Social Enterprise Solutions for Improving Livelihoods of Smallholder Farmers in Sub-Saharan Africa

Apollo – Application

Proponent details

Lead organisation legal name:	Apollo Agriculture
Primary contact person:	
Primary contact e-mail address:	
Primary contact telephone number:	

Key proposal details

Title of proposal:	Supporting the growth of Apollo
Target country:	Kenya
Target population:	Smallholder farmers
Overview of organisation and proposal concept:	Apollo Agriculture facilitates a model where smallholder farmers are provided with loans to purchase high quality inputs, agronomic advice and insurance. Apollo will procure inputs directly from manufacturers, develop and deploy remote training modules with basic farmer training, and engage specialist third parties to provide farmer insurance products. In the future, Apollo will develop the ability to recommend and commission regional optimized fertilizer blends and seed/fertilizer regimes to small scale farmers. Loans will be provided by Apollo and, in the future, by other financial institutions relying on Apollo's insights into smallholder farmer creditworthiness. Apollo's innovation is in the use of mobile technology, satellite data and advanced statistical methodologies to predict farmer credit outcomes and assess credit risk, thereby enabling the provision of loans to underserved farmers. This integrated approach to working with farmers is intended to increase yields and improve livelihoods. The process works as follows: 1. Farmer requests an input loan via SMS 2. Credit risk is evaluated based on satellite data and remote farmer registration survey 3. A field marking agent walks the boundary of the fields for which the farmer is requesting a loan. 4. Qualified farmers make a nominal downpayment via m-PESA, receive inputs and take out bundled weather insurance 5. Basic farm practice training is provided remotely (e.g. via SMS) 6. Farmer repays loan via m-PESA following harvest Linking specific farmers to unique farm plots encourages repayment and reduces the risk of fraud.
Current project status and/or details of relevant experience in sub-Saharan Africa:	Apollo's roadmap is presented below:

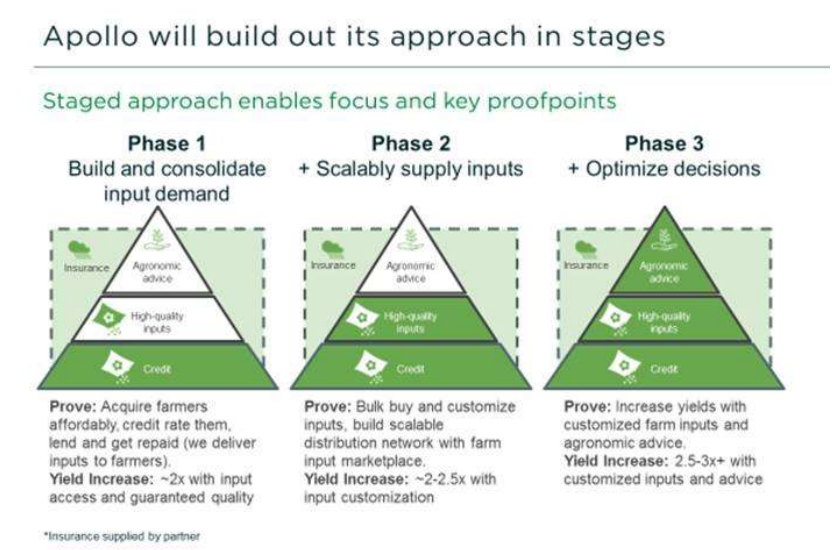

The business is currently testing whether it can acquire farmers and garner their interest in the model – this is taking place through traditional and digital marketing and advertising, and SMS referral. It will then affordably credit rate them, lend and get repaid. The first loans were provided to 1,000 farmers (with over 6,000 requesting support) covering 2,000 acres in Nakuru District, in March 2017, due for repayment after harvest (likely November 2017). Critically, this process will allow the business to test how effectively credit risk can be evaluated based on satellite data and remote farmer registration. Current assumptions on the relationship between farm/farmer characteristics and successful repayment are based on the team's background experience, but the team will not rely on assumptions to make lending decisions as it builds the credit rating model. Rather it will lend broadly and then rely on sophisticated, "machine learning" techniques to identify the factors that are most predictive of repayment. The following schematic summarises the approach to predicting which farmers will repay the loans at the conceptual level:

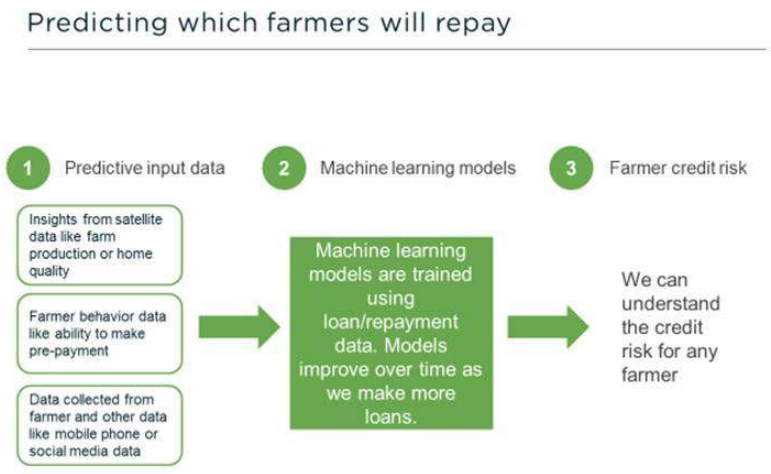

Depending on the success of this pilot, Apollo will then bulk buy inputs and develop a scalable distribution network, before fully developing and rolling out the supplementary agronomy and insurance service offering.

Experience of management team:

The lead team at Apollo previously worked at The Climate Corporation, a digital agriculture company that analyses weather, soil and field data to help US farmers determine potential yield-limiting factors in their fields. This company was recently acquired by Monsanto for circa $1bn. The team also includes staff with strong experience in working with smallholder farmers, agronomy, and insurance in Africa. As such, the team has credible experience in technology development and business commercialisation of relevance to Apollo, as well as in working with smallholder farmers in Africa.

Detail on how the proposed concept will provide cost-effective services using a social enterprise model of delivery:

The cost of serving remote areas is a key reason behind the lack of access to inputs in these areas. A key value add of Apollo's proposition is through the use of satellite data to remotely acquire data on productivity and challenges, which would otherwise require farm visits and group meetings, to predict credit outcomes. This is expected to dramatically reduce the cost of serving remote farmers.

The table below details the unit economics of the model. As shown, in Phase 1, Apollo will charge farmers KES11,500 (c. $115) per acre for inputs, remote agronomic advice, and crop and weather insurance. The programme costs equate to $101, creating a gross profit of $14 per acre, which increases as the model is rolled out over Phases 2 and 3.

Per Acre	Phase 1 - Build Input Demand	Phase 2 - Improve Supply	Phase 3 - Optimize Decisions
Farmer price *Due after harvest*	$115.00	$115.00	$130.00 *Price increases from yield increases*
Farm inputs	($81.00) *Apollo purchases inputs and delivers to farmer*	($68.00) *Bulk purchase inputs, distribute through Apollo marketplace*	($68.00)
CAC + Field Ops	($7.50) *Mobile distribution, remote credit rating*	($3.00) *Marketplace handles some field ops*	($3.00)
Capital	($9.00) *10% capital cost*	($7.10) *Reduced loan size*	($7.10) *Reduced loan size*
Defaults	($3.50) *4% with satellite credit rating*	($2.60) 3% with local marketplace support	($2.60)
Gross Profit / Acre	$14.00 / 12%	$34.30 / 30%	$49.30 / 38%

For farmers, the model during Phase 1 is expected to double yields over baseline, which increases to 2-2.5 times baseline in Phase 2 and 2.5-3 times baseline in Phase 3.

Detail on how the concept will be sustainable:

Apollo's business and revenue model is as yet untested, but is designed to become financially sustainable over time.

At $1M annual burn rate (according to financial model cash flows April '18-March'19 provided to EfD) and $13.5 per acre/farmer margin, we theoretically reach break-even at the entity level once we are serving 74,000 farmers annually. This is actually a conservative estimate, since we anticipate improved input pricing as we scale and bulk our purchasing. If we realize half of the $13 improvement that we estimate for Phase 2, then our per farmer margin increases to $20. At $20 per farmer margin, we need to serve 50,000 farmers to break even as a company. While we, of course, have a long way to go before we achieve that milestone, early indicators have only served to validate that this is a reasonable projection.

Detail on how the affordability of the proposed services to the target population:

The per acre charges are broadly in line with the charges levied by One Acre Fund, NGOs and agro-dealers in Kenya. It is acknowledged that this will exclude some marginalised groups of farmers, but has been proven by other actors to be affordable by the majority of smallholder farmers.

<table>
<tr><td>Detail on the feasibility of scaling up the concept to other populations and regions:</td><td>Taking advantage of the increasing adoption of mobile technology among even the most remote smallholder farmers, and the use of satellite data and remote agronomic advice, Apollo is designed to be a highly scalable model that can be expanded into new areas at relatively low cost. This is a key advantage over other initiatives that require face-to-face meetings with farmers, and face difficulties in serving remote areas in a cost-effective manner.

Appropriate working capital arrangements will be needed to support the scale up the business to other regions. We are looking ahead towards raising $1.35M in debt for the 2018 long rains season. Our current impact lending partner (Ceniarth LLC) has suggested they intend to grow their lending commitment alongside our growth next year. We will be formalizing this expectation in the weeks ahead. Additionally, we are cultivating additional impact lenders, including one discussion that is mid-stage and high probability. We are working to build long term relationships with lenders with greater capacity, such as OPIC and Developing World Markets (raised and is managing a $60M fund for off-grid energy finance). We are participating in a global conversation OPIC is organizing centered on “fundamentally changing the scale and speed of support brought to small holder farmers around the world,” which includes participants from AGRA, BMGF, CDC, DFID, USAID, the World Bank, and others.</td></tr>
<tr><td>Existing financial and technical partners:</td><td>Technical work is led and completed in-house, with support and oversight from senior Stanford academics on the use of satellite imagery and data science for food security.

The capital costs of setting up the business, developing the satellite technology and data collection are being borne by a mix of investors, including:
Equity partners – Accion Venture Labs, Factor[e] Ventures (linked to Shell Foundation), OEL Venture Investments I (Oberndorf Ventures), and four angel investors.
Debt – Ceniarth (Harry Davies) will provide debt to purchase inputs for the next two seasons.</td></tr>
<tr><td>Project risks:</td><td><table><tr><th>Risk</th><th>Level of Risk</th><th>Mitigation</th></tr><tr><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td></tr></table></td></tr>
<tr><td>Project costs:</td><td>TBC</td></tr>
<tr><td>Proposed EfD involvement:</td><td>EfD support is requested to support the salary for an additional employee to accelerate the pace of our research and development (R&D) program. The employee will contribute to the development of our ‘yield model’ – this is part of our core technology – by collecting additional data, such as soil samples, which could help us build more robust services in the future. The new employee will focus on improving the automatic analysis of satellite imagery, as well as the algorithms that produce seed and fertilizer recommendations for farmers. Our primary restriction on accelerating our growth is one of capacity.

More specifically, the employee will take forward a defined piece of modeling work which will be a combination of data collection, software engineering and image processing (taking raw satellite imagery and efficiently translating it into usable information – feeding into our yield model, but also yielding additional insights such as field shape and boundary angle, roof material, and so forth), and model building (either/both with more straight forward statistical modeling seeking to identify a relationship between a set of input data – yields, family size, loan size, farm practice, etc. – and output data – repayment and with</td></tr>
</table>

the machine learning/artificial intelligence approach that can identify relationships more dynamically between unstructured data and a particular outcome such as repayment).

This image reflects some of the "feature recognition" work that an additional engineering resource would help us automatically process from satellite imagery:

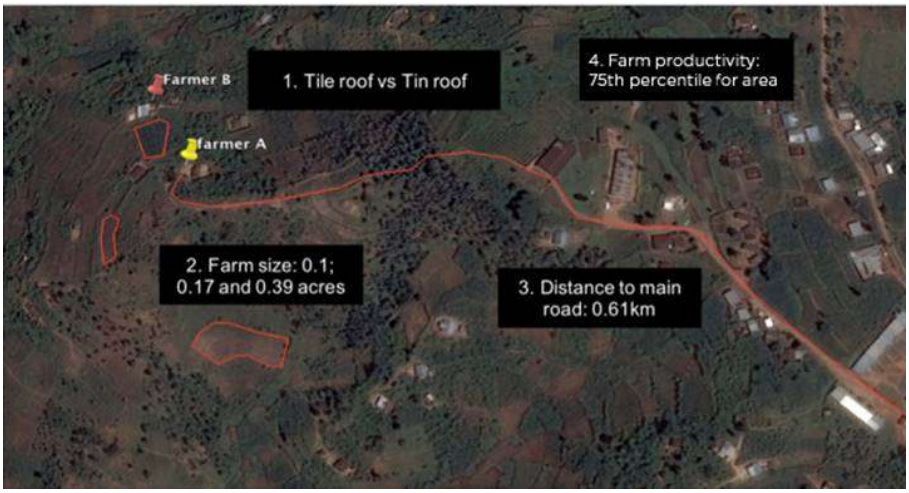

And this one the ways in which a sophisticated yield model would provide us with in-season comparative risk information:

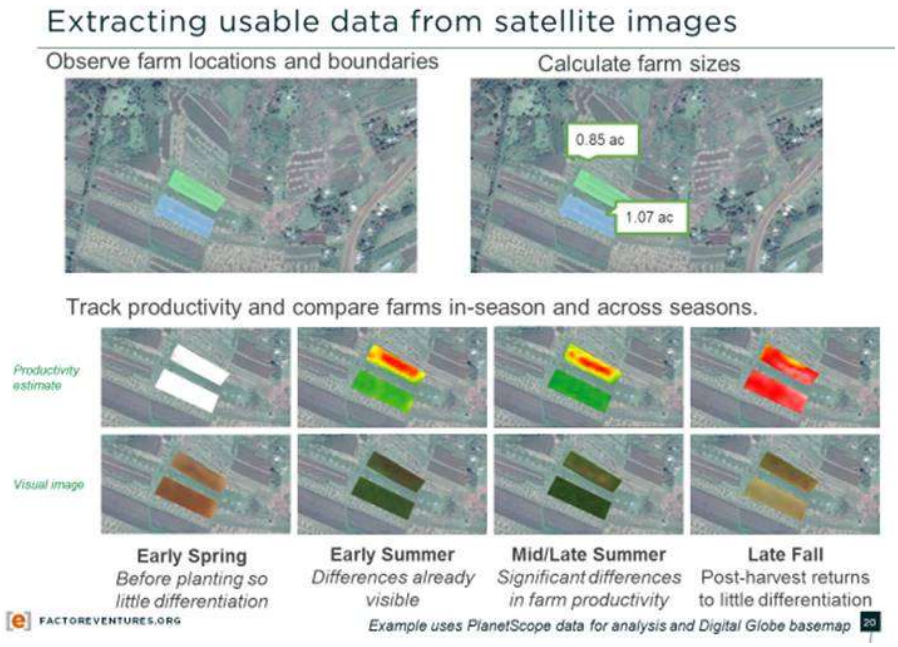

The yield model is both an exciting and novel contribution in its own right and an essential input to our eventual credit model. The yield model is important for a number of reasons not least of which is that it helps us gauge the impact of the loans we are providing. We have randomly declined loan applications from certain farmers for this season and we will be eager to gauge and understand the difference in yields between farmers we serve and similarly situated farmers we do not. With a robust yield model, we will be able to use land production history as an input into our credit score itself if we determine that there is a relationship between historical yields and loan repayment success. We are also experimenting with registration survey questions that can be cross-checked against observational satellite data: when we ask a farmer whether her harvest was better than

	average, average, or worse than average last year, we can have an externally validated cross-reference for the answer with the potential that any difference or agreement between those two sources indicates honesty, memory, or seriousness about their farming business, all of which have a strong chance of being related to repayment behavior. Over time, a good yield model will also allow us to target our marketing and service territory for both success and impact.

EfD due diligence

Feasibility of proposed solution:	• What evidence is there from other initiatives (e.g. One Acre Fund) to demonstrate that the Apollo model will work? • Is the solution economically, technically and culturally appropriate? • Does the current research conducted help to prove the model's feasibility? • Is there sufficient local capacity to implement the solution at scale?
Feedback from existing investors:	

Responses to EfD questions/ concerns

What is the cost of building an advanced yield model and who will bare this cost?	Apollo will build – and actually have already outlined the architecture for – a smallholder maize production yield model on the basis of satellite imagery. That investment is both already in the works for us and is sufficiently "mission critical" that we feel like we must build it one way or another. Where the additional resources from EfD would feature is in enhancing the sophistication and applicability of the model we build by affording us the ability to invest more in both data collection and model development. • In terms of data collection, more resource will allow us to capture more and different types of data. With your support, we would collect harvest measurements on the fields of all the farmers we serve this season, alongside all those in the control group. We have already built an efficient rural task allocation technology that allows us to farm out discrete jobs to casual rural labor on a per piece basis at a cost of roughly $2.50 per task. (This is how we accomplished the GPS field marking aspect of our farmer registration and enrollment process.) Soil samples are another example of data we could collect more comprehensively (though probably not universally) across the farmers we serve. Wet lab soil analysis is more expensive (~$5/sample plus the labor and logistical costs of collection) and we are rushing already to collect at least some samples before the rains come. This soil data could be related to the yield model and, down the road, inform soil sampling methodology/density for the purposes of optimized input design regimes. In addition to harvest box measurements and soil samples, your resource would allow us, similarly, to conduct a planting practice audit (row and whole spacing, fertilizer application, etc.), which allows for us to a) add nuance to our yield measurements, b) understand how significant planting practice is to yield results (partners/friends who have researched similar questions have partial results suggesting it is meaningful), and c) feed into planting practice and input regime recommendations down the road. As part of this enhanced data collection process, we would likely build out our off-mobile network compatible field application to collect, geotag, and deliver the data that is collected efficiently. A very rough estimate of $20,000-$25,000 should support this enhanced data collection.

	• As for model development, there is a similar logic. Our existing resource will support our hiring an additional data scientist/software engineer for our team. Based on the work we need to do to both build software to support our field operations and develop our first generation yield and credit models, this hire is essential and was part of our original investment budget. With more and richer data and yet another engineering resource we could build out these more sophisticated models that reflect the relationship between our soil samples, yield data, and repayment while building moment for the future work we want to do in generating more particular agronomic advice and input regimes. Our existing supporters are understandably focused on supporting our efforts to prove out step one: that satellite imagery can pair with other data we have collected to provide reasonable predictions of repayment risk. As we go forward, however, we need to both refine that core model and component and leverage the data and systems we will uniquely have to offer better value to smallholders through more personalized input recommendations (e.g. appropriate fertilizer blends, seed variety recommendation). In simple terms, we will not have the capacity to begin this work without an additional hire and we would love to look to EfD to cover part of the salary for this hire to enhance our capacity. Realistically, we expect to be able to find world-class engineers for our team at $80-$90k – a steep discount from market, but meaningfully higher than our management team's current compensation. You mentioned that having individuals identified will help our case. We have been courting particular individuals that Eli and especially Earl have worked with professionally in earlier chapters of their career, one or more of whom might convert in our recruiting process. We have been reluctant to "broadcast" a position until we were more confident in our decision to hire for it (and be able to support the salary in our budget). With that said, we discussed yesterday the need to move tech hiring to the top of our priority list immediately following post-delivery cleanup activities.
We are keen to understand your staffing arrangements in the short- and medium-term, and what your personal circumstances are.	Eli Pollak is our full time CEO, Earl St Sauver is our full time CTO, and Benjamin Njenga is our full time Operations Manager. All are based full time in Nairobi and travel frequently to the four wards we serve in Nakuru county. We have a small call/data center in our office in Nairobi to process farmer registrations as they come through, which is staffed by Japheth Hosanna and Elizabeth. We have three "Field Associates" – low cost, but slightly higher capacity field agents on retainer with us through the end of April – and most of our field facing work is done on a per task basis by a fleet of rural field workers. In the medium-term we expect to add a software engineer/data scientist and would like to add another, resources allowing. Chris Seifert, who is completing his PhD at Stanford basically in our work, is part of our founding group and contributes to the science elements of our R&D work. We expect Chris will join us full-time after completing his PhD, but that is not in the medium-term. Rose Goslinga also contributes in a part-time capacity focusing on building our crop insurance product and company strategy. Finally, Seth Silverman wears two hats, both at Factor[e] and as a co-founder/Director/member of the management team with Apollo. This is unique to Factor[e]'s co-development model. He does not draw a salary from Apollo, but is involved in all major operational decisions, manages operations design and, in partnership with Ben, our operations, and in particularly critical times such as input delivery acts more like a full-time member of the team.

Is there an ongoing relationship with One Acre Fund? How can/will you leverage this relationship?	We have an ongoing friendly, but informal relationship with One Acre Fund. Personally, many of my closest friends remain in One Acre Fund's leadership and our model takes inspiration from their success. They provided us with their microdose "planting scoops" for this season and we have traded notes on their soil sampling work and ways to build sophisticated agronomic modeling with the data that they have collected. Some of their grants are specifically geared toward generating large volumes of open-source data about farmer activities and we will certainly take advantage of that data whenever available. We have, in turn, shared our early indicators of success and approach to distribution through agrodealer distribution, alongside the general logic we are testing and pursuing. Though we do not have a formal relationship at the entity level, there are several future paths to leverage that relationship, including selling our insight to OAF, especially for lower farmer density areas and to identify those farmers who will not succeed if they do not have more support, working together to develop fertilizer blends for particular sub-regions, and in accessing working capital.
Have you liaised with AGRA and the Gates Foundation, both of which have invested heavily in this area? We are happy to make connections, as needed.	We have had an early conversation with AGRA related to their interest in scoping out the kind of work that they might do in digital tools for agriculture. We think there is tremendous long-term potential there, particularly in using the insights our model can generate to make national fertilizer subsidy programs more efficient and impactful. We haven't yet had the evidence from our direct delivery activities or the capacity to chase this vision, but are excited to and would be happy to have an earlier conversation with your partners at AGRA. Similarly with the Gates Foundation, we haven't yet begun serious conversations about this very early work, but if you had someone that you would recommend we connect with, we would welcome the help in making that connection.
How have you selected pilot plots (e.g. based on crops/ proximity to markets/ plot size)? How will this affect the results of the pilot?	We have assiduously avoided biasing our selection process during this first season. We have deliberately restricted our service territory geographically to ensure that we do not only serve the progressive, early adopter-type farmers. Following a range of broadcast marketing techniques (each tracked as best as possible with a different "join code"), we do basic Know Your Customer screens and then mark the farmer's field. Everyone who walked through our registration process in full in time for our deliveries (and wasn't randomly assigned to a decline control group) received a loan.
Are you able to use credit scoring data available to mobile phone companies in Kenya?	In terms of our ability to incorporate the credit insights of companies like Branch, Inventure, and First Access and Safaricom's mShwari program, where it is available, this data will be readily incorporated into our model. We are reluctant to build a model that depends on this data as it won't always be available to us and, itself, necessarily only presents shadows of their customer's economic lives. Imagery of farm-based activity, we think, can provide a clearer picture of the financial lives of farmers.
What are your initial thoughts on the likely approach to attract finance to reach scale and who from?	This is a very important and live question. We are confident in our ability to attract equity finance to keep growing our business and chasing our vision. We have had productive conversations with a number of primarily impact oriented early-stage investors and have relationships for follow-on funding that we are hopeful about through Accion Venture Labs and Factor[e] Ventures/Shell Foundation. Raising working capital for our lending is a longer term challenge, though for the coming season, we have every reason to believe that Ceniarth LLC will grow with us. In the longer-term we will have to both zero in on the model (direct deliveries and guiding other partners like fertilizer or seed companies, government fertilizer

	subsidy programs, etc.) and mobilize significant debt capital through partners like OPIC and others. This is an area we would gratefully accept advice and guidance in. Though arguably vague as a plan for the moment, our friends at One Acre Fund provide an example of an organization that has managed to find the debt/working capital to scale their lending as their operations have grown. I believe they have a portfolio of around $60M in smallholder lending this year.
Please could you confirm the different mechanisms for generating revenue.	With a per acre package (e.g. 11,500 ksh per acre – interest isn't broken out from the core package), our model generates revenue first and foremost from margin on each acre repaid in full. In the future, we see potential to generate revenue through selling the product of our credit model to partners (e.g. fertilizer/seed companies, governments) as well as finding other ways to leverage the proprietary data we will have built up.
How will you ensure crop inputs are adapted to different farmer circumstances, and are accessible to Apollo)?	We will both draw on official regionalized agricultural expertise (Kenya maize farming different from Zambia different from Nigeria) and work in subsequent phases to build our own, sophisticated crop input models to optimize input regime design for the regions where are farmers live and produce. This was discussed above and is a longer-term vision that EfD's support could help us begin building momentum towards.