there was to use tools from AI to make a positive contribution to the wellbeing of sub-Saharan Africa. He and collaborators explored several projects (security games, drug distribution) some of which fell through for a variety of reasons. Kudu was the project that stuck, though it too underwent adaptation. This kind of flexibility is necessary for working on projects that are not only interesting from a theoretical or research perspective, but can also be implemented in practice and have direct impact on people's lives.

## The Problem

Subsistence agriculture serves as the core of the Ugandan economy, but buyers and sellers of agricultural products have trouble connecting in part because development and infrastructure is low. Markets are decentralized, subject to political, natural, and social shocks, and rely heavily on social connections. This leads to sporadic food supply and fluctuating prices despite relative abundance.

As evidence of inefficiency in the Ugandan agricultural markets, Leyton-Brown et al. measured the revenue that could be made from inter-temporal or inter-spatial arbitrage of three (nonperishable) agricultural goods: matooke, maize, and beans, and found large, statistically robust evidence of arbitrage opportunities for all three commodities. We did this to quantify our notion that the market is failing.

**The goal:** facilitate transactions between farmers in rural areas with buyers at markets in cities.

**The medium:** SMS. Even though electricity penetration is not as high as you might hope in Uganda, cell phone signal is fairly easy to come by in Uganda. The phones themselves are mostly basic flip phones (feature phones). So, you can't make some app. Anything you make that you'd want the farmers to use would have to work on these phones.

Note that there were already market centers in Kampala (the capital) that people go to

with their produce and buyers and sellers connect there. There was no need to intervene there since the market was thick and it seemed to run efficiently. The markets in the outskirts, on the other hand, had much less access to information, transportation and quantity was also an issue. We wanted to serve as a bridge between small scale aggregators or large farmers on the rural side and big traders in the cities.

## How Kudu Works

We built a system called Kudu, which is basically “Craigslist running over SMS.” Kudu is conducted via USSD: a menu-based SMS system for cell phones. Seller post asks via SMS and buyers post bids. Note that the quality of products being more-or-less consistent ends up being very helpful here. The system then matches asks and bids based on location, quality, reliability, quantity, and price.<sup>1</sup>

Initially, matches were automatic; participants would be sent each others’ contact information. We now use manual matching via call centers. There is also an additional service with SMS-based price alerts. Note that these trades are fairly consequential since these are individuals whose livelihoods depend on these relatively small-scale farms and they do not trade very frequently.

Buyers are typically intermediaries. They are generally entrepreneurs with trucks. Sellers are farmers or small-scale aggregators of agricultural products. The former group often will have one or two large harvests per year, which they must sell off relatively quickly.

## Timeline

- 2010: Initial sabbatical in Uganda, where we started planning.
- 2013: We had an initial working pilot through 10k Google grant. We advertised on the radio, which is relatively cheap to do. A bunch of users signed up. We matched people and we didn’t know whether the trades actually happened or not. We paid for

<sup>1</sup>Note that there is some measure of reputation, which roughly measures the likelihood of the trade actually going through. See paper for details.



Kudu

### Sell Produce

Highlighted fields are Required

Produce  
Maize

Quantity  
1000

Price (UGX)  
800

Sell produce  
Help On Selling »  
Produce You're Selling »

(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.

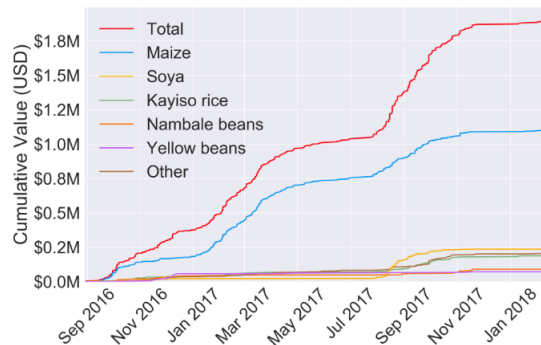
SMS and web-hosting. The grant ran out about 6 months later. (It is very difficult to get grant money for operation costs!)

- mid-2015: Reboot with Lauren and Craig. They had a big project where they wanted to test some interventions and our system looked like it would be a good fit. Neil also came on board, who made the system much better.
- Since 2015: User base has gone from 2.5k to 20k today.
- 2016: Ability to verify trades acquired via IPA. (Note, 80-90 percent of trades don't actually happen, but there's still a substantial number of trades that do go through. So far, there have been a total of 2 million USD in trades, which is substantial.)



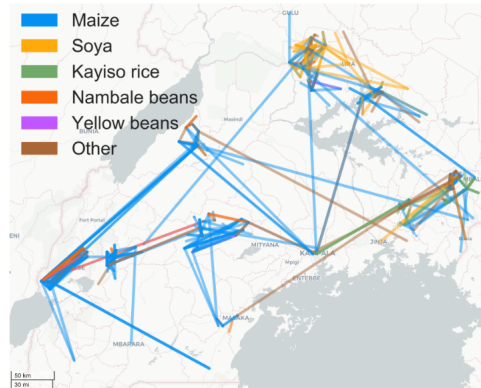
## Recent Usage Statistics

The most traded commodity by far is Maize, representing more than half of all trades as of September 2017. Other commodities include: soya, kayiso rice, yellow beans, and nambale beans. These smaller items are starting to happen more and more given that the market is becoming thicker.



Activity is concentrated in six or seven regions, with many small trades in each region, and occasional long-distance traded between regions. About a third happen in the smallest

geographic unit that we record (a parish), which do not show up on the map below. About 90 percent of all trades happen within a distance of 50 kilometers. Then, there's a long tail of distances, including driving across the country. Certain commodities are concentrated in specific areas (e.g., soya in the north). Kudu still represents a tiny fraction of the agricultural market in Uganda (perhaps single digits). We can expand to other areas of the country with enough budget to advertise there.



The data that we have now might be too sparse to see many fine-grained trends, but there are some interesting observations: e.g., over time, the quantity of trade is going down, meaning that buyers/sellers are able to find enough trades at a given point to cobble together smaller quantity trades as the market thickens. A lot of work goes into training users. There's a lot of interest; many farmers and intermediaries show up to training sessions since they want to have better trades, but it requires training people to learn to use the system before they are up and running on their own.

## Designing the Market

Bids and asks are managed via a computer system. In the old implementation, matching happened manually, which ended up being very inefficient. We even tried to get people to call a global center with their asks/bids, but they ended up reverting to the manual matching, which was unsustainable.

Our view of the system is that you have a bipartite weighted matching problem. We modeled the sellers as being non-strategic, and use VCG pricing for the buyers. Then we decide on the matching.

A lot of critical issues arise in this above framing. Location mattered more than the model gave it credit for; roads differ in safety and safety differs from region to region. We also had to make decisions like: how often should the markets clear? Once per day? Twice? There is a trade-off between speed and thickness of the market. Is it possible to aggregate supply to make trips more worthwhile for buyers? How do you handle the fact that most people have outside options which you do not know about?

Preferences are also complicated. There are various things that matter including: geography, reliability, recency of bid/ask, quantity, willingness to make multiple stops, among many other hard-to-measure metrics. Soliciting this information via SMS was especially difficult; people didn't always answer the questions they're asked or they were resistant.

The conclusion from these observations was that global clearing might not make sense at all. Our solution has been to use a hybrid system where bids and asks are guided by the computer system, but transactions occur with a human intermediary via phone who uses this system. These intermediaries are helpful for navigating a lot of challenges, including that some farmers are illiterate and cannot use the system otherwise. The process is still a lot of work for the intermediaries. There is a lot of back and forth and trades that fall through. Trades that are happening everyday are still on the order of dozens, but this is going up as trust in the system and trades goes up. In the long term, we want to make the system more and more automatic since this is not totally scalable. In the short term, we want to increase the efficiency of each call (have the AI in the back-end suggest better trades more quickly and increase the quality of suggestions). The overarching theoretical question here is: how do you do a matching when there are strong unmodeled outside options, and you don't have strong belief that any of the offers are necessarily real?

Other issues that came up in designing the market are:

- Buyers or sellers are unreliable, or try to renegotiate in person. There is not a social norm of fixed prices or electronic trade, so renegotiation is very common. (There are even cases where a seller drives very far to get to a buyer, but the buyer says that they want a better price after all.)
- How can the market identify the types of different buyers (aggregators vs. distance traders)?
- Is there a way of introducing micro-finance to increase liquidity?
- How do you guarantee the safety of buyers who are transporting cash? Can you have an escrow system? What are the costs involved to implementing this?
- Language is heterogeneous. It is sometimes a constraint that the buyer and seller do not speak the same language. The call center has to operate in many different languages.
- People can drop out of the system and rejoin with a new SIM card (hence identity) easily. How do you build a good reputation system when it is so easy to rejoin the system as a new user?

We noticed that there are different types of buyers. Some are very local. Some are distance traders, but want to end up in one region after they're done. Some are very flexible regionally. You might think that buyers would be very willing to tell us what kind of buyer they are, but it was difficult to get this information from them. There are technological and cultural challenges to these. It's difficult to collect information at scale in this setting. Overall, there's still a lot of work to do, but there has been real success as well; a lot of trades have happened because of Kudu. We're moving towards a world with continuous clearing, and this project is inspiring a lot of interesting market design questions.

## Tips for Practical Projects

Several pieces of practical advice came out of discussions throughout the talk, for those interested in pursuing similarly practical projects:

*Question:* How do you find good research projects that might be of practical use for research at this interface?

*Answer:* Don't expect cleanly specified problem. Getting a problem which is phrased in terms that we would call a problem is 90 percent of the work, and it is even ongoing in this case. We had an email thread over the weekend about should we pay attention over prices at all. Maybe geography is key here. Finding the right way to specify the problem so it matters to the stakeholders is huge. It is helpful to have a lot of conversations with stakeholders. They might not know what the problems are either, so you would need a lot of back and forth to find something that works.

*Question:* You were on sabbatical for a semester. Was that sufficient time or were you immersed in the process before?

*Answer:*

The semester was pretty good for getting started. Going somewhere and living there even for a month or two is different from having the odd Skype call. Seeing the actual conditions or talking to people casually makes a huge difference. You don't leave with a well formulated problem and things keep changing even now, but we were able to leave with a much better understanding of what might be helpful. Before talking to stakeholders, you have a lot of wrong ideas about what the problems are, what people need, and what might work. This mirrored some of my experiences with working with startups in Silicon Valley. You have some idea about what the market wants, and you need to be able to pivot again and again.

*Question:* Were you already experienced in this process of immersing yourself and talking to lots of people before getting to Uganda?

*Answer:* That kind of experience comes with time. I had worked a little bit in consulting and worked on other ill-defined projects which helps develop this skill and be comfortable with the process. This might be the first time that I worked on something that is very application-driven, and talked to a whole bunch of people without a lot of broader planning. It helped a tremendous amount that there was a local collaborator who was also a researcher. I would recommend that model. There was a strong sense from the people that we worked with in Uganda that they really valued the presence of a Ugandan connection. It helped that the problem was local and there were local partners. It's also important to understand that the local partners also didn't necessarily have a cleanly defined problem. Whenever you're working in a new domain, you have to spend a lot of time trying to understand what the most helpful problems are in this domain. But, if you have some connections or you are local, then you can make the initial contacts in a more sensible way.

*Question:* What were their connections? Did you already have contacts because they were local or did they just know where to look?

*Answer:* Each connection is different. E.g., Richard has other members of his family who have worked as agricultural traders in one form or another. He has also been able to put us in touch with microfinance people since he now works in banks. It wasn't like there was

some local organization that works on exactly this thing. You just build the connections over time. In some sense, Craig and Lauren as development economists, have had tighter connections to some of the stuff about measuring impact. But, having local connections ends up being very valuable in a variety of different ways that you might not expect a priori.

*Question:* How do you trade-off working on projects that our research communities easily recognize as research and doing something practical? A lot of the practical challenges that you've been facing are not necessarily research questions. How do you balance these competing needs?

*Answer:* I'm more focused on getting the system off the ground. I'm less concerned about publishability in part because designing markets that people actually use has to be important to us as a research community. I'm very motivated by understanding what it takes to design markets that work in practice. Engaging deeply with a real problem will surface questions that we don't yet have good answers to in our research community and can lead to theoretical questions that would be of interest to the community. It's definitely a bit of a gamble and it is not the fastest way to get publications. But, there is a definitely a desire to see markets that work in practice. So, one approach might be to play a mix strategy. This allows you to take the long view and do something risky alongside doing more traditional stuff. The more stressful aspect has been funding the project since it's very challenging to get grant money to pay for operational costs. We, as a research community, need to work on how to create these research funding opportunities, advocate for successes, and share information with each other.

## References

- [1] Neil Newman, Lauren Falcao Bergquist, Nicole Immorlica, Kevin Leyton-Brown, Brendan Lucier, Craig McIntosh, John Quinn, and Richard Ssekibuule. Designing and evolving an electronic agricultural marketplace in uganda. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, page 14. ACM, 2018.