DESIGNING INDEX-BASED WEATHER INSURANCE FOR FARMERS

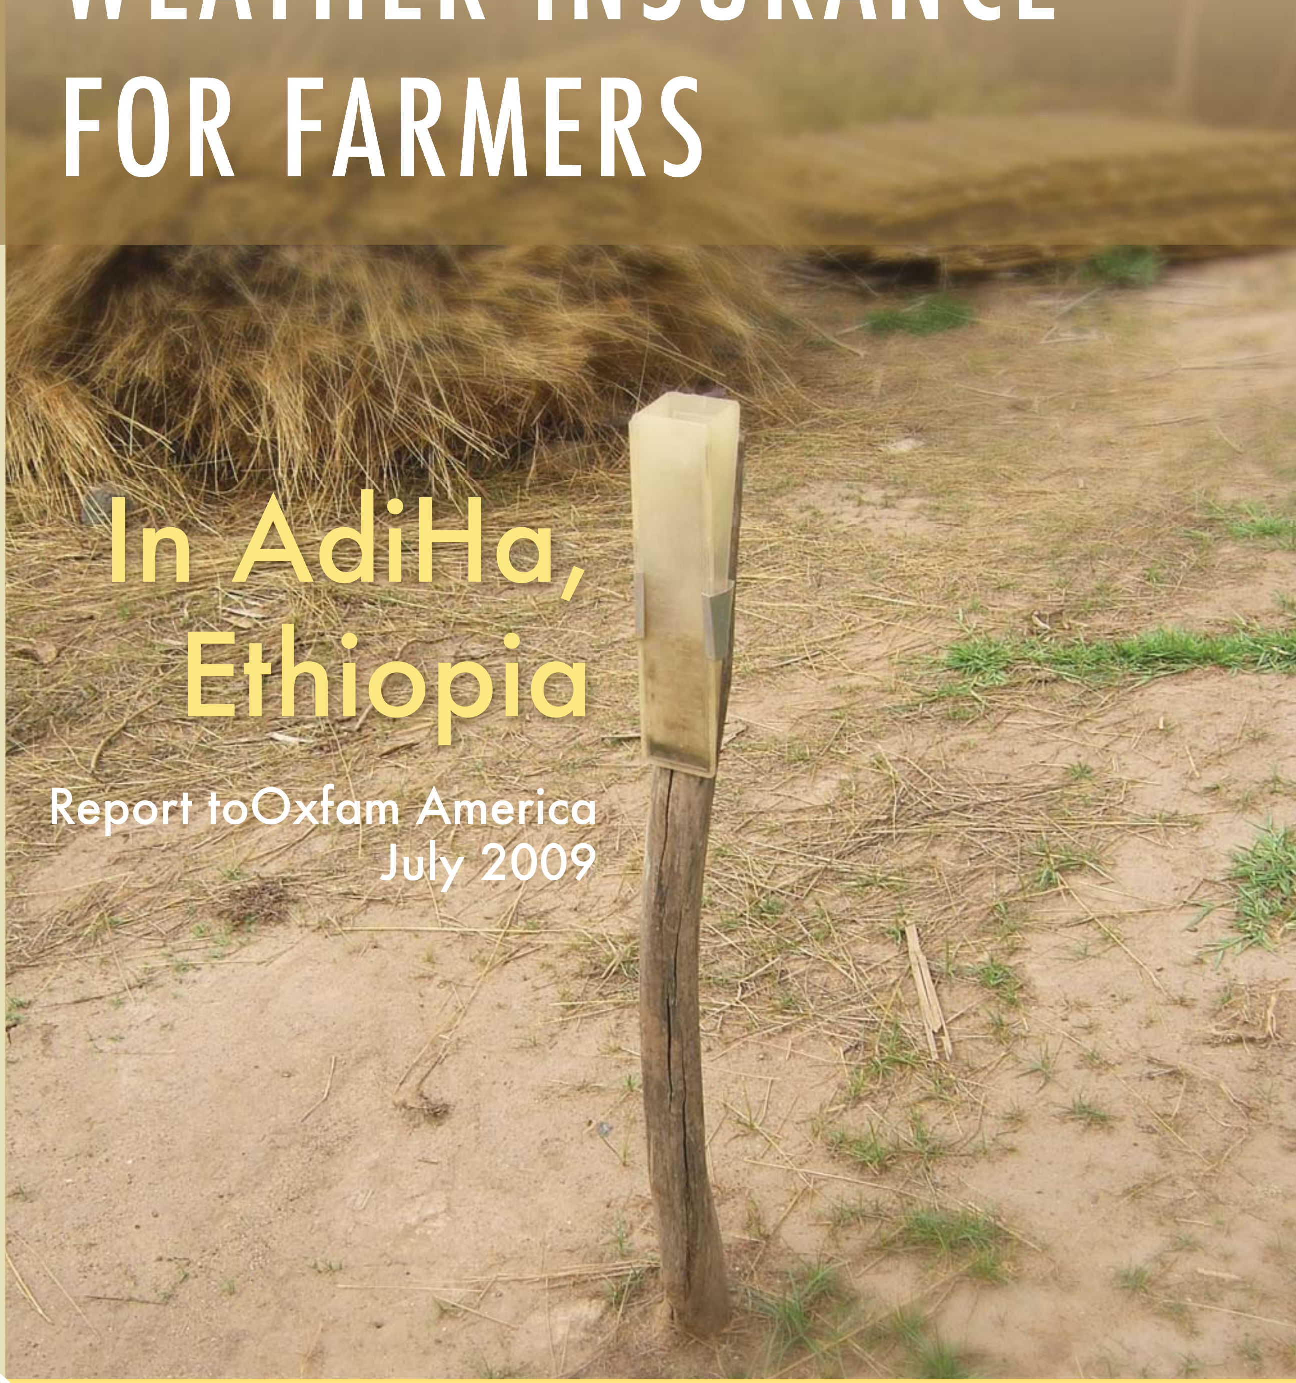

In AdiHa, Ethiopia

Report toOxfam America
July 2009

Prepared by: Tufa Dinku; Alessandra Giannini; James Hansen; Eric Holthaus; Amor Ines; Yasir Kaheil; Kristopher Karnauskas; Bradfield Lyon; Malgosia Madajewicz; Megan Mclaurin; Connor Mullally; Michael Norton; Daniel Osgood; Nicole Peterson; Andrew Robertson; Kenneth Shirley; Chris Small; and Marta Vicarelli

International Research Institute for Climate and Society
Earth Institute, Columbia University

Cover Design by Jason Rodriguez

Disclaimer

Final responsibility for the views expressed in this report lies with the authors. The views are not necessarily those of Oxfam America or NOAA.

Contact:

Daniel Edward Osgood

Associate Research Scientist in Economic Modeling and Climate

International Research Institute for Climate and Society

The Earth Institute at Columbia University

deo@iri.columbia.edu

http://iri.columbia.edu

Phone: (845) 680-4461

Fax: (845) 680-4865

This report was funded by Oxfam America and the US National Oceanic and Atmospheric Administration (NOAA), which provided its support under cooperative agreement NA050AR431104. Additional funding sources include: The Earth Institute, LDEO, and the Applied Statistics Center.

IRI Technical Report 09-04

http://iri.columbia.edu/publications/id=908.

DRAFT FOR CIRCULATION

INTRODUCTION

This report documents the process and results of the index insurance design effort leading to the index insurance contracts for Adi Ha in 2009. This report represents deliverables 1 and 2 in the terms of reference with Oxfam America, it outlines and compares analysis and design methodologies including the performance of rainfall simulators for index-based contract design. It also details contracts, methodologies, associated issues, and important lessons learned. A separate project report details the Experimental Games, deliverable 3 in the terms of reference.

IRI has delivered contracts and methodology with a significant portion of the effort provided without cost to Oxfam America. This includes contributions from several external researchers who have contributed time at no cost. Substantial funders who have provided additional sponsorship include NOAA, NSF, (through its CRED DMUU center at Columbia), the Columbia Earth Institute postdoctoral program, the Columbia University Lamont postdoctoral program, and an Earth Institute CGSD seed grant.

The insurance effort is one component of the Oxfam HARITA project, with the insurance designed to target key gaps left in the HARITA holistic climate risk management portfolio. Because this report focuses on the insurance contract design, it does not provide a comprehensive overview of the full climate risk management strategy, although it does refer to particular components in discussion of design issues. For a complete overview of the risk management and adaptation portfolio, the reader should refer to Oxfam HARITA project documentation.

Many partners have participated substantially in the contract design effort. Major partners in this project include the Relief Society of Tigray (REST), Dedebit Credit and Savings Institution (DECSI), Nyala Insurance Company, the Ethiopian Productive Safety Net Program (PSNP), the Government of Ethiopian National Meteorological Agency (ENMA), Swiss Re, Mekele University, Oxfam Horn of Africa Regional Office, and Oxfam America (project coordinator). Any successes of this project depend strongly on the outstanding effort, skill, active engagement, and level of insights from the partners. Crop modeling parameters were obtained from LEAP and Mekele University. This project benefits strongly from groundwork and ongoing index insurance projects in Ethiopia and elsewhere by WFP and the World Bank CRMG.

Project goals and strategies

The project is part of the transition into the next generation of approaches necessary to overcome the challenges facing index insurance if it is to help address poverty at large scales. Central to this challenge is to consider index insurance as the last piece of the puzzle. Instead of implementing and scaling index insurance per se, the insurance is designed to help reduce remaining risks after existing development interventions and traditional risk management activities are utilized. The goal is to design and build cost effective, robust, scalable index insurance into a package of development interventions yielding products demanded by low income farmers and workable in data poor contexts.

There have been several index insurance projects around the world, providing a foundation of technology and illustrating its potential. At IRI we have contributed to many of these efforts[1]

[1] At IRI, we have worked to overcome challenges in index insurance in Ethiopia, Ghana, Honduras, Indonesia, Kenya, Malawi, Mali, Nicaragua, Nigeria, Rwanda, Senegal, Tanzania, and Uganda. From these projects, ranging from farmer to macro level, thousands of insurance

(Osgood et al. 2007, Gianinni et al 2009, and Hellmuth, M.E. et al. 2009, Holthaus et al, in preparation). From this experience, it is clear that new developments are necessary. The next steps must now be taken in order for index insurance to meaningfully contribute to poverty reduction at large scales.

For demonstration of concept, past pilots have focused on the sites that are relatively data-rich, for example, locations for which official government met-stations exist with several decades of historical data. A key challenge not addressed by these pilots is how to meaningfully offer responsible products for the site for which there is no existing met-station and where there is very little other historical data. Although there are many potential analysis tools that may useful in these situations, they must be tested and their performance, limitations, and pitfalls evaluated. In addition, the existing index contracts themselves may not be workable in 'data-poor' settings potentially requiring new types of products.

Addressing these issues is particularly important as projects move to large scales. Since it is possible to closely monitor and address problems arising in a small scale pilot, it is a productive environment to handcraft prototypes. However, for products at large scale, the transition must be made to a factory floor, which can produce a large number of extremely robust, useful, and low cost products for which the performance and pitfalls are well known and easily communicated. This motivates the need for a new generation of streamlined, industrial strength approaches. Because the challenges vary across locations, development scenarios, and risk levels, there are many different types of approaches that must be explored to address the diversity of challenges. The Adi Ha project is a one approach intended to explore potential steps forward in addressing many of the challenges facing low income farmer level insurance projects that compliment development activities in data poor environments.

Next generation challenges:

- Integrate easily to **strengthen the portfolio** of complimentary development and risk management activities, **unlocking development potential** blocked by climate risk.
- Be built around local **bottom up demands** using local and regional design expertise.
- Be effective and responsible in a **data-poor** environment.
- Be well validated and extremely **robust**. Limitations of the product should be quantified, and the level of understanding about the product performance should be quantified.

---

contracts have been transacted. Partners include the World Bank CRMG, Oxfam America, UNDP, OI, MIA, FIDES, the Millennium Villages Project, Swiss Re, with additional national partners in each location. We have also worked to document the current state of knowledge through the Climate and Society Publication Vol. II (CSP2) process, which is designed to capture the experience, concerns, and innovations of the broader index insurance implementation as it transitions toward large-scale solutions. Partners in the Climate and Society Vol II include Swiss Reinsurance, World Bank, World Food Program, International Fund for Agricultural Development and Oxfam America. The foreword is signed by Kofi Annan. The publication was funded primarily by UNDP and NOAA. Additional partners and funders for Climate and Society II events include MCII, UNU, NSF-DMUU. Institutions from outside of Columbia University that are represented in the writing team include Cornell, Duke, FAO, IIASA, Le Centre Européen de Recherche et d'Enseignement des Géosciences de l'Environnement, NASA, Partner Re, Red Cross, Oxfam, University of Florida, UCSB, University of Reading (UK), Wageningen University (NL), WFP. Institutions from outside of Columbia University that were represented by workshop registrants at the workshop on technical issues hosted at IRI include Cornell University, GEF, Green Ink, Harvard Business School, IFAD, IIASA, Jawaharlal Nehru University (New Delhi), Liverpool School of Tropical Medicine, MIA, Millennium Promise, Mississippi State University, NASA Goddard Space Flight Center, NOAA, OECD, Oxfam America, PartnerRe, Penn State University, Purdue University, Rockefeller, U Miami/RSMAS, UCSB, UNDP, UNEP, University of Florida, USGS/UCSB Geography, WFP, World Bank, Yale University, and YellowJacket.

- Be **extremely transparent**. The clients should know exactly what the products do. Perhaps more importantly, they must know what the products will not do so that they are aware of what risks they are still exposed to. **Pricing must be transparent**. Clients must be very aware of exactly what they are receiving for the price they pay. Without these features, clients, whether farmers or development agencies, will not be able to make the necessary trade-offs to determine how, and if, the insurance can play a worthwhile role in their portfolio of risk management activities.
- Be **reproducible and adaptable at low cost** by local designers.
- Be effective for those with very **low incomes**. For very low income groups, the challenge with insurance is not to develop a 'Cadillac' product that eliminates all risks, but instead to find the product that provides **the most risk reduction for the first few dollars spent**, targeting the risk that could not have been addressed more effectively through other activities, and the risk that prevents other activities from succeeding.
- Have available a suite of effective, robust, and well understood scoping, communication, design, and analysis **tools that can be harnessed by local designers** to address the range of local challenges.
- Have **communication systems** between clients, local designers and researchers, and international designers and researchers so that new problems can be solved as they arise, with appropriate contributions from each level and knowledge can be built for the global community.

The Adi Ha project was developed to explore important initial steps towards these goals. Adi Ha was selected for this exercise because it represented many of the common challenges to be faced by index insurance, challenges that will be important to address for index insurance to be a meaningful tool over large regions of the developing world. Instead of replicating the efforts of existing pilot activities, the project was designed to directly target a site with challenges that index insurance has not yet overcome: using insurance to complete a broad portfolio of risk management and development activities for very low income farmers in a data-poor environment.

The design process was a holistic, farmer driven approach. New levels of transparency, robustness, and simplicity were sought through next generation index insurance development techniques. These next generation techniques utilized much more sophisticated modeling, design, and analysis tools than we have used before to streamline contracts. Through more sophisticated tools, we have arrived at much simpler, more effective, and more robust contracts with lower and more transparent prices.

The project has attempted to forge some of the next generation of tools and approaches that can address the challenges of low income and data poor sites such as Adi Ha, tools that could be refined and used in the future to build a new 'factory floor' that could assemble robust, cost effective, and tailored insurance products at large scales that are worthwhile components of complimentary development interventions.

As part of the next generation of index insurance tools, the contracts for Adi Ha have features that may look different from other existing index insurance solutions, and in many cases reflect different fundamental risk management strategies.

Some of the key features of the Adi Ha design process and contract are below. Some of these features are built upon the strategies of previous projects, while others are unique to this implementation.

Key features of design approach

- **Robustness to data quality in pricing and performance.** The contracts and design process involve new strategies to achieve much more robust design. Because there is **never a single source of flawless information**, we performed a great deal of evaluation and tuning using the wide range of imperfect information available. Beginning with complex contracts, models, and statistical tools the contracts were simplified so that they performed relatively uniformly for the alternate datasets and sources of information. These include farmer and expert interviews, surveys, existing historical data from regional informal raingauges, alternate types of satellite estimates of rainfall, regional official government raingauges, alternate types of satellite estimates of vegetative vigor, formal representative design processes, experimental games, informal raingauges and reporting from farmers. Any hazard for which there was not a sufficient level of agreement across data sources was cut from the contract. The contracts were redesigned to perform well given uncertainty in data, models, and parameters. New statistical tools were prototyped to formally address data quality uncertainty from several sources, explicitly building the additional uncertainty due to data quality into modeling and design optimizations. Funding permitting, these statistical tools will be finalized and packaged for broader use. The findings of these tools were replicated through manual comparisons across datasets for validation, transparency, and increased intuition.

- **Systematic democratic and scientific processes to build farmer knowledge and demands into contracts.** Because farmer input is a critical part of index insurance design, it is important to have processes and methodologies providing the farmers and local stakeholders with a more formal leadership role in product design to increase the level for which their knowledge and demands are reflected.
- **Tools not directly taken from off the shelf**. Instead of taking available off the shelf tools or datasets and directly applying them, we compared tools. We performed validation exercises and developed validation techniques. We engaged leaders in the use of the tools as well as experts with experience in their application in Ethiopia. Our goal was to understand shortcomings and develop ways to address those shortcomings. This process was particularly involved for the remote sensing of rainfall, the remote sensing of vegetation, statistical modeling of rainfall, and developing an understanding of what was known about the response of teff to rainfall.
- **High level of analysis**. Although the end contracts are much simpler than those we have been involved with in the past, they utilize a much broader and more sophisticated set of modeling and analysis tools. In addition, they are the result of an increased systematic utilization of informal farmer and expert knowledge. Many of these tools were built specifically for this project to address known failings of existing methodology. Funding permitting, we will work to complete these tools and make them easier to apply so that they may be used systematically in large scale implementation, largely by local partners.
- **Systematic democratic and scientific processes to build farmer knowledge and demands into contracts.** We have worked with partners to build processes that give the farmers and local stakeholders a more formal leadership role in product design to increase the level for which their knowledge and demands are reflected. Funding permitting, we are working to build these approaches into a toolbox with the goal that large numbers of contracts can be systematically built around local knowledge and demands without requiring unfeasible levels of resources. Not all tools will be relevant for all situations. Part of this process will be to formalize processes to flag the situations for which the more expensive tools and interaction levels must be engaged.

Key contract design goals

- **Fill missing piece of larger suite of climate risk and development activities**. The contracts have been designed to complete, enhance, and leverage a much larger suite of complimentary development interventions and traditional activities. For more information on the complete suite of interventions, please refer to the Oxfam project report.
- **Insurance designed to target climate risk driven poverty traps**. The insurance is not built with the strategy of income smoothing per se, but instead with the goal of **unlocking development activities** through strategic risk reductions, such as the loss of access to inputs or civil penalties due to an inability to repay loans.
- **Simpler but smarter**. The contracts are extremely simplified distillations capturing the performance of a wide range of more complex contracts, crop models, and analysis tools.

The simplifications are made for the sake of transparency, robustness to data and farmer livelihood uncertainty, and to dramatically reduce insurance loading costs in a data-poor environment.

- **The contracts are simple proxies of clear hazards--The teff contract in Adi Ha is designed simply to provide a major payout when the rainfall season ends early.** This is a departure from the complex contracts and directly modeled indexes that we have worked on in the past. Instead of attempting to precisely target payouts to subtle hazards, the single and simple hazard of the rainfall season ending early (during teff flowering) was the risk that was identified by farmers as being the most important for teff, and the risk for which we could obtain the most robust and cost effective contract given the information available. This risk was also one important risk to most of the other annual crops in the region, since they flower at this time as well. Note that **an additional proxy addressing the risk of a weak start of the rainfall season is likely also to be valuable for crops other than teff**.

Motivation behind the choice for a simple product in Adi Ha, foregoing the benefits of a complex index. Although contracts that are simple proxies for the occurrence of one or two well defined and simple hazards have the disadvantage reduced coverage from more comprehensive contracts, they provide a number of benefits that we feel make them very well suited for the context and goals of the project.

- The **farmer does not pay for the coverage eliminated** from the contract. The farmer is not required to purchase coverage that she may not want, coverage for risks that she may be able to respond to more effectively through other risk management approaches.
- The farmer can **easily understand exactly what is not covered** so that she can use other risk management tools to address what she remains vulnerable to, or can demand changes to the product.
- The **demands on the limited data are reduced**. For example, the basis risk question is no longer "How different is rainfall between the official amount and what I experience on my farm?" Instead it becomes "Do the years where the rainfall season ends early using the official amount agree with the years that the season ends early for me?" Although the simplified contract provides less precise coverage, the latter questions are more easily answered by farmers or imperfect datasets than the former leading to more certainty in the quality of the coverage.
- The simplified contract leads to a much greater level of **robustness** across data sources.
- The simplified contract leads to a much lower cost due to the **removal of risks for which insurance would have very expensive loading**. Data uncertainty and complex index features dramatically increase loading. This is a challenge to obtaining workable pricing even in data-rich environments. In data poor environments it has made index insurance infeasible. **The simplified contract is one way to substantially reduce insurance loading.**
- **Scalability**. The simplified contract can be **more easily adapted** to capture the simple hazard (eg. if the rainfall ends early) in a wide range of locations.
- **Basis risk is lower for covered risks**. By eliminating risks for which index based solutions might have high basis risk, the index can address the risks that it targets effectively. Because this approach reduces the number of risks covered, the disadvantage is that **basis risk from hazards not included in the contract increases**. Nevertheless, the farmer is not paying for coverage for these risks, and can be more effectively made aware that they are not covered.
- **Simple contracts often offer very similar payouts as complex contracts when tested using decades of historical rainfall data.** At IRI we have helped develop dozens of index insurance contracts. For the vast majority of these contracts, a very simple contract could have provided nearly identical payouts using forty or fifty years of historical rainfall data. Even though the payouts would have been nearly identical, the farmers would have had to pay for much more coverage, had to face much higher loading costs, and may have had the false impression that they were covered for risks that they

would have not received any payouts for in their lifetime. The simple contracts in Adi Ha were designed to offer the actual coverage that farmers would have received through complex contracts, eliminating coverage (and associated costs) that would not have been experienced.

An additional contract strategy goal that we have worked towards is to have **frequent large payouts.** Unfortunately, because of the current uncertainty in climate trends, we have only made partial progress towards this goal. Additional work leading to more precise quantification of trends will be necessary for further progress. Given the promising analysis efforts that are underway, we are optimistic that meaningful progress can be made. Reasons to work towards the objective of frequent large payouts are below.

- **Cost**. When there is a limited amount of data, a substantial portion of insurance loading costs are due to uncertainty in the frequency of rare events. In other words, if there are fifteen years of data, it is difficult to know how frequently the hundred year event will occur. This is particularly challenging due to the uncertainties associated with climate change. Given this uncertainty, insurers must secure expensive additional financing necessary to responsibly honor contract commitments. This expensive component of the premium loading represents risks that may be unlikely to occur in a farmer's life, and risks that insurance may be unable to adequately address. Large, very rare events may be beyond the ability of a low income farmer to finance requiring more substantial interventions.
- **Transparency in coverage**. If a maximum payout occurs very rarely, the farmer may not be aware that the insurance she is paying for will almost always provide much smaller payouts than the maximum coverage.
- **Clear delineation of risk layer**. By eliminating the insurance coverage from common very small payouts and very rare large payouts the role of the insurance contract is clarified. Larger events must be addressed by other mechanisms, most likely government or NGO responses, and smaller events must be covered by the farmers and communities. In Adi Ha, the strategy is that the project would turn to disaster relief programs to address the risk layers more extreme than the one in five year event. The Adi Ha microinsurance is designed to be complimented by the response systems of the Ethiopian government and the World Food Program, including macro-level index insurance in Ethiopia.
- **Transparency in pricing.** We believe it is valuable to clearly communicate the insurance pricing in terms of what the farmer is likely to experience. Large, low frequency payouts can obfuscate pricing. For example, if a major payout occurs with a five year frequency, the percentage cost of the contract will be approximately 1/5th of maximum liability plus loading costs. For that same contract, if there is a maximum payout only once in one hundred years, the percentage price will appear to be on the order of 1/100, even if the farmer paid the same premium and received the same insurance payouts as the five year frequency contract. However, the less frequent contract is likely to have much higher loading costs. Therefore it may appear that the contract with the additional hundred year payout costs only a fraction of the premium of the simpler contract, even if it actually provides less protection for a higher premium.

Initial transactions with farmers

The project has successfully reached the important milepost of actually being purchased by farmers. Although formal monitoring and analysis remains, initial information about the insurance purchases made by farmers in the initial pilot suggests that some hurdles may have been overcome. This was a new product with an unsubsidized premium, complete with reinsurance, transacted in a location for which a met station had not previously existed. The teff contract in Adi Ha was a simple contract, designed to provide a major payout only when the rainfall season ends early**.**

According to the information currently available, in the pilot transactions, approximately twenty percent of the farmers in the community purchased insurance in an approximately two day enrollment period. By design, the pilot size was constrained to this level. The index, purchased by very low income farmers, was entirely based on remote sensing of rainfall, used as a transitionary dataset until sufficient data could be gathered for accurate calibration of the newly installed met station.

Some of the farmers purchased their contracts directly using cash while others used equivalent days of labor, through the PSNP program. The insurance was priced at slightly over twenty percent of maximum liability. The transacted price was approximately two times the expected payout. Of the loading, approximately half was due to uncertainty in long term climate trends, uncertainty that is likely to be addressable in the near future. Although the goal of the index was to aid in loan repayment for teff inputs, the choice of the purchase of the insurance was voluntary, and there were no formal insurance requirements to allow access to credit. Even without a requirement to purchase insurance to gain access to credit, the average farmer selected a contract substantially larger than the minimum option.

Much of the success in the Adi Ha pilot is due to the incredibly high level of trust between project partners and farmers, the impressive marketing efforts, and the extremely high level of community involvement in building the product. Of course, those elements are always essential in a consumer product, but they are often overlooked in these kinds of projects. Clearly the strength of the partners on these fronts has been phenomenal, and it is important that they receive adequate resources as projects scale.

The scope of the project thus far has been to explore, build, and test approaches. Although there is a great deal of work left to be done, many improvements to make, and substantial hurdles to be addressed, we are optimistic that this work will be a strong foundation for effective scaling, if efforts are continued.

Of course, these encouraging preliminary enrollment results are in no way conclusive. Given the strong presence of the project partners in the village, Adi Ha may have more substantial take up than other sites. It will be critical to continue close monitoring and evaluation of the current pilot and future scaling to determine what true impacts are.

Given the possibility for large scale implementations due to the capacity of current delivery channels (currently serving millions of low income Ethiopians), and the potential for low cost contract design, there may be considerable potential for scaling. It will be important that growth from pilot be a healthily, but prudently scaling implementation with careful monitoring, and an increasing reliance of Ethiopian expertise as technologies, practices, and capacity are built.

In addition to the broader education, scaling, and monitoring related issues, there are some **extremely pressing issues** that must be addressed in the near future for the Adi Ha implementation. These include finding solutions that will allow for finding effective and transparent **transitioning from indexing on remotely sensed rainfall estimations to the new met stations** and **reductions in the insurance loading cost,** particularly developing effective and robust strategies for quantifying and pricing climate trends. It will also be valuable in future work to **complete the vegetative remote sensing work** with a field visit and associated analysis for the growing season so that the ability of remote sensing to detect crop vigor can be quantified. Additional work that is necessary to address issues for scaling will be discussed throughout this report. Perhaps the most critical action is to **take advantage of the experience of the past year to improve upon the current exploratory contract, design processes, and complimentary parts of the risk management package**. It will be essential to build formal processes for locally led and locally designed product that can adapt each year to address new lessons learned and changing situations.

This report begins with a brief summary of the specifics of the contract that was transacted. It then provides an overview of the contract design process followed by of a set of chapters that explore issues in more technical detail.

Summary of transacted index

The index is designed to proxy an early end to the rainfall season, determined by consensus to be the primary risk for Teff.

If the cumulative rainfall as reported by the ARC satellite rainfall remote sensing product for the months of August and September (11 August - 09 October, European calendar, 5 Nahase – 29 Maskaram, Ethiopian calendar) does not surpass the trigger value of 105mm (using ARC rainfall), the insurance policy will pay out. The direct source for ARC data is from NOAA[2] and IRI has made ARC data and background information available through the IRI data library.[3]

Payout is calculated according to the following formula, decreasing linearly between 105mm to 60mm (no rain):

Payout= (1 – ((Rainfall Sum – Exit) / (Trigger – Exit)))*Max Payout

Payout = (1 – ((Rainfall Sum – 60) / (105-60)))*1000

This formula is in terms of a unitless currency, with a maximum liability scaled to 1000 units so that it can be adjusted for any maximum liability desired.

Historically, in the 14 years of ARC satellite rainfall records we have available for analysis (since this particular satellite was launched), this index would have produced payouts in 1997, 2000, and 2004 in Adi Ha. The year 2000 was identified independently by numerous groups of farmers in Adi Ha as the last major dry cropping season for Teff. The contract has been designed so that there is broad agreement between rainfall simulations, surrounding stations and other satellite rainfall estimates.

Table 1 reports the payouts, rainfall, and comparison with a model of crop water stress presented as a proxy for yield losses.

[2] ftp://ftp.cpc.ncep.noaa.gov/fews/AFR_CLIM/DATA

[3] The ARC data is also directly available through the data library at the URL below: http://iridl.ldeo.columbia.edu/SOURCES/.NOAA/.NCEP/.CPC/.FEWS/.Africa/.DAILY/.ARC/.daily/X/39.089/VALUE/Y/13.736/VALUE/.est_prcp/T/(1%20Jan%202009)(30%20Oct%202009)RANGEEDGES/T+exch+table-+text+text+skipanyNaN+-table+.html. By changing the latitude and longitude in the URL directly you can access the ARC for other sites. This link is set up for Adi Ha.

Table 1 Historical burn payouts for transacted index

Harvest Year	Crop Water Stress Index– a proxy for losses	Satellite estimated Rainfall (in mm over contract window)	Losses (calculated from Crop Water Index - % of max liability)	Payout (% of max liability) 105mm trigger and 60mm exit
1995	0.5691	127.01	4.87	0
1996	0.6071	197.94	0	0
1997	0.5687	85.60	5.35	43.11
1998	0.6626	195.55	0	0
1999	0.6539	177.37	0	0
2000	0.4469	99.17	127.12	12.95
2001	0.5151	125.72	58.91	0
2002	0.5254	113.51	48.59	0
2003	0.6325	188.94	0	0
2004	0.4461	80.92	127.94	53.51
2005	0.6116	118.63	0	0
2006	0.5646	173.71	9.41	0
2007	0.6499	138.88	0	0
2008	0.5789	112.91	0	0
		Correlation (teff health vs. rainfall) 0.7104		Correlation (losses vs. payouts): 0.5388

For this contract, the mean payout is approximately 8% of the maximum liability using the raw historical rainfall data, and 11% when applied to rainfall simulations.

Much of the price of this contract is driven by the observed downward trend in the rainfall data, and the exits of this contract have been reduced to keep pricing workable when the climate trend is included. The trend is not represented in the raw data presented in the table above, although the contract has been modified in response to the trend. More discussion of this issue is in the section Trend Analysis.

To make the contract workable given the climate trend, the exit was reduced from 80mm to 60mm. With the original exit that does not account for the downward trend (80mm), the contract has a full payout approximately nearly one year in twelve. With the exit reduced in response to the trend, the contract would be expected to have a full payout approximately one year out of twenty seven.

The mean payout using either detrended historical data or simulated rainfall is approximately 15% using simulated rainfall[4]. The climate trend is included in the pricing, leading to a premium with loading and fees of approximately 24% of maximum liability.

It is likely that strategies can be brought to bear to attenuate the impacts of this apparent trend on the pricing. Further discussion is in following chapters.

The contract options for the Adi Ha project were chosen to match likely farmer needs and outcomes during the season. The minimum contract option that farmers were allowed to purchase would completely cover their inputs for the teff growing season (400Birr max liability). The maximum contract option would completely cover the average expected loss for the teff growing season (1200Birr max liability). This was calculated as a difference of the average yield in a "good season" and the average yield in a "bad season", per farmer surveys. Three additional contract options were offered as gradiations between these two. The packages are presented in

Table 1. There were options to purchase the contract using cash as well as days of labor and in the case of a payout, payments could potentially be in grain or in cash.

Table 2 Insurance packages offered to farmers

Max Payout	Approximate Premium (Birr)	Premium in days of PSNP Labor
1200	288	24
1000	240	20
800	192	16
600	144	12
400	96	8

[4] It is likely that the difference between simulated and historical rainfall is less for the detrended data because the detrending process changed many non-payout years to payouts. This is often the primary driver for differences between rainfall simulations and historical burn analysis. Fairly quantifying the potential for near payouts in historical data is one of the primary reasons that rainfall simulators are used.

Overview of contract development

Preliminary scoping

Several preparatory, scoping, and data gathering exercises were performed prior to the analysis presented in this report. These included several scoping visits, focus groups, and feasibility study efforts (Teshome et al 2008) designed to determine if there was an index product that might compliment existing activities in the village that was feasible and demanded by the farmers of Adi Ha. In addition, a data quality assessment was performed to determine the potential for development of index insurance in the context of the existing data (Block et al 2008). We direct the reader to these references for a detailed discussion of the demands and risks facing Adi Ha but quickly mention key issues here.

Drought risk was identified as the primary source of crop risk in Adi Ha. The benefits and applicability of insurance for the rainfed crops and irrigation infrastructure were investigated, and it was determined that for the best risk reduction, direct investment in the infrastructure was more cost effective than insurance products. These investments are being undertaken as part of the interventions that are complimentary to the index insurance project. For more information on the complete suite of interventions, please refer to the Oxfam project report.

Drought risk is therefore greatly reduced for irrigated crops, leaving those with rainfed crops much more exposed. Several crops and livelihoods were investigated and it was determined that Teff drought risk would be the best risk to target with the initial insurance product.

Teff is a rainfed crop that is part of the portfolio of many farmers in Adi Ha, ranging from those with irrigated land to landless laborers. It is a cereal crop (endemic to Ethiopia) grown both for home consumption and for supplemental income. It is also somewhat more drought resistant than maize, sorghum, and millet, the alternative cereal crops for the region around Adi Ha. Teff is in the process of being listed on the Ethiopian Commodities Exchange (ECX), providing the possibility of formal price risk management.

Teff has a short growing season, and can be sown in the latter part of the rainy season. When there is low early season rainfall, farmers frequently shift toward sowing more teff as a last resort. However, teff is labor intensive and requires expensive inputs that often must be purchased, often through loans. If the rains early in the season are bad, farmers have many risk management options. However, if a substantial portion of their labor, expenses, and debt is invested in teff, they are highly at risk if the rainfall season ends early, in the flowering growth phase of the crop (near the middle of its growing season). In this case, their investments and labor and loan repayment capability are all threatened. For this reason, teff was chosen for the Adi Ha index insurance pilot – as a missing piece in an existing drought risk management strategy. By removing some of the key risks in using teff as part of their risk management strategy and by reducing loan default threats, the insurance is intended to allow the farmers to make more productive investments.

The data quality study found that the available datasets for Adi Ha had a substantial level of disagreement for early season risks, leading to concern that it would be challenging to validate the data for early season risks. However, the set of nearby official met stations and remote sensing were highly correlated for risks later in the season, once adjusted for slight differences in the timing of the season.

Given that the end of the rainy season was identified as the key risk for Teff, and that it was the hazard for which the rainfall data was most robust, a prototype contract was developed for the purpose of design discussions with farmers, partners, and experts in Ethiopia.

Prototype contract

In our initial explorations of potential indexes for a drought insurance contract for Adi Ha teff growers, we worked with the multi-phase contract structure described in Osgood et al. 2007. This structure is intended to mimic crop phenology to address the series of water deficit stresses that may be experienced by a crop over the growing season. In this structure, the contract calendar is broken up in to a number of phases of several dekads (ten day periods) each. Payouts are calculated using simple piecewise linear formulas of the sum of capped dekadal rainfall occurring over the phase.

A "sowing window" is set for the contract with a start dekad and an end dekad. The contract calendar begins in the first dekad of the sowing window for which rainfall exceeds a threshold amount, the "sowing trigger." If the trigger is not exceeded during the window, a failed sowing condition is signaled, a failed sowing payment is paid, and the contract is terminated. If the sowing trigger is reached, the contract calendar begins with the dekad in which the trigger was reached. Although this sowing feature does allow the contract to more literally map to crop production, it can be quite unstable, as it can yield dramatically different outcomes if rainfall is one mm above or below the sowing threshold, reflecting and amplifying volatility to rainfall not reflected by the actual crop. With imprecise measurement of rainfall, or small differences in rainfall between nearby farms, this instability can become problematic. Analysis tools designed to validate contracts often perform counterintuitively when applied to this sowing feature, leading to problems in analysis and pricing (Giannini et al. 2009). Because of these factors, the sowing rule we have used in previous contracts has led to a great deal of additional loading on contracts. It was therefore important for us to pay particular attention to assure the potential benefits of this feature outweigh the increases in loading it can cause.

The payout function for each phase has three parameters, a trigger, an exit, and a maximum payout. If the capped rainfall total during a particular phase is more than the trigger, no payout occurs for that phase. If the rainfall total is less than the exit, the maximum payout is rewarded. Crop water stress modeling is used to proxy losses due to drought risk. We adapted parameters from LEAP for our WRSI model (described in detail in Osgood et al. 2007).

Numerical algorithms are used to optimize the triggers to provide the greatest reduction in variance for a given price constraint. In essence, the optimizer adjusts the weighting of the insurance between the different phases to obtain the best risk reduction with respect to the WRSI modeled losses. We have also utilized information from Ethiopian experts at Mekele University on Teff (Kessahun et al 2009 and Mengistu 2009).

We developed a series of optimized three phase contracts for the Adi Ha teff crop. These initial contracts were designed to mimic the growth phases as outlined in our simple crop models. As a result of the optimization, payouts for the three phase contract occurred in those years where the rainy season had begun late, leading to delayed sowing, or finished early. In all cases case, payouts occurred because of the rains ending before the crop had reached the more drought sensitive phenological phases towards the end of the crop's growth cycle.

Figure 1 Farmers preparing land in Adi Ha at the beginning of the season

To directly target this risk, we created a simplified contract, consisting of two phases focusing on the beginning and end of the growing season, identified using our initial simple crop model as the month of June[5] and the end of August through the beginning of September. The goal of this contract was not to literally replicate the full set of risks represented in the crop model, but to trigger payouts based on the key risks identified through a very simple proxy that would avoid some of the cost, transparency, and robustness problems in either the more complex, literal contract, or an index written directly on the WRSI output.

The goal of the two phase contract was to trigger a major payout when the rains ended before the crop had reached the more drought sensitive phenological phases near the end of its growth cycle, either due to the rains beginning late, or due to the rains beginning early. The phase timing, triggers, and exits for the early and late phase were set to best proxy season onset and cessation respectively, and they were optimized to provide the most similar coverage to the three phase contract as well as the WRSI based loss simulations.

Table 3 Table 4, and Table 5 present an illustration of how closely a single phase can proxy payouts from a three phase contract that has been optimized to shift coverage to the phases that best reduce risk. The single phase contract is the index that was submitted for pricing to the reinsurer, with a trigger of 105 mm and an exit of 80mm. The three phase contract is one of the options initially investigated, although it has been calibrated so that its pricing is similar to the one phase contract for easy comparison of performance. Note that the simple pricing algorithms behind the prices presented for the contracts do not include the additional loading that would be necessary for the more complex contract.

[5] As we validated our crop phenology parameters with additional field and expert information, this was later updated to early July. This updated timing and associated modeling did not meaningfully impact the trade off between single and multiple phase contracts.

Table 3 Parameters of three phase contract for Adi Ha teff

Phase	Phase Length	Trigger	Exit
1	1-4	10	10
2	5-7	65	60
3	8-10	95	70
Fixed Contract Start Dekad: June 1-10			

Table 4 Historical burn payouts of three phase contract (ARC rainfall estimates)

Year	**3 Phase Contract**			**1 Phase Contract**
	Phase 1	Phase 2	Phase 3	
1995	0	0	0	0
1996	0	0	0	0
1997	0	0	759.98	775.96
1998	0	0	0	0
1999	0	0	0	0
2000	0	0	93.4	233.03
2001	0	0	0	0
2002	0	0	0	0
2003	0	0	0	0
2004	0	0	942.57	963.22
2005	0	0	0	0
2006	0	0	0	0
2007	0	0	0	0
2008	0	0	0	0

Table 5 Summary statistics for three phase and 1 Phase contract

	3 Phase Contract	1 Phase Contract
Correlation	0.52	0.52
Payout rate	0.21	0.21
Pseudo price	0.17	0.18
Payouts in worst ¼	33.33	33.33

Two prototype contracts were developed for field discussions with farmers, partners and experts, one with a single phase at the end of the rainfall season, and one with an additional phase during sowing to proxy a delay to planting due to a late start of the season.

The simplicity of this contract structure makes it easier to communicate, more robust to data issues and more practical for pricing and implementation than the other more complex alternatives we have researched. This more streamlined structure also increases the potential for scalability to surrounding areas with similar climate and risk profiles.

Contract development and farmer involvement

Following the development of the prototype contracts, an intense set of focus groups and interviews was performed in Adi Ha, discussing the contract with project partners, local experts, and different groups of farmers. These included discussions with a group of farmers with irrigated land, a focus group discussion with female headed households, a focus group discussion with landless laborers, local extension experts, and with village elders and local leadership.

Figure 2 Contract design discussions with project partners and farmers

These meetings included discussions of the risks that farmers face in the context of their livelihood strategies. Each group independently expressed the need for insurance coverage for the rainfall season ending early for teff, primarily so that they could repay loans for teff inputs. Basis risk issues were discussed, as well as their comfort with remote sensing. One concern farmers expressed about remote sensing was that they were unsure if the detectors would remain operational into the future.

With each group we discussed the timing of the phenology of their crops, as well as cataloging the years for which there was drought stress, and what part of the season the stress had occurred in those years. Independently, each group expressed a great deal of interest in the end of the rainfall season and was not interested in coverage for a lack of rainfall for other parts of the season. Farmers explained that they sow teff based on the calendar, during what was already the middle of the rainy season, so rainfall onset was not of interest in the Teff index. They provided us with consistent updated timing for their crop development.

Although Adi Ha did not have an official automated met station onsite, it did have a manual rain gauge (Figure 3) for approximately seven years for operation of the Adi Ha nursery. This data has been utilized in the design of the product. Because this gauge was used for nursery management, metadata flagging missing data was not in the dataset, so it was utilized mostly for diagnostic purposes and testing of rainfall simulators.


Figure 3 Manual Raingauge at Adi Ha Nursery

Farmers expressed concern that the rainfall at their farm might be different from that at an official station, indicating a good understanding of basis risk. Given the contract strategy of proxying the end of the rainfall season, basis risk was also discussed in terms of if the rainfall season ended early in the same years at the site for the official index and particular farms, and farmers expressed good agreement for their farms using that metric.

Discussions of spatial basis risk included the efforts we are engaged in of monitoring of spatial rainfall patterns using remote sensing, both of rainfall and vegetation as well as the potential for the farmers themselves to measure rainfall across the region to gain an understanding of how severe the problem might be.

Following these discussions, the farmers elected a local index insurance design team, with representatives from each of the subgroups and along with the partners we formulated the design and implementation plan. This plan included using the upcoming (2008) Teff season as a 'dry run' to gain experience with how the proposed contract might perform.

Figure 4 Election of Adi Ha insurance design team

Based on the updated phenological timing, calendar-based sowing rule, and preference for late season risk coverage, we updated the WRSI model used and developed an improved index using a single phase to proxy the damage from early cessation of the rainy season.

In addition, following the initial meetings with farmers, in follow up visits, a group of approximately 20 volunteer observer farms were provided with manual raingauges in August 2008, their farms were georeferenced, and they received intensive training on how to measure rainfall. This dataset has been recorded and utilized as part of the index validation process.

Figure 5 Adi Ha farmers with manual raingauges

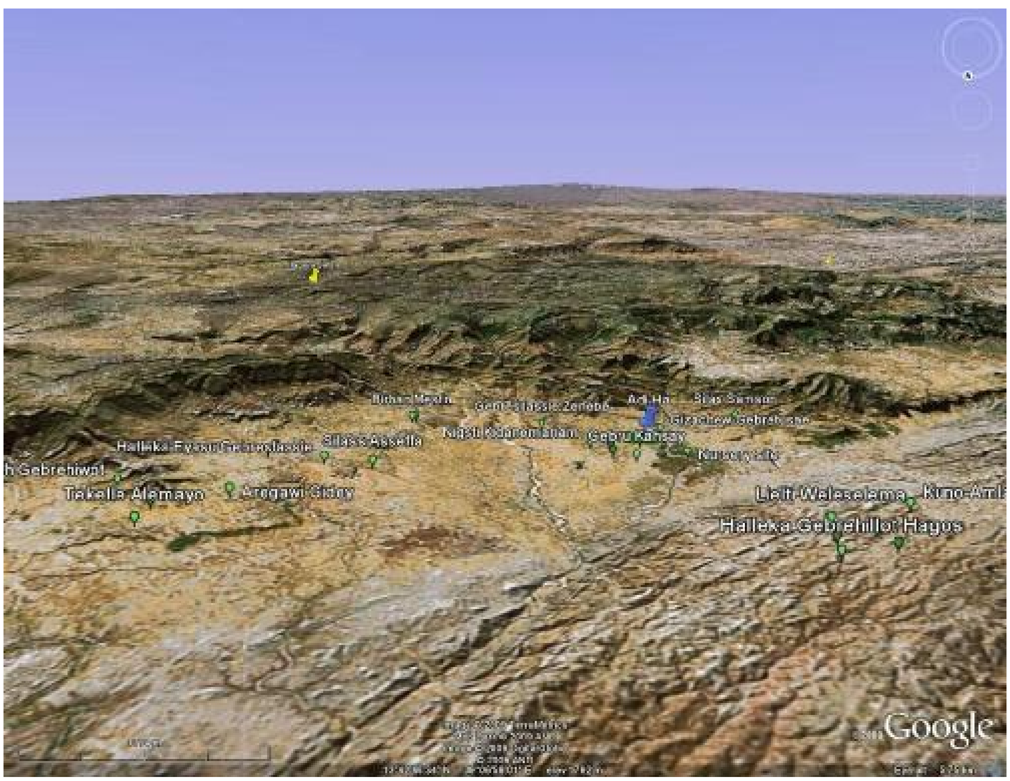

Figure 6 Locations of farmer operated manual raingauges

In addition, a new automated met station was installed, in cooperation with the Ethiopian National Met Service, who are responsible for data collection.

Figure 7 Automated met station installation

Robust contract design

In coordination with the ENMA, we compiled an array of nearby 10-20 year datasets (with intermittent missing years) to assist in defining the general climate history of Tigray. However, for Adi Ha itself, the best option was determined to be satellite rainfall estimates (varying in length from 5 to 15 years). The fully automated rainfall observing station was installed in Adi Ha on August 28, 2008 to assist in characterizing rainfall with very high precision – suitable for insurance purposes.

With the feedback and improved information another round of more methodical contract design was undertaken. The WRSI models were updated based on the revised sowing rules and phenological calendars, and the contract was set to a single phase proxy to detect an early end to the rainfall season.

Using the rainfall simulator described in Rainfall Simulators, tens of thousands of years of synthetic rainfall was generated based on the ARC rainfall dataset for Adi Ha. These realizations are intended not only to reflect the set of possible rainfall events that could likely occur in Adi Ha, but the simulator also includes years reflecting variation due to the uncertainty in the characteristics of the rainfall due to the relatively short dataset. WRSI based loss estimates were generated for each of these simulations and the single phase contract was optimized using a new, more comprehensive optimization engine.

The four parameters of the new optimizer, the 'window tuner' are the upper trigger (rainfall mm), the lower trigger (rainfall mm), the sowing dekad, and the width of the sowing window (in dekads). Given a certain contract, the objective function, for which a minimum value is sought, is the variance of the difference between the payouts vector and the losses vector. The optimization problem is bound with a set of design constraints, with bounds for payout frequency and insurance price. The window tuner employs two consecutive Nelder-Mead search algorithms. Starting from a random initial guess of parameter values, the first algorithm narrows the parameter search space by varying the values of parameters at relatively larger intervals, whereas the second one uses the resulting parameter set from the first one as an initial guess and fine-tunes it by changing the parameter values at smaller intervals.

We are in the process of developing and testing more sophisticated daily simulators that can formally map information from nearby met stations, remote sensing of vegetation, and sea surface temperatures into a set of improved realizations that reflect information from these longer data series (Robertson et al, 2008). We have experimented with an array of alternate simulation techniques in Adi Ha. Funding permitting, we will be able to finalize and apply these simulators for more robust design in the near future. These simulators should also allow for design and pricing that reflects seasonal forecasts and be the backbone of a systematic transition from the ARC data to the new met station.

In parallel, a contract was manually tuned on the ARC data with the objectives and constraints programmed into the computer optimization. Both processes arrived at very similar contracts, illustrating it as a solution robust to design technology, validated by formal optimization, but also arrived at through human intuition.

The index resulting from this design process up to this point had a trigger of 105mm and an exit of 80mm, with a cap of 60mm for each dekad. The index was designed to have major payouts approximately once every five years and a maximum payout approximately once every ten to fifteen years. The index is on the cumulative rainfall for the months of August and September (11 August - 9 October, European calendar, 5 Nahase – 29 Maskaram, Ethiopian calendar) does not surpass the trigger value of 105mm (using ARC rainfall), the insurance policy will pay out.

Historically, in the 14 years of ARC satellite rainfall records we have available for analysis (since this particular satellite was launched), this index (with a 105mm trigger) produces payouts in 1997, 2000, and 2004 in Adi Ha. The year 2000 was identified by numerous groups of farmers in Adi Ha as the last major dry cropping season. Our confidence in these payout years is increased by the broad agreement we see between rainfall simulations, consistency with surrounding stations and other satellite rainfall estimates, and local collective memory.

Table 6 Historical Burn Payout table for feedback-updated contract

Harvest Year	Crop Water Index (simple "teff health" model) – a proxy for actual losses	Satellite estimated Rainfall (in mm over contract window)	Losses (calculated from Crop Water Index - % of max liability)	Payout (% of max liability) – using 105mm as trigger and 80mm as exit
1995	0.5634	127.00	15.34	0
1996	0.6004	193.12	0	0
1997	0.5744	85.60	4.31	77.59
1998	0.6567	191.23	0	0
1999	0.6463	166.39	0	0
2000	0.4524	99.17	126.30	23.30
2001	0.5176	125.72	61.11	0
2002	0.5305	113.50	48.24	0
2003	0.6325	186.99	0	0
2004	0.4504	80.91	128.31	96.32
2005	0.6049	118.62	0	0
2006	0.5635	173.70	15.25	0
2007	0.6429	143.92	0	0
2008	0.5830	117.18	0	0
18.875% approximate premium price		Correlation (teff health vs. rainfall) 0.7058		Correlation (losses vs. payouts): 0.5275

The formula utilized to provide a rough estimate for the approximate magnitude of the pricing, we use the following simple historical burn pricing formula (this is a very simple price approximation only for the sake of illustration):
Premium = Max Liability x Percentage Price
Percentage Historical Burn Price = (Total Historical Payouts/ # of years) + (6% x Max Payout)

Assuming the above payout rate and magnitudes, the historical burn price is 18.875%.
By raising the trigger to 115mm, the burn price rises to 21.238% (2002 is included as an additional payout year in that scenario).

Using output from our rainfall simulators (which take into account uncertainty associated with short datasets), the price for each contract is a few percent higher (19.468% and 22.989%). The average payout using the simulated rainfall for the 105mm trigger is 16%, so it is worth noting that the pricing algorithm we are using leads to a loading of only about 20% above average payouts. This loading is likely to be an underestimate as it is common to have loading of average payouts that range from 50% to 100% of average payouts as well as processing fees of a few percent.

In comparing station data from surrounding locations and satellite rainfall estimates of those locations, we have found that there is a broad agreement on the driest years over the contract window in the time frame for which there is overlapping data (1995-present). Of the five surrounding stations we compared, all of them have in common two to four of the five driest years in the Adi Ha satellite record. Hagere Selam, the one exception to this, is located at a significantly higher elevation which explains the difference in climate. 1997, which would be a large payout year in Adi Ha, is the driest or second driest year over the contract window for every station dataset and satellite estimate (again, with the exception of Hagere Selam).

We have done preliminary analysis on contracts utilizing dry day counts or dry spells, as these might serve as effective proxies for early cessation of the rainy season. In addition, since remote sensing of rainfall has more skill in detecting if it rained or not than the precise amount of rainfall that fell, there is potential for increased robustness between satellite and met stations, as well as less challenges in calibration of new stations. However, our results have been inconclusive and more work is necessary before this strategy can provide useful results.

Figure 8 Locations of nearby met stations

Figure 8 illustrates the locations of the stations. A companion image (Figure 27) includes an overlay of the ARC remote sensing pixel over Adi Ha.

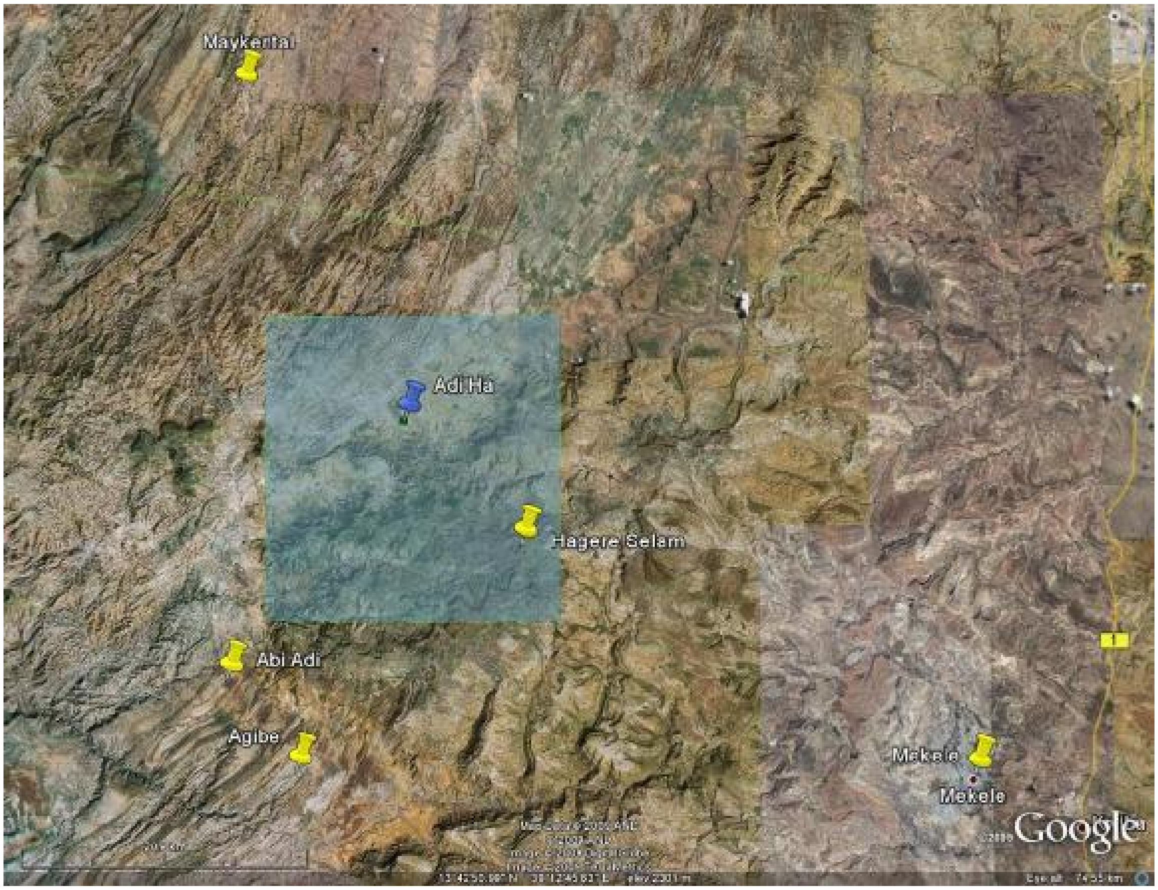

Figure 9 Adi Ha region with ARC pixel

We have also done a preliminary validation of the payout years using NDVI (Normalized Differential Vegetation Index) data which is a remotely sensed measure of biomass greenness. It is used as a means of assessing drought independent of rainfall. The NDVI dataset is also long (1981-present) in comparison to the local rainfall data available. For the Adi Ha example, NDVI for both October and November agree very closely with surface station and satellite estimates of rainfall over the contract window, when picking the driest years over their shared histories (especially 1997, 2002, 2004). Figure 10 includes the NDVI pixel used. Given the long data series, this may be of value in estimating historical payout frequency. We are very actively working on validation of the NDVI datasets as well as improved indices using more advanced satellites and techniques. Funding permitting, our goals are to validate the product in developing payout probabilities and also to map the boundary around a met station for which basis risk becomes too large. See the section on remote sensing of vegetation for more information.

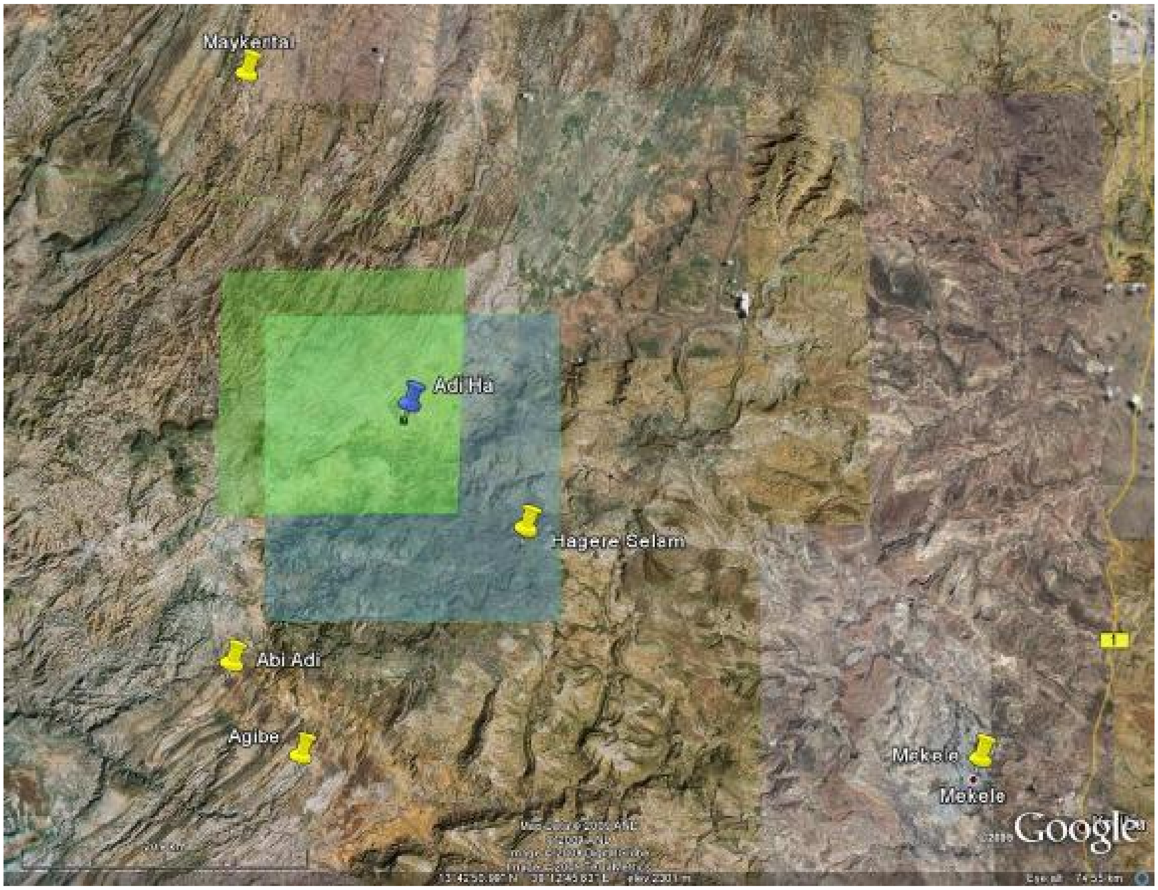

Figure 10 Nearby met stations with example AVHRR-NDVI (Green), ARC (Blue), MODIS (Small, Dark Greenunder Adi Ha 'pin') pixel

Michael Carter and Conner Mullally, economists from The University of California at Davis developed a set of analysis notes (Conner Mullally and Michael Carter, 2009; Michael Carter, 2009) exploring the proposed contracts using an existing household survey dataset, the Ethiopia Rural Household Survey (EHRS)[6]. This analysis was provided by those authors independently from the IRI project. We have attempted to utilize these notes as well as their ongoing feedback to arrive at an improved product. For a discussion of their analysis, please refer to the section entitled Contract analysis using available household survey datasets.

Experimental games with entry and exit surveys

To further validate the contracts, and to allow for a mechanism to build in farmer preferences from a quantitative source, farmers in Adi Ha participated in an experimental games exercise along with an additional associated survey.

[6] This dataset, made available by Stefan Dercon and the Center for the Study of African Economics at the University of Oxford, was collected by IFPRI, the Economics Department at Addis Ababa University, and the Centre for the Study of African Economies. Funding for the data collection was provided by Economic and Social Research Council (ESRC), the Swedish International Development Agency (SIDA) and the United States Agency for International Development (USAID). The datasets, as well as additional information can be found at http://www.economics.ox.ac.uk/members/stefan.dercon/data.htm

As part of the investigation into the attitudes of farmers to index insurance, Nicole Peterson of Columbia University and Conner Mulally of UC Davis conducted a series of games and surveys with people from the Adi Ha woreda from March 2nd to March 8th, 2009. Rest played a major role in the implementation of the effort. Many of the other authors on this report also contributed to the design of these games. The majority of the labor in the design and implementation of these games was funded through leveraged resources and provided at no cost to the project.

We present a quick overview of the exercise below. For a complete presentation of the game methodology and results, please refer to Peterson and Mullally, 2009.

Figure 11 Farmers playing experimental games

A set of questions were administered in two separate surveys in conjunction with the games. Basic demographic and exposure to insurance questions were asked before the games, and questions about yields, understanding of, and interest in insurance followed the games.

In the games, farmers played the actual contract we were considering using the ARC data and our WRSI yield estimates. The farmers played the contracts for real money and with realistic pricing. In addition to a game contract representing the index developed in the previous section, two additional alternatives were explored, to determine if farmers exhibited preferences for variations in the contracts. Figure 12 presents some of the graphics utilized in the game to represent late season rainfall and anticipated yields.

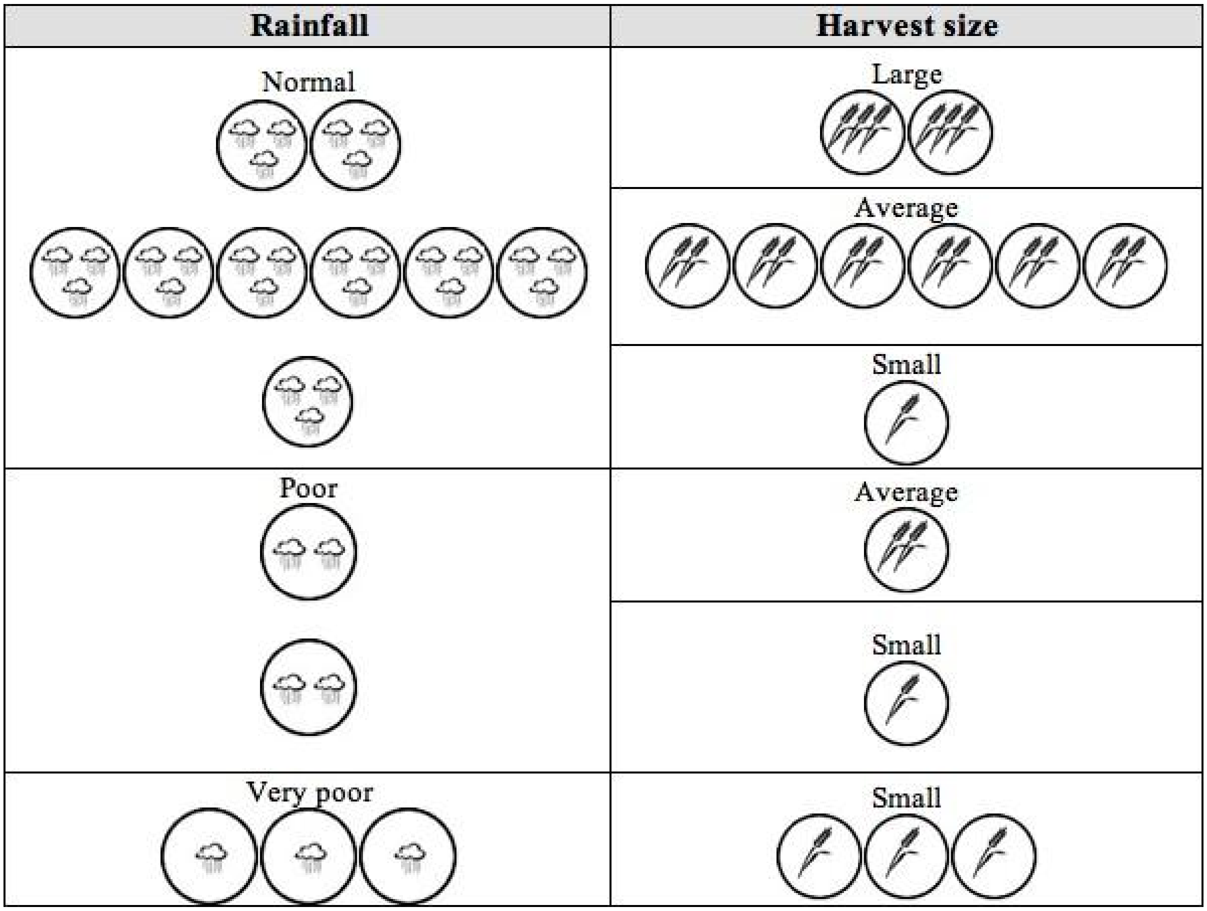

Figure 12 Graphic of rainfall and Yields used in Experimental Games

The benchmark contract is referred to as the yellow contract. The two additional contract variations included the red contract, which had more frequent payouts and the blue contract, which had larger payouts. For each of these options, the premium was increased to reflect the actual cost of the increased coverage. The intent of including the price increases was to determine if the farmers demanded increased coverage by more than the increased cost of the coverage.

In particular, the suite of games were designed to establish the

Interest of farmers in index insurance

a. Estimate farmer interest in index insurance through game play which simulates a number of years, and have monetary outcomes dependent on decisions during the games
b. Estimate farmer interest in index insurance through survey questions that address this directly

Interest of farmers in specific contract characteristics

a. Estimate farmer preferences in terms of frequency of payout and magnitude of contract by allowing farmers to play three different contracts over the course of the game, estimating demand through overall purchase percentages
b. Estimate farmer preferences for these three contracts with a game that allows them to choose among all three contracts at once.

Unfortunately, due to extremely limited time and resources, the game activity was an exploratory pilot exercise, with only a subset of the intended games instruments implemented. Given the incomplete nature of the implementation, some key questions remain to be answered through more complete game implementations. Even these exploratory games provided a great deal of valuable information, and illustrate the potential for the use of games in quantifying farmer preferences and providing education on contracts through transactions of the actual contracts involving real money.

Out of 1783 households, 240 were randomly chosen for a representative sample. Of these, 30% had women as the head of household, which is the same percentage of the overall sample. Table 7 presents a summary of the socioeconomic data of the game participants.

Table 7 Game Socioeconomic data summary

	Mean	Std. Deviation
Gender (% female)	29.17	
Age	43.46	16.713
Years of school	2.72	11.626
Years making farming decisions	22.79	17.573
Timad of rainfed land owned	3.30	11.214
Timad of rainfed land worked in 2007	3.04	1.632
Timad of irrigated land owned	.44	.433
Timad of irrigated land worked in 2007	.59	.682
Timad of teff grown in 2008	2.36	1.282
Timad of teff grown in 2007	2.30	1.665
Number of oxen owned	1.18	.895
Number of adults in family	2.10	1.145
Number of dependents in family	2.56	1.702
Reported yield in normal year (quintal of teff)	2.07	.850
Percent that took loans in 2008	.3417	.47526
Percent that took loans in 2007	.2301	.42180

As is common with these types of exercises, there was a great deal of experimentation in early rounds. Over the course of the game, the adoption choice approached an approximately fifty percent take up rate. Most interestingly, there was very little distinction between the packages, perhaps indicating that the farmers did not understand the distinctions between the packages. Note that the take up rates are not consistent with farmers playing the game totally randomly, so there was evidence of some decisionmaking involved.

If farmers did understand the different packages, the results suggest the fascinating conclusion that the **farmers were indifferent between the benchmark, the more frequent payouts, and the larger payout contracts if they had to pay for any of the increases in coverage**. That is, although they preferred more frequent payouts and larger payouts, the value that they placed on the payouts was the same as the additional cost of delivering the payouts.

Because of the very limited resources, we had to include a very limited set of options. Given the game specifications applied, we do not have game based evidence directly testing if the farmers understood the games. In hindsight, we would have liked to include contract options that were substantially higher cost for the protection offered than all of the contracts offered, both as a method to understand the impacts of loading and as a diagnostic to validate if the farmers understood the differences between the contract options.

Figure 13 Adi Ha Farmer filling out worksheet for experimental game

Some information about farmer comprehension was obtained through exit survey questions. The results of the exit survey indicate that the majority of farmers understood the basic concepts of drought insurance and the game. Some responses are summarized in

Table 8.

Table 8 Understanding the experiment.

1. In the activities today, you had the option to buy insurance. What needed to happen for you to receive an insurance payout?	[8.2%] Crop loss in your field [69.1%] Poor or very poor rainfall at the rainfall gauge [19.8%] Poor or very poor rainfall at your field [2.5%] Don't know/no answer
2. In the activities today, if you had a small harvest, but rainfall was normal, did you receive an insurance payout?	Yes [15.2%] No [80.7%] Don't know/no answer [3.7%]

The vast majority of farmers answered at least one of these questions correctly, and 58 percent gave the right answer for both. Farmers clearly understand that drought insurance is not based upon losses on individual farms. Furthermore, this high level of comprehension is not driven by wealthier and better educated farmers.

By the end of the experiments, the majority of participants understood that drought insurance was not based on compensation for losses on individual farms, and that payouts would be based on rain collected at the rain gauge in Adi Ha tabia. Farmers were asked if they would actually purchase index insurance, leading to estimates with extremely high take up rates but a very low willingness to pay. Because there are rational reasons to answer these questions strategically, it is challenging to use these results directly in estimating take up and willingness to pay. Nevertheless, they do provide valuable insights into the perceptions and negotiating positions of farmers.

Results from the questions on actual rainfall and yields were inconclusive, with comparisons to historical rainfall and focus group data suggesting that the survey questions needed to be more precisely targeted, for example to identify the season during which droughts occurred and which crops the droughts impacted. It is therefore likely that this part of the survey will need to be made more specific to provide useful information for implementation. In addition, strategies that had been eliminated due to insufficient resources are likely to be necessary if this part of the exercise is to provide accurate results.

Calibration Analysis for New Met Station

Following the contract refinement and robustness analysis, we performed a set of exercises to understand the calibration task for contracts designed using ARC data to be indexed on the newly installed met station. Since the statistical challenges are sophisticated, in this initial work we examined initial diagnostics to determine if there was evidence that the short data series from the new met station was closely linked to the ARC measurements. Given the concerns raised by the insurers, additional, more sophisticated analysis will be necessary to finish addressing this challenge.

In our investigation of the index robustness and the connection between ARC and the newly installed rainfall station, we utilize two primary datasets. The first is the remotely sensed ARC daily rainfall estimates used as the benchmark for our contract design. The second is the diagnostic experimental dataset collected by the farmers of Adi Ha. Twenty one farmers began recording daily rainfall at their farm on July 7, 2008. These farmers were trained by Ethiopian met service staff. In addition, the Oxfam project maintains 2 project run gauges in the immediate proximity of the new automated station. Although the data obtained is unofficial, it is of some use in getting an approximate sense for the level of variation of rainfall across the region. We will refer to the farmer data as the Pilot Participant (PP) observations.

Since the contract is designed to average out 'weather noise' by summing over a long period of time, we compared datasets for the time period over the contract window in 2008 to get a sense of how the remotely sensed ARC rainfall estimates for Adi Ha agree with the other measures of rainfall for the area. We also compared datasets for the time period since the installation of the automated rain gauge at Adi Ha, to get a sense of the relationship between the automated gauge at Adi Ha and the other sources of rainfall data. We then used a longer time frame of 1995-2008 to see how ARC remotely sensed data for Adi Ha relates to data from other satellite products. Previous studies of the available rainfall remote sensing products (Dinku et al) has shown that RFE (ARC) provide reasonable skill. The best product (CMORPH) has only 5 years of historical record. Note that IRI is involved in a project to develop thirty years of rainfall estimates for Ethiopia using the improved techniques of CMORPH. As soon as this product is finished, we will utilize it in scaling, design, and analysis. See the section titled Remote sensing for more background information.

To help the reader get a sense of the general patterns in daily rainfall, an animation of daily rainfall at each of the PP stations as well as the new automated gauge is provided in the supplementary materials[7] Over the contract window in 2008, data from ARC remotely sensed estimates for Adi Ha, stations surrounding Adi Ha, and averaged PP observations all show almost

[7] http://iri.columbia.edu/economics/adhiha_farmer_animation

an identical number of wet days (20 days). Figure 14 illustrates the rainfall totals for the farmer observations over the contract period. The average of PP observations and ARC estimates for Adi Ha for total rainfall during the contract window differ by only two millimeters (117 mm and 119 mm, respectively).

Observations taken at the nursery and training center by well trained staff also show similar rainfall totals for the contract window (109 mm and 116 mm). Of course, individual PP observations vary. The variation can be seen in Figure 14 as well as Figure 15, which is a histogram of the PP rainfall totals.

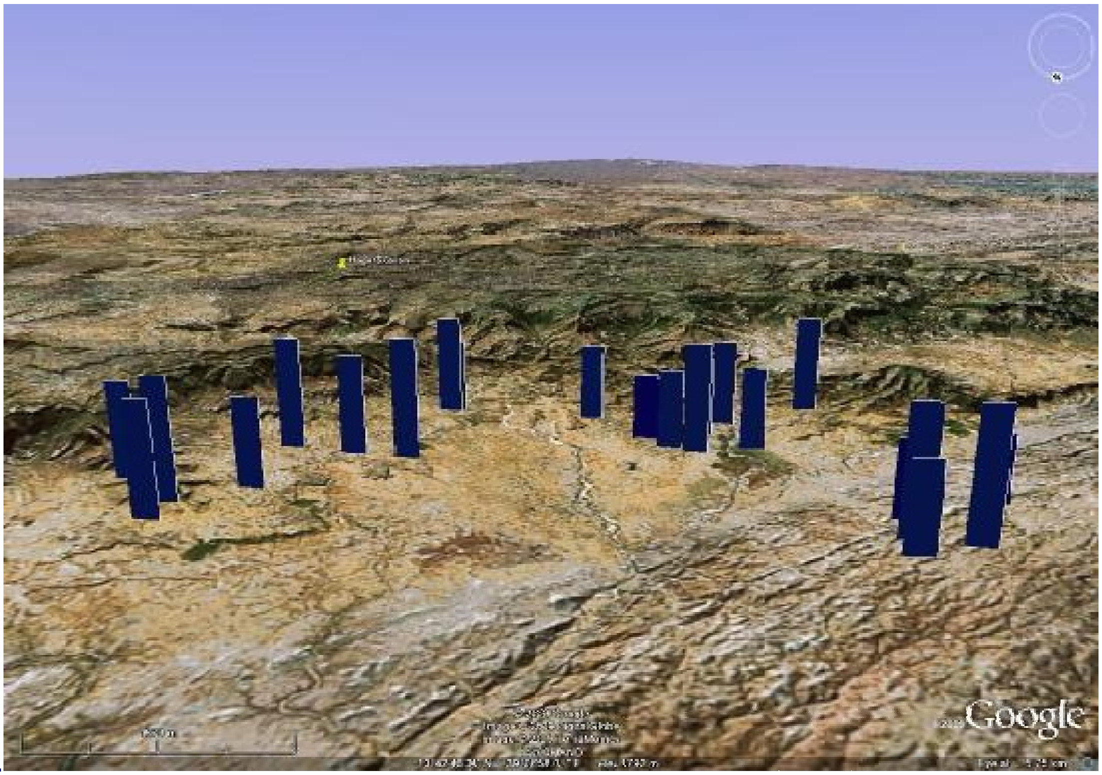

Figure 14 Total Rainfall for farmers over contract window

It is worthwhile to note that rainfall at manual stations may be reported slightly differently than that from automated stations, reflecting strategies to obtain the highest accuracy in total rainfall. Because it is difficult to accurately measure a small of rainfall in a manual gauge, readers are often reading a total that represents more than one accumulated days of rainfall. Therefore, in manual records, a two day rainfall event may be recorded with a first day total of zero and a second day total that reflects the accumulated rainfall for both of the days. Because the automated gauge can accurately measure small amounts of rainfall, it does not aggregate daily rainfall. Therefore the automated rain gauge reads approximately two times as many wet days as the “cup on a stick” station data for Adi Ha. Given this caveat in measurement strategies, observations and estimates from the measures of rainfall around Adi Ha look very similar for 2008.

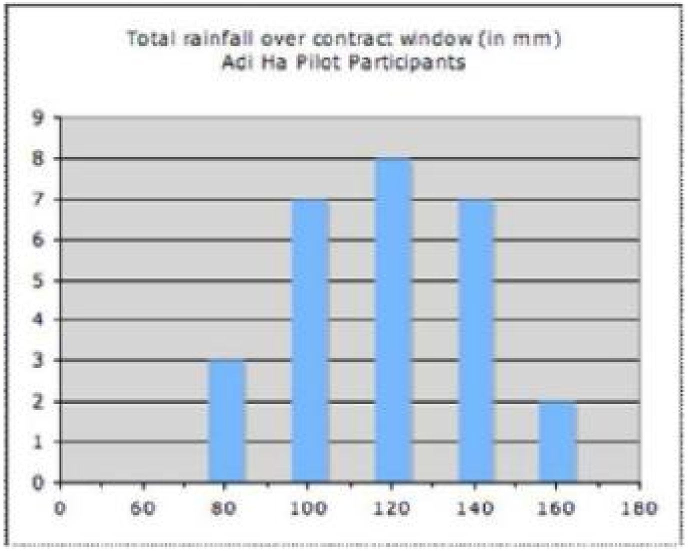

Figure 15 Histogram of farmer rainfall totals

For the period after the automated rain gauge installation, RFE, ARC, averaged PP observations and the automated rain gauge all observed rainfall totals between 80 and 90 mm. ARC rainfall estimates for Adi Ha differ by automated rain gauge observations by only 4 mm. ARC, PP observations, and automated rain gauge also show a similar number of wet days (16, 14, and 13). Note that the automated rain gauge measurements could not be intentionally emulated by the cooperative observers because the automated station data were logged into a storage device and the data was not accessible until the storage device had been transported to the met service and the information downloaded using special software. Figure 16 shows the rainfall totals for the PP stations and the new automated station for the overlap period (the new station is in red, partially hidden, near the center of the image).

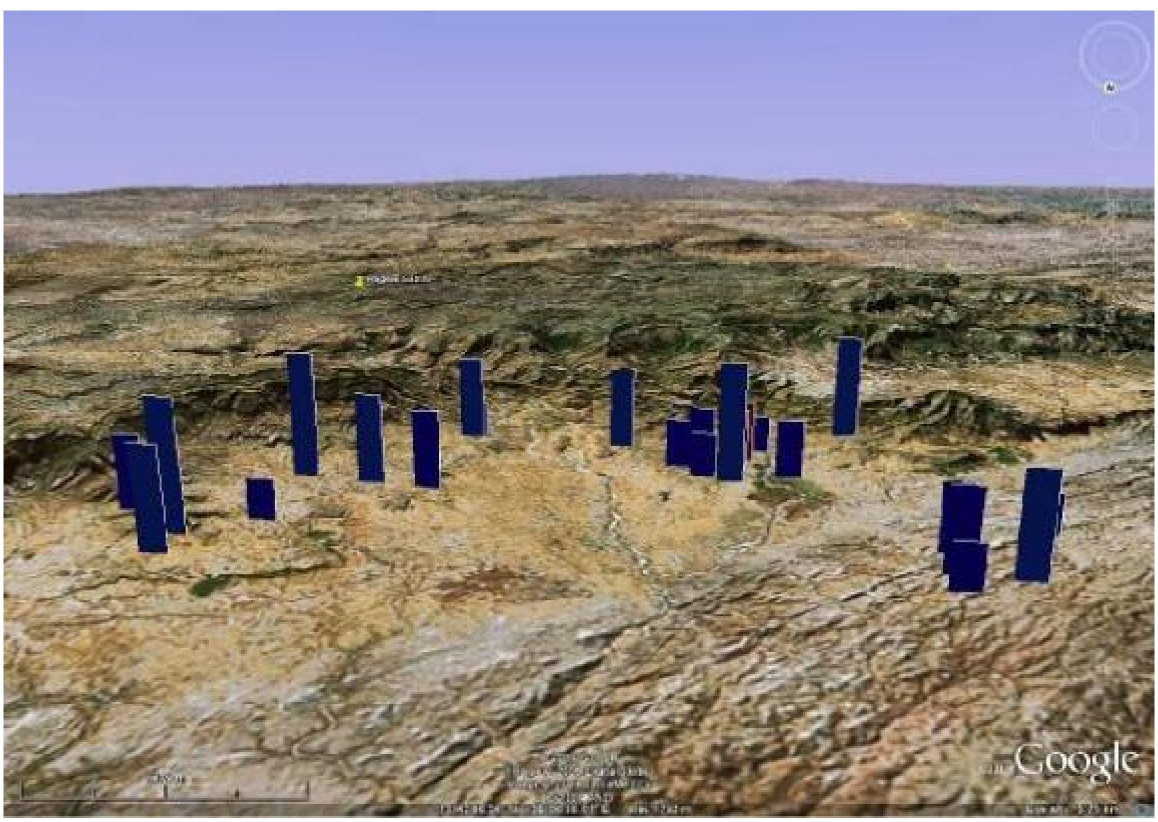

Figure 16 Total rainfall for overlap between farmer and automated station

Using averages over a longer time frame (1995-2008) to study the ARC estimates for Adi Ha, we found surrounding stations and their corresponding ARC estimates show a similar number of wet days over the contract window (26 versus 24). Adi Ha ARC and the Adi Ha “cup on a stick” data both show 27 wet days over the contract window.

Again, RFE and CMORPH overestimate wet days for this longer time period. Adi Ha ARC and RFE estimates were closer to the “cup on a stick” data, with ARC estimating 158 mm, RFE 175 mm, and the “cup on a stick” station showing 168 mm. CMORPH overestimated total rainfall over the contract window for this time period.

Figure 17 Automated met station in Adi Ha Nursery with manual raingauges in background

A comparison between the automatic gauge and the ARC satellite estimates (total rainfall since gauge installation) reveals a discrepancy of only about 4%. When we factor this difference into our pricing formulas (using the actual ARC data as well as the 490-yr rainfall simulation), payout frequency increases by less than 3%, and percentage premium price increases by less than 2%. The automatic station was installed during the contract window in 2008, on August 28, leaving only 39 days of overlap possible with the ARC satellite estimates. During this overlap period, the two recorded somewhat different rainfall totals, leading to an additional uncertainty loading in the final pricing process, which would be passed on to farmers.

To get a sense of the impact of differences of this magnitude, we recalculated our simple pricing and payout frequency for the 105mm trigger contract after scaling both the historical ARC rainfall and the simulated rainfall by that amount, finding that the price for the historical data would increase slightly to 20.2% (from 18.8%) with the payout rate being unchanged at 21.4%. Using the simulated data we see the price increase to 21.3% (from 19.4%) and the pay rate increase from (22.6% to 25.5%).

During 2008, the manual gauge colocated with the automatic gauge recorded slightly lower amounts of rainfall than the average of the 21 farmer gauges, so we consider this site to be a conservative estimate of total rainfall over the Adi Ha region. However, any individual manual rainfall observation may have an underreporting bias as rain gauges are likely not emptied every day, leaving time for accumulated rainfall to evaporate.

Also, we have evidence of this from an observed depressed frequency of rainy days in individual farmer reports - however when this data is aggregated in an average of all farmer reports - the number of rainy days closely matches that observed at the automatic gauge. Therefore, we think the automatic gauge will provide a comparatively consistent, accurate measurement of rainfall in Adi Ha - both in frequency and magnitude, very closely matching the average of all farmer reports. Naturally, conclusive statements cannot be made based on a single season of rainfall.

Discussions and changes with reinsurer

At this point of the design and analysis process, the reinsurer was consulted so that the contract could be further refined to lead to the least amount of loading possible, and adjusted to arrive at a workable price.

We presented a suite of possible contracts with our approximations of historical burn prices somewhere in the neighborhood of fifteen to twenty percent of maximum liability insured and with payouts in the range of approximately one in four to one in five years. The goal was to understand the trade offs in design and insurance loading.

The benchmark contract was the single phase contract presented in Table 6, with the window from August 11-October 9, the 105mm trigger and an 80mm exit. The average payout using historical data for this contract is 16% of maximum liability. Our calculations of premiums produced a price of a little less than twenty percent of maximum liability, or a modest total loading of about twenty five percent of expected payouts, using a pricing formula that reflect only a subpart of the loading components typically charged by insurers. The formula is simplified because it is used not to arrive at official pricing, but instead to characterize risk cost tradeoffs in design. We are always certain to communicate with design partners that final prices will be higher than our estimates.

Several cost saving strategies were suggested by the reinsurer, such as setting a realistic cap on maximum liability to help protect against unnecessarily high cash premiums due to data quality uncertainty, and simplifying the contract structure. Many of these strategies agreed with strategies we felt would reduce loading costs, and were strategies we had been working with in our design process.

There were additional simple modifications to the contract structure that they recommended. For example the initial definition of dekads allowed for 11 day dekads at the end of long months, which kept the dekads aligned with the monthly calendar. The reinsurer suggested that we define all dekads as ten days. This subtle modification was trivial to implement, impacting one day of our window.

Our primary interest was to determine if the benchmark contract would yield workable pricing as well as if more aggressive contracts would yield workable pricing. In addition, we were interested in determining if the first few months of data from the newly installed met station would be sufficient to allow the index to be transacted using data from that station. Our goal was to work interactively with the reinsurer to determine what solutions could be found, both for the current year and for the near future.

Given the calibration analysis for the Adi Ha met station, the reinsurer was interested in performing an additional diagnostic, how similar the rainfall levels were between ARC and the new station for brief period (39 days) that they overlapped in during the contract window. For this specific window, the newly installed station recorded rainfall that was approximately thirty percent less than the ARC estimate. Given the evidence from the other calibration analysis, it is very possible that this was due to random noise in the short time period. There is insufficient data at this point to know if the difference would have averaged out during the complete contract or not. With the more sophisticated statistical techniques we are working on we could address this question more systematically, providing a quantification of the risk of calibration uncertainty that could be used for pricing.

Given the substantial loading that the difference between the ARC and met station contract overlap had, it was determined that the only workable option was to directly index the ARC satellite rainfall in the first year instead of the new met station, in order to have more time to gather a complete dataset of automatic station data, especially over the overlap period in the contract window. It is important to note that although there was a great deal of concern about the use of remote sensing as the data for the contract, it was a decision reflecting the inputs of all partners. In the follow up surveys, approximately ninety percent of the farmers were comfortable with using remote sensing for the contracts this year. Nevertheless, it will be critical for the project to complete analysis so that the transition to the met station to occur as soon as possible.

The primary pricing issue in the contract was the existence of a downward trend in the rainfall data. When the data is detrended to account for this decrease, the average payout and associated loading increase substantially, leading to unworkable prices for any of the contract proposals.

Given 14 years of data, it is very difficult to know if any trend observed is due to a true process, or simply the small number of years observed. The trend was not significant, meaning that the errors bounds on the trend would include possible trends of increasing rainfall as well as decreasing rainfall. However, if the trend does exist, it is important to secure the proper financing to be able to honor payouts, which leads to increased loading costs.

It is important to note that we believe these pricing issues do not reflect unfairness from the part of the reinsurer. The contracts under consideration involve a very small amount of money, through which the reinsurer does not stand to make any noticeable profit. The reinsurer could very easily simply offer a very small token premium for publicity reasons yielding a product that was not realistic or scalable. With this product, there would be no understanding of how or if products could be modified to be workable at large scales. We believe the reinsurer has a very sincere desire to develop a credible, workable, and robust product through responsible pricing.

Given the uncertainty involved, the most cautious approach was taken. Pricing included the trend, but the upper trigger was not adjusted to assure that the important payout years would be included regardless of if the trend existed or not. If there was no trend, it would have paid out in the years of the historical burn analysis. If the trend was real, it would have paid out more frequently. To make the pricing workable, the lower trigger was adjusted at our request by Swiss Re (at the request of IRI) targeting a loaded price near twenty percent of maximum liability.

Through this process, the contract exit was lowered from 80mm to 60mm, leading to the contract presented in the section
Summary of transacted index. In this contract, the cost of the premium is driven by three components. The first component is the no-loading, no-trend cost, which is about half of the total premium. About one quarter of the premium is due to loading with the remaining quarter being

due to the climate trend observed. In addition to these costs, the local insurer charges a few percent of maximum liability for processing costs. Swiss Re also provided prices for the other, more aggressive options, which ended up being prohibitively expensive.

Product packaging

A roundtable of all local partners was convened in Mekele in May to discuss options and finalize which insurance package would be accepted and presented to the farmers for purchase. This involved discussions with the elected design team and farmers about issues such as packaging, delivery, product size, and their level of comfort with satellite based products.

Once the final index parameters were decided upon by the stakeholders, we presented the option to the farmers in Adi Ha. This required a training on satellite remote sensing. A training was also conducted on the scalable distribution mechanism developed by Oxfam with partners. For details of the delivery and distribution arrangements, please refer to the Oxfam project report.

The basis of the distribution/implementation model is the ability of farmers to choose whether to pay their insurance premium in labor or in cash. Also, 5 tiers of insurance were made available for purchase, ranging from covering pre-season inputs at the lowest level, to the average total yield loss during drought years at the highest level. These tiers are presented in Table 1. The selection of the maximum and minimum tiers involved a great deal of negotiation, as there are concerns associated both with offering farmers too much insurance as well as too little insurance.

The contract options for the Adi Ha project were chosen to match likely farmer needs and outcomes during the season. The minimum contract option that farmers were allowed to purchase would completely cover their inputs for the teff growing season (400Birr max liability). The maximum contract option would completely cover the average expected loss for the teff growing season (1200Birr max liability). This was calculated as a difference of the average yield in a "good season" and the average yield in a "bad season", per farmer surveys. Three additional contract options were offered as gradations between these two, to provide some flexibility to the farmer, and to allow us to observe what packages farmers preferred.

Farmers were able to purchase the contracts either through cash or through labor, with a close link to the PSNP program. Purchase of the insurance was entirely voluntary and although the insurance was encouraged by lenders, and lenders played an important role in the distribution, they had no requirement of insurance purchase for access to credit.

For details about the contract packages and distribution arrangements, please refer to the Oxfam project report.

Figure 18 Adi Ha village chief providing input to the index insurance distribution mechanism

Technical sections

Contract analysis using available household survey datasets

Michael Carter and Conner Mullally, economists from the University of California, Davis, developed a set of analysis notes (Conner Mullally and Michael Carter, 2009; Michael Carter, 2009) exploring the proposed contracts using an existing household survey dataset, the Ethiopia Rural Household Survey (EHRS)[8]. This analysis was provided by the authors independently from the IRI project. The discussion below is our interpretation of the notes, and may not represent the opinions of the authors.

The dataset did not directly include the Adi Ha village. The nearest households included in the survey were approximately 50 miles from Adi Ha. Unfortunately, these households were located in a region in which Teff was not a major crop. Therefore, the researchers utilized villages in the survey for which Teff was a major crop. The closest of these villages were approximately 280 miles from Adi Ha, and the furthest was approximately 460 miles. The villages were clustered close to Addis Ababa, with a distance of about 220 miles between the furthest villages. The locations of these villages with respect to Adi Ha are reflected in Figure 19. The survey information for the villages nearby to Adi Ha is anticipated to be of particular value in the next phase of the project because they are located around one of the sites selected for future scale-up.

[8] This dataset, made available by Stefan Dercon and the Center for the Study of African Economics at the University of Oxford, was collected by IFPRI, the Economics Department at Addis Ababa University, and the Centre for the Study of African Economies. Funding for the data collection was provided by Economic and Social Research Council (ESRC), the Swedish International Development Agency (SIDA) and the United States Agency for International Development (USAID). The datasets, as well as additional information can be found at http://www.economics.ox.ac.uk/members/stefan.dercon/data.htm

Figure 19 Location of Adi Ha (Yellow), Villages used to examine Teff in Survey (Green), and farmers in survey closest to Adi Ha (red)

Conner Mullally and Michael Carter, 2009 describes the initial statistical analysis. In order to use the survey data to be able to evaluate the contract, it was necessary to map the contracts being considered for Adi Ha to the survey villages in the south. The researchers did this by selecting the 60 day window for which rainfall reflected the highest correlation with Teff yields. This resulted in a window beginning on October 1, as opposed to the August 11 start of the Adi Ha window.

The household survey included the years 1996, 1998, and 2003. The Adi Ha contract would have had payouts in 1997, 2000, and 2004, years that were not included in the survey. For evaluation of the contract performance, it was necessary to extrapolate to the 14 year time period for which rainfall estimates were available. To do this, the researchers utilized the relationship between rainfall and household teff production to predict out of sample household level Teff yields for the 14 year period. Rainfall explained about half of the variation in yields, the other half of the variation being unexplained. This illustrates both the importance of adequate rainfall to the farmers as well as the importance of alternate risk management strategies when implementing index insurance, such as community risk pooling to address the remaining idiosyncratic risk.

A series of alternate triggers and exits were set, tailored to the rainfall levels of each village. The contracts varied primarily in terms of their payout frequencies, with the most infrequent contract paying out 5 out of the 14 years, the second contract paying out whenever the rainfall was below average, and the third contract paying out in each year. In the extrapolation, all of the contracts reduced the variation in teff-related income. The least frequent contract was effective at targeting the worst years, although as would be expected it provided much less reduction in standard deviation of expected yields (on the order of 5 percent) than the expensive contract that triggers in all years (on the order of about 50%). It is encouraging to us at IRI that the contract with infrequent payouts was effective at targeting the worst years, since the contract design goals for Adi Ha were to provide a contract that would provide a substantial payment in most of the worst years to help farmers be able to repay loans for inputs. Given that the data was extrapolated from years that would not have been payout years in Adi Ha, at IRI we also find it encouraging that even the extrapolated contracts provide substantial risk reduction so that there may be scope for effective scaling from Adi Ha to other regions. The authors of the note conclude that the contracts have strong potential to reduce risk. They also note that although contracts with infrequent payouts were effective at targeting the worst years, it would be important to compliment these contracts with alternate risk management strategies for the more minor events.

Carter, 2009 utilized this analysis to explore insurance certainty equivalent benefits, setting triggers to roughly reflect those from Adi Ha, with the trigger set at the 1 in 4 payout frequency, and the exit set to the driest year in the record, as well as a contract with increased coverage. He provided a table of certainty equivalent comparisons between farmers with the insurance and without the insurance for a series of alternate loadings, insurance levels, and prices, assuming a coefficient of relative risk aversion of 1. He also included a dual strike contract in the comparison. He illustrated that the margins are very tight for the insurance to be able to provide reductions in certainty equivalent income given the risk aversion assumed. He also illustrates that the margins decrease for lower teff prices, and increased levels of coverage, and of course, with insurance loading. In the illustration, the dual trigger contracts do not provide improvements over the single trigger contracts for the same loading level. In past contracts that IRI has transacted, or been informed of, dual trigger contracts have very dramatically increased loading costs. Since Adi Ha has substantially less historical data to characterize less common events, dual trigger contracts are likely to be particularly expensive in that context.

In his simulations, although all of the options provide some improvements with no loading (price equal to the average payout), at 50% loading (a premium price of 1.5 times the average payout size), only one of the packages illustrated yields a certainty equivalent benefit. This highlights the urgency of discovering methods to reduce the loading of the contracts. In addition, it illustrates the danger of a farmer overinsuring, either because the farmer is forced to purchase a package larger than optimal, or because the farmer does not sufficiently understand the insurance to make the optimal decision.

Most importantly, it illustrates that although there may be scope for use of insurance as a purely income smoothing device driven by risk aversion, the potential for such products is extremely limited. Therefore it is likely that the bulk of the value of insurance will be due to production and development constraints, such as access to productive assets or other poverty-trap related issues and insurance is most likely to have its major benefits as the key to unlock climate risk related development constraints. The insurance product has been designed to use insurance to enable alternate activities to increase the average income in non-drought years through productivity increases as opposed to reducing income variance per se.

In Adi Ha, based on the surveys and farmer discussions, the insurance was designed to increase a farmer's ability to repay loans for Teff inputs to prevent the farmers from losing access to credit or criminal penalties after failing to repay a loan. By addressing this constraint, it is hoped that farmers will be able to obtain improved inputs allowing them to have improved production in non-drought years. In addition, since shifting to Teff production is a traditional ex-ante risk response to early season drought[9], the goal of the insurance is to reduce the risk of this risk management strategy due to the rainfall season ending early, which would impact teff as well as most other annual crops. The very tight margins that the illustrations have for utility based benefits from reductions in income variance suggest that for index insurance to have substantial impacts, it is likely that these benefits will need to arise through activities enabled by the insurance as opposed to the insurance as a stand alone product. That is, it is likely that the insurance will not only need

[9] This was identified in the scoping surveys and was validated repeatedly in focus group discussions with farmers and experts.

to reduce variance in income but will also need to improve the livelihoods of farmers to have meaningful impacts. Of course, much work must be done to definitively establish if the insurance is most useful directly for reduction in variance of income or for enabling production related increases in average income, and its role may differ for farmers in different situations. It is therefore strongly recommended that this question be studied empirically in the implementation and scale up.

Rainfall Simulators

When pricing and design analysis is primarily based on historical burn analysis, products and prices are sensitive to particular features of one or two historical events, making it possible to overemphasize the importance of the specifics of these events. However, it is important to remember that most rainfall simulators have limitations in what they can accurately reflect, and unless utilized with caution, may lead to an inaccurate understanding of the performance of a contract. An overview of issues with rainfall simulators in index insurance is provided by Shirley 2008.

In the design process for the Adi Ha contract we worked with both dekadal and daily rainfall simulators. When feasible, daily simulators are preferable. However, because a dekadal simulator has the less challenging task of simulating a ten day sum of rainfall, they have some advantages. They run much more quickly, are relatively accurate for modeling ten day sums, and do not require the manual supervision by a trained operator that high quality daily simulators need. In addition to the simulator presented below, for the Adi Ha project we are exploring the insights offered by a range of daily simulators (Robertson et al 2008). We are exploring rainfall simulators from the DSSAT suite, including modified versions designed to condition on climate processes. We are working with hidden markov model based simulators that attempt to estimate hidden states that drive rainfall processes. We are also working with a hierarchical Bayesian daily rainfall model designed to incorporate information from nearby stations, remote sensed vegetation, and climate driving processes. Finally, we are investigating how well simulation spatial interpolation techniques perform in the context of index insurance. Note that a central challenge in utilizing advanced simulators is in understanding the performance of the simulator in different conditions, developing clear knowledge of when the simulators work and when they do not work.

The dekadal rainfall simulator described below is utilized for the contract analysis presented in this report because it is a robust simulator and can be reliably run unsupervised, is not processor intensive, and can therefore be run in a loop under software control. It has the useful feature that statistical uncertainty due to short data series is represented in the simulations but does not have mechanisms to estimate and utilize links between datasets. However, the simulator as written may produce unrealistically large dekadal sums. We are currently addressing this problem by capping these rare simulated events.

In the future, a simulator with the appropriate capabilities and performance will be necessary to be able to systematically transition from satellite datasets to new raingauges, to identify the best places to locate new raingauges, to interpolate to new locations for scaling, to build seasonal forecasts into pricing and design, and to formalize statistically the robustness testing that we have done manually for the Adi Ha implementation.

The rainfall simulator presented below is available for use as code we have developed in the freely available R statistical package. It also has been implemented in graphical user interface through the online Weather Index Insurance Educational Tool developed by IRI in partnership with the World Bank CRMG. This tool will be externally released in the near future.

Designing a Dekadal Rainfall Simulator

To illustrate the potential benefits and challenges in applying rainfall simulation, we have been working on a simple rainfall simulator. This simulator represents many of the strategies utilized by industry standard products. In this section, we present the model for dekadal rainfall, we fit it and evaluate the fit for two example data sets: Adiha (Ethiopia), and Lilongwe (Malawi), where these data sets consist of 7 and 44 years of observed dekadal data, respectively. The ultimate goal is to use the model to provide realistic dekadal rainfall simulations for any site in the world.

To be useful, we want the model to have the following qualities:

1. It must not take a long time to fit the model to data. Ideally we would be able to simulate a few thousand years of data within 1-2 seconds. This restricts the class of models we can consider.
2. It will be parametric, so that we can simulate values that did not occur in the observed data, which is impossible using certain resampling strategies.
3. It will be sensitive to the amount of data with which it is fit, so that less observed data results in more uncertainty in simulated rainfall. To do this conveniently, we use the Bayesian framework.
4. Lastly, we want the procedure to be automatic, so that as little supervision and data set-up as possible is required. This might be the greatest challenge of all, because different amounts of observed data (5 years vs. 50 years, for example) and different rainfall-generating processes throughout the world may require different types of models.

Here is some general notation, and the basic outline of how we will fit and simulate rainfall from our models. Let Y_{dt} be the total amount of rainfall during dekad *d* of year *t*, for dekads *d* = 1,...,36, and years 1,..,*T*. We will fit a parametric model, $p(q|Y)$, where *Y* denotes the vector of observed values of Y_{dt} for a given site and time span. In Bayesian statistics, the statistical distribution assumed (the prior distribution) is updated using the data to arrive at an updated (posterior) distribution. Given our assumed prior distribution for q (which requires careful construction) we will sample values of q from its posterior distribution, $p(q|Y)$. The posterior sample of *G* values is denoted $(q^{(1)}, q^{(2)}, ..., q^{(G)})$, where *G*=1000 is typical. To simulate new dekadal rainfall observations, we do the following:

1. Draw a value of $q^{(g)}$ from its posterior distribution.
2. Draw a value of $Y_{\text{pred}}^{(g)} | q^{(g)}$ where Y_{pred} denotes a new, simulated realization of *Y*.

Finally, we compare the distribution of Y_{pred} to the observed values of *Y*, to make sure that the simulations are similar in character to the observations. We also check to see that we don't overfit the model by performing a cross-validation exercise. This Bayesian framework is a convenient way to simulate rainfall in a way that reflects not just the natural uncertainty of the rainfall process, but also the uncertainty associated with our parameter estimation.

The Basic Model and Adiha Data

Let Y_{dt} be the total amount of rainfall during dekad *d* of year *t*, and let X_{dt} be the indicator of whether there was any rainfall during the specified dekad and year, such that $X_{dt} = 1\{Y_{dt} > 0\}$.

The model we fit has two components: a model for the frequency of rainfall, and, when rainfall occurs, a model for the intensity of rainfall.

The frequency model is

$$X_{dt} \sim \text{Bernoulli}(p_d),$$

for $t = 1,...,T$ and $d = 1,...,36$, and the intensity model is

$$Y_{dt} \left| (X_{dt} = 1) \sim \text{Gamma}(a_d, b_d), \right.$$

where a_d and b_d are the shape and rate parameters of the gamma distribution with mean (a_d / b_d). The notation $Y_{dt} | (X_{dt} = 1)$ is a reminder that this is the model for intensity of dekadal rainfall conditional on the fact that the rainfall amount is non-zero for that dekad. The full distribution of dekadal rainfall depends on a_d, b_d, and p_d.

The probability density function of a single observation can therefore be written:

$$\text{P}\ (Y_{dt} = y_{dt}) = \begin{cases} 1 - pd \\ p_d \dfrac{b^{a}_{D^d}}{T_{(a_d)}} y_{dt}^{(a_d - 1)} \exp(-b_d y_{dt}) \end{cases}$$

if $y_{dt} = 0$

if $y_{dt} > 0$

Since the occurrence of rainfall in a given dekad is independent across years, the frequency model can be simplified in the following way: Let $Z_d = \sum_{t=1}^{T} X_{dt}$, the number of years out of T in which a non-zero amount of rainfall occurred. Then

$$Z_d \sim \text{Binomial}\ (T, p_d),$$

for $d = 1, \ldots 36$.

Example: Adiha Cup-on-a-stick

The data for the cup-on-a-stick site in Adiha consist of dekadal sums for each of the 36 dekads in the year for the years 2000 - 2007, with the exception of 2004 (whose data were missing); hence for this data, $T = 7$.

Table 9 Annual Rainfall data from manual raingauge at Adi Ha

	2000	2001	2002	2003	2005	2006	2007	mean	sd
1	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA
2	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA
3	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA
4	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA
5	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA
6	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA
7	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA
8	0.00	14.0	0.0	0.0	0.0	0.0	0.0	14.0	NA
9	0.00	2.0	0.0	0.0	7.0	0.0	0.0	4.5	3.5
10	25.00	0.0	0.0	0.0	0.0	20.0	0.0	22.5	3.5
11	0.00	0.0	0.0	2.0	70.0	17.0	0.0	29.7	35.7
12	6.00	0.0	0.0	0.0	22.5	15.0	0.0	14.5	8.3
13	35.00	0.0	0.0	0.0	0.0	47.0	0.0	41.0	8.5
14	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA
15	1.07	0.0	0.0	0.0	0.0	36.0	0.0	18.5	24.7
16	24.30	14.0	2.0	0.0	0.0	0.0	0.0	13.4	11.2
17	0.00	46.5	11.0	16.0	0.0	5.0	19.0	19.5	16.0
18	18.00	41.0	61.5	81.0	33.5	43.0	89.5	52.5	25.9
19	45.00	92.0	60.5	14.5	72.0	46.5	52.0	54.6	24.2
20	56.00	28.0	45.5	29.0	122.5	32.0	104.0	59.6	38.4
21	64.50	107.5	82.0	100.0	57.0	81.0	129.5	88.8	25.3
22	65.00	6.0	58.5	51.0	76.5	67.0	103.0	61.0	29.4
23	57.50	0.0	72.5	81.5	43.5	59.0	94.5	68.1	18.4
24	42.00	75.0	26.0	74.5	101.5	132.0	39.5	70.1	37.7
25	31.00	10.0	40.0	41.0	0.0	0.0	44.0	33.2	13.8
26	29.50	0.0	31.0	10.0	55.0	0.0	0.0	31.4	18.4
27	4.00	0.0	4.0	0.0	27.0	0.0	0.0	11.7	13.3
28	6.00	0.0	3.0	0.0	8.0	0.0	0.0	5.7	2.5
29	17.50	0.0	0.0	0.0	0.0	0.0	0.0	17.5	NA
30	38.00	0.0	0.0	0.0	0.0	0.0	0.0	38.0	NA
31	23.00	0.0	0.0	0.0	0.0	0.0	0.0	23.0	NA
32	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA
33	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA
34	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA
35	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA
36	0.00	0.0	0.0	0.0	0.0	0.0	0.0	NA	NA

Let's look at some graphical summaries of the data. First, Figure 20(a) plots the percentage of years in which the rainfall sum was non-zero for each dekad. Second, Figure 20(b) plots the mean amount of rainfall by dekad, for the dekads in which there was non-zero rainfall in at least one of the seven observed years. Lastly, Figure 20(c) plots the standard deviation of rainfall by dekad, this time for only the dekads in which there was nonzero rainfall in at least two of the seven observed years (so that the observed standard deviation could be calculated). In all three plots, it is clear that there is dependence between dekads, due to seasonality, that the current model doesn't include.

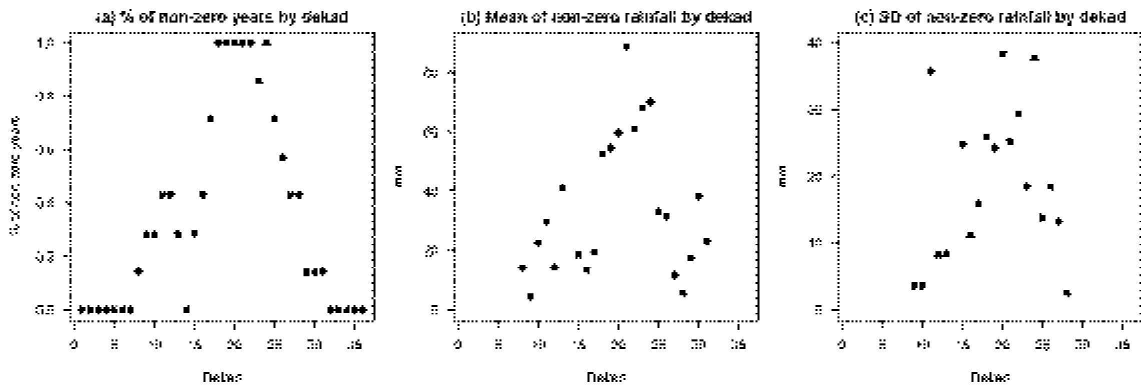

Figure 20 Graphical summaries of Adiha cup-on-a-stick data

A simple version of the model

A very simple model would use non-informative priors for the parameters of the binomial and gamma distributions. Ultimately, we would like it if the same priors provided realistic simulations no matter how large the data is. This is difficult, though.

The Frequency Model

For the frequency model, A beta distribution, p_d ~ Beta(1,1) for *d* = 1, …, *D,* is a conjugate prior for the binomial likelihood, and leads to a beta posterior:

$$p(p_d | Z_d) \sim \text{Beta}(Z_d + 1, T - Z_d + 1).$$

This posterior actually leads to appropriate inferences and simulations for small and large data sets, so we use it to model the frequency of rainfall by dekad without modification.

The Intensity Model

In order to fit and simulate from the intensity model quickly, we seek a conjugate prior for the parameters of the gamma distribution so that the posterior distribution is available in closed form. Unfortunately, the gamma distribution has a convenient conjugate prior only when the shape parameter, a_d, is known. In such a case, if we use a gamma prior for the rate parameter, b_d, then the posterior distribution of b_d is also a gamma distribution. The questions, then, are how do we choose values for a_d for each dekad, and what prior distribution do we use for the rate parameter, b_d?

In our first attempt at fitting the intensity model, we assume a_d is the same for each dekad, and we use the data to estimate an average value. The value of the shape parameter, a, of a gamma distribution determines the relationship between the mean and the standard deviation of the distribution. To be precise, if Y_{dt} ~ Gamma (a_d, b_d), then Sd(Y_{dt}) = Mean (Y_{dt})/$\sqrt{a_d}$. In order to set a constant value of a_d for all *d* =1,...,36, we regress the observed standard deviations of dekadal rainfall against their means, and compute the estimated slope via weighted least squares, where the intercept is fixed at zero. Figure 21 shows this relationship and the weighted least squares regression line for Adiha, which has a slope of 0.44, which equates to a value of a_d = $1/0.44^2$ = 5.1. In other words, on (weighted) average, the standard deviations of dekadal rainfall sums by dekad are about 44% as large as the means.

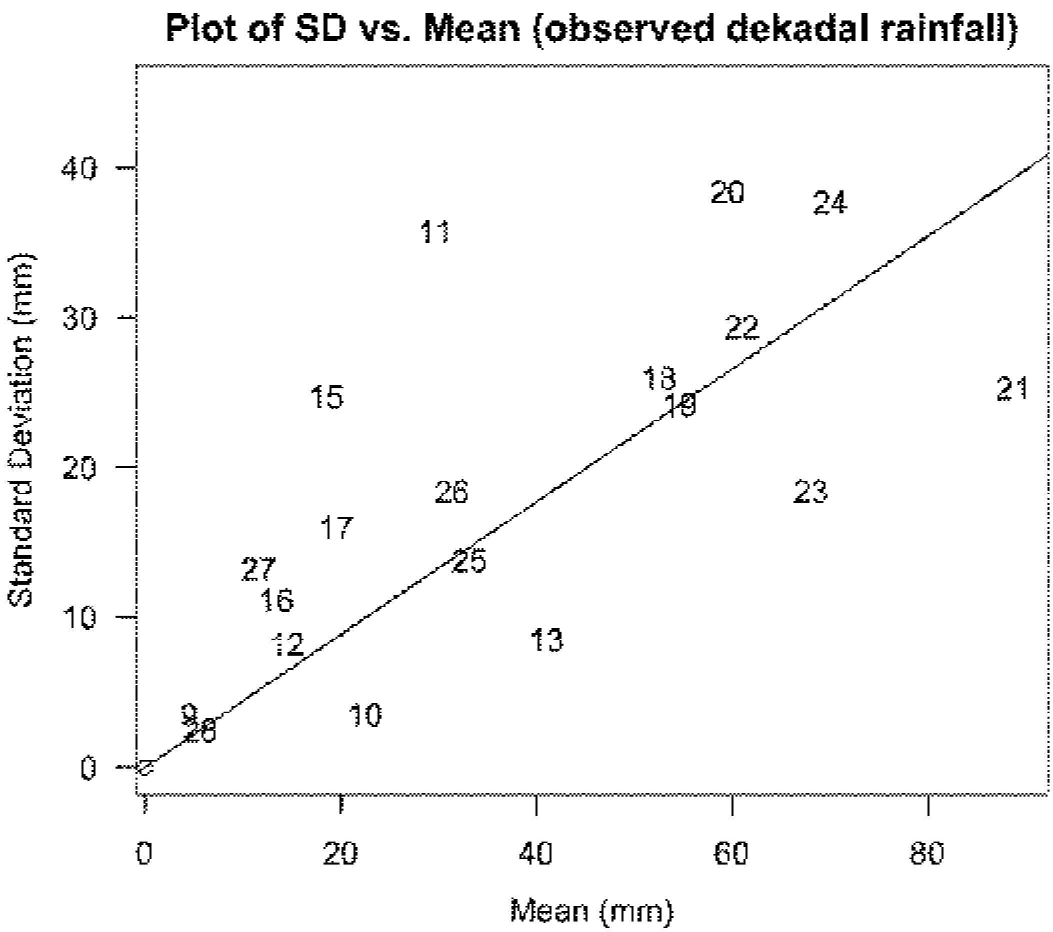

Figure 21 Plot of standard deviation vs. mean of observed non-zero rainfall, for dekads in which there were at least 2 years of non-zero rainfall, where the points are plotted using their dekad number (1-36), and the line is the weighted least squares regression

If we use a gamma prior for the rate parameter,

b_d = Gamma(a_0, b_0),

then the posterior distribution of b_d is given by

$$p(b_d | Y_d) \sim \text{Gamma}(a_0 + Z_d a_d, b_0 + \sum_{t=1}^{T} Y_{dt}$$

where Y_d denotes the vector ($Y_{d1},...,Y_{dT}$). We choose $(a_0, b_0),$ = (1.5,1.5), so as to be relatively noninformative. For dekads in which no non-zero rainfall sums are observed, this prior results in a rate parameter with a mean of one, and an sd of about 0.8.

To see the results, consider Figure 22. This figure displays samples from the posterior distributions of the gamma distributions from which dekadal rainfall sums are drawn. The procedure is as follows: For posterior simulations *g* =1,...,*G*,

1. Draw $p_d^{(g)}$ (for *d* =1,...,36) and b_d^g (for all *d* such that dekad *d* had at least one non-zero observation) from their posterior distributions.

2. Draw $X_{dt}^{(g)}$ from the Bernoulli($p_d^{(g)}$) distribution for *d* = 1,...,36, *t*=1,..,*T*; Draw $Y_{dt}^{(g)}$ from the Gamma($a_d = \hat{a}, b_d^{(g)}$) distribution for *t* = 1,...,*T*, and all *d* for which there was at least one non-zero observation, where $\hat{a}$ is the fixed value of a_d for all *d* that was computed using weighted least squares. The new simulated rainfall amount for iteration *g* is $Y_{\text{pred}(d.t)}^{(g)} = X_{dt}^{(g)} \times Y_{dt}^{(g)}$.

Figure 22contains the densities of the gamma distributions that correspond to different draws of b_d for four dekads of the year.

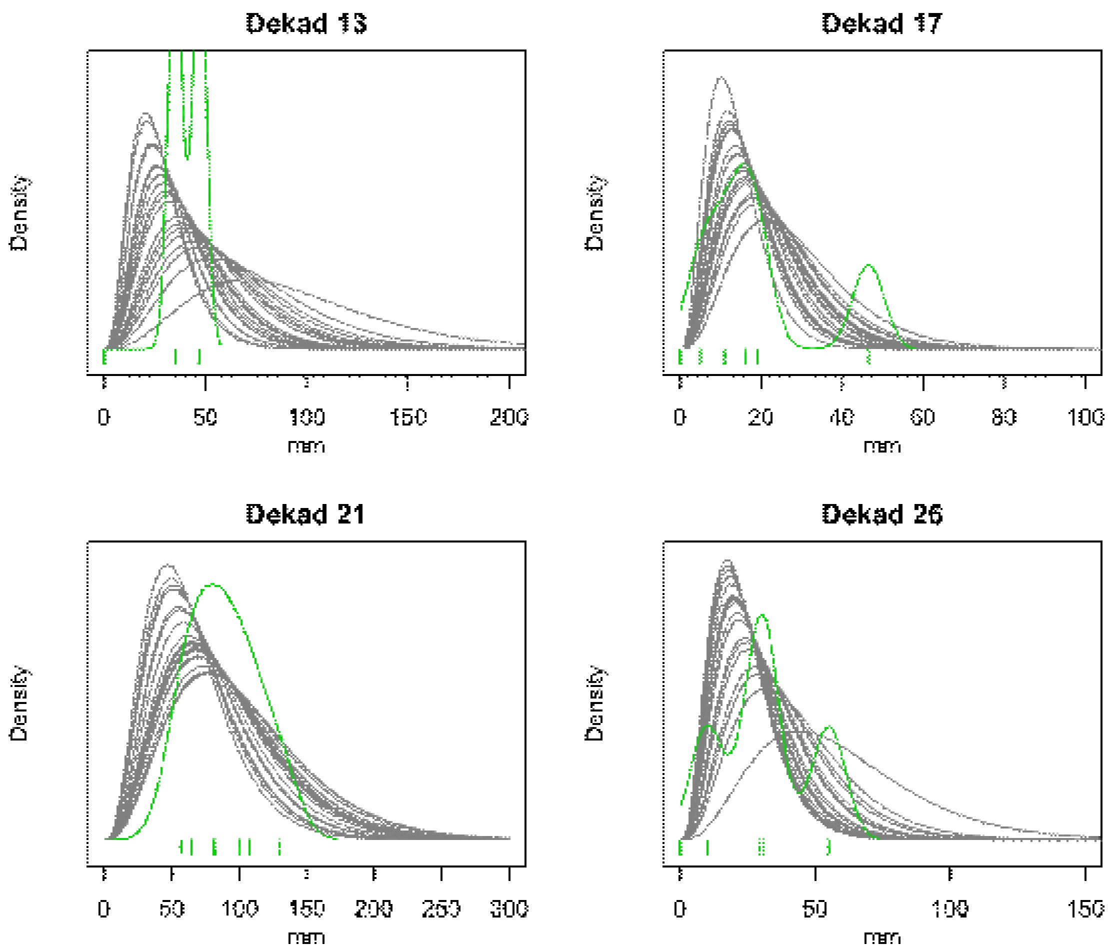

Figure 22 Densities (gray) of the gamma distributions that correspond to different draws of from their posterior distributions, compared to kernel density estimates (green) of the densities of

Y_{dt} from the observed values, plotted as green tick marks along the x-axis, for dekads *d*=13,17,21,26.}

Comments on Figure 22

- The green kernel density estimates are not smooth enough. For example, for Dekad 17, the kernel density estimate is bimodal, with virtually zero probability mass between 30 and 40mm. This is why we use a parametric model - because we know that the distribution of dekadal rainfall should be somewhat smooth, and probably unimodal.
- The gamma distributions have long right tails. For Dekad 13, for example, one of the gamma distributions is very flat, and puts probability mass above 150 mm, even though the only two non-zero observations for this dekad were 35 mm and 47 mm. The data can't rule out the possibility that this gamma distribution produced these two data points,

but observations from surrounding dekads suggest that this gamma distribution is not appropriate.

- For Dekad 21, the observed mean rainfall amount is much larger than the observed sd, but setting the shape parameter of the gamma distribution a_{21} =5.1 restricts the standard deviation to be just under half mean. From Figure 22, it looks like there is too much mass at low values, and too long a tail on the right.

Two alternative versions of the model

To find a better-fitting model, we modify the way that we estimate the shape parameter for each dekad. Instead of setting it to a fixed value for each dekad, we will allow it to vary across dekads. Ideally, we would estimate it simultaneously with the rate parameter in a Bayesian way, resulting in a fully flexible model for the intensity of rainfall. Unfortunately, there is no simple way to do this and simulate thousands of years of data in a matter of seconds (i.e. there is no joint conjugate prior for both the shape and rate parameters that yields a closed form posterior).

We propose two different alternative methods for estimating a_d for each dekad. First, we estimate a smoothed version of a_d by fitting a model to the method of moments (MOM) estimates that assumes a yearly cycle. Second, we simply use the MOM estimates themselves. The first alternative is more robust to outliers, but is also less flexible and less able to pick up patterns beyond those that follow a yearly cycle. Figure 23 shows the smoothed estimates and the MOM estimates of the shape parameter a_d for each dekad for the Adiha data.

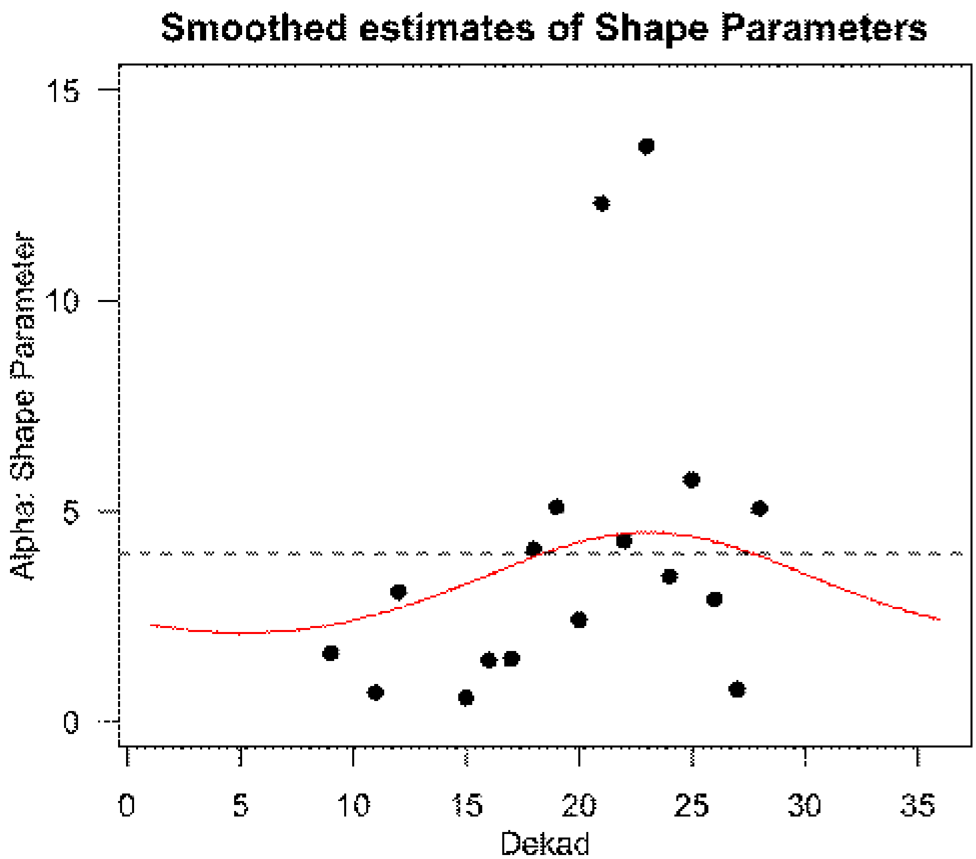

Figure 23 Smoothed estimates (red curve) and MOM estimates (black points) of the shape parameters for Adiha, where two of the MOM estimates (= 40.5 for Dekad 10, and = 23.5 for Dekad 13) are very far away from the others, and lie outside the bounds of this plot

Results

We evaluate the three models

1. Model 1: Fixed $a_d = \bar{a}$ for each dekad
2. Model 2: Smoothed a_d (see Figure 23)
3. Model 3: MOM a_d

in two different ways. First we use posterior predictive checks. Second, we use cross-validation to compute out-of-sample errors.

Posterior Predictive Checks

Posterior predictive checks consist of comparing the value of an observed statistic to the posterior predictive distribution of that statistic. For example, suppose we are interested in the mean and standard deviation of rainfall in Dekad 21. The observed values of these statistics are 89 mm and 25 mm, respectively. Figure 24 contains histograms of the posterior predictive distributions of these statistics for the three types of models:

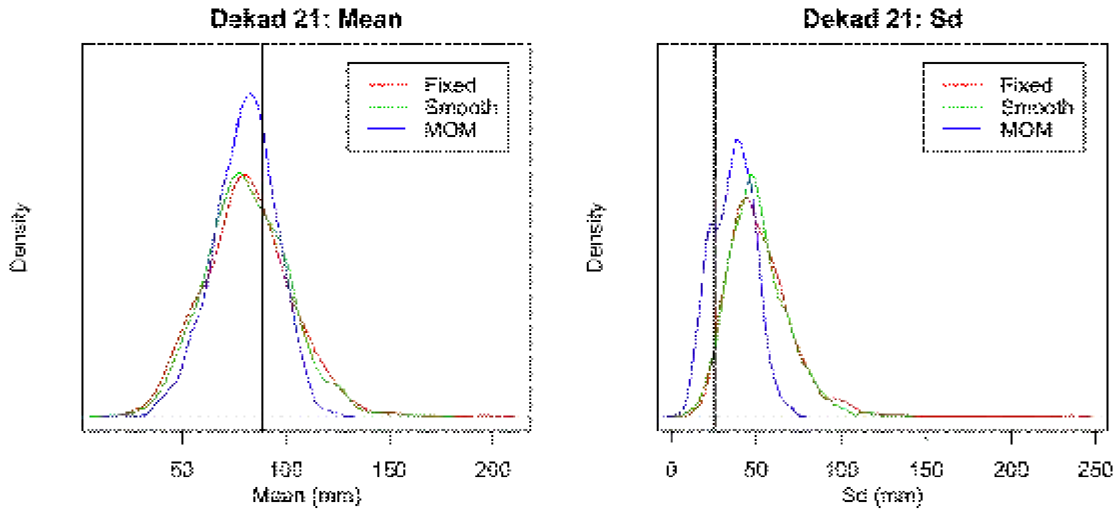

Figure 24 Posterior Predictive checks of the mean and sd of dekadal rainfall for Dekad 21 for all three models, where the black vertical lines indicate observed values of these statistics.

All three models appear to be a reasonable fit, although Models 1 and 2 both tend to overestimate the standard deviation of rainfall in this dekad, as evidenced by the fact that the posterior predictive distribution of this statistic for both of these models puts much of its mass in the range 30 mm - 80 mm, whereas the observed value is about 25mm.

For another posterior predictive check, consider the maximum dekadal rainfall value for each dekad. Figure 25 plots the observed maximum values (x-axis) against the posterior mean maximum values for each dekad (y-axis) and model. Two points in Figure 25 stand out. First, in the upper right corner, there is a green and a red point; these both reflect the fact that for Dekad 21, the gamma distributions for Models 1 (red) and 2 (green) have right tails that are too long - they overestimate the maximum dekadal rainfall. The other point that stands out is not actually in the plot, because it is too far away. In Dekad 15, the observed maximum dekadal rainfall was 36 mm, but the mean maximum dekadal rainfall in the posterior predictive distribution using Model 3 is 245! The reason for this is that the MOM estimate of $a_{15} \approx 0.56$, which results in a very long-tailed distribution (a_d =1 is the exponential distribution), allowing for a few very large values of dekadal rainfall to be simulated.

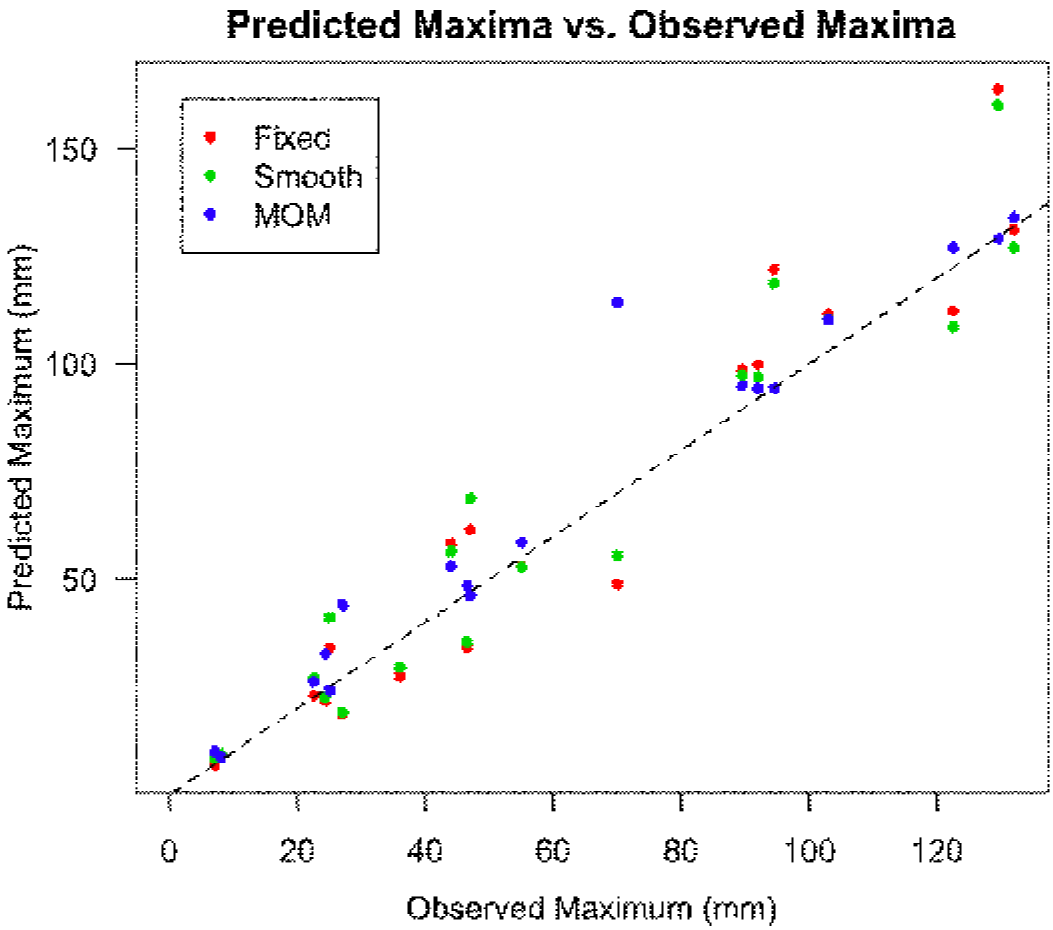

Figure 25 Posterior Predictive checks of the maximum dekadal rainfall by dekad and model, where the y-axis provides the scale for posterior means of the maximum rainfall for each dekad and model.

In summary, the posterior predictive checks can provide detailed accounts of exactly where a model is fitting the data well, and where it is fitting the data poorly. If we are concerned about simulating a few huge dekads of rainfall, we would avoid Model 3, or at the very least, cap the dekadal rainfall simulations that come from this model using a reasonable maximum. For the Adiha data set, checking means and standard deviations for all dekads shows that Model 1 and 2 contain more uncertainty in the simulations than Model 3. None of the three models is clearly superior, although there is reason to be concerned about Model 3 because of some unrealistically large simulated values.

Cross-Validation

Another way to evaluate models is by using cross-validation to simulate out-of-sample predictions. We perform 7-fold cross validation and display the results in Figure 26. The basic idea of cross-validation is to divide the data into a training set and a test set, fit the model to the training set, and then evaluate the model's predictions on the test set. 7-fold cross-validation consists of partitioning the data into 7 sections, and for each section, using the 6/7 section of data as the training set, and using the 1/7 section of data as the test set. This results in each data point being part of the test set exactly one time.

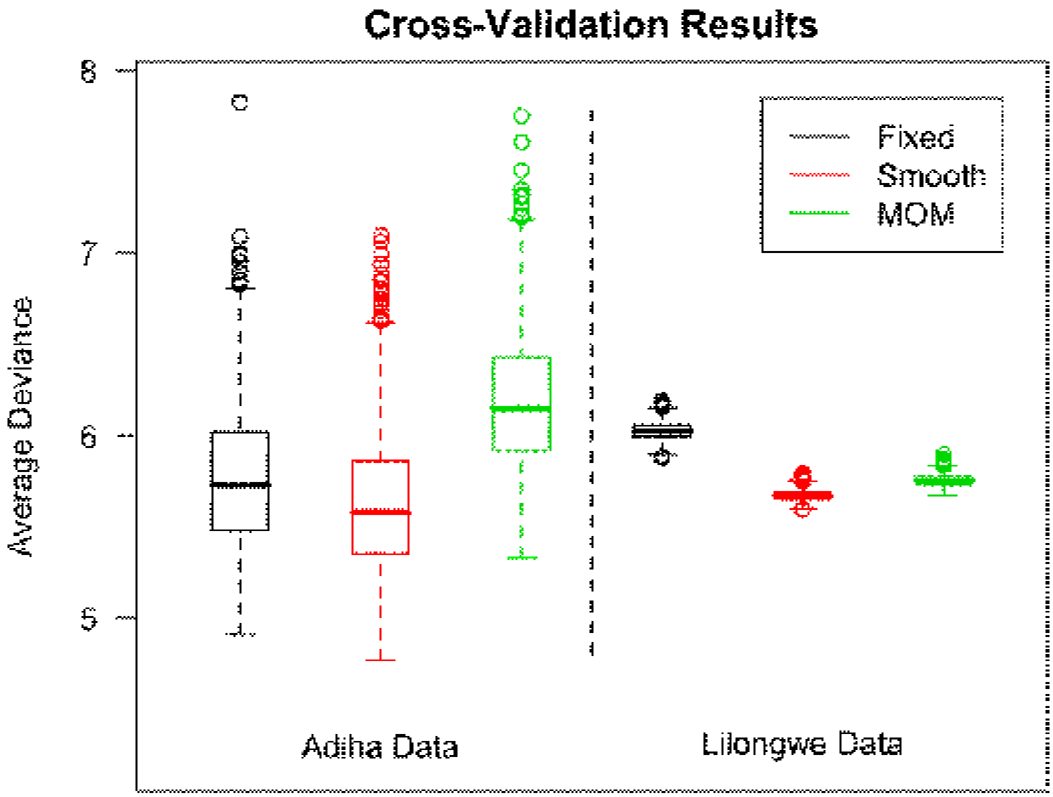

Figure 26 Cross-validation results for 3 models, and 2 data sets. The leftmost 3 boxplots contain results from the Adiha data set, and the rightmost 3 boxplots contain results from the Lilongwe data set.

In our final procedure, we perform 7-fold cross-validation 1000 times, for each of the 3 models, and each of the 2 data sets:. That is, for each model and data set,

1. For *g* = 1,...,1000:
   a. For data sections *i* = 1,...,7:
      i. Fit the model to 6/7 of the data (Adiha or Lilongwe), leaving out section *i* of the data. For the Adiha, the left-out section is one year. For the Lilongwe data, it is 6 or 7 years.
      ii. Sample one value from the posterior distribution of the parameters: $q^{(g)} = (p_1, p_2 ..., p_{36}, b_1, b_2, ..., b_{36}, a_1, a_2, ..., a_{36}),$, where given the data and the model, a_d is determined for all *d* = 1,...,36 and doesn't vary from iteration to iteration. In other words, it doesn't depend on *g*, whereas p_d and b_d do vary from one iteration, *g*, to the next, *g* + 1.

iii. Compute the deviance of each data point in the test set (section *i* of the data), where the deviance of a data point is $d_{dt}^{(g)} = -2 \times \mathrm{P}(Y_{dt} | q_{(g)})$..

b. Compute the average deviance across the whole data set,

$$\bar{d}(g) = \frac{1}{36T} \sum_{d=1}^{36} \sum_{t=1}^{T} d_{dt}^{(g)}$$

Recall that for each iteration, *g*, each data point Y_{dt} is a member of the test set exactly once.

Figure 26 contains boxplots of $\bar{d} = (\bar{d}^{(1)}, \bar{d}^{(2)}, ..., \bar{d}^{(1000)})$ for each of the three models and two data sets, where each boxplot contains *G* = 1000 points. Lower values of the average deviance indicate a model that makes better predictions. For the Adiha data set, Model 2 (smoothed a) does the best, Model 1 (fixed a) is the second best, and Model 3 (MOM a) is the worst. This isn't surprising - the MOM estimates for a small data set are likely to be bad, compared to the fixed a model results, because the fixed a is more robust to outliers, whereas MOM estimates are sensitive to outliers.

For the Lilongwe data, the results are slightly different: the ranking of the models from lowest average deviance to highest is 2,3,1. In this case, the data set is relatively large, and the MOM estimates aren't that bad, whereas the fixed a model is too conservative, and fails to pick up important differences between dekads, that the MOM estimates have the flexibility to model.

Ultimately, however, for both data sets Model 2, smoothed a , was the best, and we recommend it going forward. It is quick and easy to compute, and its most important advantage is that it models dependence between neighboring dekads, because the smoothing function is a regression of dekadal means and sds on the periodic elements $x_1 = \cos(2pt/36)$ and $x_2 = \sin(2pt/36)$, and then the shape parameter is estimated from these smoothed functions. The denominator in the sine and cosine terms contains a 36 because that is the number of dekads in one year, which is the most dominant cyclical component of rainfall in virtually every site in the world. It is essentially a compromise between MOM estimates of the shape parameter for each dekad and a constant value of the shape parameter for each dekad.

Remote sensing

Remotely sensed satellite products show great promise for use with weather index insurance, either to write an index or as a supplemental source of data. They are virtually immune to tampering, often have a longer series of data than ground-based gauges, and are easily available in almost real time over the internet. Because of these advantages, it is tempting to think of satellite data to be a convenient catch-all solution to the lack of data in poor rural areas of Africa, but in reality each product has its advantages and disadvantages. Satellite observations can only be considered a proxy and each separate product must be carefully evaluated for pitfalls.

There are two types of products that can be used to write indexes – rainfall and vegetation measurements – as well as products intended for visualization of ground features as different resolutions, such as Landsat and Quickbird.

Remote sensing products have failings that require calibration to the specific problem and the ground conditions. Remote sensing is a rough approximation of what is happening on the ground. It is currently not known to what extent it can be relied upon for index insurance applications at scale or exactly how it should be used. We are aware of this challenge firsthand from our development of prototype remote sensing indices that have been used in the Millennium Villages Project, and through our collaboration with other projects involving remote sensing.

For example, remotely sensed estimates of daily rainfall rarely capture the actual daily rainfall amounts accurately. However, in many more situations, they are very good approximations of monthly totals or number of rainy days. For remote sensing of vegetation, there is biweekly data available for thirty years. However, this data is of very coarse resolution, with a single pixel covering kilometers of farms, mountains and forests, making it difficult to understand if greenness changes reflect crop behavior. Figure 27 illustrates the relative pixel sizes of some remote sensing products over Adi Ha. High quality data is only available recently, or from rare individual snapshots, preventing it from being directly useful in understanding historical behavior.

Therefore the remote sensing product utilized will need to negotiate a careful series of compromises to be useful. Depending as it does on cutting-edge technology, one tradeoff that occurs with satellite data is one between length of time and accuracy of estimates. In other words, the data that goes back the farthest could be the least reliable. This is especially important for calculating an insurance premium based on historical frequencies, as some of the newer products have existed for less than ten years, hardly enough time to provide an accurate picture of drought risk on the ground. For instance, two of the satellite rainfall products available over Africa are ARC (African Rainfall Climatology) and CMORPH (CPC morphing technique) The ARC data offers the best tradeoffs in terms of length of time series and spatial resolution, it has a pixel size of one arc minute (approximately 11 km) going back to 1993. On the other hand, CMORPH has a spatial resolution of 25 km and goes back to 2003. ARC had the added advantage of including the three woredas of Adi Ha into the same pixel, which simplifies the contract creation process. But CMORPH is superior to ARC in terms of accuracy (Dinku et al., 2008). Currently there are efforts underway at IRI to generate long rainfall time series with ARC's spatial resolution and CMORPH's accuracy. This will be accomplished through local calibration and merging of ground measurements with satellite estimates. Both these products use infrared imaging to estimate the height of rainclouds, on the assumption that taller, colder rainclouds produce more surface precipitation. Satellites also exist to measure passive microwave radiation, which measure actual rainy areas and are therefore more accurate. However, there are no satellites in a geosynchronous orbit that have this capability, so less information is available. Future efforts may include a blend of these two products and data gleaned from surface rain gauges to produce a composite index. We are currently involved in research collaborations that will extend the remote sensing estimates of rainfall in Ethiopia to approximately thirty years. Although these extended chronologies will not be available in time for the current contract, they should be available for more robust pricing of contracts in the future.

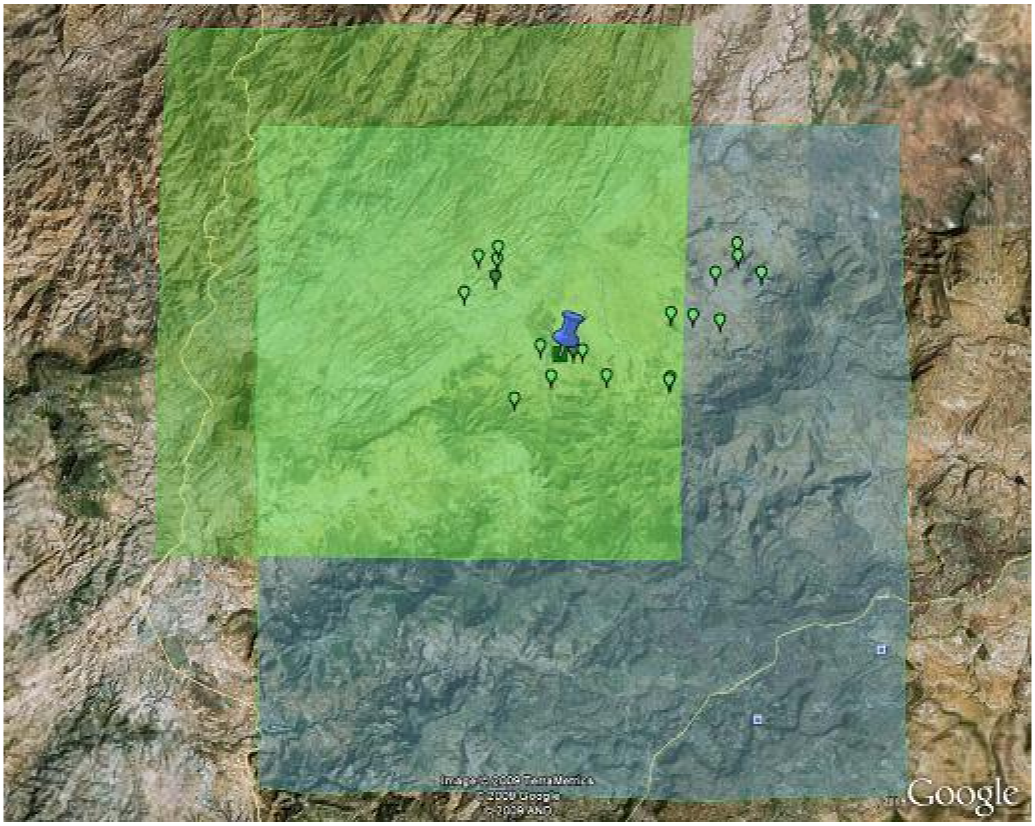

Figure 27: Relative size and placement of pixels for ARC (blue), AVHRR-NDVI (light green), and MODIS-NDVI (dark green/very small) . Also pictured: farm sites participating in rainfall measurement experiment

Vegetation indexes also show some promise. An overview of issues in vegetative sensing and index insurance is presented in Ceccato et. al 2008. NDVI is the most commonly constructed greenness index, which uses the ratio of infrared and near-infrared light to determine the amount of plant cover based on the observation that plants absorb infrared light for photosynthesis but reflect the potentially harmful near-infrared light. NDVI values range from -1, where land cover like desert, rock, or tundra reflecting all infrared light, and +1, where plant cover absorbs all light in these wavelengths. To measure NDVI, we have two remote sensing options to use – AVHRR (Advanced Very High Resolution Radiometer) and MODIS (Moderate-resolution Imaging Spectroradiometer). AVHRR is the more established of the two, with a time series dating to 1978, but is comparatively crude to the MODIS product. AVHRR captures five wavelengths of radiation, while the MODIS captures a full twenty. The tradeoff occurs in the length of the time series, as MODIS is much shorter, with scarcely ten seasons of data.

NDVI is used as a surrogate for drought stress. Insurance contracts will indemnify producers for deviations from the normal NDVI index, which can be at a coarse scale because droughts are often highly spatially correlated, afflicting large areas equally. NDVI works best in countries that practice industrial agriculture, like the United States or Canada, and encounter estimation problems in countries in which natural vegetation is present, like Ethiopia. There is approximately a 2-3 week lag between the rainfall and the appearance of increased plant cover on the NDVI index. NDVI also requires variations in greenness to be effective, and only provides useful information for locations with the right mix of characteristics. For example vegetative 'greenness' indices appear to work well in grasslands and regions where all plants brown during droughts. However, greenness indices are nearly useless in regions where there is a thick coverage of trees that remain green even in dry years. It is important to note that NDVI includes sources of variation that are not related to crop production such as solar angle, atmospheric attenuation, satellite angle, sensor ageing, differences between different satellites, confusion due to bare soil, and the response of vegetation that was unrelated to crop losses.

The use of NDVI has been explored in index insurance contracts in the United States, Canada, and Mexico. In 2007, IRI has designed transacted NDVI-based, dual trigger contracts in Kenya, Mali, and Ethiopia (for Koraro, approximately 30km from Adi Ha) for the Millennium Villages project. Agricultural economists like Michael Carter at UC Davis have explored the crop yield/vegetation index relationship in Burkina Faso (Carter 2009). We are aware of ongoing activities exploring the use of NDVI for index insurance for livestock in Kenya . Because NDVI does have serious limitations, it is important to continue work to better understand these limitations in terms of index effectiveness and develop improved solutions.

In addition to their potential use as an index, vegetative remote sensing estimates may be of great value building, analyzing, and validating an index based on other sources of data. We are particularly interested in determining if vegetation indexes can be developed to address the particular challenges for index insurance as well as quantifying their limitations in these tasks.

- **Mapping the area around a met station for which an index product can be applied.** ***This challenge is one that we are explicitly investigating as part of the Terms of Reference commissioning this report.*** Techniques have not been established to determine the maximum distance from a rain gauge at which one can design a responsible contract, i.e. how many gauges are required and where they need to be. Implementers have largely followed an arbitrary rule of thumb of 20km radius around the gauge, although in some places 20km is far smaller than necessary, and in others the boundary is dangerously large. For example, in a large valley, a single station may represent the rainfall experienced over a great region. However, when there are mountains between a met station and a farm, these mountains may cause rainfall amounts to vary substantially. This issue is of particular concern in Ethiopia with its mountainous terrain.
- **Establishing the historical frequency of major events.** Some NDVI products provide data from the early 1980s. If vegetative sensing can be validated as reflecting years with major crop loss, these products may be extremely useful in proxying the frequency of major events, to aid in the design of insurance products using much shorter datasets.
- **Validating information from other sources.** Vegetative sensing may be used as a method to validate other sources of information, such as crop modeling or satellite based estimates of rainfall to determine if there is agreement in the identification of loss through very different sources of information.
- **Understanding and improving vegetative sensing for direct use as an index.**

As part of our project for Oxfam, Dr. Chris Small of Columbia University completed a field reconnaissance in Adi Ha in March 2009 to determine the possibility of using a vegetation index, which he concluded was "challenging but ... feasible." Although his focus was to explicitly investigate the potential for the use of remote sensing in mapping the area around a met station for which the index could be responsibly offered, the research provides products that will be useful in understanding the other questions we are interested in. His analysis requires that he visits the field at two key time periods. The first time period is before the crop has been sowed, and the second is when the crop has full leaf cover so that he can compare the satellite imagery (and associated ground truthing) between the two periods to determine the level to which the satellite imagery reflects the presence and vigor of the crop. To date, the first field visit has been completed, and it is intended that he will follow up with a field visit in approximately August, if permitted by continued funding.

The purpose of the field reconnaissance was to determine the feasibility of satellite monitoring of agriculture in different environments. Specifically, to investigate the distribution and abundance of indigenous vegetation that would need to be distinguished from agriculture when both occur simultaneously. The spatial scale of agriculture and the diverse environments present in northern Ethiopia present a number of challenges to satellite monitoring. Field reconnaissance is essential to determine, prior to analysis, what environmental variables need to be considered when designing the analysis. The preliminary conclusion based on the field reconnaissance is that the task will be challenging but appears to be feasible.

Agriculture monitoring will be done with Landsat TM/ETM+ imagery as it is available free of charge, generally within days of acquisition. While Landsat's 16 day revisit period is less than ideal for monitoring, the combination of 30m spatial resolution and the deep (27 year) archive of imagery compensate for the temporal resolution. While daily coverage would be desirable for monitoring, the spatial resolution of wide-field sensors like MODIS is inadequate for distinguishing fine scale agriculture from indigenous vegetation. In contrast, the 30m spatial resolution of Landsat, combined with its Shortwave Infrared (SWIR) sensitivity allows for the possibility of distinguishing agriculture from indigenous vegetation on the basis of both spatial distribution and temporal phenology. However, in order to accomplish this, it is necessary to establish validation sites in a variety of different environments where the abundance and type of indigenous vegetation is known *a priori*. One of the primary objectives of the reconnaissance mapping is to establish such sites. A complementary objective is to determine the type and variability of soil reflectance properties in the study area. Knowledge of soil reflectance can be important to estimation of vegetation abundance because of the spectral effects of exposed soil on the aggregate reflectance measured by the sensor. The Landsat sensors' 30m resolution approaches the scale of individual agricultural plots but will inevitably result in mixed pixels containing both indigenous vegetation and agriculture. In order to understand the effect that varying proportions of indigenous vegetation have on aggregate reflectance at the scale of the 30m Landsat pixel it is necessary to know the type and abundance of indigenous vegetation present in different environments. This is accomplished through the combined use of Quickbird imagery (2.4 m spatial resolution) and field observations.

Field reconnaissance in March 2009 allowed for six days of observations in Tigray. A total of 3.5 days of 4WD travel and 2.5 days of hiking resulted in adequate dry season coverage and sampling at both regional and local scales. A total of 1866 photographs, 35 GPS stations, 128 GPS tracks and 14 substrate samples were acquired. Maps, photographs and satellite imagery are available online at: http://www.ldeo.columbia.edu/~small/Ethiopia2009.
Thanks to excellent support from colleagues at REST, we were able to obtain 4WD coverage from Mekele to Adi Ha with northward extension as far as Aksum. Despite limitations imposed by high mid-day temperatures and sandy soil, it was also possible to hike to key locations in the vicinity of Adi Ha to obtain substrate samples well beyond the irrigated areas.

Figure 28 below illustrates one of the reconnaissance sites. The high resolution image is the Quickbird data. The site is at the center of the Quickbird image. The green square illustrates the size of a MODIS pixel, the highest resolution vegetative image available at frequent timescales.

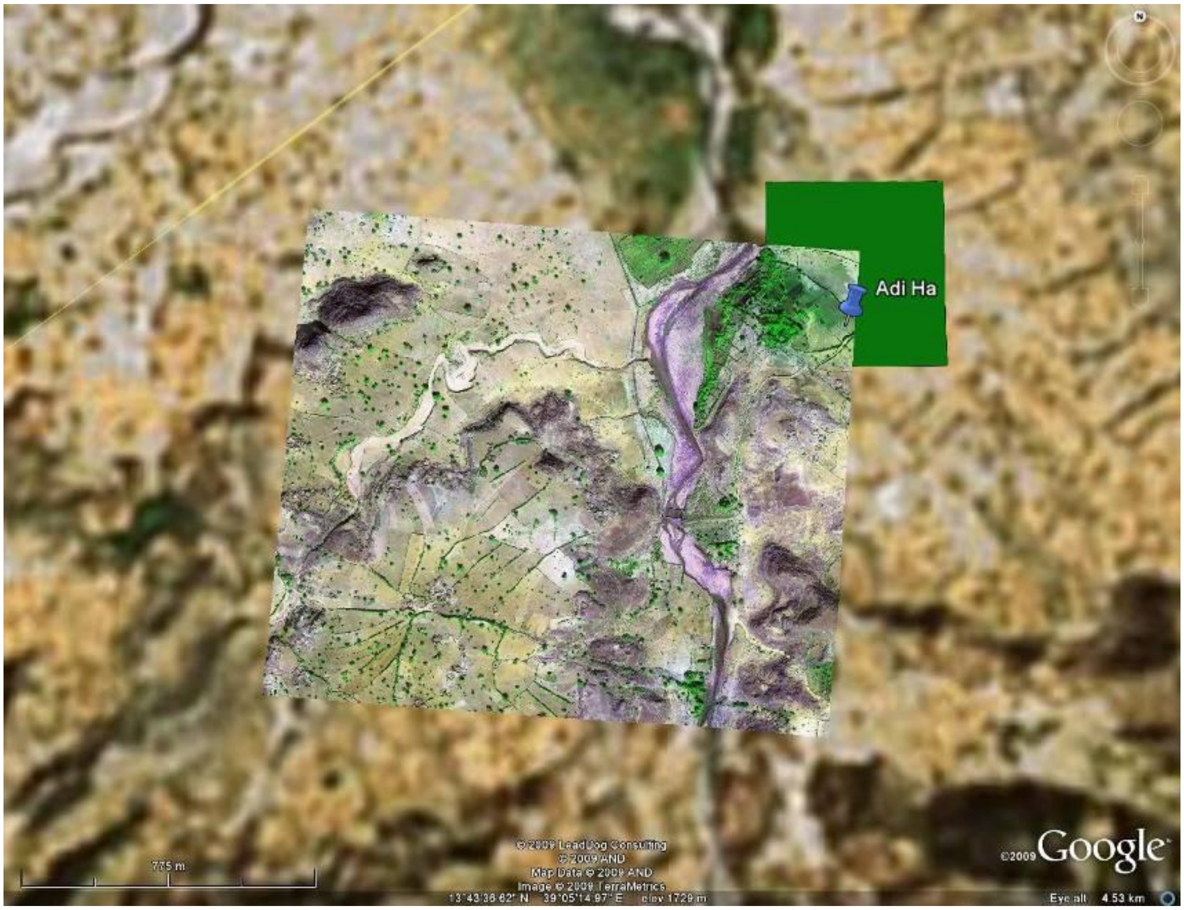

Figure 28 Field Site Example

Figure 29 is a perspective view of this site. Figure 31 is a photo of the site from a similar perspective, to help give a sense for what is depicted in the perspective satellite rendering, and Figure 30is a landscape composite photograph generated for the site, which was part of the groundtruthing data gathered for each site.

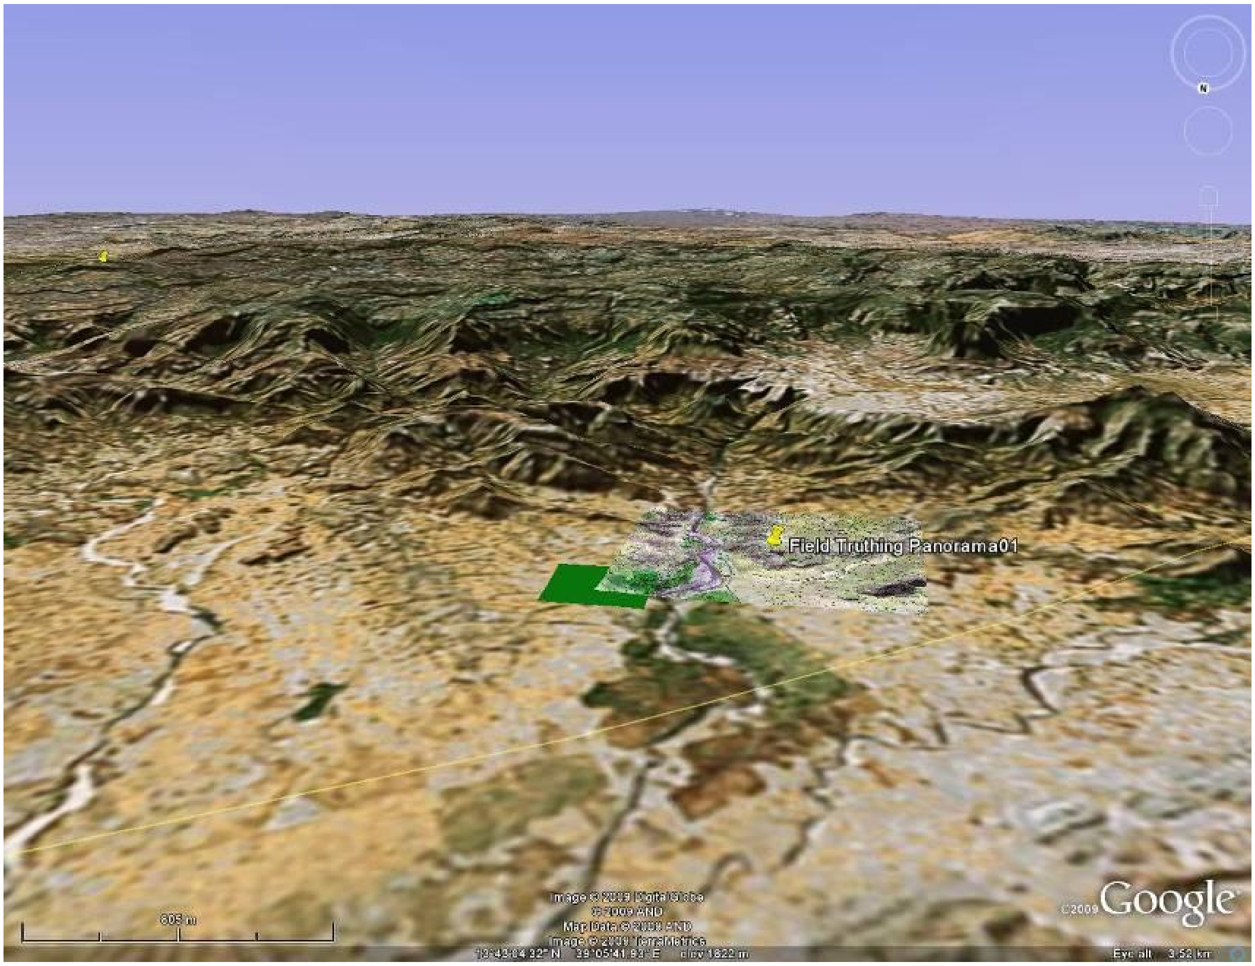

Figure 29 Perspective view of a site

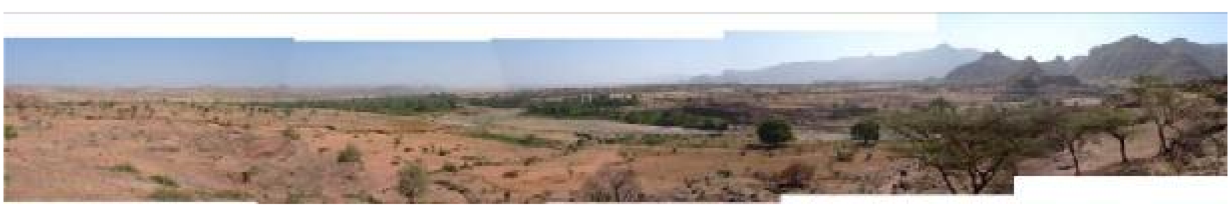

Figure 30 Example panoramic image at field site

Figure 31 View of Adi Ha

Field observations confirm the diversity of geologic terrains and substrates inferred from the Landsat imagery prior to arrival. The recent extension and uplift of the flank of the East African Rift results in the superposition of rugged topography on a wide range of rock and soil substrates. While the environs of Adi Ha village are dominated by very sandy iron-rich soils derived from erosion of the adjacent Jurassic and Miocene highlands, the adjacent PreCambrian metamorphic belt produces a wide variety of substrates and soil textures. The metamorphic belt extends northward to Miocene volcanic terrains near Adwa and Aksum.

The ubiquitous presence of terraces, stone walls, and maintained treelines reveals a far greater extent of terraforming than is apparent in the Landsat imagery. The higher spatial resolution Quickbird imagery gives some indication of the extent of terraforming as indicated by linear discontinuities in substrate reflectance marking boundaries of agricultural plots. However, Quickbird imagery has been acquired only for the area surrounding Adi Ha so multiscale analyses will necessarily be limited to this area. While our colleagues from REST report that Kiremt season agriculture is extensive, preliminary interpretation of Landsat imagery from Sept. 2008 suggests that considerable land area remains fallow or sparsely vegetated around Adi Ha.

The next phase of the analysis would involve development of a multi-scale spectral mixture model to quantify the spatial abundance and distribution of vegetation as a function of time. A multiscale mixture model applied to both Quickbird and Landsat imagery will allow us to develop accuracy estimates for the Landsat mixture fractions in different environments and to quantify the ambiguity that indigenous vegetation will impose on identification of agriculture. This will provide a basis for determining where indigenous vegetation can reliably be distinguished from agriculture and provide a starting point for development of a temporal phenology for vegetation assemblages in different areas. This, in turn, could provide a basis for a vegetation classification that may distinguish between agriculture and indigenous vegetation. If the two classes of vegetation assemblage can be distinguished, this may provide a basis for monitoring crop evolution throughout the growing season. Given sufficient Landsat coverage of past years, it may then be possible to distinguish successful harvests from premature termination of the growing season on the basis of reflectance signature.

While a vegetation index is technically possible, some challenges would remain even if we could construct a perfect representation of agricultural activity. First, clouds can easily obscure infrared reflectance of land cover, especially in a period of high rainfall which naturally contains the agricultural season. Second, it is not clear how to write a contract on a vegetation index. In our current contract structure the payout increases linearly with decreasing millimetres of observed rainfall. To produce a similar contract for a vegetation index would likely require some negotiation and analysis between partners. Third, farmers in Adi Ha have shown a willingness to purchase insurance based on satellite measurements, but relating the more abstract concept of a vegetation index could be a challenge.

Remotely sensed data might be considered at this point in time to be a reliable fallback option, but increasing refinement in measurements and interpretation is likely to provide a reliable source of data in coming years, if the necessary analysis is performed.

Trend Analysis

A non-statistically significant downward trend was noted in the rainfall sum indexed for the Adi Ha contract. For the initial year of the contract this trend has lead to an increase in price of about one quarter of the insurance price, or approximately the same size as the insurance loading. Building upon the initial first year solution it is important that work be done to understand this trend and to build pricing mechanisms that are appropriate and are not vulnerable to instability due to year to year variability. Given our exploratory analysis to date, we are highly optimistic that simple strategies can be found for substantially reduced premiums and price volatility concerning this trend. To do this, we will need to work closely with partners, particularly insurers and reinsurers to agree upon a workable and scalable solution.

It is important to keep in mind that temperature trends are a much more straightforward topic. There is a consensus that climate change will lead to higher temperatures in Ethiopia, as in much of sub-Saharan Africa. This anthropogenic temperature increase will increase the water needs of crops, leading to more severe and more frequent water stress. However, higher temperatures could also lead to higher yields in years with adequate amounts of rainfall. Thus, on top of any impacts climate change has on rainfall, its temperature impacts are expected to increase the variability faced by farmers and hence also increase the value of insurance.

The precipitation trends in Adi Ha are much more complex to understand, and have more intimate quantitative impact on insurance design decisions. It is possible that the trend is due to climate change, decadal processes, or is a spurious result due to overfitting a short dataset. In any case, for the insurance to be effective and affordable in the future, we must develop a better understanding of this trend and how to responsibly build it into the product design, product pricing, and broader risk management.

Research on this issue has fundamental implications for the design of index insurance and for the broader policy environment. If the drying is due to climate change, it could signal the importance of an eventual transition out of the insured activities, or of fundamental changes that enable these activities to be viable with less rainfall. In addition, if the drying is driven by greenhouse gas emissions, the burden faced by the farmers is directly due to those emitting the gases. However, if the drying is part of a decadal cycle, it will be more important for contracts to be built into climate risk strategies that protect against alternating decadal periods of drought and high rainfall. Fortunately, this issue need not be disentangled all at once. If we put effort into observing and adapting our strategies as our information improves, it is likely that we can arrive at worthwhile solutions in time for the decadal or longer timescales that these trends operate at. More discussion of this issue can be found in Hellmuth et al, 2009, and Greene et al.

Implications on the contract if the trend is an accurate representation of the climate

We begin with an illustration based entirely upon the historical data for the Adi Ha index. Figure 32 presents the indexed ARC rainfall during the contract window, along with a line fit to the fourteen years of data. The slope of the trend is negative, that is, the trend reflects a decrease in rainfall over time. In percentage terms, the decrease is approximately 1.27% per year.

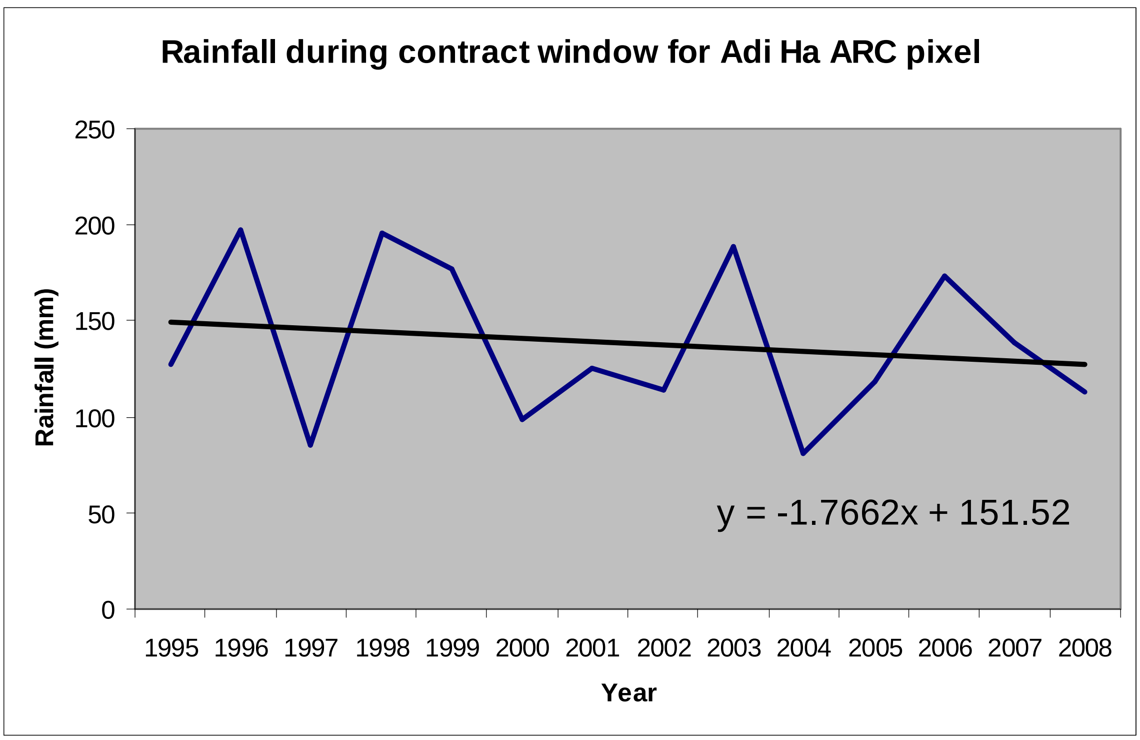

Figure 32 Adi Ha precipitation with trend

Figure 33 includes the trigger level of the contract, illustrating the level below which there would have been a payout.

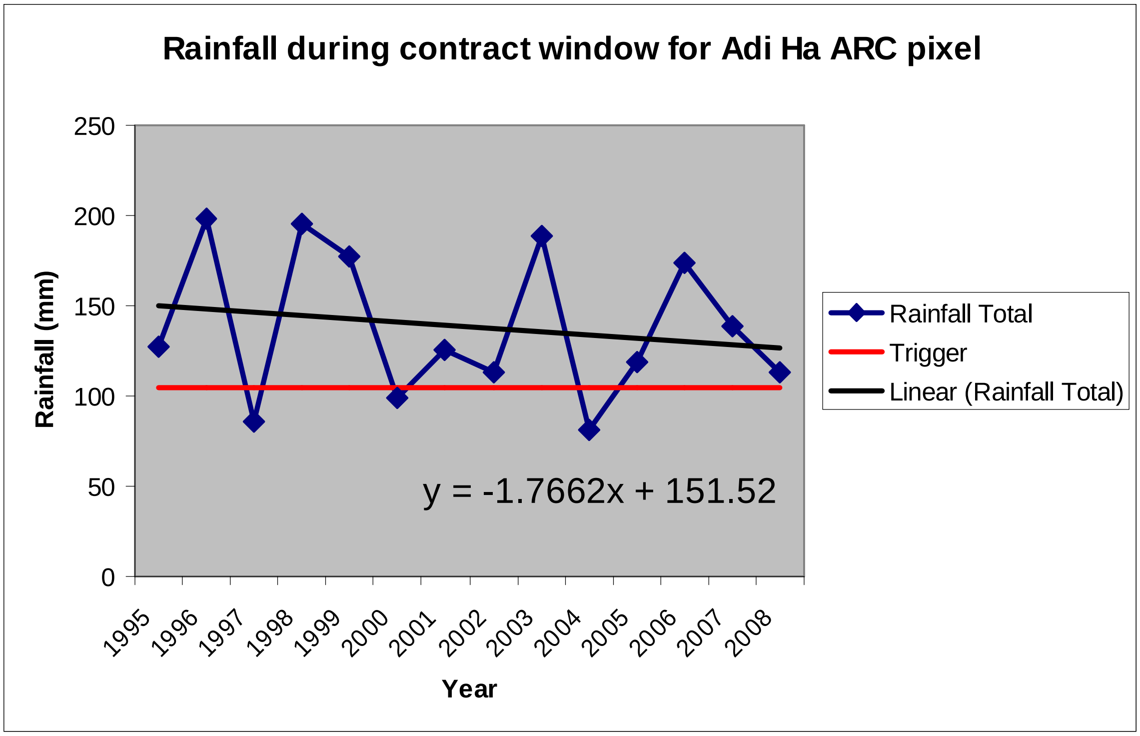

Figure 33 Precipitation with trend and trigger

The use of historical data in contract design is to approximate the performance for the coming year, using the restrictive but transparent assumption that the coming year will be draws from previous years. If the trend is assumed to be a true process in the rainfall, for an understanding of next year, it would have to be removed from the historical data. That is, data from previous years should be made drier to reflect the current climate. Adjusting the historical data to reflect this drying trend yields the precipitation presented in Figure 34. From this figure, one can see that if the drying trend is accounted for in the historical data, payouts would have been much larger and more frequent for the trigger used in the Adi Ha contract.

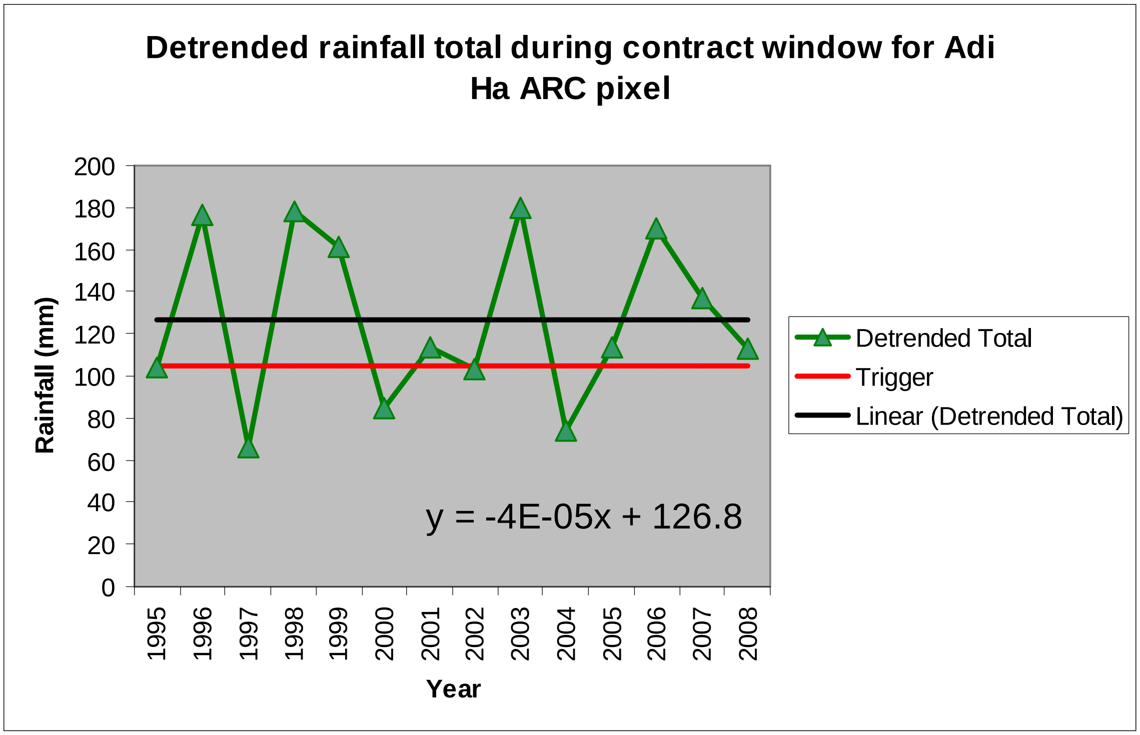

Figure 34 Detrended precipitation

The number of payouts increases from three to five, as 1995 and 2002 register small payouts, and the existing payouts become more severe. For example, the 2000 rainfall is 99.17 mm, for a payout of 13.0% of maximum liability, to 85.04 mm, 44.4% of maximum liability.

Table 10 presents what payouts would have been for observed and detrended rainfall.

Table 10 Historical Burn Payouts with and without trend as a percentage of maximum liability

	Observed	**Detrended**
1995	-	2.1%
1996	-	-
1997	43.1%	86.3%
1998	-	-
1999	-	-
2000	13.0%	44.4%
2001	-	-
2002	-	4.7%
2003	-	-
2004	53.5%	69.2%
2005	-	-
2006	-	-
2007	-	-
2008	-	-
Mean	**7.8%**	**14.8%**

If the detrended data truly reflects the rainfall distribution to be expected in the coming year, the zero-loading price for the contract increases from 7.8% to 14.8%, approximately doubling. As will be discussed later, it is not clear if the detrended data adequately reflects the rainfall distribution to be expected.

Applying the rainfall simulator (section Rainfall Simulators) to the observed data and to the detrended data, we can better understand the implications of this trend[10], if it does exist. The simulator was run to generate 10,000 possible versions of the coming year in each case. We present statistics for the contract designed without taking the trend into account (exit of 80mm) and the contract adjusted for workable pricing on the detrended data (exit of 60mm). Note that two averages of payouts are presented in the payout category. The first, entitled 'Mean' is the average of all payouts including both payout and non payout years. This is identical to the no-loading insurance cost. The second is the average of payouts excluding zero payout years, entitled Mean Payout (Non-Zero), which gives the reader a sense of the size of a typical payout. Also included in the table is the percentage of years for which maximum payouts would have occurred (% Maximum Payout).

The severe implications of the trend on pricing are evident, with the trend leading to substantially lower rainfall with substantially more frequent payouts, and substantially higher prices even without loading. Interestingly the costs of the trend appear to be driven by more frequent payouts, as opposed to larger payouts since the maximum payout frequency does not change much, and the average payout for the years in which payouts occurred is also nearly unchanged. This suggests one strategy that may make the contract robust to the trend.

[10] Note: for purposes of the rainfall simulator, the contract window was considered to be the six dekads from August 11th to October 10th, which is one day different from the six standard 10-day windows used in the actual contract that end on October 9th.

Table 11 Simulations calibrated by observed vs detrended rainfall

	Observed	**Detrended**
Total Simulations	10,000	10,000
Rainfall - Year		
Mean (mm)	451.28	412.38
Std. Dev.	103.39	91.53
Rainfall - Contract Window		
Mean (mm)	126.62	119.57
Std. Dev.	32.94	32.15
Frequency		
% Payout	25.42%	32.58%
% No Payout	74.58%	67.42%
Payouts - 60 mm exit		
Mean	11.26%	14.74%
Std. Dev.	24.96	27.88
Mean Payout (Non-Zero)	44.29%	45.25%
% Maximum Payout	2.51%	3.61%
Payouts - 80 mm exit		
Mean	16.09%	20.95%
Std. Dev.	32.65	36.09
Mean Payout (Non-Zero)	63.28%	64.30%
% Maximum Payout	8.65%	11.61%

Lowering the trigger to pay out for events that are approximately one in seven year events might make the contract have very similar prices between observed and detrended data. This strategy can also be seen by comparing Figure 33and Figure 34, in which it is evident that there are several non-payout years that would have become payout years if they had been a few percent drier. Lowering the upper trigger to approximately 85mm would have eliminated these years from the payout series, making the contract less expensive and more robust. If the trend truly exists, the three payout years from the historical dataset would still be reflected in the coverage. If the trend does not exist, then the contract would have only protected against the two most extreme years.

Because we do not have much statistical confidence that the trend does exist, and the climate science does not provide us with evidence that it is occurring, we were uncomfortable in suggesting that solution. We were particularly uncomfortable with assuming the trend exists through reduced upper triggers because the year 2000 was noted uniformly by farmers as a year in which Teff production was damaged due to an early cessation of rainfall.

Evidence demonstrating the need for improved methodology in estimating the trend

As a temporary and very cautious response, it was important to build the initial pricing around the trend. However, as we have performed additional analysis, there is a great deal of evidence that the approach used will **not be an adequate solution for the future**.

The first concern is for extreme insurance price volatility due to the algorithm used to determine the trend. The estimated trend has a slope of -1.7662 and a standard error of 2.774, giving us a 97.5% confidence interval of [-7.810, 4.277]. Note the differing signs on the confidence interval - the direction of the trend cannot be confidently ascertained. Using the current approach, **the sign and size of the trend is extremely sensitive to a single year of data**.

An observation this coming year in excess of 195 mm (which happened 2 times in the 14 years) would in fact reverse the trend's sign and indicate that the village is getting wetter. Likewise, if there is a payout, meaning that rainfall is below 105 mm, the slope of the trend will increase in severity to at least -2.267 mm/year. If we use the same methods to adjust data to price a premium, we get a premium of at least 18%. Thus, because of the small amount of data**, instead of quantifying the long term change, a simple fitting of the trend will make the insurance price change dramatically simply based on the year to year variation. Although the occurrence of a single payout or wet year does not alter long term climate trends, in the current algorithm, it could greatly increase, or reverse the estimated trend.**

Because of the short time series, relying entirely on the single set of data is not adequate to understand what trends exist. It is therefore essential to include our understanding of climate processes to help identify the extent to which a trend actually exists.

For the case of Ethiopia, and much of the Horn of Africa there is currently a heated discussion about if long term climate change is leading to lower precipitation. The IPCC models do not provide a consensus for what will occur for precipitation in Ethiopia due to climate change. If anything, the models suggest that there may be a slight increase in precipitation (see Figure 35and Figure 36).

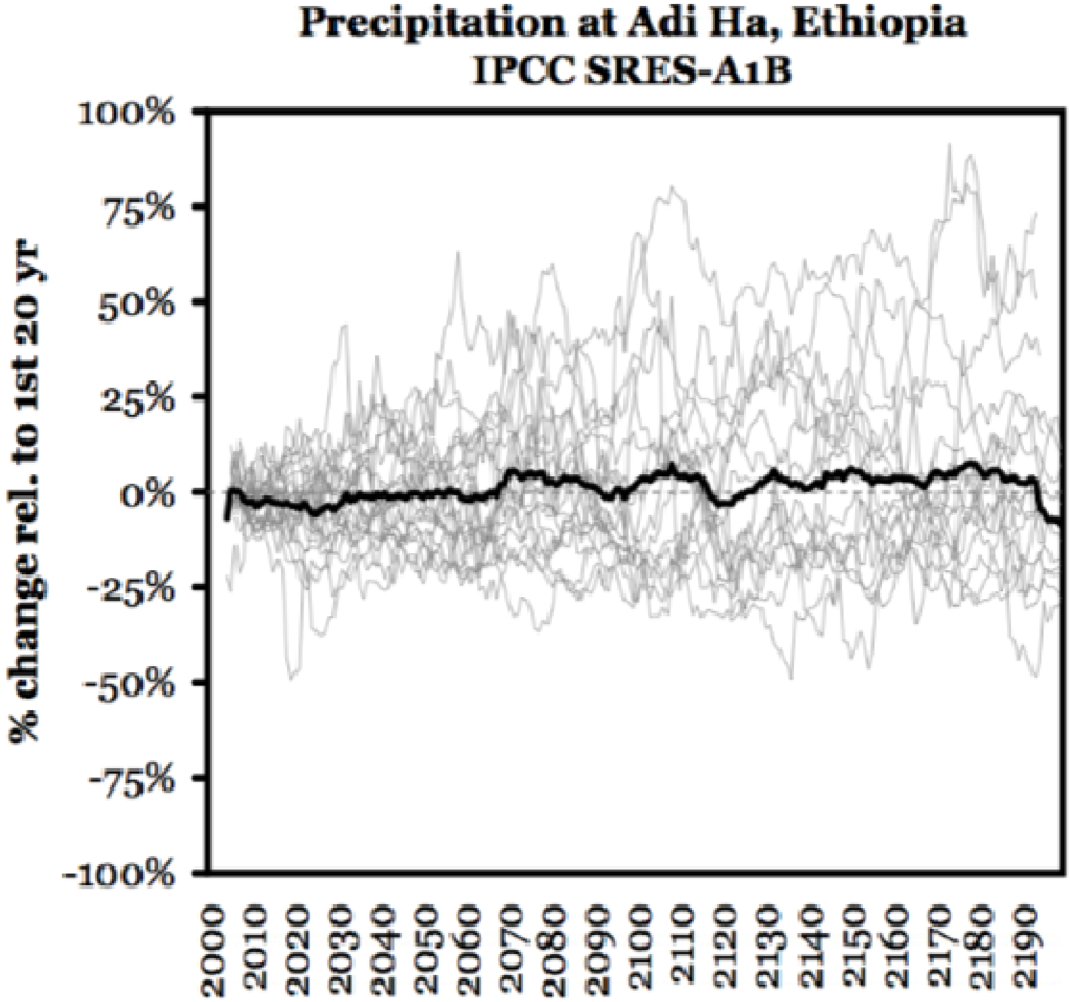

Figure 35 IPCC climage change model precipitation forecasts for Adi Ha

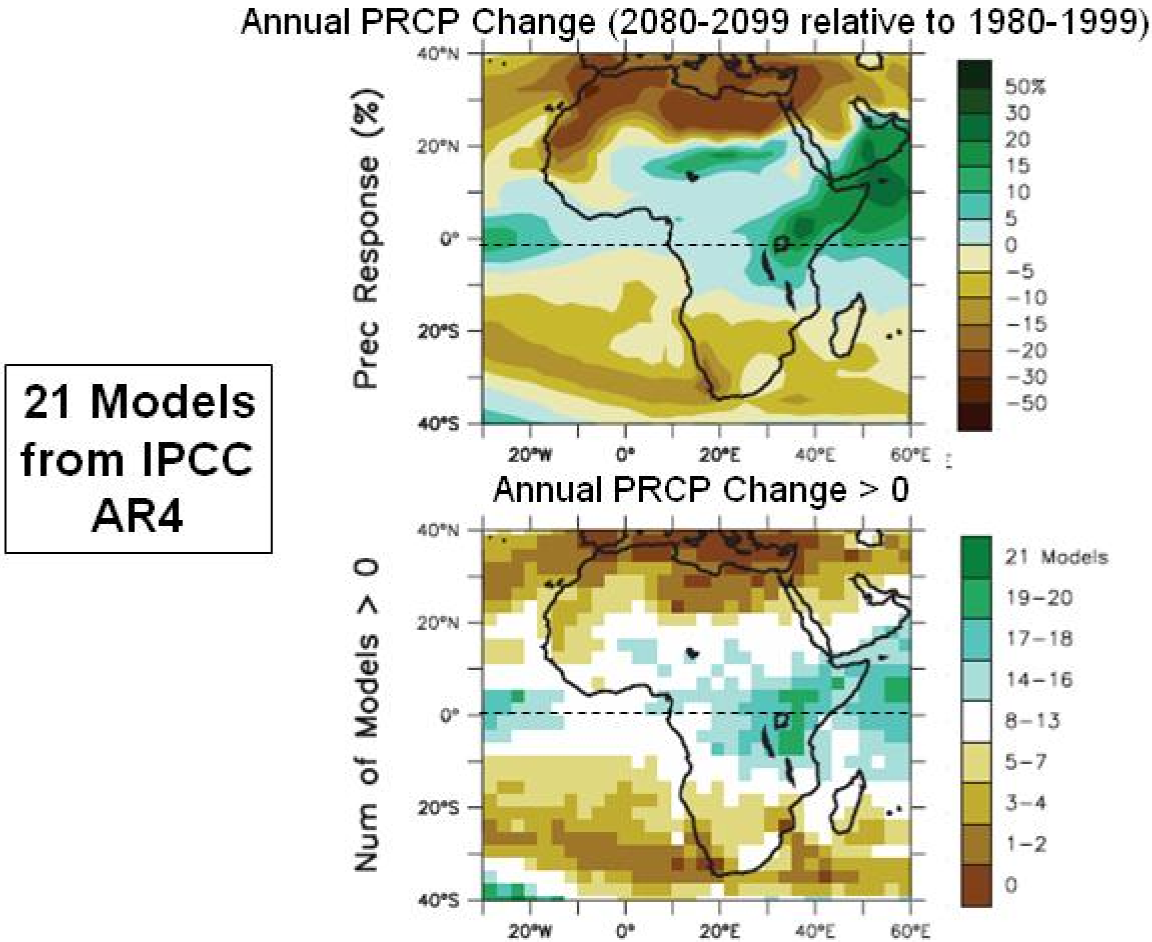

Figure 36 IPCC forecast Annual Precipitation Change

However, for the spring rains, the observed data suggests a recent drying trend. Because of this paradox, the climate science community has been recently exploring if the IPCC models are failing to accurately predict rainfall for this part of the world (Funk *et al*., 2005), if the recent observed drying trend is simply a spurious statistical feature, or if it is driven by a non anthropogenic decadal process similar to (and possibly related to) the decadal cycles observed in the Sahel. Independent of the index insurance project, there is ongoing work at IRI on near term climate trends. Decadal timescale work is led by Lisa Goddard and the Ethiopia effort is driven by Bradfield Lyon, Alessandra Giannini and Tufa Dinku. Arthur Greene is involved in parallel efforts for other regions. This work is being informed through discussions with related ongoing work by other groups, including the Ethiopian National Meteorological Agency.

For the IRI effort, some preliminary work has been performed, including:

- Comparing recent trends in eastern Africa across multiple rainfall datasets based on a) station data only, and b) combined station data and satellite rainfall estimates.

- An analysis to separate out precipitation patterns over eastern Africa and the Indian Ocean associated with ENSO, the dominant source of interannual variations, from the patterns more closely associated with the recent observed trend (see Figure 37).

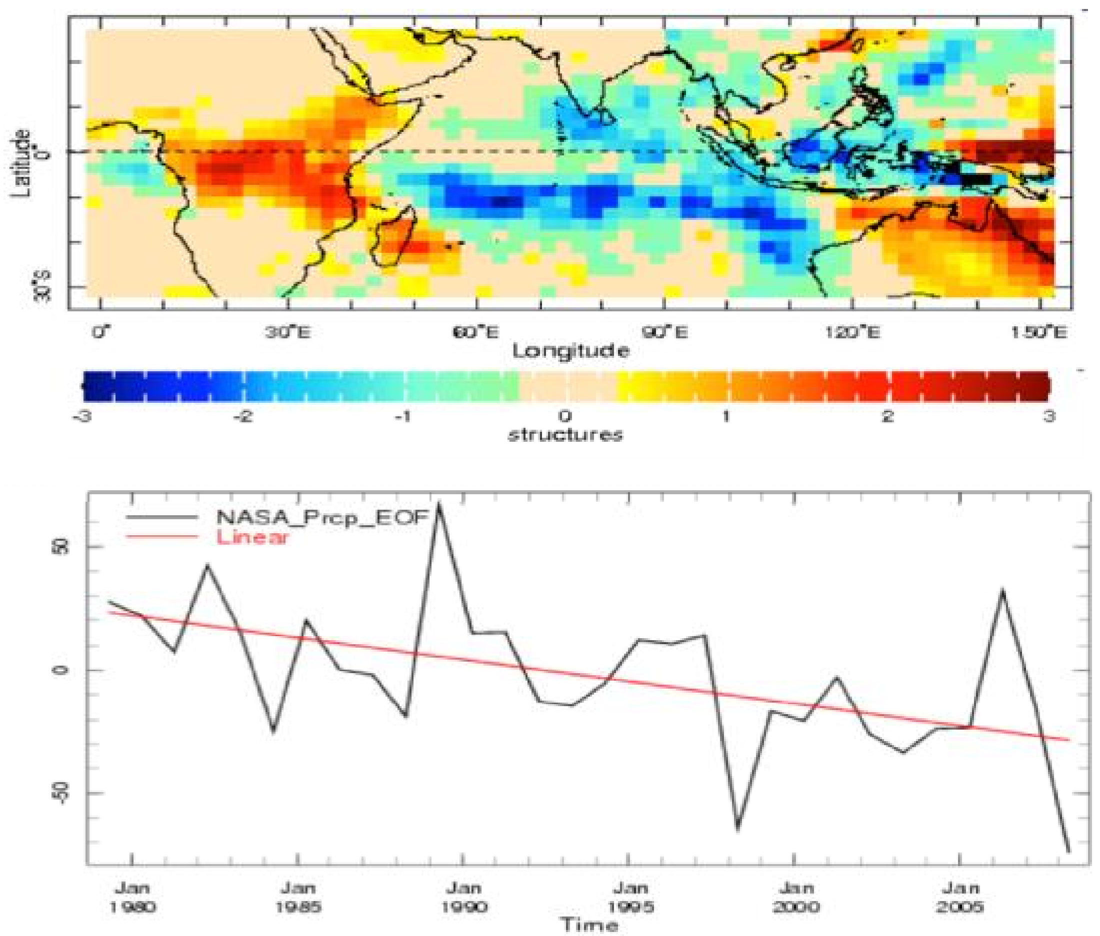

Figure 37 Pattern of Mar-Apr-May rainfall (top) and its associated trend (red line, bottom) based on EOF analysis.

The opposite-signed values in the top plot over eastern Africa and the Indian Ocean indicate opposite trends, with precipitation decreasing over eastern Africa and increasing over the Indian Ocean.

Future work will focus on

- An analysis of observed precipitation variability, and trends, as a function of season.
- Investigating whether the observed trends are associated with anthropogenic climate change (long-term climate change) or whether they alternatively might represent variability that is occurring on shorter, decadal time scales and thus will not be expected to continue.

For the spring rains, the observational record suggests a recent drying trend across East Africa. Because an individual dataset may show a great deal of random noise, IRI researchers have averaged rainfall over a large box, spanning from Ethiopia to Tanzania (see red box in Figure 38) to explore if there are systematic large scale changes in rainfall over time.

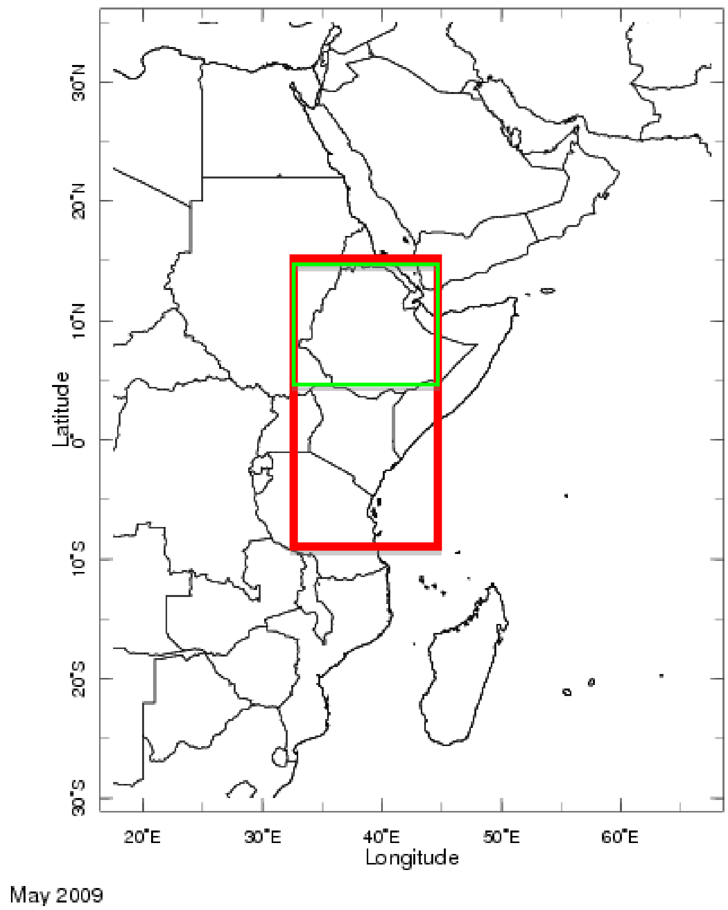

Figure 38 Rainfall Averaging boxes: Regional (Red), Ethiopian (Green)

Much of the work to date has been performed for the March April May season, which appears to reflect a drying trend. To illustrate this, Figure 39 presents the average rainfall over March April and May in this box for a wide array of satellite and ground based datasets.

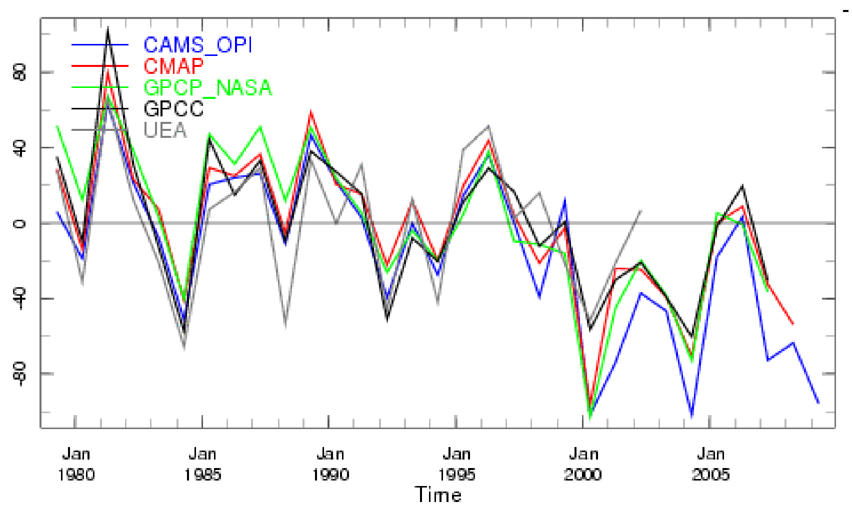

Figure 39 March April May Average Rainfall for several satellite and ground based datasets

In support of the index insurance project, these researchers have performed some preliminary analysis for the contract period.

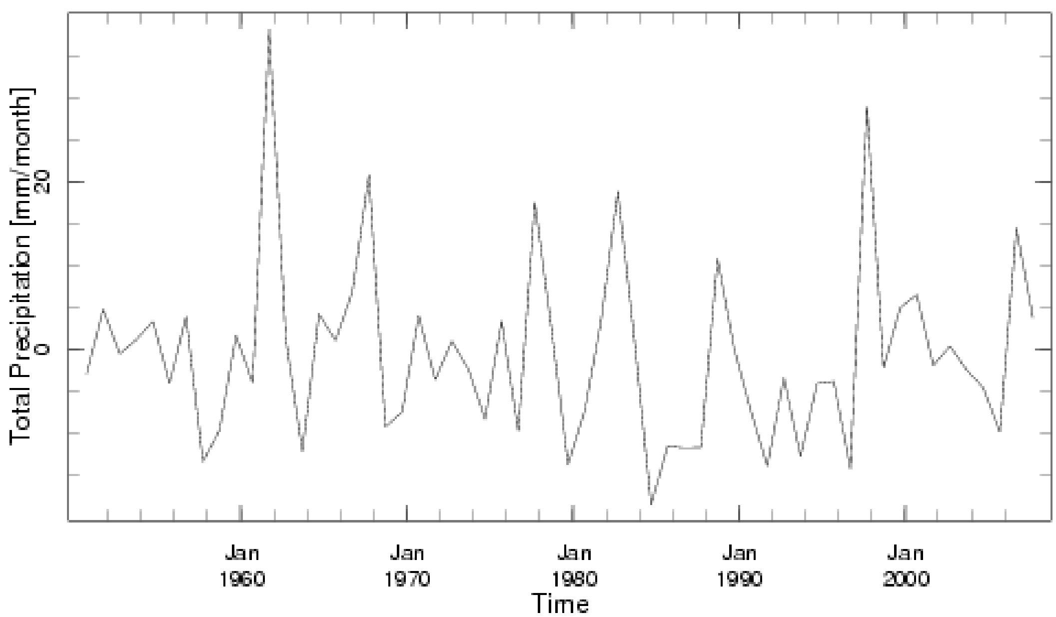

Figure 40 Aug-Sep-Oct rainfall anomaly for larger (red) boxed region

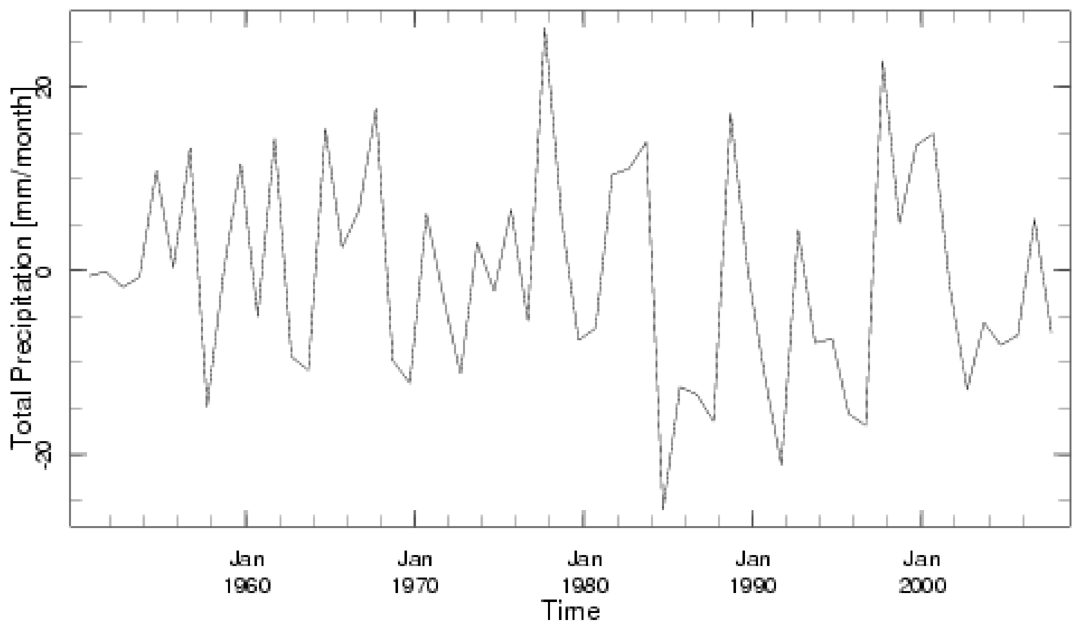

Figure 41 Aug-Sep-Oct rainfall departure from 1961-1990 average for box roughly covering Ethiopia (Green. 4N-14N, 35E-45E).

Figure 40 presents the precipitation average across the entire box for the Aug-Sep-Oct season, showing no evidence of a downward trend, and slight evidence of a modern increasing trend in rainfall. Figure 41 presents the averages from a box approximately cover Ethiopia. In the Ethiopia rainfall averages for the contract period, there is no evidence of a recent decreasing trend, and perhaps a trend towards increasing rainfall in recent years.

Although this analysis is very preliminary, it does suggest that a range of strategies may be very effective at reducing the additional cost of insurance due to trends observed in the data. One approach would be to utilize updated strategies for estimating the trend. **Another avenue that would likely reduce the premiums is to pool multiple sites together, as they would likely average out the potential trend observed in the Adi Ha data.**

Next steps

Given the possibility for large scale implementations due to the capacity of current delivery channels (currently serving millions of low income Ethiopians), and the potential for low cost contract design, there may be considerable potential for scaling. It will be important that growth from pilot be a healthily, but prudently scaling implementation with careful monitoring, and an increasing reliance of Ethiopian expertise as technologies, practices, and capacity are built.

The coming year of the project will likely include advances in the Adi Ha implementation, taking advantage of the experience of the past year to improve upon the current exploratory contract, design processes, and complimentary parts of the risk management package. It will also include continued design of a broader set of implementation strategies and solutions utilizing the broader set of scenarios reflected by the additional sites. It will be essential to build formal processes for locally led and locally designed product that can adapt each year to address new lessons learned and changing situations.

The project has attempted to forge some of the next generation of tools and approaches that can address the challenges of low income and data poor sites such as Adi Ha, tools that could be used in the future to build a new 'factory floor' that could assemble robust, cost effective, and tailored insurance products at large scales that are worthwhile components of complimentary development interventions. This would be necessary for further large scale up in future years.

We suggest some key next steps below concerning the concrete activities that IRI has been heavily involved in. These are selected both to address immediate issues in the Adi Ha implementation as well as to make further progress on next generation challenges listed in the introductory section.

- **Improvement of the prototype index and complimentary risk management package based on initial experiences.** Given that this is the first year of an exploratory product, it is likely that substantial improvements in the product and design process will be found in analysis and discussion of the initial year. In addition, the refinement of the risk management activities that compliment the index insurance will be important. Examples of potentially useful innovations that are becoming evident include explicit tuning of the rainfall cap to assure that dry spells are weighted appropriately in the index and the possibility of a community controlled risk/savings pool that could be utilized for idiosyncratic risks, or risks smaller than what the index would be a cost effective solution for.
- **Methodology formalizing the transition from indexing on remotely sensed rainfall estimations to the new met stations.** We have been developing statistical tools to complete this process. It will be important for us to work with partners to finalize those tools, validate their performance, and utilize them to transition between datasources through a transparent, consensus driven, and scalable process. This task relates to future work on rainfall simulation and modeling.
- **Solutions leading to reductions in insurance loading costs.** This will be driven by climate analysis and development of increasingly effective robust, validated, consensus statistical analysis tools, heavily leveraging future work on rainfall simulation and modeling as well as climate analysis.
- **Developing effective and robust strategies for quantifying and pricing climate trends.** As discussed in this report, we believe that through an improved understanding of climate trends it and associated improvements in insurance design, hedging, and pricing, it is likely that the impacts of climate trends on loading may be reduced. As it stands, strategies that were effective for the initial implementation may to lead to premium instability in the future without adequately reflecting true climate trends. We

believe that working with partners we should be able to arrive at meaningful improvements.

- **Completion of the vegetative remote sensing work**. It is necessary to perform additional validation and analysis in order to complete the initial exploratory vegetative remote sensing work. In the further future, it will be valuable to build upon this exploratory work to arrive at scalable, industrial strength vegetative sensing tools that can be used reliably in index insurance projects.
- **Leverage and catalyze improvements in remote sensing of rainfall**. We are involved in efforts to have Ethiopian researchers extend estimates of remotely sensed rainfall to thirty year histories. It is possible to explicitly tailor remote sensing estimates of rainfall to optimize their performance for specific insurance implementations utilizing existing capacity of Ethiopian scientists.
- **Healthy and prudent scaling**. The Adi Ha exploration should be followed with scaling to a broader range of crops, locations, livelihood portfolios, and contexts so that the necessary tools can be designed and validated for a more complete set of insurance implementation typologies.
- **Experimental Games**. It would be worthwhile to build upon the implementation of the initial pilot experimental games to address issues raised as well as additional design questions in a wider range of contexts.
- **Completing and extending exploratory research in rainfall simulation.** For the Adi Ha project, we have initiated research on several fronts. Much of this research is near completion and it we hope to finalize it in the near future. For example, we have developed and begun testing of advanced rainfall simulators that formally utilize information from multiple sources. These include a daily Bayesian rainfall model, hidden Markov (illustrated in Figure 42), and other models. If validated, these simulators may allow contracts to be optimized and tested through formal statistical inclusion of a range of different types of information, including forecasts, sea surface temperatures, vegetation remote sensing, multiple nearby stations, farmer measured rainfall, alternate sources of satellite rainfall, and tree ring data. The successful performance of these simulators could allow the development of contracts that are conditional on seasonal forecasts, and offer formal statistical tools to transition from remote sensing to contracts on newly installed met stations. These tools could formally incorporate warning signs from long historical records or nearby stations that we have currently been investigating manually, through handcrafted analysis by experts. It will be valuable for us to complete and publish this research, as well as investigate avenues to build the research models into tools that could be harnessed in more applied design and scale up situations.
- **Completing and extending climate science research.** With partners, we have been performing exploratory climate research on long term trends (discussed in the section on climate trends), onset and cessation of rainfall in Ethiopia (eg Moron and Robertson, 2008), and on seasonal forecasting in Ethiopia. The continuation of this research will be extremely important for informing index insurance projects.
- **Completing and extending economic research**. We have been developing research on economic issues concerning the role of insurance in credit, livelihood improvement, and risk management. It will be valuable for us to complete this research and utilize it to better inform insurance project decisionmaking.
- **Monitoring and Evaluation** will be critical, for understanding project performance, design trade-offs, and signaling when approaches are inappropriate. It is essential that monitoring be utilized to improve the exploratory products over time.
- **Capacity building and tools to allow local designers to lead the insurance design process**. We have begun developing processes (and relevant training and design software tools) to allow designers to systematically identify and validate effective index insurance strategies to support development interventions. It would be worthwhile to work with Ethiopian partners to develop formal processes, communication networks with national and international researchers and software education and design tools. Figure 43 illustrates an educational software tool on index insurance design that we have

developed in partnership with the World Bank CRMG that could be tailored (or expanded) for use in Ethiopia.

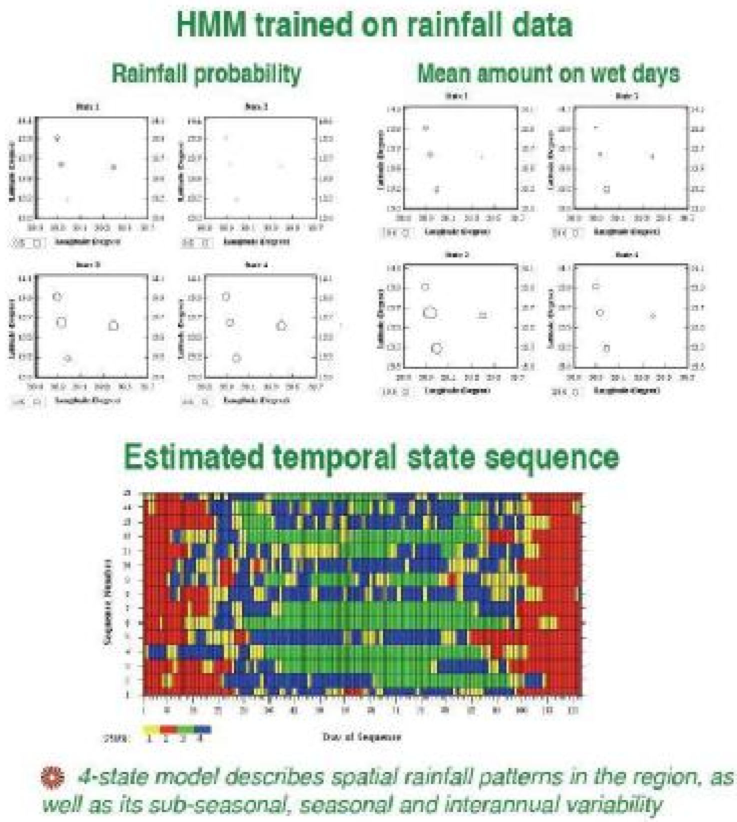

Figure 42 Initial Explorations for Rainfall Simulation using HMM models

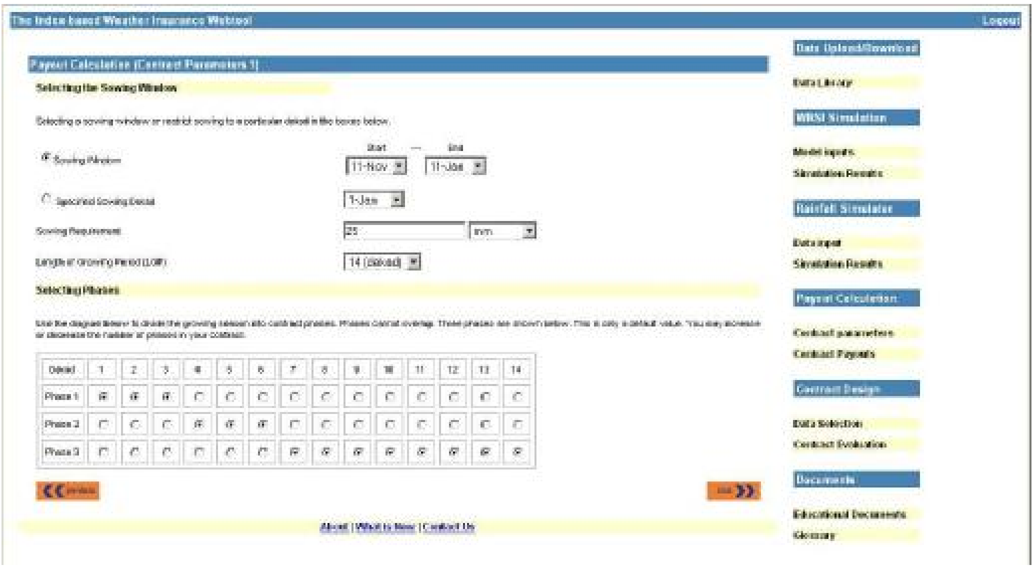

Figure 43 Software tools for index insurance

Bibliography

Osgood, D.E, McLaurin, M. Carriquiry, M., Mishra, A., Fiondella, F., Hansen, J., Peterson, N., and Ward,N. (2007). Designing Weather Insurance Contracts for Farmers in Malawi, Tanzania, and Kenya, Final Report to the Commodity Risk Management Group, ARD, World Bank. International Research Institute for Climate and Society (IRI), Columbia University, New York, USA.

Peterson, N. and C. Mullally. (2009) Index Insurance games in Adi Ha Tabia, Tigray Regional State, Ethiopia. Report prepared for Oxfam America.

Block, P., Dinku, T., McLaurin, M., Osgood, D. (2008). Designing a Weather Insurance Contract for Farmers in Adi Ha, Ethiopia Report to Oxfam America Precipitation-Data Quality Assessment and Verification.

Carter, M. (2009). Using Satellite Imagery as the Basis for Index Insurance Contracts in West Africa. BASIS CRSP memo to Rebecca Nelson, Director, CCRP, Cornell University.

Ceccato P., Brown M., Funk C., Small C., Holthaus E., Siebert A. and Ward N. (2008). Remote Sensing Vegetation. Paper presented at a workshop on 'Technical Issues in Index Insurance', held 7–8 October 2008 at IRI, Columbia University, New York. Available at http://iri.columbia.edu/csp/issue2/workshop

Dinku T., Funk C. and Grimes D. (2008). The Potential of Satellite Rainfall Estimates for Index Insurance. Paper presented at a workshop on 'Technical Issues in Index Insurance', held 7–8 October 2008 at IRI, Columbia University, New York. Available at http://iri.columbia.edu/csp/issue2/workshop

Funk C., Sanay G., Asfaw A., Korecha D., Choularton R., Verdin J., Eilerts G. and Michaelsen J. (2005) Recent Drought Tendencies in Ethiopia and Equatorial-Subtropical Eastern Africa. Famine Early Warning System Network, USAID, Washington, DC.

Giannini, A., Hanson, J., Holthaus, E., Ines, A., Karnauskas, K., McLaurin, M., Osgood, D., Shirley, K. and Vicarelli, M. (2009). Designing Index-Based Weather Insurance for Farmers in Central America Final Report to the World Bank Commodity Risk Management Group, ARD, World Bank. International Research Institute for Climate and Society Earth Institute, Columbia University.

Greene, A.M., Goddard, L., Ward, N., Siebert, A., Holthaus, E., Hellmuth, M., and Baethgen, W. (2008) Climate change, one decade at a time. Paper presented at a workshop on 'Technical Issues in Index Insurance', held 7–8 October 2008 at IRI, Columbia University, New York. Available at http://iri.columbia.edu/csp/issue2/workshop

Hellmuth M.E., Osgood D.E., Hess U., Moorhead A. and Bhojwani H. (eds) (2009). Index insurance and climate risk: Prospects for development and disaster management. Climate and Society No. 2. International Research Institute for Climate and Society (IRI), Columbia University, New York, USA.

Holthaus, E. M., Siebert A., Ward M.N, Baethgen W., Brown M., Osgood D., Indeje, M (in preparation). Exploring practical climate impact index development to support drought risk management across 10 agroecological zones in Africa. International Research Institute for Climate and Society (IRI), Columbia University, New York, USA.

Kassahun, D., Assefa, D. and Gebrekidan, D. (2009). Teff (*Eragrostis tef*) agronomic research in Adi Ha, Central Tigray, Ethiopia. Oxfam America Horn of Africa Risk Transfer for Adaptation (HARITA) Progress Report no.01.

Moron and Robertson (2008). Preliminary analysis of the daily network of Ethiopia (1975-2004). Research note. International Research Institute for Climate and Society (IRI), Columbia University, New York, USA.

Mengistu, D.K. (2009). The influence of soil water deficit imposed during various developmental phases on physiological processes of tef Agriculture, Ecosystems & Environment Volume 132, Issues 3-4, August 2009, Pages 283-289 http://www.sciencedirect.com/science?_ob=ArticleURL&_udi=B6T3Y-4W9S2N1-5&_user=18704&_coverDate=08%2F31%2F2009&_alid=958383849&_rdoc=1&_fmt=high&_orig=search&_cdi=4959&_sort=r&_st=4&_docanchor=&_ct=6&_acct=C000002018&_version=1&_urlVersion=0&_userid=18704&md5=2206f5c31b9198d83d0853087b76b777
doi:10.1016/j.agee.2009.04.013

Mullally, C. (2009) A Note on Risk Reduction and Drought Insurance in Ethiopia.
Peterson, N. and C. Mullally. (2009) Index Insurance games in Adi Ha Tabia, Tigray Regional State, Ethiopia. Report prepared for Oxfam America.

Osgood, D.E, McLaurin, M. Carriquiry, M., Mishra, A., Fiondella, F., Hansen, J., Peterson, N., and Ward,N. (2007). Designing Weather Insurance Contracts for Farmers in Malawi, Tanzania, and Kenya, FinalReport to the Commodity Risk Management Group, ARD, World Bank. International Research Institute for Climate and Society (IRI), Columbia University, New York, USA.

Robertson, A. W.; Baethgen, W. E.; Block, P.; Dinku, T.; McLaurin, M.; Osgood, D.; Shirley, K.; Ward, M.N. (2008) Modeling weather-within-climate for index insurance contract design. Poster presented at the Eighth EMS Annual Meeting and the Seventh European Conference on Applied Climatology (ECAC), 29 September- 03 October 2008, Amsterdam, Netherlands.

Shirley, Kenny (2008). Rainfall Modeling and Simulation. Paper presented at a workshop on 'Technical Issues in Index Insurance', held 7–8 October 2008 at IRI, Columbia University, New York. Available at http://iri.columbia.edu/csp/issue2/workshop

Teshome, W., Peterson, N., Gebrekirstos, A., Muniappan, K. February 28, 2008. "Microinsurance Demand Assessment in Adi Ha Tabia, Tigray Regional State, Ethiopia." Working paper, Oxfam America.