State of Knowledge Report — Data Requirements for the Design of Weather Index Insurance[1]

Innovation in Catastrophic Weather Insurance to Improve the Livelihoods of Rural Households

June, 2010

GlobalAgRisk, Inc.
1008 S. Broadway
Lexington, KY 40504
859.489.6203
Jerry R. Skees, President

[1] This report is based on research funded by the Bill & Melinda Gates Foundation. The findings and conclusions contained within are those of the authors and do not necessarily reflect positions or policies of the Bill & Melinda Gates Foundation.

Acknowledgements

This report was prepared by Benjamin Collier, Barry Barnett, and Jerry Skees of GlobalAgRisk, Inc. Collier is a research analyst at GlobalAgRisk and a doctoral student in the department of agricultural economics at the University of Kentucky. Barnett is president of Acacia Economic Consulting, LLC, and professor of agricultural economics at Mississippi State University. Skees is president of GlobalAgRisk, Inc., and H.B. Price Professor of Agricultural Policy and Risk in the department of agricultural economics at the University of Kentucky. Appendix B was developed by Mario Miranda, Andersons Professor of Finance and Risk Management in the department of agricultural, environmental, and development economics at The Ohio State University. Appendix C was developed by Upmanu Lall, Alan and Carol Silberstein Professor of Engineering in the department of earth and environmental engineering at Columbia University.

The authors gratefully acknowledge assistance and comments from Grant Cavanaugh, Marcia Rapchak, Michele Vollaro, Jock Anderson, Nadezda Nikolova, Jason Hartell, Anne G. Murphy, and two anonymous reviewers from the reinsurance industry. Editorial assistance was provided by Celeste Sullivan.

Table of Contents

Figures

Tables

Boxes

Acronyms and Abbreviations

ARMA	*Autoregressive Moving Average*
AVHRR	*Advanced Very High Resolution Radiometer*
CAPM	*Capital Asset Pricing Model*
CFSRR	*Climate Forecast System Reanalysis and Reforecast*
CHARM	*Collaborative Historical African Rainfall Model*
CMAP	*CPC Merged Analysis of Precipitation*
CMORPH	*CPC Morphing Technique*
CPC	*Climate Prediction Center*
CV	*Coefficient of Variation*
DRR	*Disaster Risk Reduction Programme*
ECMWF	*European Centre for Medium-Range Weather Forecasts*
ENSO	*El Niño Southern Oscillation*
FEWS NET	*Famine Early Warning Systems Network (USAID)*
GCM	*Global Circulation Model*
GRP	*Group Risk Plan*
IBLI	*Index-based Livestock Insurance (Mongolia)*
MODIS	*Moderate Resolution Imaging Spectroradiometer*
NAO	*North Atlantic Oscillation*
NASA	*United States National Aeronautics and Space Administration*
NCDC	*United States National Climatic Data Center*
NCEP/NCAR	*National Centers for Environmental/National Center for Atmospheric Research*
NDVI	*Normalized Difference Vegetation Index*
NMHS	*National Meteorological and Hydrological Services*
NOAA	*United States National Oceanic and Atmospheric Administration*
PERSIANN	*Precipitation Estimation from Remotely Sensed Information using Artificial Neural Networks*

RFE	*Rainfall Estimate*
SAR	*Synthetic Aperture Radar*
SD	*Standard Deviation*
SKR	*State of Knowledge Report*
SST	*Sea Surface Temperature*
TRMM	*Tropical Rainfall Measuring Mission*
USAID	*United States Agency for International Development*
UCAR	*University Corporation of Atmospheric Research*
WMO	*World Meteorological Organization*

Executive Summary

In recent years, weather index insurance has gained significant international attention. Multilateral agencies and donors are supporting the development of index insurance products and practitioners are venturing into this work for the first time. However, designing index insurance products is quite challenging and requires strong analytics, significant research, access to sufficient quantities of relevant data, capacity building among local stakeholders, legal and regulatory expertise, etc. Based on experiences to date, some questioning has begun about the scalability and sustainability of weather index insurance — in some cases, these questions are motivated by beliefs that data limitations may cause significant geographic constraints to offering weather index insurance, and in other cases, that the potential of current index insurance models on poverty reduction is doubtful. Thus, the theory and implementation of index insurance is at somewhat of a crossroads. In recognition of these emerging questions, this is the first of several documents that describes the current state of knowledge on key aspects of index insurance. This State of Knowledge Report (SKR) is on the data component of developing weather index insurance.

We originally envisioned this SKR as a review of the quantity and quality of data needed to support weather index insurance offerings — a general guide for practitioners and those in the development community. Yet, as we considered basic questions about how much data are required or whether the data quality is sufficient, the answers were almost universally, “It depends.” In assessing risk, data needs are always contextual. One simply cannot address questions about data in isolation from broader questions about the type of weather index insurance product being developed, its target market, and its application. Based on this consideration, we approach our analysis of data requirements by looking at index insurance designs that tend to be most robust for very limited data. In other words, what models are most scalable and sustainable in terms of data limitations? Another standard for scalability and sustainability is in regard to economic development — how does the insurance product contribute to poverty reduction?

Our conclusions for overcoming data constraints and contributing to poverty reduction converge to three recommendations. From our analyses and field experience implementing index insurance projects, we conclude that weather index insurance programs should focus on: 1) **consequential losses** from extreme weather events that extend beyond crop yields; 2) **catastrophic losses** rather than moderate losses; and 3) **risk aggregator products** instead of, or in addition to, household products.

This document is grounded in both our academic economic research and experiences developing and implementing index insurance programs in lower income countries. Our current interest in weather index insurance has its roots in analytical work by Skees and Barnett in the 1980s and 1990s in regard to the U.S. Federal Crop Insurance Program. Along with J. Roy Black, Skees and Barnett developed the first agricultural index insurance product in the United States — an area-yield insurance product called the Group Risk Plan (GRP) that is still offered today. Based on our experience in developing GRP, our interest turned to how index insurance could be used to protect against agricultural losses in lower income countries where traditional loss-based crop insurance was not feasible. Nonetheless, area-yield data are quite sparse in lower income countries but weather data, at least in some countries, are available. Peter Hazell (then working for the International Food Policy Research Institute) had considered this possibility and worked with us to introduce these ideas to the World Bank via a project on weather index insurance in Nicaragua, in 1998. Since 2001, GlobalAgRisk has been developing and implementing index insurance programs in lower income countries. Currently, GlobalAgRisk has projects in Mongolia, Peru, and Vietnam, and examples from our experiences are

used throughout the SKR. As academics, we also frequently write articles on index insurance for scholarly journals — sharing our evolving understanding of these products and encouraging the research efforts of our colleagues to further this work.

We share our background to demonstrate our extensive experience with weather index insurance and our long-term commitment to the fundamental principles on which index insurance is based. This journey of discovery has been an iterative process through which theory and practice synergize to advance our understanding of index insurance product development. Developing this SKR is another step in that process. Writing this document has allowed us to consider our recent experiences, lessons learned from other index insurance programs, premises from economic theory and research, and the reality of data constraints to advance and formalize new thinking on weather index insurance. We freely admit that, as a result, we now critique some of the ideas and methods that we once helped develop. Such is the evolution of knowledge.

Chapter 1 provides introductory material on weather index insurance. We remind the reader that evaluating weather risks is a very different process than many scientific endeavors undertaken by economists. Whereas much economic research discounts outliers, extreme values are the most important observations for weather risk analysis. Thus, when sparse data suggest that an event may be an outlier, insurance underwriters will typically use any information they can find to learn more about that event. Even so, the available data typically consist of small samples, which can cause large estimation errors. For this reason, insurance underwriters try to understand more than simple statistical relationships. They may work with scientists who understand the underlying physical processes of weather to evaluate patterns and any potential non-stationarity of data.

In Chapter 1, we also present a conceptual model of how insurance can facilitate poverty reduction. We review why traditional loss-based insurance is infeasible in rural areas of many lower income countries and present weather index insurance as a potentially viable alternative. The chapter also develops a conceptual framework for evaluating potential index insurance contracts under idealized conditions — with sufficient access to relevant, high-quality data. In actuality, data constraints create a contrast between how index insurance contracts are assessed in theory and how they are assessed in practice. Basis risk, an inherent constraint of index insurance, is discussed in detail, and we critique some of the methods practitioners use to demonstrate that they have reduced the basis risk associated with weather index insurance products.

Chapter 2 transitions from the conceptual model presented in the introduction to real-world data constraints facing practitioners developing index insurance products. In particular, Chapter 2 describes how, in a data-constrained environment, one can use qualitative information to determine relationships between potential weather indexes and realized losses of potential insureds. We argue that to gain an understanding of a potential for index insurance products, qualitative data obtained from scientifically based risk assessments with participating key stakeholders may provide superior insights to sparse quantitative data.

Chapter 3 describes data needs for weather index insurance. Historical data are needed to evaluate and price the risk, and real-time data are needed to make payments. It is here that the contextual nature of data requirements becomes most apparent. Rather than proffering absolute minimum data requirements, Chapter 3 describes how various contextual elements (e.g., spatial and temporal presentation of the weather risk; perceptions regarding the validity, security, and credibility of the data source; access to alternative measures of the underlying weather phenomenon; and sophistication of the target market) can affect data requirements.

Chapter 4 examines the real data constraints associated with using weather stations for index insurance. Scaling up their current efforts for weather index insurance, many of the poorest countries lack or are unable to maintain this infrastructure. We review these issues using data from Africa in addition to cost data obtained from providers of weather measurement instruments. We demonstrate why these data constraints are less binding for risk aggregator products. Next, Chapter 4 develops our assessment of the current state of satellite-based technologies. While we are enthusiastic about the promise of these technologies, we conclude that, with the exception of some limited applications for pastoral settings using the Normalized Difference Vegetation Index (NDVI), further developments are needed before satellite-based data sources can support weather index insurance offers. We predict that major breakthroughs are likely once combined information sources used to develop index models begin to demonstrate their value in developing insurance products to protect against catastrophic risk.

Chapter 5 presents important lessons derived from preceding chapters and from experiences to date developing and implementing weather index insurance in rural areas of lower income countries. We begin by evaluating selected aspects of some current index insurance programs. We then present our recommendations for advancing weather index insurance given data limitations. As mentioned previously, those recommendations are: 1) expand the focus beyond just crop yields to consequential losses; 2) transfer catastrophic rather than moderate losses; and 3) initially target risk aggregators rather than, or in addition to, households.

Chapter 6 presents a number of outstanding research questions related to key challenges to demand for weather index insurance. We also develop some research questions on both data availability and how data can best be used in developing weather index insurance products. This research agenda is supported by two technical appendixes. Appendix B, contributed by Dr. Mario Miranda, develops a research agenda to test the properties of extreme events relative to moderate weather events. For some time, we have hypothesized that the covariance of weather events and losses is likely not linear throughout the distribution. Given the correlated nature of weather risk we believe that when extreme events occur, classic diversification strategies will break down as the strategies used (e.g., having a number of farm enterprises) will all suffer losses at the same time. In short, we will test if the variance-covariance matrix changes given extreme weather events. If the answer is "yes," as we suspect, this research agenda will add rigor to our recommendations on catastrophic insurance to cover consequential losses. Dr. Upmanu Lall contributes Appendix C, which reviews the potential for developing weather index insurance based on global teleconnections using sea surface temperature (SST) measures to learn if we can replicate our work on "forecast insurance" in Peru in other areas. Appendix A details our use of an SST measure as the basis for an index insurance product that protects against extreme flooding in the northern coastal region of Perú.

Thanks to our grant from the BMGF, we will continue working on research topics described in this document. We view this SKR as a work-in-progress. We welcome comments and critiques of the ideas presented. Our second SKR focuses on legal and regulatory component of developing weather index insurance, in particular, the legal and regulatory challenges of creating index insurance products designed to protect against consequential losses. We have experience designing consequential loss index products in two very different jurisdictions: Peru and Vietnam. The third SKR focuses on evaluating the scalability and sustainability of index insurance products. The fourth, final SKR is a synthesis of the previous SKRs, with a more thorough discussion of index insurance as a tool for reducing poverty and supporting rural economic development. In each case, our research and inquiry is informed by economic theory and empirical analysis that includes our field experiences in Peru, Mongolia, and Vietnam, and what we have learned from others' experiences.

Chapter 1 Introduction and Conceptual Model

This document presents our analysis of the current state of knowledge regarding data systems to support weather index insurance. Part of that analysis identifies the limits of current theoretical understanding and empirical practice in a developing country context. Additional work is needed to make weather index insurance more effective, efficient, sustainable, and scalable. This state of knowledge report (SKR) does not address all the challenges and constraints to developing index insurance markets. Our goal is to provide a better understanding of how index insurance can be developed within existing data constraints with a long-term perspective that identifies important research needs.

We assume that readers of this SKR have at least a basic understanding of index insurance and its advantages and disadvantages relative to other forms of insurance. Nonetheless, readers may want to refer to more detailed documents to gain a deeper understanding (e.g., Skees et al., 2007; Barnett and Mahul, 2007; or GlobalAgRisk, 2006, 2009). The focus of this SKR is market-based products — products sold to individuals or firms operating in the private sector. Still, many of the key data issues discussed in this document are also relevant to products for non-market institutions such as governments and donors. This introductory chapter frames elements important to the body of the SKR.

Designing index insurance in lower income countries involves expert judgment in an environment of significant data constraints. Practitioners are increasingly recognizing the importance of the contextual nature of weather risk and are adapting their methodologies for evaluating data.[2] For example, reinsurers are moving away from set, arbitrary data standards — requirements for at least 30 years of data, or an insured must be within 20 km of a weather station, etc. — to more analytic, context-specific approaches that consider how a variety of sources can be used to evaluate weather risk. Thus, while we hope that this document is straightforward and pragmatic, it will not be a checklist of requirements for data systems that support index insurance. Instead, our intention is to discuss key issues that will help practitioners understand, identify, and address data constraints and also motivate additional important research that may lead to innovation in developing appropriate weather index insurance products.

1.1 Methodologies for Insurance Analysis

Our approach to analyzing data issues is, in many ways, similar to that of an insurance underwriter. Readers with scientific or statistical backgrounds will be familiar with many of the techniques described in this document. For example, many scientific disciplines are interested in modeling relationships between variables. Thus, the coefficient in a regression analysis is interpreted as, "A one unit change in the independent variable results, *on average*, in an *X* unit change in the dependent variable." As with many statistical analyses, the emphasis is on the average or *central tendency*. In contrast, insurance underwriters are generally more concerned with what happens in the extremes. They want to know what variables cause extreme losses for insureds (e.g., drought for insured crop farmers), the probability that these extreme events will occur, and the relationship between the extreme events and insured losses.

[2] In this document, "practitioner" refers to those individuals developing and implementing an index insurance program such as insurers, reinsurers, international development agencies and firms, etc.

Insurance underwriters evaluate risk using both statistics and some knowledge of the underlying physical processes that cause extreme events. Therefore, insurance practitioners can benefit from scientific findings of other disciplines (e.g., plant growth analysis, climate circulation models, household livelihood survey data, etc.) even though the objectives of insurance practitioners may be quite different from those of the scientists collecting and reporting raw data. For example, extreme events that scientists in other disciplines may commonly omit as outliers tend to be among the most important sources of data for insurance practitioners (see Collier, Skees, and Barnett (2009) for further discussion of central tendency, outliers, and insurance). Likewise, many statistical procedures are helpful for understanding relationships around the most frequently occurring outcomes, but are not very helpful for describing the extremes. Relationships between variables may be quite different under extreme conditions than under typical conditions. Thus, different statistical methodologies are also needed for evaluating risk than those typically used in many scientific disciplines. Recognizing the distinction between insurance underwriting and typical scientific methods is important for scientists and other practitioners who venture into developing index insurance products as the misapplication of scientific findings and statistical procedures can lead to inaccurate conclusions about the risk.

A specific problem for insurance practitioners is that extreme events tend to occur very infrequently. Fewer data observations of the event mean that 1) the underlying scientific processes tend to be less well-understood; and 2) statistical results are much less precise. Given the lack of data for extreme events, insurance underwriters turn to numerous other sources of information when evaluating insurable risks. For example, to understand extreme rainfall patterns, insurers may compare data from weather stations, satellites, airplanes, weather balloons, tree-ring analysis, lake-bed analysis, etc. Different data sources tend to agree around the most frequently occurring values, but may diverge greatly at the extremes, and insurance practitioners must determine how much weight to give each of these sources when estimating risk. Thus, insurance analysis must rely on both strong scientific and statistical analysis as well as expert judgment.

1.2 Two Types of Indexes

Data supporting index insurance products can be classified in two broad categories: indexes that aggregate losses over a group and weather-based indexes. Aggregate loss data describe losses across many individuals, typically in the same geographic region. The index of group losses serves as a proxy for the losses of individual members of the group. The Group Risk Plan (GRP) in the United States and the Index-based Livestock Insurance (IBLI) Program in Mongolia are examples of index insurance programs using aggregate loss measures. The GRP uses county-yield data for specific crops as the index for determining compensation (Skees, Black, and Barnett, 1997). The Mongolia IBLI uses government-developed estimates of mortality by species for concentrated geographical areas as the index for indemnities (Mahul and Skees, 2007). An aggregate loss index insurance contract can be considered as a type of valued policy.[3] With these products, the aggregate data are on a large enough scale to reduce the likelihood that any individual insured can significantly influence an indemnity. Thus, these products also have lower moral hazard and adverse selection than traditional insurance products.

[3] For a thorough discussion of the insurance classification of index insurance contracts please see GlobalAgRisk. 2010.

Weather-based indexes use measurements of weather events highly correlated with losses of the insured as the basis for an insurance payment. The objective of the index is not to serve as a direct proxy for loss, but rather as a predictor or proxy for the insured event itself, e.g., flood or drought. Weather-based index insurance can be likened to contingency insurance in that a specific event (e.g., death, loss of leg, etc.) can trigger an insurance payment. A commonly used weather-based index is rainfall data from local weather stations; however, other measures serve as weather-based indexes, as well., For example, the Normalized Difference Vegetation Index (NDVI) is a measurement of vegetation density and has been used to provide index-based drought insurance (Box 8 in Chapter 4). Another index insurance product uses sea surface temperatures (SST) as a predictor of extreme flooding in northern Peru. The SSTs used are indicative of extreme El Niño events, the primary cause of catastrophic flooding in that region (Appendix A).

Both types of indexes have their relative merits and shortcomings. Aggregate loss indexes are generally easier to develop and scale up than weather-based indexes. However, in lower income countries, weather data tend to be more readily available than aggregate loss data. Also, weather data are often easier to collect and may be less prone to tampering than, for example, subregional yield data, which have sometimes been adjusted to support political agendas. While aggregate loss indexes may be feasible in some regions of the world (Carter, Galarza, and Boucher, 2007), this SKR focuses on data issues related to weather-based indexes.[4]

To date, many weather index insurance pilots have been based on weather station data. For some of our examples, we use weather stations as a helpful point of reference for readers, but we want to emphasize that in most regions of the developing world, weather station infrastructure is insufficient to support index insurance. Thus, practitioners need a broader vision for what data can be used to support index insurance for many regions, including satellite and other forms of remotely sensed data. We develop this vision as the document progresses and describe alternative data systems.

1.3 Two General Types of Products

Two general classes of products have been developed for market-based index insurance programs: those for households and those for risk aggregators. Household index insurance products have often been designed with the intention to protect against crop-yield losses due to adverse weather risk; however, other designs are also possible such as contracts that protect a household's livelihood portfolio more generally from a specific, severe weather risk. *Risk aggregator* refers to firms such as lenders and agricultural value chain members who are negatively affected by the correlated production risks in a geographic region. For example, given the correlated nature of drought risk, lenders are affected by the drought exposure of their agricultural borrowers. If a drought occurs, many borrowers are likely to experience repayment difficulties concurrently. Products designed to protect these risk aggregators are intended to protect the solvency of the firms and improve access to their services. As we discuss below, the target market for the index insurance product has significant implications for which data sources can potentially be used to support the insurance offer.

[4] Unless otherwise stated in the SKR, we use the terms “index insurance” and “weather index insurance” synonymously to refer to insurance products using weather-based indexes.

1.4 Background and Historical Context: Practice and Theory

In the 1980s, the development community essentially gave up on crop insurance being workable in developing countries (Hazell, Pomareda, and Valdes, 1986). Given the characteristics of agriculture in developing countries (small farm sizes, intercropping, etc.), developed country models of crop insurance were simply not working in developing countries. The potential use of index insurance in developing countries started to gain attention in the late 1990s (Skees, Hazell, and Miranda, 1999). The history of these developments is available from Skees (2008). Much of the literature and experience since then has focused on using index insurance to replace crop insurance with an underlying belief that data would be available to develop effective indexes. The data requirements for index insurance tend to be less than those for traditional crop insurance. Copious data are required for underwriting and rating traditional crop insurance. Additionally, while the contract is in force, it may be necessary to send personnel into the field to collect data on the activities of insureds and/or data for settling claims. In contrast, weather index insurance uses established weather information systems to provide all the data required for underwriting, rating, and settling claims, and there is no need to monitor the activities of individual insureds.

Most of the scholarly literature on weather index insurance is based on a traditional agricultural insurance framework because weather index insurance was developed to overcome problems with traditional crop insurance in developed countries (e.g., Skees, Hazell, and Miranda, 1999). As with the traditional crop insurance literature, most weather index insurance studies focus on smoothing income from a single crop in a single year (Wright, 2006). This framework does not fit for developing countries where small households are engaged in a host of livelihood strategies including farming. Instead, much of what is developed in this document is framed with a broader consideration of households protecting their wealth positions over time. By framing the insurance decision as a portfolio problem for small households striving to protect wealth over time, the focus of linking index insurance to a single crop within a single year becomes less important. Instead, the focus is on protecting household wealth from catastrophic events that have multiple consequences.

The risk of catastrophic events can also cause risk aggregators to ration the services they provide to poor households. Index insurance products targeted to risk aggregators can increase the likelihood that service providers throughout the value chain will provide poor households with access to their services. In fact, if the objective is to improve the lives of the working poor, products targeted to risk aggregators may be the most effective place to start. As will be explained later in this document, data constraints are less problematic for index insurance products that are targeted to risk aggregators than for those that are targeted directly to households.

1.5 Conceptual Framework

This chapter describes the conceptual underpinnings of index insurance. Deconstructing index insurance allows one to identify critical points where data limitations create challenges for successful product development. Chapter 2 builds on this background by discussing strategies for addressing these data challenges.

1.5.1 Why Insurance?

Consider a decision maker (a household or a business) with initial wealth W_0 who invests an amount K $(K \leq W_0)$ in various activities which are expected to generate positive net returns so that wealth will increase over time. Each activity $A_i (i = 1, 2, \ldots n)$ generates a periodic expected net rate of return $E(r_i)$. In any given period, the realized net rate of return r_i may differ from $E(r_i)$ due to a number of factors including, but not limited to, variability in weather conditions. For example, one of the activities may be production of a particular crop for which realized yield depends critically on adequate rainfall.

The realized periodic net return on the entire portfolio of activities is $R = \sum_{i=1}^{n} \tau_i K r_i$ where τ_i is the proportion of K that is invested in activity A_i and $\sum_{i=1}^{n} \tau_i = 1$. The variance of net returns for the portfolio is calculated as $\sigma_R^2 = \sum_{i=1}^{n} \sum_{k=1}^{n} \tau_i \tau_k \sigma_{ik}^2$ where σ_{ik}^2 is the variance in returns on the single activity when $i = k$ and the pairwise covariance in returns between activities when $i \neq k$ .Thus, the overall variability in net return for the portfolio of activities depends on the variance in net return for each of the activities, the proportion of the overall portfolio that is invested in each activity, and the *covariances in net returns across the different activities.* As long as returns are not perfectly, positively correlated, engaging in more than one activity provides some diversification benefit, reducing the variability in the net returns for the overall portfolio — i.e., because the activities are not perfectly correlated, the variance of the portfolio is less than the weighted average of the variances for each of the activities. Negative or small positive covariances provide significant diversification benefits, while large positive covariances create little diversification benefit. A fundamental challenge for rural areas in many lower income countries is that many of the available wealth-generating activities are susceptible to the same extreme weather events. Thus, increasing the number of activities in the portfolio may provide little protection against very low portfolio net returns when extreme weather occurs.

The impact of extreme weather events is not limited to high variability in single period net returns. Extreme weather also has long-term impacts on wealth. Consider a simple two-period model for the evolution of wealth

$$W_1 = W_0 - K + (K + R) - AL$$

where the decision maker's ending wealth W_1 is equal to W_0 minus the level of investment K plus the realized return on that investment ($K + R$, where R may be positive or negative) minus any asset losses AL that occur during the period. Simplifying this expression yields

$$W_1 = W_0 + R - AL$$

The occurrence of an extreme weather event may affect ending wealth W_1 by reducing the realized return R and/or causing losses to assets such as buildings or livestock (i.e., $AL > 0$). This evolution of wealth model could be generalized from two periods to an infinite number of periods. Note that if $W_1 < W_0$, the decision maker will have less to invest in subsequent periods, reducing the *level* of wealth in future periods. Extreme events may also reduce the *growth rate* of wealth. If households must reduce consumption or sell livelihood assets to cope with the losses of a catastrophic event, it may reduce their expected returns in future periods (Barnett, Barrett, and Skees, 2008).[5] Thus extreme weather events affect not only single period net returns *but also the long-run accumulation of wealth*.

Risk-averse decision makers have utility functions that are increasing in expected wealth and decreasing in risk (variability in wealth).

That is,

$$U = f\left(E(W), \sigma_W^2\right) \text{ with } \frac{\partial U}{\partial E(W)} > 0 \text{ and } \frac{\partial U}{\partial \sigma_W^2} < 0 .$$

As risk-averse decision makers recognize their vulnerability to extreme weather events, they are likely to allocate their portfolios to activities that are less susceptible to extreme weather (e.g., plant cassava instead of maize). But since lower risk activities generally offer lower expected rates of return, this decision also affects the long-term trajectory of wealth accumulation. Thus, exposure to extreme weather events reduces wealth accumulation whether directly by destroying assets and reducing net returns in the period of the shock or indirectly by encouraging low-risk, low-return portfolio allocations across all periods.

Insurance purchasing can be conceived as another activity in the decision maker's portfolio. Insurance purchasing reduces expected wealth $E(W)$ because the premium paid by the insured must exceed the expected indemnity to compensate the insurer for taking on the risk. However, insurance purchasing also reduces the variance in wealth σ_W^2 because, the decision maker receives an indemnity only when R is lower than expected and/or when asset losses AL occur. Risk-averse decision makers will purchase insurance only if the utility gained from reducing the variability in wealth exceeds the utility lost from having a lower expected wealth.

Consider an example where we assume that a decision maker manages a portfolio that consists of only one activity — crop production. Also, assume that there are *n+1* possible weather outcomes — *n* types of bad weather, each occurring with some probability π_i plus the possibility of good weather occurring with probability $\left(1 - \sum_{i=1}^{n} \pi_i\right)$.

These assumptions simplify the presentation without loss of generality. Also assume that with good weather there are no yield losses, the realized return is R_g and no asset losses occur. Bad weather events can cause crop-yield losses YL_i and asset losses AL_i . When bad weather occurs, returns are presented as the return on investment in good years minus yield losses ($R_i = R_g - YL_i$).

[5] We could model this change in the growth rate using a household production function that changes functional forms depending on the levels of physical and human capital available to the household. Intuitively, if a household sells a productive asset such as a plow or livestock, it is likely to reduce their farming productivity. As another example, if households must reduce their caloric consumption, it can affect both the physical and intellectual development of children in the household, reducing the future labor productivity of the household (Grantham-McGregor et al., 2007).

Without insurance purchasing, the decision maker's expected utility is

$$E(U)_{NI} = \left(1 - \sum_{i=1}^{n} \pi_i\right) U\left(W_0 + R_g\right) + \sum_{i=1}^{n} \pi_i U\left(W_0 + R_g - YL_i - AL_i\right)$$

The decision maker can also purchase an insurance policy at a premium rate *p* that provides a sum insured (maximum possible indemnity) I. Thus, the premium cost is pI. With insurance purchasing, the decision maker's expected utility is

$$E(U)_I = \left(1 - \sum_{i=1}^{n} \pi_i\right) U\left(W_0 + R_g + b(I) - pI\right) + \sum_{i=1}^{n} \pi_i U\left(W_0 + R_g - YL_i - AL_i + b(I) - pI + q_i I\right)$$

where q_i is an indemnity function that determines the magnitude of the indemnity conditional on an indemnity being triggered and $b(I)$ is the monetary value of any ancillary benefits associated with insurance purchasing reducing risk exposure (e.g., improved access to credit), thus $\frac{\partial b}{\partial I} > 0$.

Define p^{max} as the premium rate that would cause the decision maker to be indifferent to being insured or uninsured. Then the decision maker would be better off by purchasing the insurance if the premium rate is less than p^{max}. Obviously, p^{max} depends on the decision maker's utility function (i.e., how risk averse is the decision maker?), the magnitude of potential yield losses YL_i relative to R_g, and the magnitude of potential asset losses AL_i relative to W_0.

As will be developed later, it also depends on the covariance between q_i and YL_i, the covariance between q_i and AL_i, the magnitude of any ancillary benefits associated with insurance purchasing (which likely also depends on the covariances between q_i and YL_i and q_i and AL_i), and the decision maker's subjective assessment of the probability of loss π_i and associated magnitudes of losses YL_i and AL_i.

While we have modeled the above in a household-level expected utility framework, the model is generalizable to risk aggregator firms in many contexts. In neoclassical economic models, firms are often modeled as risk neutral. These models are built on assumptions that firms can fulfill their demand for physical capital and labor and that these factors can easily be replaced. In reality, these assumptions often do not apply. Catastrophic weather events can severely disrupt rural risk aggregators in developing countries (Box 1). Because of these business disruptions, it seems likely that many risk aggregators may be risk averse, especially for catastrophic, correlated weather risks.

Box 1 Risk Aggregator Portfolios

Consider a bank that holds a portfolio consisting of *n* loans. Each loan $L_i(i=1,2,\ldots n)$ generates a periodic net return $r_i = Int_i - C$ where Int_i is the interest earned on the loan and C is the bank's cost of capital on the funds that have been loaned (which, for simplicity, we will assume does not vary across loans). In a given year, the realized net return for a specific loan is largely a function of whether the borrower is able to repay the loan at the agreed terms. In some cases borrowers may experience losses that make it difficult or impossible for them to repay their loans. They may be able to repay only part of the loan or perhaps they are able to repay the entire loan but only after renegotiating the terms so that the net return for the lender is reduced.

Assume that for each loan L_i there are m possible discrete levels of net return for the lender. The variance of net returns for loan L_i is calculated as $\sigma_i^2 = \sum_{j=1}^{m} \alpha_{ij}\left[r_{ij} - E(r_i)\right]^2$ where α_{ij} is the probability of net return level j for loan L_i and $E(r_i)$ is the expectations operator. The periodic net return on the bank's entire loan portfolio is $R = \sum_{i=1}^{n} w_i r_i$ where w_i is the proportion of the total value of the portfolio that is invested in loan L_i and $\sum_{i=1}^{n} w_i = 1$. The variance of net returns for the portfolio is calculated as $\sigma_R^2 = \sum_{i=1}^{n}\sum_{k=1}^{n} w_i w_k \sigma_{ik}^2$ where σ_{ik}^2 is the variance in returns on the single loan when $i = k$ and the pairwise covariance in returns between loans when $i \neq k$ with $\sum_{i=1}^{n} w_i = 1$ and $\sum_{k=1}^{n} w_k = 1$. Thus, the overall variability in net returns for the bank's loan portfolio is a function of the variance in net returns for each of the loans, the proportion of the overall portfolio that is invested in each loan, and the *covariances between the net returns for each loan*.

If the net returns generated from the loans are not highly correlated (the covariances are low), the variance in net returns for the bank's loan portfolio will be greatly reduced. However, if the loan net returns are highly correlated, the variance in net returns for the bank's loan portfolio will be quite high. This is the problem faced by banks that lend to borrowers who are exposed to spatially correlated, catastrophic, weather risks. A single catastrophic weather event could affect a large proportion of the bank's borrowers and result in very low net returns for the loan portfolio.

1.5.2 Probability Distributions of Loss, Cause of Loss, and the Index

The above model describes the accumulation of wealth and household decision making in the presence of extreme events. Here, we develop a conceptual framework that describes how practitioners would approach index insurance product design in an ideal world, if they had sufficient quantitative data. However, it is important to remember that in many lower income countries, statistical analyses described herein will not be possible due to data insufficiencies. Methods to overcome these data constraints are described in the next chapter.

This framework is an important benchmark for practitioners seeking to design sustainable insurance products. It is motivated by two principles. First, the insurance should be priced based on the risk being transferred to the insurer.[6] Second, the indemnity should be highly, positively correlated with losses of insureds. Otherwise, purchasing the insurance is not likely to benefit the insured as it would not effectively reduce the volatility in ending wealth.[7]

Regarding the first principle, pricing the risk, practitioners need data based on the mechanism used for paying insurance indemnities. For weather index insurance, the mechanism determining indemnities is the weather index (e.g., rainfall levels at a specific weather station); for traditional forms of insurance, the mechanism for determining indemnities is an on-site assessment of losses (see Box 2 for a comparison of traditional and index-based insurance).

Box 2 Data Constraints Create the Need for Index Insurance

In theory, a loss-based insurance product could be offered that makes payments based on losses experienced by the insured. In many respects, this would be the most straightforward way to insure against losses caused by extreme weather events because payments are based directly on the losses experienced by the insured. But this direct connection between the loss experienced by the insured and the payment received by the insured also causes significant problems. Some potential insureds will have greater loss exposure than others. To offer a loss-based insurance product, the insurer must be able to accurately estimate the loss distribution for each potential insured and charge a premium rate that accurately reflects the potential insured's loss exposure. So those with higher (lower) loss risk will be charged higher (lower) premium rates. But the data required to estimate a loss distribution for every potential insured are often not available. If the insurer is unable to accurately classify potential insureds according to their loss exposure, the pool of insurance purchasers will be disproportionately composed of those who have been offered premium rates that understate their actual loss exposure. This problem, known as *adverse selection*, will undermine the long-run sustainability of the insurance product.

Another problem with loss-based insurance products is *moral hazard* — a reduction of insureds' incentives to reduce their exposure to losses since insurance payments will at least partially compensate for any realized losses. For example, in flood prone areas, those who purchase loss-based insurance may be less likely to invest in building levies or elevating buildings. Moral hazard can be controlled to some degree by policy provisions requiring the insured to utilize specific risk mitigation strategies, but the cost of monitoring and enforcing these policy provisions can be excessive.

[6] The risk being transferred to the insurer, called the "pure risk" is but one component of the cost of providing insurance (e.g., administrative costs are another). For market-based insurance products, all costs of the insurance must be passed on to those purchasing the insurance. The important point regarding pure risk is that the price of the insurance is directly related to the amount the insurer is likely to pay in indemnities. Please see Collier, Skees, and Barnett, 2009, for more on insurance pricing fundamentals. Returning to our expected utility model, assume that an insurance product is created to protect against losses from the m worst weather outcomes ($m \leq n$). The premium rate p comprises the pure risk g and other costs oc such as administrative costs (i.e., $p = g + oc$). For insurance to be priced sustainably, $g = \sum_{i=1}^{m} \pi_i q_i$ where π_i is the probability of the insured event i and q_i is the rate of insurance indemnity for event i. In other words, the pure risk premium equals the (expected) rate of indemnity.

[7] Insurance is priced based on the pure risk as g and other costs oc , and therefore, has a negative expected return. Generally, insurance is beneficial to decision makers when 1) indemnities have a large, positive covariance with losses, and 2) decision makers are risk averse. Insurance may also create additional benefits as described in the expected utility model.

Deductible and co-payments are also often used to help control moral hazard.

A final problem with loss-based insurance is the very high administrative and delivery costs. As indicated earlier, the insurer must assess the loss exposure (estimate the loss distribution) for every insurance applicant. This often requires traveling to the exact location where any insured losses would occur. While the policy is in force, it may be necessary to travel to the location again to make sure that the insured is abiding by all relevant policy provisions. Finally, in the event of an extreme event that triggers an insurance payment, the insurer would again have to travel to the location to assess the magnitude of loss and determine the payment due the insured. These administrative and delivery costs are quite high even in developed countries where transportation infrastructure is good, insurers have access to the latest computer and communications technologies, and the insured value for a single policy may be quite large. In developing countries transportation infrastructure is often not good (especially in rural areas), insurers often do not have access to the latest in computer and communications technologies, and the insured value for a single policy may be quite small.

High administrative and delivery costs, along with adverse selection and moral hazard, make loss-based insurance infeasible for insuring against extreme weather events in rural areas of most developing countries. Index insurance is designed to address each of these problems. With index insurance there is little potential for adverse selection or moral hazard because the payment is based on the realized value of the index rather than on the insured's realized loss. Administrative and delivery costs are greatly reduced because there is no need to assess each potential insured's loss exposure, no need to monitor for violations of policy provisions by insureds, and no need to assess the actual losses experienced by insureds. Thus, the lower data requirements for index insurance make it feasible in some regions where traditional insurance is not.

In principle, index insurance could be tailored to each insured, but in practice, sufficient historical data are not available and the data that are available lack the spatial specificity, i.e., the spatial resolution with which a data system records measurements that would be required to estimate a unique probability distribution for each insured. For these reasons, index insurance is based on a generalizable index. For example, Figure 1 presents a probability distribution for an index based on November–December average sea surface temperatures (SST) measured off the northern coast of Peru in the composite ENSO Regions 1 and 2, called ENSO 1.2. High SSTs are correlated with flooding in northern Peru, so this index is being used as the basis for an index insurance product that protects against flood losses (Khalil et al., 2007).

Figure 1 November–December Average Sea Surface Temperatures in ENSO Region 1.2

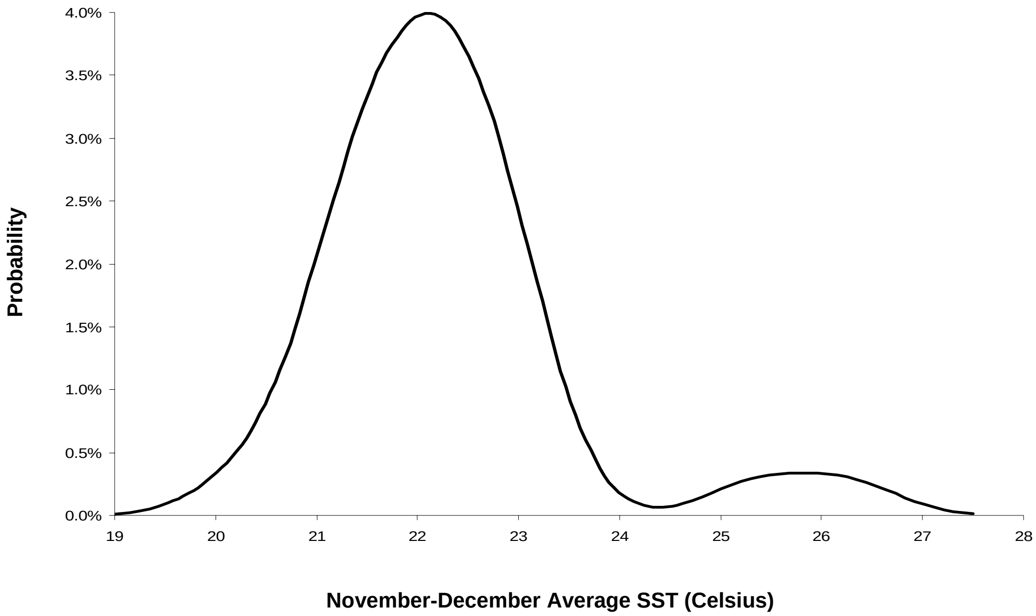

Source: Authors using NOAA historic data from 1979 to 2008

Whatever the mechanism, the principles are the same. Practitioners need to know the expected indemnity of the insurance contract they design, which depends, of course, on the probability of an insured event occurring. For example, if an insurance product made payments based on flooding, the insurer would need to know the probability of flooding. Practitioners need data on historical flood events to assess this probability. Flooding data would be organized based on the frequency and severity of flood events into a probability distribution, as shown in Figure 2. From the probability distribution, insurers can estimate the expected level of indemnities for the insurance product and can, thus, identify a sustainable price for the insurance.

Figure 2 Hypothetical Probability Distribution of Flooding

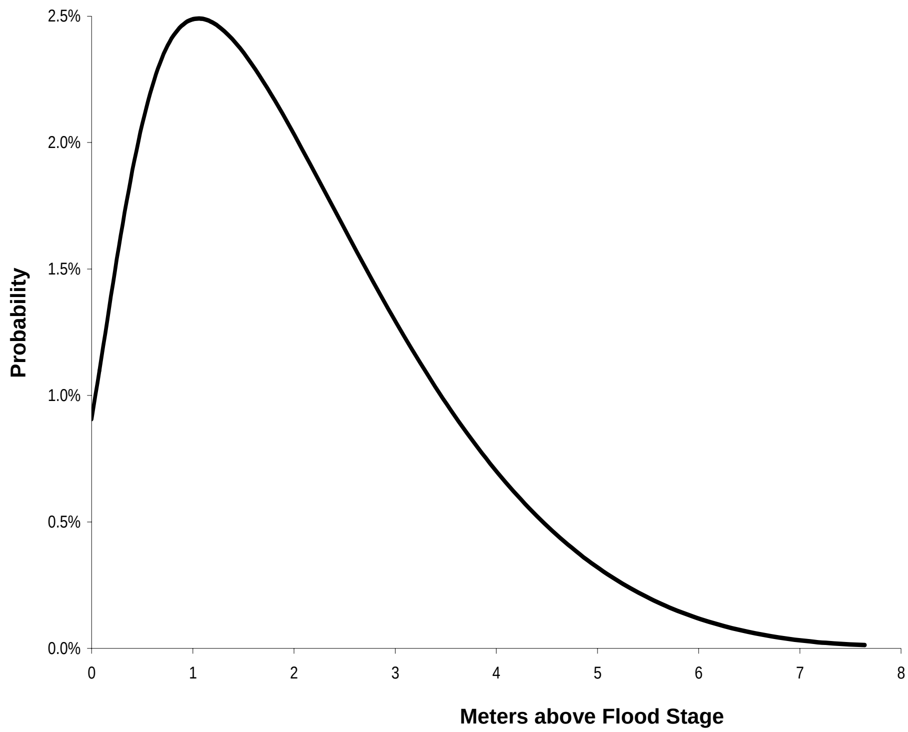

Source: Authors

In practice, accurately estimating the probability distribution for a weather index can be quite difficult. Many data are needed to accurately estimate the statistical characteristics (or moments) of the distributions that are of interest to a practitioner (Box 3). In many cases, these data are simply unavailable in developing countries.

Box 3 Estimating the Moments of a Distribution

Sample data are used to fit a probability distribution. For example, a sample of 30 years of historical cumulative rainfall data for the month of June collected at a specific weather station can be used to fit a probability distribution. The shape of a parametric probability distribution can be summarized by a few standard characteristics called *moments*. The basic moments of a probability distribution are as follows:

First moment – central tendency: the mean of the distribution.

Second moment – variance: describes how potential outcomes are positioned relative to the mean. It is small when potential outcomes are narrowly distributed around the mean and high when they are widely distributed around the mean. The standard deviation (sd) is the square root of the variance.

Third moment – skewness: characterizes the symmetry of the distribution. Skewness is equal to zero if all the data are symmetrical around the mean as with a perfect bell curve. It is negative if there is a fat tail (many low probability events far from the mean) on the right and positive when the fat tail is on the left.

Fourth moment – kurtosis: measures how data are stacked over the distribution. With a high kurtosis, there is a distinct peak near the mean that declines rapidly. Low kurtosis is relatively flat near the mean so that a larger proportion of the events in the population are found in the tails.

An approximation used in many applications is that a sample of at least 30 observations is required to estimate the central tendency and variance with an acceptable degree of accuracy. But this is just an approximation. More specifically, the accuracy with which sample data can estimate the true central tendency of a distribution is shown by the square root of *n* rule.

$$\text{standard error of the estimate} = \frac{\text{sd}}{\sqrt{\text{n}}}$$

This shows that the accuracy with which sample data estimate the true central tendency increases with the sample size and decreases with the standard deviation of the distribution. Accurately estimating higher moments of the distribution requires even larger sample sizes. The sample size needed to estimate skewness and kurtosis is much higher than what is needed for mean and variance. Whereas the mean and the variance deal with the bulk of the distribution, skewness and kurtosis are essentially measures of extreme, rare, events which may be underrepresented or overrepresented in a small sample. As a rule of thumb, for a given sample size used to estimate the mean and variance of a distribution at some desired level of accuracy, a sample size 6 times larger is required to estimate skewness at the same level of accuracy and a sample size 24 times larger is required to estimate kurtosis. For example, if 30 years of data are sufficient to estimate the mean and variance of the distribution of a weather variable with some desired level of accuracy, approximately 180 and 720 years of data are needed to estimate skewness and kurtosis, respectively, with the same level of accuracy.

In sample vs. out of sample. When sufficient data are unavailable, practitioners must estimate the distribution given the data they have. Regarding the skewness and kurtosis in particular, this practice relies heavily on the few extreme values (i.e., catastrophic events) available in the data. In this case, *in sample* values may be very poor predictors of *out of sample* values (e.g., extreme events in the future). In other words, the *in sample* distribution may not accurately reflect the actual but unknown distribution. Ignoring this limitation can lead practitioners to believe that they have accurately accounted for catastrophic risk exposure when, in reality, this may not be true. A probability distribution fit from limited in sample data may greatly underestimate or overestimate catastrophic risk exposure.

Additionally, if the risk is changing (e.g., due to climate change, multi-year weather cycles, hydrological engineering developments on rivers, etc.), practitioners must adjust for this in estimating the distribution. Chapter 3 discusses these difficulties more fully; here, we want to note that accurately estimating a single probability distribution for a weather risk can be challenging because of data constraints (see Box 4 for an illustration of how limited data can easily lead to misestimating the probability distribution).

Box 4 Working with Small Samples

To further illustrate the sensitivity of getting the correct expected value for the underlying index, we demonstrate the standard error of the estimate with a Monte Carlo simulation. Practitioners who strive to develop index insurance products would be well-served to consider sampling issues in this fashion. What is presented below is a relatively straightforward process that is designed to put some of these problems into context.

Under the strong assumption that a variable is normally distributed, the square root of N rule can be extended to demonstrate errors. If the true distribution is known to have a mean of 100 and a standard deviation of 50, one can quickly gain perspective on the sampling errors that are possible. Having this level of risk is not uncommon for rainfall in many regions of the world. Table 1 provides the lower and upper bound for this distribution at the 95 percent confidence interval, given different sample sizes.

Table 1 Range Boundaries for Confidence Intervals

Sample Size	Lower Bound	Upper Bound
10	53	147
15	61	139
20	66	134
25	70	130
30	73	127

Confidence intervals are a measure of the precision of an estimate. Of course, these errors can be even greater when the distribution is not normal or if the distribution is not stationary.

Working with small samples also creates opportunities for misestimating the pure risk. For illustrative purposes, once again we assume that the index is distributed normally and that it is stationary. With a normal distribution, two parameters can be manipulated, the mean and the standard deviation. In this example, assume the risk is constant as measured by the coefficient of variation CV (CV=sd/mean). The only unknown variable is the mean — the sample statistic that requires the least amount of data to estimate.

We test the sensitivity of the pure risk to errors in estimating the mean using a Monte Carlo simulation of 1,000 draws given a sample size of 30 and an algorithm to develop the pure premium risk from a normal distribution. For a rainfall insurance contract that pays for losses below 80 percent of the expected value, the pure premium from the given distribution is 11.5 percent. Running the Monte Carlo simulation, about 10 percent of the pure premium values are less than 8.5 percent and another 10 percent are greater than 14.9 percent. These are quite large errors in the estimates of the underlying risk, especially when one recognizes that the only estimate that is allowed to vary is the expected value and that we assume the index is stationary and normally distributed (i.e., in this case we perfectly know the true parent distribution) — a rather strong assumption! When the sample size drops to 20, 10 percent of the pure premium values are less than 7.9 percent and 10 percent are above 15.7 percent.

Regarding the second principle, data requirements are much higher to ensure that the index insurance indemnities are highly correlated with the losses of the insured. Practitioners need several types of data from the same time period. The data needed, particularly data on the specific risk exposure and losses of individuals in the target market (e.g., households or firms), are unavailable in many developing countries. Because of data constraints, practitioners will typically have to rely on alternative methods to assess how indemnities relate to losses for the insured. In the next chapter, we describe these alternative methods. Here, we discuss how practitioners would accomplish this, if they had sufficient data, pointing along the way to the types of challenges practitioners are likely to face in data constrained environments

Practitioners developing index insurance products would like to be able to demonstrate a strong relationship between the index and the losses of the insured. To do so, practitioners must work across three types of data for index insurance: losses of the insured, the cause of loss, and the index.

The insured's *losses* can be conceptualized in many ways: losses in income or consumption, reduced profits, yield losses, asset losses, etc.

The *cause of loss* is the specific natural phenomenon (typically, a weather risk such as drought) that causes losses for the insured during a particular period of time. The index insurance product is designed to provide protection against the financial losses caused by this natural phenomenon.[8]

The *index* is a measure on which the index insurance indemnities are based. It is defined by a number of specific characteristics such as: what is being measured, how it is being measured, where it is being measured, and over what period of time.

Each of these variables (loss, cause of loss, and the index) has a probability distribution. Ideally, a practitioner would like to have sufficient data to accurately estimate these three probability distributions. Additionally, practitioners need data on each variable to have been collected during the same time period to estimate relationships between the variables. In particular, one would like to understand the relationships between these variables when an extreme event occurs. Practitioners want to identify the relationship between losses and the cause of loss (e.g., low rainfall). We write this as losses ←→ cause of loss, for shorthand. Also, with index insurance, the index is an approximation for the cause of loss for the insured so practitioners also want to identify the relationship between the cause of loss and the index (cause of loss ←→ index). For example, rainfall at the closest weather station may differ from the rainfall a household experiences on its farm.

As an example of these three variables, consider our experience in Peru. We assessed solvency risks to rural lenders in Peru, and our analyses indicated that regional floods create borrower default losses, liquidity constraints, and increased operational costs for the lender. Furthermore, extreme El Niño events are the primary cause of flooding in the region, and the best measurement of El Niño is sea surface temperature (SST). Therefore, we designed an insurance contract based on SST (the index) that we intend to reduce exposure to catastrophic flood risk (the cause of loss) as a means to address bank losses. Figure 1 represents the probability distribution of SST. Because this variable is used as the index, insurers price the risk using this distribution. For example, if the insurance contract paid when index values were above 24° C, the insurer would estimate the probability of the event using the area under the curve for temperature values above 24° C. To design the contract so it is relevant to lenders in the region, the practitioner would like to know the flood level that is likely to occur for a specific SST and the losses that lenders are likely to experience given that specific flood level (Appendix A). With this information, the characteristics of the index insurance product could be tailored to optimize the protection offered to the insured.

In statistical terms, practitioners are attempting to identify a conditional distribution. A conditional distribution is the probability distribution of one variable given a specific value of another variable. For example, we could potentially identify a probability distribution of losses based on a specific level of flooding. Suppose from Figure 3, that we are interested in the level of losses experienced by a household or firm for a flood that is 3.0 meters above flood stage. Figure 3 is a conditional distribution. It shows the distribution of losses given this level of

[8] It is also possible to construct more complex weather index insurance products that cover losses from more than one cause of loss. Assuming a single cause of loss simplifies the conceptual presentation without any loss of generality.

flooding. The expected level of loss is roughly USD 4,000 for this level of flooding; however, losses are represented as a distribution because the relationship is not deterministic. A given flood level will not always generate the same magnitude of loss. Consider this variation around the mean as error, which we call *basis risk*. This error occurs because many variables may influence losses from a specific level of flood (e.g., whether individuals had time to prepare for impending floods, whether flooding as due to heavy rains or some other factor such as river overflow). These other variables that influence flooding affect the variance of this distribution — how precisely flooding can be used to identify losses. Thus, the closer the distribution is around the mean of the distribution, the lower the basis risk.

Figure 3 Hypothetical Probability Distribution of Losses Given Flooding of 3.0 Meters

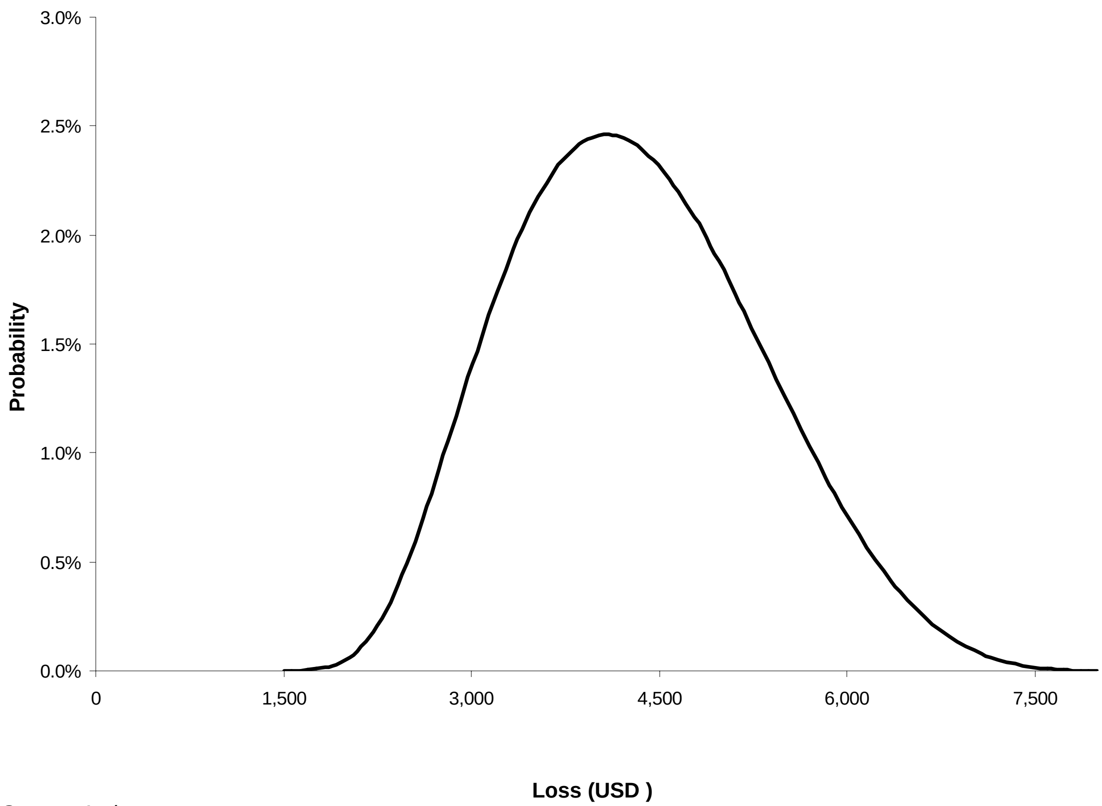

Source: Authors

For index insurance, practitioners want to estimate the conditional distribution of losses for a specific value of the index — i.e., given a specific value of the index (e.g., SST) what are the losses experienced by the insured. Practitioners can then design the insurance contract based on the expected value of losses for any given value of the index. The conditional distribution will include the errors associated with the ability of the index (SST) to estimate the cause of loss (flooding) and the errors associated with the ability of the cause of loss to predict losses of the insured (borrower default rates).

In practice, data limitations create two major challenges with this framework. First, the conditional distribution of losses given a specific index value is somewhat unique to each individual. **Index ←→ cause of loss** changes depending on the physical location of the individual (e.g., how close a household farm is to the weather station used as the data source for the index). Also, **cause of loss ←→ losses** is affected by differences in factors such as business or livelihood activities and risk mitigation strategies. For example, some business or livelihood activities may be less prone to losses from flooding and some insureds may have invested in levies or elevated buildings to reduce their exposure to flooding. This is especially likely with household products because each household farms different soils, uses different inputs, plants different crops, manages different livelihood portfolios, etc.

Second, practitioners will not have sufficient data to properly estimate the distributions and conduct these analyses. As Boxes 4 and 5 illustrate, many data are needed to accurately estimate the moments of a probability distribution. However, cause-of-loss data for individual insureds are generally not available at all. Some individual loss data may be available for a sample of potential insureds but these data are unlikely to include enough observations to allow for accurate estimation of probability distributions. Furthermore, to develop a conditional distribution of losses for a specific index value, practitioners need several observations of losses for each level of the index. Because the insurance products of interest here are designed for low frequency, high severity events — the types of events that tend to occur no more frequently than say 1 in 15 years — even data for the index may, at best, tend to include only one or two high-severity events.

So, in reality, practitioners can typically only estimate the probability distribution of the index (and, as will be discussed later, it is often not easy to identify a variable with sufficient data to serve as the index). Because there are insufficient data, *it is generally not possible to statistically estimate the relationships between the index, the cause-of-loss magnitude experienced by the insured, and the losses experienced by the insured. Instead, practitioners will typically make inferences based on qualitative sources and limited amounts of available quantitative data*.

Before we proceed with describing how contract design would be approached in practice, it is essential to examine in more detail the concept of basis risk that was briefly introduced earlier. Basis risk is the foremost limitation of index insurance; therefore practitioners want to estimate and minimize basis risk to the extent that is practically feasible. If access to data were not an issue, practitioners would learn about the sources and magnitude of basis risk by examining the relationship between the probability distributions of the index, the cause-of-loss, and loss. However, as noted earlier, these statistical analyses may not be feasible in a real-world setting and practitioners will generally have to rely on qualitative information to estimate basis risk.

1.6 Basis Risk

In an insurance context, basis is the difference between the loss incurred by the insured and the indemnity received. Basis can occur due to factors such as contract characteristics (e.g., deductibles or co-payments) or errors that occur in the process of establishing the sum insured or in conducting loss assessment. If basis is relatively small and predictable, as would be the case with a modest deductible, it is generally not a major concern for an insurance purchaser.

Variability in basis, or *basis risk*, on the other hand, can be a major concern and is the primary limitation of index insurance. Basis risk creates the possibility that indemnities will not be highly correlated with the losses of the insured. As described above, a source of basis risk is the imperfect relationships between the index, the cause of loss, and the loss. Consider the insurance product equation that is described in Section 1.5.1. If the indemnity $q_i I$ is not highly positively correlated with yield losses YL_i and/or asset losses AL_i, insurance purchasing will not reduce the variability in ending wealth.[9]

Basis risk describes the precision with which the index can be used to estimate losses of the insured. It can be represented in part by the variance of the conditional distribution of losses given a specific value of the index. Because data are insufficient to capture the conditional distributions described above, practitioners have used correlations (or covariances) between losses and the index to estimate basis risk.

Miranda (1991) demonstrated how basis risk affects the efficacy of index insurance by using a modified version of the Sharpe-Lintner Capital Asset Pricing Model, or CAPM (Sharpe, 1964) — a model that is widely used in finance to describe the relationship between returns for a given asset (e.g., a stock) and returns for the aggregate market (e.g., the Standard and Poor's 500 Index). Specifically, Miranda describes the relationship between yield outcomes on a given farm and a spatially aggregated regional yield as

$$(y_{it} - E(y_i)) = \beta_i (Y_t - E(Y)) + \varepsilon_{it}$$

where *y* is farm-level yield, *Y* is the regional yield, *i* represents different farms, *t* represents different crop years, and *E* is the expectations operator. The parameter β_i shows how the *i*th farm's yield deviations from their expected value vary with regional yield deviations from their expected value and is defined formally as

$$\beta_i = \frac{cov(y_{it}, Y_t)}{var(Y_t)}$$

The error term ε_{it} captures that part of the *i*th farm's yield deviations from expectation that are not explained by regional yield deviations from expectation. By assumption $E(\varepsilon_{it}) = 0$ and $cov(\varepsilon_{it}, Y_t) = 0$.

Thus, Miranda's model decomposes farm-level yield deviations from expectation, $(y_{it} - E(y_i))$, into a spatially covariate component, $\beta_i(Y_t - E(Y))$, and an idiosyncratic component ε_{it}. Miranda used this model to demonstrate that for a specific farm *i*, the efficacy of an area yield index insurance would depend on the farm's β_i, which, in turn, depends on the covariance of the farm yield and the regional yield. The higher (lower) the covariance between the farm and regional yield, the higher (lower) the value for β_i and thus, the more (less) that an area yield insurance policy would protect against farm-level yield losses. Said differently, the higher (lower) the value for β_i, the lower (higher) the basis risk.

Miranda's model can easily be extended to other types of index insurance. Instead of regional yields, the deviations on the right-hand side of the model could be for a weather variable or a combination of weather variables. Likewise, the deviations on the left-hand side of the equation need not be limited to yields. Instead they could reflect deviations in gross or net revenue from

[9] Loss-based insurance products may also entail basis risk due to errors in estimating expected values (or sum insured) and losses (Barnett et al., 2005). Additionally, other financial contracts used to manage risk such as commodity futures also have basis risk. Nevertheless, the value of these risk management mechanisms even in the presence of basis risk is well-demonstrated in the literature and also applies to index insurance. Therefore, our discussion focuses on methods used to conceptualize, estimate, and manage basis risk that occurs with index insurance.

a number of different livelihood strategies or they could reflect deviations in net worth. Regardless of the index used on the right-hand side or the losses used on the left-hand side, β_i is a simple and convenient measure of the extent to which deviations in the index explain losses. The higher (lower) the covariance between the index and the losses, the lower (higher) the basis risk associated with the index insurance product.

To better understand the primary sources of basis risk and relate these to product design, it is helpful to break down the covariance between the index and losses into two main components: 1) the covariance between the cause of loss and the loss; and 2) the covariance between the index and the cause of loss. With regard to the first component, if a specific cause of loss (e.g., flooding) is responsible for most of the realized losses and a relatively predictable relationship exists between the measure of the cause of loss (e.g., severity of flooding) and realized losses then the covariance between the cause of loss and losses is likely to be high. If many different causes of loss can generate large losses or if the relationship between a specific cause of loss and realized losses is highly random, then the covariance between any specific cause of loss and realized losses is likely to be low. This covariance between the cause of loss and loss likely differs across individuals.

With regard to the second component, the covariance between the index and the cause of loss is affected by various factors. For example, deficit rainfall on the farm of an insured household (a cause of loss) is likely not perfectly covariate with deficit rainfall measured at the closest weather station (an index). Flooding in northern Peru (a cause of loss) may occur for reasons other than El Niño cycles as reflected in SST (an index).

A limitation of using covariances is that these estimates describe the relationship of two variables across all values of each variable. The β_i above is the same as a coefficient in a regression analysis. In the regression context, we would interpret the value of β_i as “on average a one unit increase in the independent variable will cause an *X* unit increase in the dependent variable.” Thus, the β_i value describes the general relationship between the index and losses. As we described in the opening paragraphs of the introduction, insurers are specifically concerned with the relationship of these variables in the tails of the distribution, which may differ from the general relationship between these variables. These differences may be a result of underlying physical processes. Consider a crop growth example. When rainfall is around the optimal level for a crop, many other important factors affect crop yields (e.g., soil quality, fertilizer use, pesticide use, etc.); therefore, around this level, the correlation between rainfall and crop yields would likely not be very strong. When rainfall is extremely low, however, the relationship between rainfall and yields is expressed more strongly. At low levels of rainfall, other variables such as fertilizer use have very little effect on yields. Therefore, statistical methods for estimating basis risk are needed that go beyond simple covariances. Appendix B, developed by Miranda, provides this methodology. In brief, Miranda proposes estimating these relationships using copulas, which is a statistical technique that can be used to compare covariances in the tails of the distributions.

Basis risk cannot be completely eliminated from index insurance but careful product design can reduce basis risk. Additionally, proper marketing of index insurance in light of its limitations is critically important to reduce misunderstandings in the target market about basis risk. In Chapter 2, we describe the risk assessment process for developing index insurance. The risk assessment chapter serves as a reality check to the academic discussion above. As will be clear, this estimation of distributions and the statistical relationships between distributions is only possible in a qualitative way in many lower income countries. However, before moving to the discussion on risk assessment it is useful to review two important efforts that are sometimes made to reduce basis risk. First, when practitioners have yield and weather data, they can work extensively to fit models that will provide the best fit as a way to reduce basis risk. As is developed below, working with limited data can result in overfitting models that explain the in-sample farm-yield data. Second, when developers do not have yield data, there have been attempts to use plant growth simulation models with local weather data as a means for compensating for the lack of data. While this is also a good practice in principle, it can lead to wrong conclusions about the quality of the weather index insurance product for protecting farm yields if the limitations of this approach are not kept in mind.

1.6.1 Fitting Models with Limited Data to Reduce Basis Risk

As indicated earlier, index insurance products are designed around some perception of how the index relates to losses. If loss data are available, practitioners may utilize statistical techniques such as regression analysis to better understand the relationship between weather variables and losses.

Practitioners have sometimes developed very complex statistical models of the relationship between various weather variables and crop yields *within the available data*. They then use these models to create insurance product designs. Practitioners feel more comfortable if the underlying index for the index insurance product is based on a statistical model that explains much of the variability in losses.

The problem is that more complex models tend to overfit the in-sample data. In other words, while complex models will generally fit the in-sample loss data better than simpler, more parsimonious, ones, simpler models will often perform better in predicting out of sample events. If an index insurance product is created using the parameters of a complex statistical model, it is likely that out of sample indemnities will not match losses as well as the model would suggest. The complex, overfit, model will underestimate the basis risk, giving practitioners a false sense of confidence in the index insurance product. As a result practitioners are likely to "oversell" the benefits of the insurance so that insureds believe that they are much better protected than they actually are.

1.6.1.1 Product Designs with Specific Temporal Characteristics

Some practitioners have also designed their models using quite specific, discrete time intervals to determine indemnities. For example, index insurance products have been designed using a dekadal (10-day) rainfall measurement as the basis of indemnities. These measurements provide a specific probability distribution across the growing season with which one can estimate the pure risk of the contract. The difficulty with these estimates is if the dekadal intervals are started two days later, remarkably different indemnities may occur in the historical data.[10] In other words, this is just another example of in-sample overfitting. Such a contract lacks external validity and increases basis risk relative to a simpler contract.

1.6.2 Using Crop Growth Simulation Models to Compensate for Missing Yield Data

Practitioners have also used crop growth models to compensate for missing data. If loss data are not available, crop growth models have been used to simulate the relationship between various

[10] This finding was shared with our team by professionals from an international reinsurer.

inputs (including weather variables) and yields for specific crop varieties and regions. For example, practitioners have used these models to estimate the effects of low rainfall values on crop yields. They then use this information to design the indemnity structure for an index insurance contract.

Various data sources are used to construct crop growth models. Data on weather variables are often collected by creating a test plot at or near a weather station. Crops are planted in this test plot and their growth is observed by researchers. Each year, the researchers try to control certain inputs (e.g., soil quality, fertilization, pesticide applications, etc.) and observe the relationship between the uncontrolled variables (e.g., rainfall, temperature, and sunlight) and yields. Over the years, they obtain observations that can be used to estimate how different combinations of inputs, including weather variables, affect yields.

A major concern about building index insurance products around relationships inherent in crop growth models is that these models are parameterized for very specific crop varieties and regions. One cannot simply assume that the parameters contained in the models are generalizable to other crop varieties, regions, or farming practices.

Another concern is that while crop growth models are quite useful for estimating the effects of a change in a variable around the central tendency of the distribution, they are much less useful for predicting the effects of extreme weather events on yields. A crop growth model is, in essence, a complex regression model that specifies relationships between various inputs (including weather variables) and yields. Regression coefficients are interpreted based on the relationship across all values and are most accurate near the mean since most observations occur around the mean. Regression estimates tend to be least accurate for extreme values because extreme values occur much less frequently. When there are very few observations of low yields, the regression will be much less accurate in predicting the relationship. Also, with few observations of extremely low yields (e.g., a typical model may have one or two cases) it also becomes quite difficult to isolate the effects of one weather variable versus another.

There is a common theme to our concerns about building weather index insurance by overfitting available data or relying on exclusively on crop growth models. That theme is the danger of "overselling" the potential benefits of index insurance. Both models that overfit in-sample and crop-growth models are likely to underestimate the true basis risk that will occur out of sample. If index insurance is sold based on unrealistic representations about the true magnitude of basis risk, practitioners will lose credibility with increasingly frustrated insureds — potentially undermining any efforts at long-run scalability and sustainability.

1.6.2.1 Implications

In response to problems associated with overfitting and reliance on crop-growth models, some development economists have questioned the value of index insurance investments citing that basis risks seems unacceptably high, development costs for these programs are significant, and opportunities for scalability and sustainability are minimal — especially given the complicated statistical approaches of some practitioners. However, the problem is not index insurance. Index insurance is built on sound economic principle, and new applications for this class of products continue to emerge. Generally, problems emerge when practitioners ignore the consequences of the data limitations by: 1) relying too heavily on a small data sample to provide an accurate estimation of the underlying probability distributions and the relationships across distributions; and 2) using scientific models intended to explain physical processes under general conditions to identify relationships during extreme events.

Chapter 2 Working under Real-World Constraints: Data, Product Design, and Risk Assessment

The conceptual model presented in the introduction describes the underlying framework for index insurance product design. In practice, product design is quite challenging, marked by uncertainty, and often, qualitatively approximates this theoretical framework. Vast data constraints in lower income countries limit the techniques available to practitioners. Consider, for instance, the task of estimating the loss distribution. Historical, quantitative data on losses in many lower income countries simply do not exist, especially data on household losses. Sparse or missing quantitative data greatly limit what can be done in estimating probability distributions. Mapping relationships between imprecise distributions results in more error. Thus, practitioners are left with insufficient quantitative data to determine how measurements of a weather event translate into losses for potential insureds.

Index insurance developers rely heavily on what little information may be available and employ several approaches to overcome these data constraints. In this chapter, we describe an approach that we believe is most likely to create products that are in line with the priorities of the target market.

The basic requirements for an index insurance product are 1) an index that is highly correlated with the insured risk; 2) historical data on the index to establish the pure risk being insured; and 3) some indication of how the index relates to the consequential losses and costs of the potential insureds in the target market. The index must be based on a secure and objective data source because indemnities are based on this measure. Accurate premium rating requires sufficient historical data to estimate the probability distribution of the index. The first two items in the list above are discussed further in Chapter 3. The focus of this chapter is on the third item — determining how the index relates to losses given that limited quantitative data in many lower income countries make it particularly challenging to estimate this relationship statistically.

The dearth of loss data in lower income countries leads to the question of how to begin product design. Because weather index insurance bases indemnities on the measure of a weather event, an appropriate starting point is identifying what weather risks are of greatest concern in the region. This approach focuses on the types of risks that can be insured with weather index insurance. It also allows for the possibility that the weather risk causes more than one type of loss. Thus when developing weather index insurance, the following is the logical sequence of questions:

1. What weather risk (cause of loss) causes the greatest/most disruptive losses?
2. What losses are associated with this weather risk? What risks should the insurance target? and
3. How well can the causes of loss for different insureds in the target market be approximated with a generalizable weather index?

To address these questions the practitioner utilizes any available quantitative data but must also usually rely on qualitative data collected from local stakeholders through risk assessment.

2.1 Risk Assessment

As a critical step of product design, we recommend a risk assessment process that is informed by scientific inquiry into weather risk in the region and enhanced by local knowledge collected through focus groups or other interview techniques to supplement any available quantitative data. Scientists with specialized knowledge of the region provide an invaluable starting point for practitioners. The findings of these researchers may be particularly important for identifying a suitable index, targeting vulnerable populations, and guiding focus groups. Local stakeholders who have lived through previous extreme events tend to have a clear picture of how households and businesses in the region have been affected and often have assessments regarding how vulnerable their community is to future catastrophic events.

Risk assessment identifies the major risks affecting households and businesses and systematically develops a model for understanding the risk. The risk assessment process operates under the recognition that weather risk and resulting losses occur in a larger system affected by many components: household livelihood strategies, geography, weather patterns, population dynamics, industry growth, cultural values, etc. Risk assessment attempts to estimate the cost of a specific risk in this context. The process includes assessing how households and businesses currently "pay" for this risk. Households pay *directly* when a catastrophic risk event causes yield, revenue, or asset losses or increased costs; however, households also pay *indirectly* for catastrophic risk by foregoing business opportunities because the risks are too high. For example, a bank may ration lending in a region exposed to flood risk, or a household may avoid higher-risk, higher-return production strategies because it deems drought risk too great. Often, weather events result in many concurrent consequences. For example in Vietnam, when coffee farmers in the Central Highlands are exposed to drought, they suffer losses in yield and quality, increased irrigation costs, and the death of coffee trees.

Risk assessments identify where existing risk management strategies are ineffective and/or inefficient for catastrophic risk, and where index insurance might be appropriate. As practitioners develop an understanding of risk in the local context, themes are likely to emerge that guide priorities in product development. For example, risk assessments will identify critical periods in which the target market is most vulnerable to specific weather risks. In our work in Vietnam we learned that rice farmers in the Mekong Delta are particularly vulnerable to flooding during the June and July rice harvest. An index insurance product was designed for lenders in the region based around a two-week critical window early in this harvest period (see Chapter 3 in GlobalAgRisk, 2009).

Another common theme that arises during risk assessments is the impact of catastrophic events on risk aggregators. Risk aggregating firms (e.g., banks, agricultural input suppliers, output processors, etc.) are exposed to the disaster risks of the communities they serve. As a result, catastrophic risk exposure can limit the role of these firms in the region and limit access to their services (Skees and Barnett, 2006). For example, a theme that emerged for our team during risk assessments in northern Peru is that, despite significant growth in financial services in the region, El Niño significantly limits access to credit for agricultural borrowers.

Risk assessment is likely to begin an ongoing process of product development with the target market. Practitioners can return to focus groups to assess the feasibility of using a specific weather index as the basis for the insurance. For example, the target market may not trust data collected by the national meteorological association, or they may not be willing to pay for an insurance product that makes payments based on satellite data such as NDVI. An important consideration regarding the index is potential insureds' perceptions of the basis risk. The target market can also be consulted regarding how they might use the proposed insurance product, which can have important implications for design and delivery. For example, index insurance products targeted to households in Malawi have been designed to protect against events (low rainfall) that create losses on the principal of an agricultural loan; in India, to protect household income from losses occurring to a specific crop; and in Mongolia, to protect household livelihood assets, such as livestock (Hellmuth et. al., 2009).

In this fashion, local knowledge can be used to partially overcome the data constraints that exist in many lower income countries. Index insurance is designed using qualitative data regarding how the index maps onto losses for the target market. While the practitioner may be unable to fully validate this aspect of the index insurance design with quantitative data, the process of developing the index insurance with local stakeholders increases the relevance of this product for the target market and reduces the likelihood of significant misunderstandings about the purpose of the index insurance.

Given all the challenges to stakeholders revealed in the risk assessment, practitioners must decide whether scarce funds should be used to develop index insurance or used to address other development priorities. In our experience, there are many positive externalities of index insurance market development beyond the development of a specific insurance product. The market development process can improve the risk management strategies of households, firms, and governments through education; strengthen legal and regulatory institutions to facilitate broader insurance market development; develop capacity among local partners such as insurers and banks; and advance other economic development efforts through the findings of risk assessments. Still, in some situations where weather risks are not severe, funds are likely better spent elsewhere. Also, if conditions are not amenable to index insurance (e.g., if data systems are severely underdeveloped), other investments are likely better. Thus, the opportunity cost should certainly be considered before practitioners pursue index insurance product development efforts — especially if significant infrastructure investments would be required to facilitate offering index insurance.

Chapter 3 Data Needs for Indexes

To estimate the loss distribution of the insured and to determine how losses tend to relate to the cause of loss (the weather risk), practitioners supplement available quantitative data with risk assessments conducted through focus groups and consultation with local experts, as described in Chapter 2. However, to actually develop an index insurance product, sufficient quantitative data are required to: 1) develop the price of the insurance based on the expected frequency and magnitude of indemnities; and 2) serve as the index for calculating the insurance payout.

It is necessary to review the desired attributes of data systems and place them within the context of challenges facing practitioners in the field. Data requirements for estimating the pure risk and calculating insurance payouts are contextual, in that they are influenced by the characteristics of the weather event and the physical environment in which the event occurs. In particular, data needs are largely determined by the spatial and temporal presentations (patterns) of the weather risk. Different weather risks tend to follow different spatial and temporal patterns that vary by region. These patterns have a direct influence on data measurement needs.

3.1 Data Needs Are Influenced by the Weather Risk

Data needs for both estimating the pure risk and for settlement of the insurance are largely influenced by the characteristics of the weather risk itself. The spatial and temporal presentations of the weather risk are the most salient features. Spatially, it is important to understand how large an area an extreme weather event tends to influence as this will determine how geographically precise data measurements must be. A basic precondition for index insurance is that the weather event must result in correlated losses; therefore, some spatial correlation must exist for any weather event suitable for index insurance. Still, these patterns may differ by type of weather event and by region. Excess rainfall will usually present a different spatial pattern than drought. Topography is an important determinant of the spatial presentation of a weather risk. For example, flood risk of an area can depend on its elevation relative to a flood risk source, such as a river, but also on the aspect and slope of the land such that heavy rainfall can produce inundation. Mountains alter weather patterns over short distances or generate microclimates. Such regions are often poorly suited for index insurance because the spatial correlations of weather events are so low.

Weather risks tend to follow temporal patterns as well. In some regions, daily rainfall in the afternoon is common; in others, it rains for most of the day every few days; and in others, several days or weeks of sustained rainfall are followed by prolonged periods of no rain. Seasonal patterns also occur throughout the world. Some regions near the Indian Ocean receive almost all their rainfall during the monsoon season from roughly June to September. August is considered the start of the hurricane season in the Caribbean. Finally, interannual patterns also occur. The Sahel, a region in Africa south of the Sahara, experiences multi-decade cycles of drought. El Niño also follows an interannual pattern in that atmospheric conditions make it unlikely for two extreme events to occur in consecutive years. In conclusion, identifying appropriate data systems for index insurance requires careful consideration of the many contextual influences to the presentation of weather events.

3.2 Estimating the Pure Risk

To price insurance sustainably, practitioners must accurately estimate the expected value of the payouts, also called the *pure risk* or *pure premium*, of an insurance policy. To do this, practitioners first fit a probability distribution to the historical weather data during the cover period, which describes the likely frequency and severity of the event in probability terms. The probability distribution can sometimes be estimated using a known distribution, such as a Gaussian, or by using kernel smoothing procedures of the empirical distribution of the weather observations. The pure risk is found by integrating the probability distribution times the payout rate determined by the contract thresholds and finally multiplied by the sum insured. Thus, the pure risk is a function of the distribution of the weather event during the period of interest, the thresholds and limits of the index where payments begin and end, and the total amount of coverage desired. The payment structure may take any form but is most frequently a linear function of the index. Figure 4 shows the probability distribution for a hypothetical index insurance contract. For example, Figure 4 could be the estimated probability distribution of rainfall at a specific weather station with the index insurance designed to protect against excess rain above the predetermined threshold of 18. The shaded area for values greater than or equal to 18 represents the expected value of the index insurance. In calculus, the area is integrated to take the value of the area relative to the entire area under the pdf. The expected value is also referred to as the pure risk.

Figure 4 The Pure Risk of a Hypothetical Index Insurance Contract

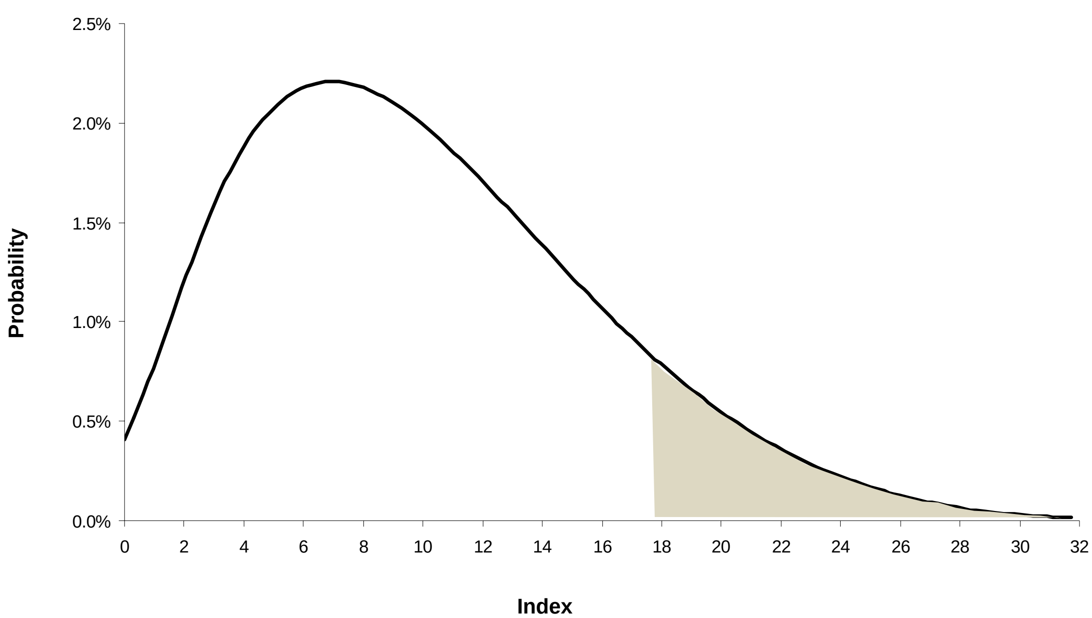

Source: Authors

When estimating the probability distribution, practitioners must also assess the historical data for trends and other systematic changes in the data over time. Figure 5 shows historical average rainfall for June, July, and August in the Sahel and is a dramatic example of trends in weather data, as average rainfall changes significantly over the time series. Data for June, July, and

August are used because of the importance of these months in this region for agricultural production. These data were collected from the rain gauges mapped in Figure 6. Practitioners are particularly worried about trends that are likely to increase the pure risk (see Box 5, which describes how these trends can affect the pure risk).

Trends in the data are a signal that the underlying probability distribution is non-stationary (i.e., the risks are changing). As a result, the probability distribution developed from the historical data must be adjusted based on this trend. How practitioners adjust the probability distribution relies heavily on their perceptions of the underlying physical process that is causing the change in weather risk. For example, practitioners may assume a stronger, more permanent trend if they believe the changes are due to climate change than if they believed the changes were largely due to multi-year cycles (as in the case of the Sahel).

Figure 5 Sahel Rainfall for June, July, and August, 1900–2006

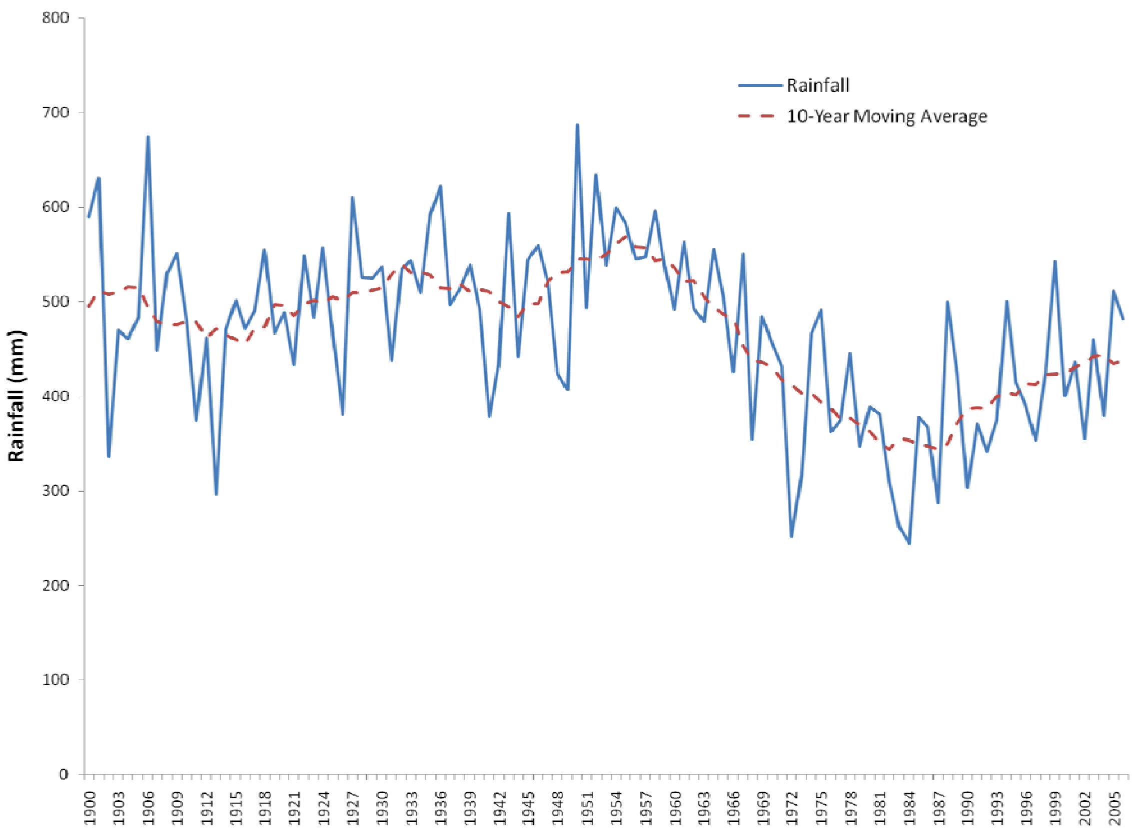

Source: Authors, based on data provided by International Research Institute for Climate and Society, Columbia University

Figure 6 Location of Rain Gauges Used in the Sahel Dataset

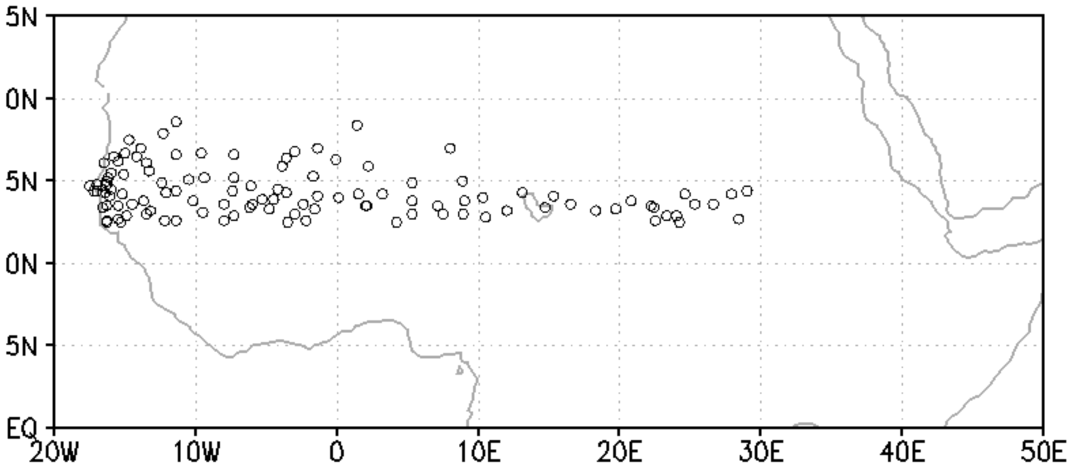

Source: International Research Institute for Climate and Society, Columbia University

Box 5 Potential for Misestimating the Pure Risk: A Sahel Example*

This case illustrates the potential consequences of misestimating the underlying probability distribution of a weather risk. It also demonstrates that probability distributions are heavily conditioned on a snapshot in time. The historical data used at that point in time can significantly misrepresent the current or future weather risk.

Consider the case of a practitioner developing an index insurance product for drought with an indemnity triggered at 425 mm. For levels of rainfall below 425 mm, indemnities increase in a linear fashion until 200 mm, below which the insured receive 100 percent of the sum insured. Suppose that the practitioner was designing such a contract in 1962, using the 1900–1961 Sahel data shown in Figure 5. Assume that the practitioner does not expect any trends in the risk. The green, dashed probability distribution in Figure 7 is based on this historical data (for simplicity we use a normal distribution for this illustration). For the contract described above, the pure risk would be 2 percent. The actual rainfall experience in the Sahel was much different in the following 26 years. While the variance of rainfall did not change, the average level of rainfall fell significantly. A probability distribution using data only from 1962–1989 (the blue, solid distribution in Figure 7) would estimate a *44 percent pure risk* for this insurance contract. Insurance is neither an efficient nor sustainable solution for managing such a large portion of the insured's risk. From 1962 onward, insurers would incorporate the new rainfall experience each year and adjust the pure risk for trend. Given how different the estimate of pure risk is from the actual risk experienced during the following 26 years, whatever adjustments the insurer made would likely be insufficient and the insurer would continue to lose money until it decided to stop offering the product.

**See Collier, Skees, and Barnett, 2009 for more details of this example.*

Figure 7 Comparing Rainfall Distributions and Indemnities, Created Using the Vantage Points of 1962 and 1990

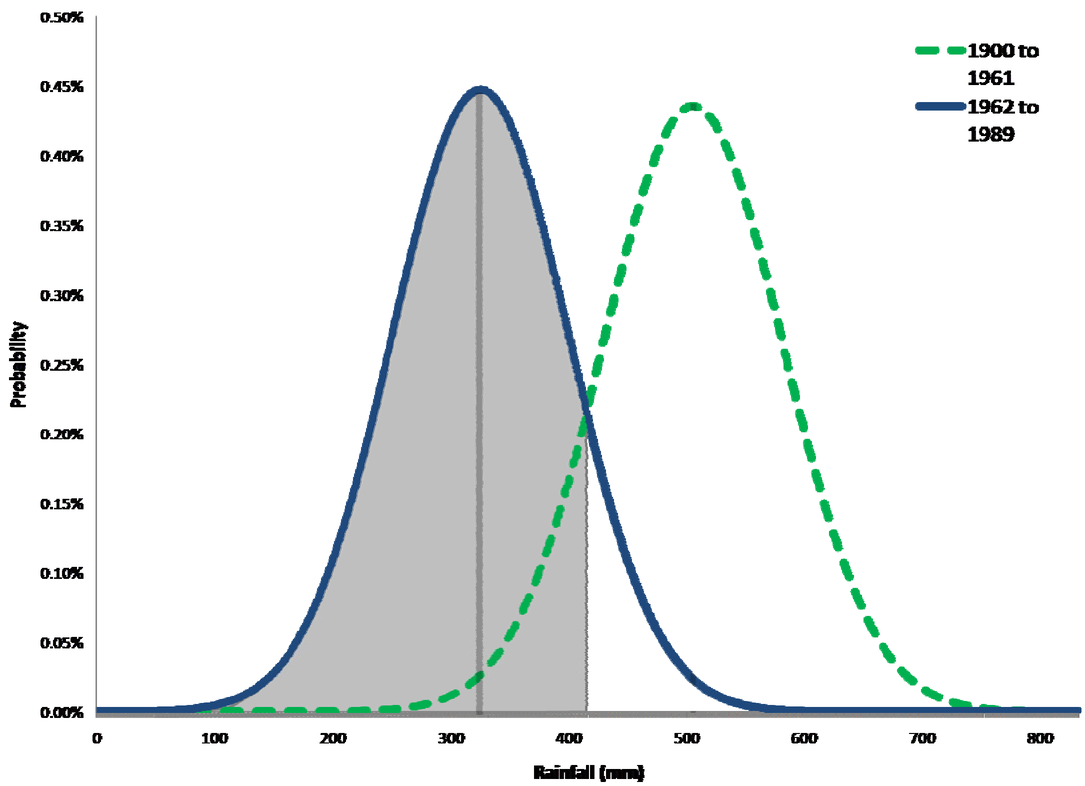

Vantage Point	Data	Central Tendency	Standard Deviation	Pure Risk
1962	1900–1961	510	77	2%
1990	1962–1989	328	76	44%

Source: Collier, Skees, and Barnett, 2009; The figure is based on data provided by International Research Institute for Climate and Society, Columbia University.

Practitioners also examine the data for clustering. Clustering describes the tendency of some catastrophic events to occur in a temporal pattern. For example, in some regions it may be that if drought occurs in one year, it is more likely to occur in the next year as well. If such interannual cycles occur, individuals will use this information when making the decision to purchase insurance, which creates an adverse selection problem for insurers. In some cases, insurers can adjust for this by offering multi-year contracts or by pricing the insurance assuming that individuals will use these signals in making a purchase decision. Both of these alternatives can cripple the insurance market: multi-year contracts require long-term commitments that may discourage potential buyers while pricing the insurance based on the assumption that people will only buy when the risk is high will increase the price of the insurance, also reducing insurance purchasing.

In assessing the pure risk, it is best to consider an estimated distribution as a snapshot at a particular point in time. The following year, new observations will emerge and may change the distribution. Practitioners are in the difficult position of interpreting the implications of new observations. Emerging trends are particularly difficult to identify. For example, if two extreme events occur in consecutive periods, it may be an unusual sequence of events, or it may indicate a new regime for the weather risk. Practitioners can adjust the price of the insurance each year based on new experiences. Still, such uncertainty often leads insurers to add an ambiguity load to the cost of insurance to protect against the possibility that they have misestimated the pure risk due to misestimating the moments of the distribution or clustering. In the next chapter, we describe some of the key characteristics for choosing a data source to estimate the pure risk.

3.2.1 Key Characteristics for Estimating the Pure Risk

Five characteristics of the data source help practitioners evaluate if it is suitable for estimating the pure risk: 1) historical length; 2) spatial specificity; 3) temporal specificity; 4) completeness; and 5) validity.

Historical Length. The length of time series determines how well the distribution can be estimated. A general benchmark for index insurance has been at least 30 years of data. While this standard is somewhat arbitrary, it has some implications for understanding the distribution of a weather variable (Box 2 in Chapter 1).

A growing number of alternative data systems have emerged to replace or supplement data of short time series. Some forms of satellite-based data, such as NDVI, were developed in the late 1970s and early 1980s and so have generated a length of time series data that may be suitable for estimating distributions. Reanalysis data (a term used to describe products that combine weather data from many sources) are often used to supplement short time series of weather station data. Chapter 4 contains more detailed descriptions of some of these alternative data sources.

Spatial Specificity. Spatial specificity describes the level of detail with which the data system can assess the weather risk in the region. As discussed above, spatial requirements for the weather risk depend on the spatial correlation of the event. For data sources on the ground such as weather stations, spatial specificity refers to the distance between weather stations. For satellite-based data, spatial specificity is the level of resolution of the index. For example, NDVI is often measured using pixels representing areas of approximately 1 km^2 (Box 8 in Chapter 4).

Temporal Specificity. Temporal specificity, the frequency with which a data system records measurements, has important implications for estimating some weather events and for mapping the index on to losses. The temporal specificity of some data systems may be inadequate to capture certain weather risks. For example, potential insureds in the target market may report sudden extreme rainfall that causes flash flooding within a few hours. There may be weather stations in the region but if the data are only collected on a weekly or 10-day basis, they will lack the temporal specificity needed to identify sudden extreme rainfall.

Completeness. Missing values in a data series can occur for many reasons — civil unrest, loss of funding for meteorological services, human error, etc. Missing values can be estimated using statistical methods such as interpolation or using other data sources. Sometimes data are missing due to a catastrophic event (e.g., a flood washes away a weather station). Dealing with missing data in these cases is problematic. Observations of extreme events are very important for estimating the tail of the probability distribution and the statistical methods used to replace missing values tend to underestimate extreme values.

Validity. Having trust in the validity of data is important for both practitioners and target users. Ideally, historic data would have been collected by an institution that is not likely to be pressured into altering data values. Typically, national meteorological associations can fulfill this role sufficiently, but sometimes the target market may prefer an alternative data source.

Practitioners must also consider possible effects of changing technology on data values. For example, new rain gauges may provide data more regularly and more accurately than older technology. Likewise, many of the satellite-based technologies are regularly updated and care must be taken to ensure the data used throughout the time series are consistent.

3.2.2 Supply Priorities

Practitioners working to supply index insurance are concerned that historic data provide an accurate estimation of the pure risk. They are particularly concerned with the possibility of underpricing the risk which can create large losses. As a result, reinsurers typically add ambiguity loads to premiums when limited data increase the likelihood that the pure risk will be estimated incorrectly. In some cases, practitioners may decide the potential for misestimating the pure risk is so great that they are unwilling to offer index insurance.

3.2.3 Demand Priorities

While the target market may not be exposed to the methods used to price the pure risk, it is in the interest of the target market for the data to be a long and accurate time series. When practitioners can confidently price the pure risk, premium loads will be lower, reducing the cost of the insurance product to potential insureds.

3.3 Settlement Index

The index used to settle the insurance payment is the core element of index insurance. It is the contractually binding mechanism used to determine insurance payments. It is also the index that the insured must evaluate when deciding whether to purchase the insurance. Therefore, insurers want to choose this index carefully. In many cases this index will be based on the same data source that is used to estimate the pure risk, such as if weather station data are used for premium rating and as the index for indemnities. But this need not always be the case. For example, the index used for settlement of payments may be based on a relatively new weather station while the pure risk has been estimated using longer series of historical data from surrounding weather stations.

3.3.1 Key Considerations for the Index

Following are some of the key characteristics that help practitioners evaluate if a data source is suitable to serve as the index.

Relationship between the Pure Risk Estimate and the Index. Given the importance of accurately pricing the pure risk, it is crucial that practitioners understand the relationship between the data source(s) used for estimating the pure risk and the index used for loss adjustment.[11] This is especially important for extreme values.

[11] If the index used for loss adjustment is the same as the index used for calculating indemnities, no action is needed in this regard.

Spatial Specificity. Spatial specificity has significant implications for basis risk, and is probably the biggest constraint to index insurance scalability. The two most prominent index insurance programs that rely on weather stations, India and Malawi, have both been constrained by insufficient weather station infrastructure. Many regions of the world, especially much of Africa, have even less developed weather station infrastructures.

Temporal Specificity. Many regions may also be constrained by weather stations that report data too infrequently (e.g., biweekly or monthly) to be useful for many weather risks. It is worth repeating that the spatial and temporal specificity demands on an index depend specifically on the spatial and temporal presentation of the weather event and the type of contract designed.

Validity/Security/Credibility. Both the target market and the suppliers of index insurance must rely on the index used for loss adjustment. Potential insureds are not likely to buy the insurance if they do not trust the validity and objectivity of the data source on which the index is based. This can occur because potential insureds do not understand the data source. It may also be that potential insureds understand the data source but believe that the data may be manipulated for political or other purposes.

It is also important to note that basing indemnities for an insurance product on a particular data source may create new threats to the security of the data source. Obviously this is a bigger concern with weather station data sources than it would be with satellite-based data sources. Regardless, threats to the security of the index should be anticipated and addressed proactively.

Completeness/Permanence. Practitioners want to choose an index that is certain to be available for the upcoming season. For example, if an insurance product uses rainfall gauges managed by a national meteorological institution as the index and the national government reduces funding for data collection and maintenance, the credibility of the index insurance program is challenged. Additionally, practitioners commit substantial resources to designing a product so they typically want to use an index that is likely to be available for many years into the future.

3.3.2 Demand Priorities

The core concern of the target market evaluating the product is how well it is likely to insure their risk. Basis risk is a major component of the consideration, but as we have described, basis risk is a complex issue. Many researchers have estimated basis risk in terms of household-level weather and crop-yield correlations. Obviously this is not an appropriate measure of basis risk for a risk aggregator product. But we would argue that it is also not necessarily an appropriate measure for a household product. Our experience has demonstrated that extreme weather events affect stakeholders in many ways. Because of the diversity of risk management strategies of households in rural areas, it is not convincing that weather/yield correlations adequately capture the ground-level consequences of a catastrophe. Farmers may save their yields but lose quality, or they may irrigate at substantial cost. No researcher or practitioner can hope to identify the effects of all of these management and coping strategies, especially for heterogeneous households.

Potential insureds consider the effects of extreme weather events on a variety of factors including the long-term health, safety, and well-being of family members. In short, they evaluate an insurance product simply in terms of whether or not they are likely to be better off having purchased the product. This process may be a complex evaluation of their risk exposure and any alternative risk management strategies, and can also include the recognition that, in many cases, an imperfect insurance instrument can be better than no insurance at all.

Credibility of the index is also a concern of the target market. We have already mentioned the possibility that potential insureds may not trust data provided by the national meteorological service. In other cases, they may not be willing to rely on data collected from satellite-based platforms. Field research in Kenya regarding an index-based livestock insurance (IBLI) pilot using NDVI has included assessments of potential buyer response to this product. (Chantarat, 2008; Chantarat et al., 2009). In early 2010, some 2,000 Kenyan herders purchased the IBLI.

3.3.3 Supply Priorities

The primary concern of those supplying index insurance is the validity and security of the data source as it serves as the basis of the insurance. Practitioners will want to dialogue with reinsurers during product development to make certain the index of interest is acceptable to the reinsurer. A second priority is the permanence of the data source. Insurers will not invest in building a market for an index insurance product unless they are confident that the underlying data source for loss adjustment is likely to be available into the future.

Chapter 4 Real-World Data Constraints: Limited Weather Station Infrastructure and Opportunities for Satellite-based Technologies

This chapter extends the conceptual framework of Chapter 1 and the real-world focus in Chapters 2 and 3 to identify some pragmatic implications regarding index insurance product development. It is divided into two subsections. First, we note that weather stations are underdeveloped in many regions and describe considerations for evaluating if weather station infrastructure is sufficient to support an index insurance product. As part of this discussion, we consider the potential cost of populating a region with weather stations and maintaining this data system. Second, we evaluate practices of some current index insurance programs. The findings in each of these subsections motivate our conclusions for advancing weather index insurance, given the prevalent data limitations.

4.1 Data Availability: Weather Stations

Weather station (or rain gauge[12]) data have been the primary data source for weather index insurance programs thus far. Yet the weather station infrastructure in many regions of the world is underdeveloped, and may be insufficient to support index insurance products based on weather station data. The U.S. National Climatic Data Center (NCDC) archives weather station data as part of the World Meteorological Organization (WMO) World Weather Watch Program according to WMO Resolution 40 (Cg-XII).[13] This data source reveals very limited weather station coverage in some regions of the world, in particular for parts of Africa. For example, for stations reporting daily values NCDC provides data on 37 stations in South Africa, 12 in Sudan, 4 in Botswana, and 1 in Congo. Still, the NCDC archives tend to underreport the actual number of weather stations in some countries. Many reasons may motivate countries to underreport to international archives. Low data quality from some stations is one. Civil and political unrest and the very limited budgets of many African countries would contribute to low quality data. But also countries may prefer not to publicly share all of their weather data. In some cases, country meteorological offices maintain a policy of charging for data. Thus, from archives such as that held by the NCDC, it can be quite difficult to estimate what data the country has actually collected. Another source reports that in Africa, there are roughly 1,000 quality-controlled weather stations (Funk et al., 2003). If weather stations were uniformly distributed, it would indicate a weather station density of roughly one weather station every 60,000 km^2 (Funk et al., 2003). These weather stations are not uniformly distributed, of course, yet their density and the NCDC archives are stark indicators of how severe data constraints are in Africa. Moreover, the

[12] A rain gauge is a device used for measuring rainfall and can be much smaller and less expensive than a weather station. Weather stations may also measure rainfall (through a rain gauge) but may measure other weather variables as well, such as temperature or wind speed.

[13] NCDC maintains at least two sources of data collected from weather stations around the world. Global Summary of the Day (GSOD) surface data can be accessed at
ftp://ftp.ncdc.noaa.gov/pub/data/gsod/
Global Historical Climatology Network (GHCN) daily data can be accessed at
http://www.ncdc.noaa.gov/oa/climate/ghcn-daily/index.php?name=data
These data archives are among the most comprehensive sources of publicly available weather data in the world.

existing weather station infrastructure in some regions, such as the Sahel, has deteriorated since the 1980s due to the costs of data collection and systems upkeep (Ali et al., 2005). For example, in recent work performed in Mali, we learned that, of some 85 weather stations that were in service at some level during the time period of 1951 to 2007, only 10 were operational in 2007 (Hartell and Skees, 2009).

Estimates on data that are may be particularly telling for index insurance. Consistent with our emphases on the reliability and validity of data sources, practitioners should question whether data that do not meet the following standards are a suitable for supporting an index insurance product: that the data are quality controlled, reported based on some minimum standards, and publicly accessible. When the data used to generate the index for an index insurance product must be purchased, these ongoing costs must be passed on to the insured, and if the data are not publicly accessible, the insured will not be able to verify the index value unless they too pay for the data.

4.1.1 Evaluating the Sufficiency of the Weather Station Density

Given that we know the density of weather stations in a region, the logical next question becomes: is this density of weather stations sufficient to support an index insurance product? Of course, this is a difficult question and motivates much of the work in this document. It is ultimately a question of basis risk. Returning to our conceptual model, practitioners are working with three probability distributions (losses, cause of loss, and the index) and two mapping functions (losses ←→ cause of loss and cause of loss ←→ the index). Chapter 2 describes how to qualitatively assess the distribution of losses and its relationship to the cause of loss. . Chapter 3 discusses considerations for selecting an index — that is, it discusses characteristics of a data system that would improve the estimation of the probability distribution of the index. The current question, whether the density of weather stations is sufficient, is a question of the relationship between the cause of loss and the index data (cause of loss ←→ the index). Specifically, it is a question of how close a geographic point of interest has to be to a weather station for the weather station to adequately measure the weather phenomenon at the point of interest. Of course, the answer to this question will depend on the specific weather phenomenon being measured. We use rainfall to illustrate these principles but other weather phenomena (e.g., temperatures) are likely more spatially correlated and would thus require less density of weather stations. Box 6 provides an empirical example for a weather index insurance pilot in Vietnam.

Box 6 Vietnam: Weather Station Infrastructure and Product Offerings

Our ongoing work in Vietnam provides a good illustration of how the lack of rainfall stations can influence what is possible in terms of product design. Vietnam generally has a strong infrastructure for weather stations. Authorities have been particularly diligent about putting weather stations near major crop production regions. As a centrally planned economy, Vietnam policy makers have influenced where certain crops are grown. The project focuses on the important Robusta coffee producing region of Dak Lak Province. Station information and rainfall observations were obtained from eleven weather stations throughout the coffee growing area. Of these, five stations were selected for the initial pilot because their available rainfall data were considered adequate to support a drought insurance product. The insurance is designed to compensate coffee growers for the consequential losses of severe drought occurring at the beginning of the normal monsoon season.

Concentric circles are drawn around the pilot stations at 5, 10, and 15 kilometer radii. Given the contract design and our analysis of the spatial presentation of the risk, the insurable area would ideally encompass an area no more than 10 kilometers from a weather station in order to minimize basis risk. However, the resulting "islands" of insurable areas were also considered to be problematic from an insurance operation standpoint. Consequently, the insurable area is expanded to +/− 15 kilometers depending on the lay of commune boundaries to develop a single contiguous insurable area.

Insurable Zones in Dak Lak Province

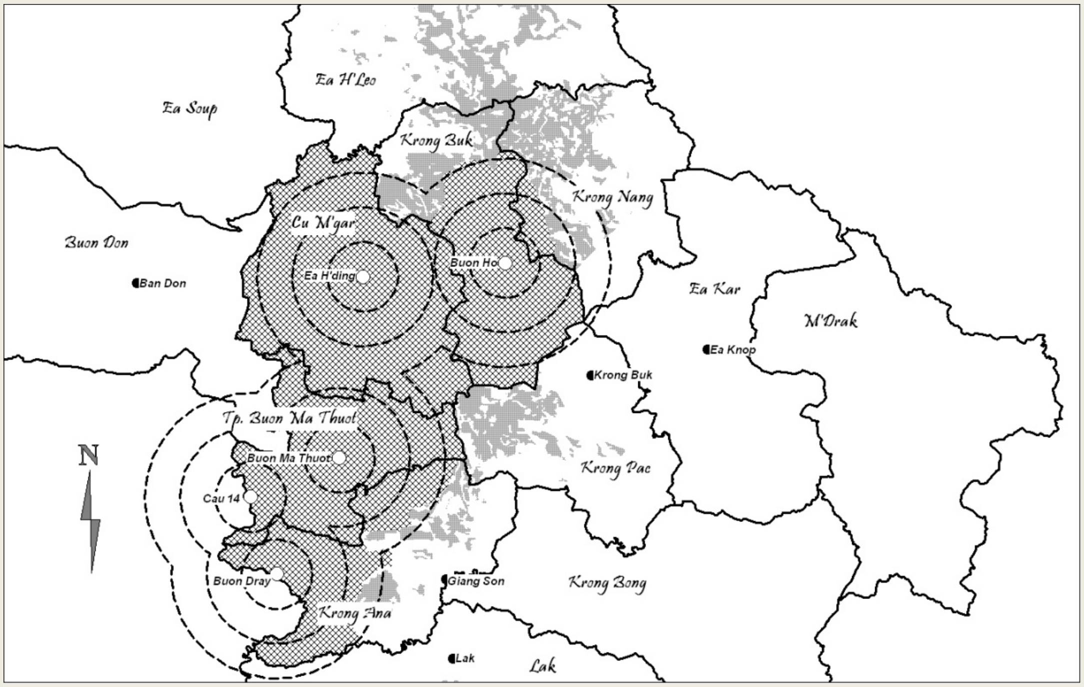

Source: Authors
Note: Hatched area indicates where insurance is feasible in the initial pilot

This area represents some of the densest coffee plantings in the region, approximately one half of the total area planted. To cope with the possibility of increased basis risk, insured growers in areas that intersect two stations at the 15 kilometer distance are given the choice of which station to associate with, based on their knowledge of local weather conditions. In addition, the threshold for payouts is set at significantly catastrophic levels to help avoid the possibility of payments without loss. This project demonstrates that, even in a country that has a strong weather observation infrastructure such as Vietnam, difficulties are frequently encountered due to data quality and density of observation points that limit insurance possibilities. The project also demonstrates how it is often necessary to find innovative solutions to technical limitations in order for an insurance to be operationally feasible.

An important area of meteorological research is comparing the accuracy of alternative data sources in estimating rainfall values. Thus, an important question for this research is how well a data system (e.g., a grid of weather stations) is likely to predict rainfall values in the region. In this context, the basis risk associated with differences between the cause of loss and the index is described as estimation error. Consistent with the theme of this of this document that weather is context specific — it depends on topography and other factors influencing the spatial and temporal presentation of a weather event — the meteorological literature identifies many variables that hinder rainfall estimation. Lebel and Amani (1999) report that the following factors contribute to this estimation error:

1. Size of the area of estimation;
2. Number of weather stations;
3. Geometry of how the weather stations are dispersed;
4. Spatial presentation of rainfall;
5. Temporal presentation of rainfall (e.g., rainfall accumulation);
6. Mean rainfall depth;
7. Type of rainfall;
8. Meteorological conditions; and
9. Season.

These researchers develop an error function that captures many of these variables (Ali, Lebel, and Amani, 2005). They use this error function for making comparisons across rainfall data sources (e.g., weather stations to satellite data sources, Ali et al., 2005). It should be noted that Ali, Lebel, and Amani (2005) develop the error function to be a general form for comparisons across regions and data sources; however, their research is in applications in the Sahel and thus the function may have unforeseen regional influences. As noted above, the presentation of rainfall can vary greatly across regions of the world, and even across seasons in the same region. We provide their error function here to illustrate the relationship between this form of basis risk (error) and these important variables.

$$e\left(A, N_g, K_T, P_T\right) = \frac{C_1}{N_g^{0.5} K_T^{0.3} P_T^{0.2}} \left[C_2 + C_3 log\left(\frac{A}{N_g} \right) \right] + C_4$$

where *e* is the error, *A* is the area in km^2, N_g is the number of weather stations, K_T is the number of rain events during the period of measurement, P_T is total rainfall during the period of measurement, and the C elements are parameters fitted based on local conditions (e.g., topography, geometry of the weather station dispersion, etc.).

Many of the relationships in the error function are not surprising but are worth noting. We consider the marginal effect of each variable — that is, the effect of increasing this variable holding all other variables constant. First, increasing the area assessed A increases the error term. The number of rain events and total rainfall are inversely related to error so rainfall estimations in regions that experience generally few rainfall events and low rainfall tend to have higher error than rainfall estimations in regions generally experiencing more events and higher rainfall (Ali, Lebel, and Amani, 2005). Perhaps the most important implications are that the number of weather stations N_g is inversely related to the error term and the error term is convex in N_g.[14] This implies that increasing the number of weather stations decreases the magnitude of rainfall estimation errors but each additional weather station contributes less and less to decreasing the magnitude of the error.

[14] That is, $\frac{\partial e}{\partial N_g} < 0$ and $\frac{\partial^2 e}{\partial N_g^2} > 0$.

Lebel and Amani (1999) estimate the error associated with using 10-day cumulative rainfall values for a 100 km by 100 km area with 10 weather stations in the Sahel. If weather stations are distributed to provide equal coverage across the 10,000 km^2 area, any point in the area should be within approximately 18 km of a weather station.[15] Estimation error tends to be 10 to 15 percent for this scenario. In other words, under these assumptions, the actual 10-day cumulative rainfall at any point within 18 km of the weather station may differ from the rainfall measured at the weather station by as much as 10 to 15 percent. Lebel and Amani also note that weather stations are of much lower density than this for almost all of the Sahel.

4.1.2 Estimated Costs of Populating and Maintaining Weather Stations

What would it cost to increase the density of weather stations on which index insurance could be based. Instruments for measuring weather phenomena come in a wide range of prices. As an example, consider that a simple plastic rain gauge can be purchased for as little as USD 5.00 while more advanced automatic weather stations with rain gauges cost between USD 12,000 and USD 15,000.

At a minimum, a basic automatic rain gauge would cost several hundred dollars (US). In addition, several add-ons would likely be required. Data loggers are not standard equipment on most automatic rain gauges. Data loggers come with a variety of features that influence the price. The most simple data loggers, that do not have remote access, cost around USD 130 but they also require software that costs anywhere from USD 55 to USD 125. To access the data logger remotely, more complex data loggers are required. These are priced from USD 435 to USD 1,440. The user may also need to purchase additional software that can cost as much as USD 600. Remote access to the data logger also requires that the data logger be equipped with either a landline phone, mobile phone, radio, satellite capability, or some other mechanism for transmitting data from the data logger to a remote computer.

Automatic rain gauges also require a power source. Some rain gauges are battery powered, though battery life can be limited to one year. The more advanced gauges are electric or solar-powered with backup batteries. Solar panels cost around USD 135 for 5 watts, which is enough to power some rain gauges, but others might require more power and more expensive panels. Other add-ons include mounting plates and brackets (USD 40–USD 90), wind screens (USD 465), stands (USD 100–USD 300), calibrators (USD 125), and various types of cables and adapters (USD 100–USD 300).

For most automatic rain gauges, regular maintenance is required to clean out any debris (leaves, sticks, etc.) from the rain gauge and wipe out mud and dirt. Sensors may need to be replaced on a regular schedule (e.g., every six months or annually). Regularly scheduled recalibration is also required. While parts (such as washers and bearings) will occasionally need to be replaced, rain gauges typically last between 5 and 20 years. The lifespan of a rain gauge depends on the location; a location with many storms (sandstorms, windstorms, rainstorms, etc.), will wear a rain gauge more than one with only occasional storms.

Based on this information, it seems that a basic automatic rain gauge with remote access to the data logger would cost, at a minimum, between USD 1,500 and USD 2,000. A weather

[15] Each of the 10 weather stations would cover an area of 1,000 km^2. Any point within the 1,000 km^2 area of a circle around the weather station should be within approximately 18 km of the weather station; $area = \pi \times radius^2$ so if area = 1,000, the radius is 17.84.

observation station that measures weather phenomena other than just rainfall would cost significantly more. This estimate does not include the cost of installing the rain gauge, the power source, the remote access mechanism (e.g., internet or mobile phone), or any security measures (e.g., fencing). Nor does it include the cost of shipping the equipment and materials to the site. Also not included are the recurring costs for routine maintenance. These recurring costs can be quite significant. Many lower income countries have weather stations that have been abandoned because of an inability to pay for the recurring costs of maintenance and upkeep.

So, what quality of rain gauge is required to support an index insurance offer? For example, would all the rain gauges used for index insurance offers in a given area need to be fully automatic with data loggers that support remote access? Unfortunately, there are no straightforward answers to these questions. While reinsurers would like weather index insurance to be based on the best possible weather measuring instruments, they may be willing to utilize data from something other than the most advanced weather stations if cross-verification can be done using data from nearby fallback (or “buddy”) stations. Alternatively, satellite or reanalysis data can sometimes be used for cross-verification.

The answer to the “how good is good enough” question may also depend on various characteristics of the insurance contract. If the contract triggers indemnities for relatively small rain shortfalls (e.g., 10 percent below the expected value), reinsurers are likely to insist on more sophisticated rain gauges. If the contract triggers indemnities only for the most extreme droughts (e.g., 40–50 percent below the expected rainfall value), reinsurers may accept less sophisticated rain gauges because they can cross-verify such catastrophic events using other data sources. If the contract is for a relatively short period of time, e.g., cumulative rainfall over a 10-day period, reinsurers will want better rain gauges. If the contract is for cumulative rainfall over a seasonal period of 60–90 days, reinsurers may accept less precise instruments.[16]

Thus, there is no clear cut answer to the question of what it would cost to populate a region with weather stations that are adequate to support index insurance offers. Instead the answer depends to the characteristics of the insurance contract and the availability of alternative data sources that can be used for cross-verification. For automatic weather stations with data loggers than can be accessed remotely, one must also consider the costs of shipping, installation, providing a power source, providing a remote access mechanism, and any security measures. It is also critically important to recognize that this is not just a one-time cost. Any weather station must be maintained. Fully automatic weather stations must be maintained by highly skilled professionals. The recurring cost of performing maintenance on automatic weather stations located in rural areas of lower income countries is likely to be quite significant.

4.1.3 Spatially Interpolated Weather Station Data

The discussion in the previous chapters suggests that the availability of weather index insurance will be severely limited if each index must be based on actual data from a single weather station. In areas where a weather station is not available, an alternative might be to create indexes based on spatially interpolated data from surrounding weather stations.

[16] For temperature-based index insurance, reinsurers will require very sophisticated weather stations if insuring against minimum or maximum temperatures during a window of time. Less sophisticated instruments may be acceptable for contracts based on average temperatures or cumulative temperatures (e.g., growing degree days).

Spatial interpolation describes the process of using data from weather stations to estimate weather variables (e.g., rainfall) at a specific point between these stations. The most basic spatial interpolation models assume that the available data are distributed independently according to the normal distribution, meaning that taking an average of the observed data points should give a reliable estimate of the value at the unobserved target point. However, weather variables are not always normally distributed or independent. Cumulative rainfall, for example, is frequently modeled using a Gamma distribution and during certain times of the year the amount of rainfall in one area may be highly correlated with the rainfall in a neighboring area.

Kriging is a spatial interpolation technique designed to correct for these correlations and give a weight to each observed data point based on its expected relationship to the unobserved point whose value is being estimated — this expected relationship depends on geographic distance, topography, etc. (see Bohling, 2005, for a general introduction to kriging). Kriging is a linear regression model, which provides not just an estimate of the unknown value but also a rough approximation of its accuracy. While there are many types of kriging, each distinguished by the way it assigns these weights, the easiest way to think about kriging is to imagine that you have the set of observations (lettered A–E) and are trying to measure a value at some unobserved target point X as illustrated below:

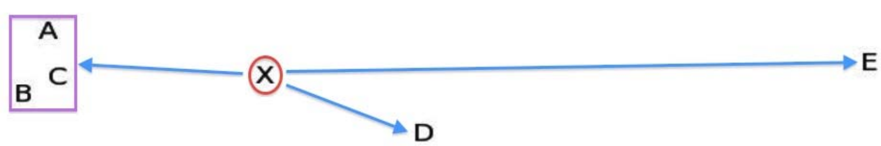

In this case, rainfall data from point E would have a low weight in the equation (Ali, Lebel, and Amani, 2005; see the equation in Section 4.1.1) because it is far from X. Points A, B, and C will also be discounted because they are located so close to each other that there is likely a high covariance between the observations taken at each of these points — meaning that they are, to some degree, redundant. Point D will likely have the highest weight because it is both close to X and is not overlapping with other values from the same neighborhood (low covariance with other observed values).

There are a few basic types of kriging that may prove useful in constructing weather indexes. *Simple kriging* assumes a single mean for the entire area of interest (similar to estimating based on the normal distribution). In *ordinary kriging*, the weights that are assigned to various data points reflect the assumption that individual regions each have a local mean. The boundaries of those regions are determined by a variogram, a statistical measure specific to the spatial distribution. In *universal kriging*, these local means are replaced by linear functions that estimate how values change as they approach the estimation target point. Finally, *indicator kriging* uses some underlying variable that can be measured at the target point as well as the known data points and tries to create an index relating the easily observed variable to the data-scarce target variable (rainfall). Sophisticated kriging models sometimes also use the general pattern of the spatial and temporal presentation of the weather event to enhance point estimates.

While kriging is a very robust estimation process in many circumstances, giving estimates with low mean squared errors, it suffers from one particularly important drawback. Kriging systematically underestimates the likelihood of extreme values at unobserved target points. It is by definition a weighted average of observed results. Consequently, its estimates tend to be drawn down by averaging the extreme observations with less extreme observations in the data. This is an important shortcoming for potential index insurance applications.

4.2 Reaction and Looking Forward

The data available from the NCDC (and others) indicate that weather station infrastructure is underdeveloped in many regions, especially regions of Africa, to a degree that likely prevents weather index insurance products based on weather stations or at least limits the types of products that could be offered. Our analysis of the cost of purchasing, installing, and maintaining new weathers stations indicates that these costs can seem quite high if the sole purpose is to support a weather index insurance program. A several thousand dollar per station investment could take years to recuperate based on the premium payments from poor households in the rural regions that the weather station would serve. Nonetheless, it is certainly worth noting that systems which provide publicly available weather data are public goods that can provide countries with many positive benefits beyond the potential for a weather index insurance program.

It is the ongoing mission of WMO and others to help governments appreciate the many social benefits of strong weather data, for example, for improved disaster response. WMO (2008) developed a report highlighting the challenges of, and rationale for, developing adequate weather information systems in developing countries. We describe some of the WMO findings in Box 7.

Box 7 World Meteorological Organization (WMO) Perspective: The Needs of National Meteorological and Hydrological Services (NMHS) Providers in Africa

WMO developed a report highlighting the challenges of adequate early warning systems in developing countries. Records of accurate and frequent weather observations provide countries with the ability to predict and plan for hazardous weather events. Natural disasters often result in casualties and resource losses, and 90 percent of the disasters from 1980 to 2005 were weather or climate related (WMO, 2008, p. 2). Weather related disasters are particularly devastating to lower income countries where early warning systems also tend to be underdeveloped and underfunded (p. 18). Instead of risk assessment and efforts that prevent casualties and economic disaster, these countries tend to focus on actions to take after a weather crisis (p. 16). Early warning systems include weather analysis, forecasting, and warning capacities and data collection (e.g. from observations, radar, satellite, etc.) and analysis by skilled workers, data storage through computer systems and information technology, infrastructure maintenance, and information dissemination. The WMO reports that ongoing funding is needed for these services provided by the National Meteorological and Hydrological Services (NMHS). Having a coordinated system that has clear technical guidelines implemented by a well-trained staff and information technology to support the communication and data transfer needs can protect the people and properties of every country (p. 17).

If they have suitable technology and funds, NMHS providers can provide essential services such as warnings for weather related hazards, assistance to emergency response organizations, data analysis for development and crisis preparation, education of the public on potential disasters and the actions to take before, during, and after such disasters, and recommendations to improve emergency notification and response (p. 6). These services can also help scientists monitor trends to guide investment policies, provide methods for improving crop production by indicating prime sowing and harvesting periods, indicate areas of land that are vulnerable to flooding, and prevent the loss of life and resources through emergency preparedness and response data. However, less-developed areas, including many of the countries in Africa, are not able to fund these services.

Underdeveloped early warning systems exist in almost all lower income countries; however, since 33 of the 50 countries the United Nations considers "least developed" (p. 37) are in Africa, the most significant NMHS needs are there. In a survey given by WMO to African NMHS providers, 100 percent of the respondents report that technical improvements to the warning systems would enable them to improve their disaster prevention capabilities (p. 30). Ninety-two percent of respondents also say that a better prepared staff would also improve the risk assessment ability of NMHS providers (p. 30). Most (92 percent) of the responding NMHS providers in Africa also think that government needs to recognize the importance of the NMHS in reducing weather risks (p. 36). According to the WMO report, 96 percent of respondents claim that they are unable to provide the best services because of a lack of resources and a weak infrastructure, with limited funds and well-trained staff cited as the biggest problems (p. 37). The report concludes that many African NMHS providers need to improve record-keeping, policy guidelines, organizational partnerships and memberships, observation networks, telecommunications, warning systems, quality control, and risk assessment tools, among others (pp. 39–42). These improvements cannot be made without continued assistance from organizations and governments. Most projects that have attempted to assist African NMHS providers have a limited timeframe, but the improvements needed in Africa require long-term support that typically must come from government (p. 43).

Source: WMO, 2008

Given the high costs of purchasing, installing, and maintaining new weather stations, it may be difficult to scale up weather index insurance offers into areas that have sparse weather stations. This is particularly true for products that require estimating rainfall at a specific point (e.g., household products). In response, we arrive at two conclusions to guide future index insurance product development. First, in many regions, due to data constraints the development of risk aggregator products will be more feasible than for household products. While some regions may have data systems sufficient to support the development of scalable household products, these are likely to be the exception rather than the rule. Instead, we encourage the development of risk aggregator products, which tend to have much lower data requirements. Second, since weather station infrastructure is likely inadequate to support scaled-up index insurance offers in many lower income countries, it is important to investigate the potential for using data collected from alternative sources, such as satellites. We discuss special considerations associated with remotely-sensed data, and consider the potential index insurance applications of these data.

4.2.1 Data Constraints Are Less Binding for Risk Aggregator Products than Household Products

The spatial specificity demands of index insurance depend greatly depend on product design — risk aggregator products require regional data (e.g., each data point represents an area of, say, 50 km × 50 km or 100 km × 100 km) while household products tend to require location-specific data. Risk aggregator products and household products tend to be quite different in terms of the cause of loss, which in turn has implications for the index. As an illustration, consider two insurance product designs in the same region — a household product and a risk aggregator product, both protecting against deficit rainfall. For the household product, the cause of loss is deficit rainfall at the site where the household lives and works (for simplicity, suppose households tend to own and live on small farms). Thus, the household product needs an index that approximates deficit rainfall at a specific, geographic location (e.g., the site of the household's farmland). The risk aggregator product is designed for a bank that lends to households and firms in a rural region. The bank is exposed to extreme deficit rain that results in high levels of loan defaults and savings withdrawals. This, in turn, leads to liquidity constraints and large losses for the bank. Thus, the cause of loss is extreme deficit rainfall over the whole region where its clients live and work. For the risk aggregator product, the index needs to approximate rainfall in the region, *not* at a specific location within that region.

This difference in data needs result in much less challenging data requirements for risk aggregator products. There are several reasons for this. First, it is easier to estimate regional values than to estimate the value at a specific geographic point in that region. Second, as indicated earlier, using several weather stations tends to lower estimation error compared to using a single weather station (Ali, Lebel, and Amani, 2005; see the equation in Section 4.1.1). Ali, Lebel, and Amani (2005) analyze two scenarios of comparable weather station density in the Sahel, one weather station in a 1° × 1° (latitude by longitude) area and six weather stations in a 2.5° × 2.5° area. They find that using the six weather stations results in lower error than using the single station. Third, risk aggregators tend to be affected by the most extreme events and so risk aggregator products are likely to be designed around extreme events. We propose a research agenda for the SKR based on our hypothesis that losses from extreme events tend to be more highly spatially correlated than for moderate losses (Chapter 5 and Appendix B).

In sum, risk aggregator products are likely to be a "low hanging fruit" for practitioners facing significant data constraints. Many regions of the world simply do not have the data capabilities to support household products, and the cost of developing weather station infrastructure is likely to be so expensive that the opportunity costs of weather index insurance development may exceed the benefits. As will be developed in a subsequent SKR, risk aggregator products have the potential to provide immediate positive benefits through enhanced performance of risk aggregators — providing important services to the region as well as the longer-term benefit of insurance market development, which may lead to a broader range of insurance products in the future.

4.2.2 Alternatives to Use of Weather Stations

Given the constraints associated with existing weather stations both in terms of the sparsity and the high maintenance cost, alternative data sources are needed to improve the scalability of weather index insurance. One promising source is the variety of observational measurements obtained from satellite-based or aircraft mounted remote sensors. In some cases there is as much as 30 years of data covering much of the globe. Other practical advantages of data collected from satellite platforms are the uniformity of the measurements and the systems, and the ability to standardize the data acquisition contracts. These factors improve the potential scalability of index insurance products using remotely sensed data. One limiting factor in the past was the computational power and data storage requirements needed to effectively and routinely use remotely sensed data. As these constraints ease, and as the science of these information systems advances, there may be significant opportunities to advance index insurance product development. As with any climate observation system, remotely sensed data has its own set of limitations and constraints.

4.2.2.1 Additional Considerations for Satellite-based Technologies

The considerations discussed in Chapter 3 also apply to remotely sensed data. However, because of challenges associated specifically with using satellite-based data for assessing the pure risk and for making contract settlement, we provide several considerations here that extend the discussion in Chapter 3.

Length and Continuity of Time Series. Some of the initial satellites collecting ongoing meteorological data went online in the late 1970s and early 1980s. Thus, they have an approximately 30-year length of available historical data. Satellite classification is related to the potential length and continuity of an observational data series. Net continuity in a data series, which can be related to satellite failure risk or the end of a specific program, must be evaluated during the design phase of an index program based on remote sensing. Understanding operational capacity, of having a continuity plan in place in case of satellite or sensor failure, will help in assessing the risk that index settlement data will not be collected. There are four classes of satellites (Hipple, 2010):

- Experimental: characterized by a limited time series and operational life.
- Research: data series may be lengthy or short although data cost may be low. Usually no continuity plan is in place so there exists insurance settlement risk from satellite failure. All NASA satellites are research orientated although they may feature lengthy data series, such as those from MODIS.
- Non-commercial operational: government operated systems, such as the NOAA-N Series (AVHRR), characterized by low cost data with continuity plans which reduces insurance settlement risk.
- Commercial operational: characterized by good continuity providing there is sufficient demand, and the possibility of high resolution imagery. Data are priced to incorporate operational costs and capital recovery which may be too expensive for some index insurance programs that do not have scale or adequate external support.

Calibration between and within Sensor Systems. Zenith angle changes due to drift and orbit changes, aging of the sensor, and other variables can create variations in the observations of a satellite-based device over time and must be detected and compensated for in the algorithms used to interpret the radiometric readings. Such challenges are especially true for data collected in the 1980s and early 1990s (e.g., see Box 8 for a review of AVHRR). Newer sensor arrays correct for some of these problems (Fensholt et al., 2009). A much investigated challenge is matching data collected from one sensor system with data collected from another (e.g., Ali et al., 2005; Fensholt et al., 2009; Uppala et al., 2005). Comparing across satellites using the same technology also requires calibration; however, the most challenging calibration is comparing across technologies. While newer technologies overcome problems with older satellite-based technologies, these changes can also create discontinuities in the data. While multiple systems created data redundancy, one must also ask how much variation is acceptable in index values across different technologies. In a related matter, the algorithms used may differ in subtle ways such that when applied to the same raw data they produce different results. Changes in the algorithm may also not be routinely reported, particularly from research satellites so understanding and documenting the algorithms used is an important part of the insurance design process. As Chapter 3 discusses, accurately assessing the pure risk is crucial to the longevity of the index insurance program so practitioners will want to make certain their data are consistent across time.

Spatial Specificity. Increasingly, new applications for satellite-based data are emerging. Much of the increase in data demand is associated with climate change analysis. While climate change and weather index insurance overlap in terms of important weather variables such as temperature and rainfall, climate models can often use data with low spatial specificity, e.g., 1° latitude × 1° longitude or greater. One degree of latitude or longitude is roughly equivalent to 111 km at the equator. Thus, a 1° x 1° spatial resolution provides one data point about every 12,300 kmPP2PP. Newer sensor systems, particularly those of commercial vendors, offer much finer spatial resolution, as small as 1 m.

Calibration to Ground Level. Similar to calibrating across sensor systems, researchers conduct perhaps even more analyses matching satellite-based data with ground-level data. Such calibrations can depend on the climatic and topographical features of a region. Thus, researchers make these comparisons for specific locations. Perhaps because they are calibrated based on the mean, satellite-based technologies tend to measure extreme events poorly (Kahel, 2009), an unfortunate feature for index insurance.

4.2.3 Evaluating Satellite-based Technologies

We consider several satellite-based data sources in the context of index insurance. While we have already discussed some of these data sources, this section provides a further review of the potential usefulness of these technologies. However, this is meant as an example of some of the main sources and is not an exhaustive review of instrumentation or potential applications. In particular, vendors of satellite-based data products and vendors of commercial operational satellite services are not reviewed.

4.2.3.1 Vegetative Indexes: NDVI

Vegetative indexes comprise a whole class of optical measures of the photosynthetic potential of the observed vegetative canopy as a result of total leaf area, chlorophyll levels, cover and plant structure. The measurements are used as proxies in estimating these and other plant canopy state variables which can be understood as an estimate of plant health (Heute et al., 2006). There are in excess of twenty different vegetative indices with most derived from sensor readings of either an Advanced Very High Resolution Radiometer (AVHRR) or a Moderate-resolution Imaging Spectroradiometer (MODIS) (Yang, Z. Willis, P., and R. Mueller, 2008). Other indexes can be created using a combination of optical and other sensory data such as moisture or thermal readings. The choice of vegetative or related index to use as a proxy for loss in an index-based insurance product will depend on how well the index correlates with the phonological cycle during critical time frames for a particular crop. As with any index, validation with in situ historical data is necessary, particularly for those years known to have generated significant losses.

The Normalized Difference Vegetative Index (NDVI) is one of the most commonly used satellite-based vegetative indexes for weather index insurance (Box 8). It has high spatial specificity, and a relatively long time series. It is frequently used as a proxy for drought conditions based on the estimate of plant health. An index insurance pilot based on NDVI is currently being developed to insure against drought risk for pastoralists in northern Kenya (Chantarat, 2008). The United States Risk Management Agency has since 2007 offered a pilot Pasture, Rangeland and Forage

(PRF) insurance product which is an area based vegetative index using NDVI data. In addition to the PRF product that is currently offered in nine states, a pilot apiculture insurance product has been developed that also utilizes NDVI data (Alston, 2010).

Box 8 Normalized Difference Vegetation Index (NDVI)

NDVI is a measure of vegetation density. Data are collected by Advanced Very High Resolution Radiometer (AVHRR) sensors on polar orbiting satellites managed by the U.S. National Oceanic and Atmospheric Administration ([NOAA]; NationalAtlas.gov, 2009). AVHRR measures both the visible light spectrum and the near infrared light spectrum. While plants absorb visible light for photosynthesis, they reflect near infrared light. Thus, discrepancies in the reflection of visible light to near infrared light are an indication of plant life (Weler and Heming, 2010). Researchers use an algorithm that transforms light wavelengths into estimates of vegetation density. Less healthy plants absorb less visible light so NDVI is also an indicator of plant health. By comparing historical NDVI values to present values, NDVI is being used to assess drought in some contexts (Bayarjargal et al., 2006; Peters et al., 2002). Satellite data also provide estimates of rainfall and temperature, which have been used in conjunction with NDVI data to create other drought estimation models; however, these different models yield differing results. Determining which NDVI-based models are most appropriate depends on the potential application of the model (Bayarjargal et al., 2006; White and Walcott, 2009).

The Moderate Resolution Imaging Spectroradiometer (MODIS) is a newer sensor technology than AVHRR and is implemented on two polar orbiting satellites, Terra and Aqua, which are managed by the U.S. National Aeronautics and Space Administration (NASA, 2010c). MODIS has a higher spatial resolution and overcomes some problems experienced by AVHRR (Weler and Heming, 2010). Studies comparing NDVI data from AVHRR and MODIS find important differences in these data sources (e.g., Fensholt et al., 2009). Newer sensor technologies (e.g., the Visible Infrared Imaging Radiometer Suite) are being advanced to replace MODIS in 2013 (NASA, 2010c).

Spatial specificity:	AVHRR, 1 km^2
	MODIS, 250 m^2
Temporal specificity:	Daily (both AVHRR and MODIS)
Length of time series:	AVHRR, 1981 to present (NationalAtlas.gov, 2009)
	MODIS, 1999 to present (NASA, 2010a)

AVHRR and MODIS have several data collection problems: imaging can be blocked by cloud cover and aerosols in the atmosphere, glare from the sun can saturate the color spectrum, and the satellites can malfunction (Weler and Heming, 2010). Because of these problems, researchers often use composite NDVI data that combines data from several days (e.g., a 10-day index; Chen et al., 2004) and several data pixels (e.g., 8 km resolution; Tucker et al., 2005).

Application of NDVI should be done with care. Despite its high spatial resolution, NDVI may not be useful for identifying drought conditions in a specific location. The method for predicting drought with NDVI depends on comparing the current value of NDVI for that location and time of year to NDVI values in previous years, which collectively is termed “normal.” On a small scale this approach can be problematic because land use may change. The biggest challenge in this regard is that different crops have different light requirements. Comparing across different crops, Thenkabail, Smith, and DePauw (2000) found crop growth was optimally modeled using different light bandwidths for different crops. NDVI values in a specific location may change if farmers plant a different crop than in previous years, or intercrop in some years and not others. One could envision a scenario in which a certain commodity price increases and the crop profile for a whole region alters, affecting NDVI values. Likewise, if farmers anticipate drought and plant different crops, it may change the NDVI assessment of drought in that location.

NDVI may also be problematic in situations where there is significant upper tree canopy or brush cover such that it becomes difficult to accurately estimate crop plant health. NDVI may also not be the most appropriate choice for an index product that focuses on yield outcomes when photosynthetic capacity is not the main yield determinant.

As a result, NDVI is perhaps best used as a gross indicator of plant health in a region. This lends itself well to the use of NDVI for pastoralists who are concerned about grassland vegetation in regions such as northern Kenya. Index insurance using NDVI is likely most appropriate for insuring against extreme droughts that are likely to have widespread effects. Additionally, it may be most useful for risk aggregators that are particularly vulnerable to regional effects.

4.2.3.2 Satellite Measurement of Rainfall

Over the past 10 to 15 years, several high-resolution products for estimating rainfall from orbiting satellites have been developed and implemented (Box 9). Satellite rainfall products use algorithms that combine data from multiple satellites to estimate ground-level rainfall. It is important to remember that these observations are made from orbit, looking down on and through clouds. To identify how much it rains from this perspective requires identifying both the intensity and duration of rainfall, yet because satellites are in orbit they are sometimes unable to observe a cloud for the duration of a storm.

Satellites estimate rainfall using infrared and passive microwave radiation data. Infrared temperature observations of the tops of clouds provide an estimate of the intensity of rainfall. Passive microwave radiation can be measured through clouds, identifying the energy emitted by rain drops to estimate the intensity and vertical distribution of rainfall. Infrared rainfall estimates tend to be less accurate than the passive microwave data; however, infrared data can be collected by more satellites. Given the coverage of orbiting satellites, infrared data can be measured continuously across the globe and are used to “fill in the gaps” when passive microwave data on rainfall are unavailable (NASA, 2010d).

Comparisons of satellite estimates of rainfall to ground-level observations show that satellite estimates contain significant error, but these estimates are improving. Generally, satellite rainfall products tend to correlate better to gauge-based data in warm seasons than in cold seasons, and correlate better in wetter regions than drier regions (Shen, Xiong, and Xie, 2008). Complex, varying terrains also challenge satellite rainfall products (Dinku et al., 2008).

Box 9 Satellite Estimates of Rainfall

We profile four satellite rainfall products that use algorithms to combine infrared, microwave, and (sometimes) rain gauge data to estimate rainfall. One of the satellites, TRMM, is the only satellite with weather radar and so has additional rainfall estimation capabilities (NASA, 2010b). One of the other satellite rainfall products, PERSIANN, also includes TRMM as part of its algorithm.

Product*	Spatial Specificity	Temporal Specificity	Time Series	Geographic Area	Data Type
CMORPH	0.07° x 0.07° (8 km at the equator)	30 minutes	December 3, 2002 to present	Global 60°N–60°S	MW, IR
TRMM 3B42	0.25° x 0.25°	3 hours	January 1, 1998 to present	Latitude 50°S–50°N	MW, IR, RG
PERSIANN	0.25° x 0.25°	30 minutes accumulated to 6 hours	1997 to present	Global 50°S–50°N	IR, MW, RG, TRMM
Africa RFE 2.0	0.25°	6 hours	January 1, 2001 to present	20°E–55°W 40°S–40°N (Africa)	IR, MW, RG

Intercomparisons of these rainfall products indicate that CMORPH tends to perform best (Sapiano and Arkin, 2009). However, CMORPH tends to significantly overestimate rainfall during warm seasons (Zeweldi and Gebremichael, 2009). In other studies, CMORPH underestimates rainfall. For example, Shen, Xiong, and Xie (2008) find that in China, CMORPH and PERSIANN underestimate rainfall, while TRMM 3B42 overestimate rainfall. In that study, the mean biases of satellite rainfall products ranged from -10 to +5.7 percent, depending on the product — the least biased product is one of the TRMM products and was -3.7 percent. A comparison of rainfall products in Africa indicates RFE 2.0, CMORPH, and TRMM satellites perform comparably overall (Dinku et al., 2008); however, on complex terrain, RFE 2.0 and TRMM 3B442 perform best and PERSIANN performs worst (Laws, Janowiak, and Huffman, 2004). Laws, Janowiak, and Huffman (2004) found satellite rainfall products tended to be positively biased for low daily rainfall values (values under 11 mm) and negatively biased for high daily rainfall (values greater than 20 mm).

Satellite-based estimates of rainfall are expected to improve in the future, especially as passive microwave sensors become more prevalent among orbiting satellites. For example in 2004, 6 passive microwave sensors were used by CMORPH (Joyce et al., 2004); as of December 2009, CMOPRH used 9 passive microwave sensors (CPC, 2009).

**This table was developed from product descriptions on NASA and NOAA websites. In the last column IR is infrared, MW is passive microwave, RG is rain gauge, and TRMM is the satellite Tropical Rainfall Measuring Mission.*

4.2.3.3 SAR

Synthetic aperture radar (SAR) is used to map contours of geospatial environments such as fault lines in earthquake-prone regions (Box 10). More importantly for index insurance, it can be used to map the contours of extensive flooding. Satellite-based systems for collecting SAR data share all the advantages of other satellite-based data systems (e.g., low marginal cost and large coverage areas). However, SAR has the unique advantage of being able to penetrate cloud cover.

Satellite-based SAR data also share a problem of other satellite-based data systems. There are gaps in the data collected because the satellite is only in a position to collect images for a given location for a few days out of every orbit cycle. This means that there are necessarily gaps in the historical record of images that may include key moments in weather or natural disaster events (Holecz, 2009). The question of SAR data coverage is further complicated by the fact that SAR satellite images are made to order. While a SAR satellite may be in position where it could be taking images of flooding in a given delta in Vietnam, for example, it may instead be recording data from other regions within its field of view that have a higher priority. This means that there is no guarantee that index insurance professionals can access a continuous historical record of SAR data for any given point on earth (MDA, 2009).

Researchers have begun using existing SAR systems for new purposes such as satellite-based measuring of wind speed in hurricanes. Such measurements remain in development and are not yet widely available (Schiermeier, 2005).

Box 10 Synthetic Aperture Radar (SAR)

Just as ship navigators use radar to map the contours of their environment regardless of fog or inclement weather, geoscientists use radar to monitor geospatial change when other imaging technologies would be obscured by cloud cover. In the past, the resolution of radar-based images was limited by the size of the antenna used to send and receive radar signals. The roughly hundred meter long antennas needed to produce high resolution images simply could not fit on an aircraft or satellite. The Synthetic Aperture Radar (SAR) process solves this problem, simulating long antenna by recording the echoes of radar beams emitted at regular intervals and compiling the resulting data as though it came from one long antenna (Sandia National Laboratories, 2008).

While areas of particular interest to geoscience, such as the major fault lines of California, are now monitored by SAR using unmanned aircraft, most SAR data relevant to index insurance come from one of the handful of SAR equipped satellites (Radarsat-2, Envisat, ALOS, TerraSAR-X, among others) orbiting the earth (NASA, 2010e; Lotsch, Dick, and Manuamorn, 2009).

Spatial specificity: Varies by satellite and mode

- Wide beam mode: image field covers 500 km^2 with a pixel size of 100 m^2
- Spotlight mode: image field covers 50 km^2 with a pixel size of 1–3 m^2 (Schiermeier, 2007; MDA, 2009)

Temporal specificity: The orbit cycles of existing satellites range from 11 to 41 days. While most of the newer satellites can record images of a given location during multiple days in every orbit, there remains some significant percentage of each orbit for which images of a given area are not available. With newer satellites that black out period is as short as 4 days (Holecz, 2009; Schiermeier, 2007).

Length of time series: Various SAR equipped satellites have been in continuous operation since 1991. However, as mentioned above, not all regions of the globe were under continuous monitoring during this time so historical coverage varies by location (Schiermeier, 2007).

Many additional SAR equipped satellites are planned for launch in the upcoming years. In particular, the European Space Agency Sentinel-1 satellite, scheduled in 2011, promises to make a significant contribution to the problem of sporadic data availability. It will acquire SAR images covering the whole planet every 6 days (Holecz, 2009).

4.2.3.4 Reanalysis Data

Reanalysis data describes a class of data products that combines and calibrates observations from many sources — weather stations, satellites (infrared and microwave imagers), weather balloons, ocean buoys, aircraft, and ships (Uppala et al., 2005). Each data source has its own time series so there are certain points in these models where there are fewer observations than at other points. Reanalysis models use periods of overlap across data sources to “assimilate” the data, i.e., to calibrate across data sources.

First- and second-generation reanalysis models (Box 11) do not seem suitable to act as the index for an index insurance program. First, the spatial specificity of these models tends to be too coarse (generally 120 km to 210 km). Second, these models tend to inaccurately estimate rainfall, especially first-generation models such as NCEP/NCAR (Funk et al., 2003). NCEP/NCAR does not use rain gauge data, which likely decreases the accuracy of its rainfall estimates (Funk et al., 2003). These first-generation models have discontinuities and biases throughout the time series. Poccard, Janicot, and Chamberlin (2000) analyzed NCEP/NCAR for Africa. They found an abrupt shift in the data in 1967 that affects data for almost all of tropical Africa. NCEP/NCAR tended to underestimate rainfall during the peak of the rainy season across regions in Africa. These authors conclude that NCEP/NCAR is useful for studying “large-scale climate dynamics” but significant problems restrict the use of NCEP/NCAR for studying “regional long-term variations.” Newer reanalysis models seem to be much improved including higher spatial specificity. However, accurately estimating rainfall is quite difficult for these models, too. For example, ERA-Interim tended to overestimate rainfall during the rainy season for eight tested regions in Africa (Sylla et al., 2009).

Specialized data products that incorporate reanalysis models (e.g., CMAP and CHARM) tend to show some improvement over the pure reanalysis models, but also seem too inaccurate to supply data for an index insurance product. Using data from the Sahel, Ali et al. (2005) show that the reanalysis models used in that study (which included CMAP) tend to show regression to the mean for both types of extreme events (an underestimating of excess rainfall and an overestimation of deficit rainfall). Additionally, Funk et al. (2003) compared CHARM to interpolated rain gauge data for two regions, one in southern Mali and another in western Kenya. CHARM tended to underestimate extreme events in Mali and Kenya. CHARM showed positive bias for rainfall in Kenya, and negative bias in Mali. Funk et al. (2003), who developed CHARM, conclude that “neither the CHARM nor a [pure satellite] product is likely to be skillful at a mesoscale resolution.” (p. 59)

Reanalysis models have come a long way in a short period of time, and reanalysis models developed in the not too distant future may be sufficient to generate data for index insurance products. Models under current development are reaching a level of spatial specificity (approximately 35 km for Climate Forecast System Reanalysis and Reforecast {CFSRR]) that may be sufficient to support some index insurance products. Climate modelers are likely to continue to improve the accuracy of reanalysis models, making them more feasible for index insurance applications.

For index insurance, reanalysis products are currently most useful for analyzing the historical weather risk. Since even newly developed reanalysis models tend to misestimate rainfall, these models would probably not be accurate enough to act as the sole data source for estimating the pure risk.

Box 11 Reanalysis Data

Reanalysis models typically develop values for their weather-related variables for the entire world. Many weather related variables are included in reanalysis models including upper-air and surface wind, temperature, humidity, sea surface temperature, land surface temperature (at 2 meters), soil temperature, snow depth, infrared and microwave radiances, surface pressure, and oceanic wave height (Uppala et al., 2005). The following are the main atmospheric reanalysis models: [A]

	Spatial Specificity (Approximate)[B]	Temporal Specificity (Hour Intervals)	Time Series	Vintage	Status
NCEP/NCAR	210 km	6	1948–present	1995	Ongoing
NCEP-DOE	210 km	6	1979–present	2001	Ongoing
CFSRR (NCEP)	35 km	6	1979–present	2009	In progress
C20r (NOAA)	220 km	6	1891–present	2009	In progress
ERA-40	125 km	6	1957–2002	2004	Done
ERA-Interim	80 km	6	1989–present	2009	Ongoing
JRA-25	120 km	6	1979–present	2006	Ongoing
JRA-55	60 km	6	1958–2012	2009	Underway
MERRA (NASA)	55 km	1–6	1979–present	2009	In progress

The NCEP/NCAR is one of the first-generation reanalysis models developed in the late 1990s (Kalnay et al., 1996; Trenberth et al., 2009) and is often included as a component of more recently developed reanalysis models. NCEP/NCAR includes over 80 variables (UCAR, 2010). These first-generation models have been widely used but experienced many problems (Trenberth et al., 2009). These models have biases that change in magnitude and direction over time. For example, Funk et al. (2003) found for Africa the NCEP/NCAR tended to consistently overestimate rainfall in the tropics from 1961 to 1996; however, in the northern and southern subtropical regions, the pattern of bias in the 1960s and 1970s differed from that in the 1980s and 1990s. Second-generation models such as ERA-40 address some of the problems of first-generation models like NCEP/NCAR and tend to outperform these models (Trenberth et al., 2009). ERA-40 incorporates data from NCEP/NCAR (Uppala et al., 2005). Instead of updating ERA-40, ECMWF has developed ERA-Interim for 1989 to present, a period for which many more data sources are available (ECMWF, 2010). New reanalysis models, such as CFSRR under development by NCEP, will continue to address problems of previous models.

Many data products use one of these main reanalysis models as a component but add other important variables, depending on the purpose of the data product. For example, CMAP is a data product specializing in rainfall, and CHARM is a regional rainfall product for Africa.

	Spatial Specificity (Approximate)	Temporal Specificity	Time Series	Vintage	Main Reanalysis Model Used
CMAP	275 km	5-day	1979–2006	1997	NCEP/NCAR
CHARM	Unclear[C]	Daily	1961–1996	2003	NCEP/NCAR

CMAP provides full global coverage and is intended to improve large-scale rainfall estimates. It uses rain gauges, satellite data, and NCEP/NCAR (Xie and Arkin, 1997). CHARM uses interpolated rain gauge, topographical, satellite imagery, and NCEP/NCAR reanalysis data. The U.S. Agency for International Development Famine Early Warning Systems Network (FEWS NET) uses CHARM to predict the development of famine in several African countries.

There are other reanalysis data series and many other data for the types of weather-related variables described above. The University Corporation of Atmospheric Research (UCAR) houses an inventory with basic descriptions of these data products at (http://dss.ucar.edu/).

[A] *Table adapted from Trenberth et al. (2009).*

[B] *Reanalysis models are spectral models and so are measured in terms of spectral resolution, not kilometers; however, considering the spatial specificity of these models in terms of approximate kilometers is most useful for our purposes. For more on spectral models see Krishnamurti et al. (2006).*

[C] *Funk et al. (2003) note that CHARM data are "generated with a 0.1° [10 km] resolution to facilitate integration with satellite-based [rainfall estimates]…the actual resolution of CHARM data are considerably courser than 0.1°."*

4.2.4 Summary of Satellite-based Data Sources

With the exception of some specific perils in particular locations (e.g., NDVI in some large rangeland areas), satellite-based data sources are not currently exploited for index insurance products. However, because these technologies continue to improve and because satellites have the potential for providing a consistent source of data that covers much of the globe, satellite applications could become more prevalent. Two advancements in these technologies would be particularly important for index insurance offers. The first is higher spatial specificity in the data routinely collected and in the methods to use this data for insurance purposes. For certain types of risk aggregator products, high special specificity may not be needed depending on the environment. The second is improvements in calibrating data collected from satellites with ground-level weather or other *in situ* data — especially for extreme events. The latter will likely be a significant challenge. Calibration is conducted using available data and far more data are available for typical events than for extreme events.

More opportunities exist for using satellite-based data and reanalysis data for estimating the pure risk. While reanalysis data may not perfectly match ground-level data, they can provide a much longer time series, making these data especially useful for identifying trends. One reinsurer reported that his firm uses reanalysis data regularly for this reason. These alternative data sources can also be useful when extreme events disrupt ground-level data sources (e.g., if a flood washes away a weather station). Satellite-based and reanalysis data sources can also serve as a quality check if weather stations provide an unusual data value, can be used to roughly price a weather station with little history, and to assess the extent of basis risk around a weather station during an extreme event. In sum, these data sources provide more information to practitioners, which can reduce ambiguity about the risk and allow insurers and reinsurers to price the pure risk more sustainably. Because insurers and reinsurers charge a premium for ambiguity, this added information can also lower the cost of the insurance.

If the market for weather index insurance in lower income countries is going to expand greatly beyond current pilot programs, offers will need to use cost effective data sources other than just weather stations as the index for settlement. While the satellite-based alternatives will no doubt continue to improve, each will have strengths and limitations. Thus, it seems likely that indexes will need to be based on data collected from multiple sources — similar to reanalysis data. As the market for weather index insurance evolves, commercial firms are already investing resources in developing and improving such measures — much as the catastrophe bond market stimulated the development of private-sector firms that provide hurricane modeling services.

Chapter 5 State of Practice: Evaluation of Current Index Insurance Program Practices in Meeting Customer Demands and Overcoming Data Constraints

To this point, the focus is on providing an overview intended to be an objective assessment of the state of knowledge regarding data. In this chapter, we provide our synthesis of how the current state of available data systems influences future development of index insurance. Our analysis is based on knowledge of practices that have been employed in developing index insurance products in recent years. As we are likely not fully informed regarding all aspects of every index insurance product implemented in recent years, there are limitations that must be acknowledged as well as acknowledging that our assessments are influenced by our own experience in developing new index insurance products.

Most of the index insurance applications have been targeted to low-income households. Nonetheless, we believe that many of our observations about these applications can be applied to risk aggregator products. The conceptual model presented in this SKR provides the theoretical framework for evaluating index insurance given real-world constraints. Based on the tension between that theoretical framework, our analysis of where index insurance may best fit in lower income countries, and our knowledge of current practices, we begin by critiquing two common practices: 1) insuring against moderate losses; and 2) designing index insurance products solely to protect against crop-yield losses for a single crop.

5.1 Insuring against Moderate Losses

In recent years, at least two developments have led practitioners to design products that insure against moderate losses.[17] First, some practitioners have come to the conclusion that households will only maintain interest in insurance if the insurance makes frequent payments. This concern relates to research suggesting that individuals have difficulty understanding the probability of, or the potential magnitude of, catastrophic events (Kunreuther, 1996, 1976; Kunreuther and Slovic, 1978; Tversky and Kahneman, 1973). Thus, in some cases, individuals make decisions by essentially assigning zero probability to low-probability, high-consequence events. This *cognitive failure* almost certainly contributes to the challenge of maintaining household interest in a product that pays infrequently. Experience with the specific weather index insurance products being offered in India seems to confirm this finding (Giné, Townsend, and Vickery, 2008). Second, some practitioners have insured against moderate loss events to facilitate the bundling of index insurance with other services. For example, some index insurance programs are bundling insurance with loans that facilitate the purchase of yield-increasing inputs. Typically, the sum insured by the index insurance is for the value of the loan and the lender has first rights to any indemnity payment to cover the outstanding loan. Thus, the index insurance acts as a type of loan guarantee. In some cases, lenders, who are concerned that even moderate losses could cause loan defaults, have encouraged practitioners to offer high levels of index insurance coverage relative to the expected value of the index.

We have a number of concerns with offering index insurance protection against moderate losses. First, we have some expectation that basis risk problems may be more problematic when insuring against moderate losses. As a simplistic example, assume that there is a 95 percent

[17] We consider products that pay as frequently as 1 in 3 or 1 in 5 years as paying for moderate losses.

chance that an index can estimate losses within plus or minus 5 percent of the actual value (i.e., the 5 percent error can be referred to as the basis risk). For this illustration, we make an additional simplifying assumption (which we will relax below) that the 5 percent error does not change depending on the size of the loss.[18] Consider two contracts, one that protects against extreme losses when the insured has experienced a roughly 25 percent loss, and another that protects against moderate losses when the insured has experienced a roughly 10 percent loss.

$$\frac{\textit{Basis risk}}{\textit{Insurance trigger}} = \textit{Misestimate of losses}$$

For insurance against the extreme loss, the misestimate of loss is ${5\%}/{25\%} = 20\%$. For insurance against the moderate loss, the misestimate of loss is ${5\%}/{10\%} = 50\%$. Because the moderate risk contract is attempting to insure smaller deviations from the mean, opportunities for misestimating the loss are much greater. Additionally, because moderate risks occur more frequently, inaccurate indemnities will occur more frequently. As a result, practitioners developing products that pay frequently may feel that it is more important to reduce basis risk and may engage in practices such as overfitting, which as we outline in Chapter 1, can actually increase basis risk.

Second, in contrast to our simplifying assumption in the example above, we believe that basis risk may change depending on the magnitude of the insured event. Specifically, underlying physical processes of both weather events and specific types of losses may result in lower basis risk for insuring against extreme events than insuring against moderate risks. It may be the case that basis risk changes depending on the severity of the weather risk due to changes in the spatial presentation of the weather risk.

In support of this hypothesis for drought, we cite the work of Bravar and Kavvas (1991) who describe the physical processes of rainfall to demonstrate that when regions experience drought, it becomes much less likely to rain in that region. In brief, evaporating soil moisture increases humidity. At a certain level of humidity, the air becomes saturated causing passing clouds to rain. If evaporating soil moisture is insufficient to elicit rain, then this soil moisture is not replaced, resulting in lower humidity and a decreased chance of rainfall. Thus, drought causes a positive feedback loop, which must be broken by a weather front with sufficient moisture that develops in another region. This positive feedback loop is a physical process much different than those in patterns of moderately low rainfall.

A research agenda motivated by this SKR is to examine how spatial correlations of drought and excess rainfall change depending on the severity of the risk (Appendix B). It is also the case that extreme values of rainfall tend to be a much better predictor of losses than rainfall values close to the mean, for specific types of loss events. Two examples may make the point. First, in flood-prone regions, extreme flooding is a much better predictor of extreme losses than moderate flooding is of moderate losses. Second, extreme drought is a much better predictor of extreme crop failure than moderate drought is of moderate crop failure. Moderate shortages in rainfall will affect crops in the same region in different ways depending on the soils, crop varieties, and other input variables. However, beyond certain thresholds of depleted soil moisture,

[18] More fundamentally, we are assuming that the variance of losses is constant across all values of rainfall. With a constant variance, the relative risk increases as the expected value of rainfall goes down.

photosynthesis slows and plants stop growing. Additionally, when plants do not meet water satisfaction requirements, inputs such as fertilizer have little impact on plant growth. Thus, under many conditions extreme rainfall deficits are a better predictor of extreme crop losses than moderate rainfall deficits are of moderate crop losses.

Third, due to basis risk, index insurance provides a rather imperfect loan guarantee — especially for moderate losses. By incorrectly suggesting that index insurance can provide a loan guarantee against moderate losses caused by more frequent but less extreme weather events, practitioners run the risk of eventually losing credibility with lenders and insureds. We strongly support the concept of bundling index insurance with other services as a more efficient delivery mechanism. Our concern is that, instead of advancing an agenda with lenders about the optimal use of weather index insurance given its limitations, practitioners who design contracts that pay frequently with the intent of protecting individual loans, may be overlooking the data limitations and giving potential users a false sense of security about the protection that these products provide.

Fourth, due to high transaction costs, insurance is a rather expensive financial instrument and is designed to protect against low-probability, extreme losses, while savings and credit are generally more economically efficient mechanisms for managing small to moderate losses.[19]

Fifth, it can take years for households and firms to recover from extreme catastrophic events. For example, our analyses with data from lenders in northern Peru indicate it took roughly five years for households and firms to recover from the 1997–98 El Niño (Collier, Katchova, and Skees, 2010). Experience with Hurricane Mitch in Honduras indicates that, for many years afterwards, some households continued to struggle due to losses from that event (Carter et al., 2007). The literature related to poverty traps suggests that in some cases the working poor may not be able to recover from these events (Barnett, Barrett, and Skees, 2008; Sachs and Arthur, 2005). The most efficient use of insurance is to protect against extreme catastrophic events that can threaten long-term wealth positions. When, instead, weather index insurance is designed to protect against more moderate losses, the price of insurance is high compared to a catastrophic policy. As a result of the higher price, households and firms purchase less insurance and, when a catastrophe occurs, are not as well protected. Thus, insuring moderate risks tends to divert resources away from the most effective use of weather index insurance — transferring catastrophic weather risk.

5.2 Emphasis on Crop Yields

Another common practice among many researchers and practitioners is to think of weather index insurance as a form of crop insurance. To those in the insurance industry, this may not be surprising. Much of the innovation that led to weather index insurance was motivated by

[19] It is important to note that while market-based insurance products are not social programs, indexes can be used to finance social programs that protect against catastrophic weather events. However, it is critical that such social programs be designed so that they crowd-in, rather than crowd-out, complementary market-based insurance products. For example, if governments or donors wish to subsidize weather index insurance offers, they should consider doing so by funding a social program that protects against the most extreme layer of risk —those extremely rare loss events (e.g., a frequency of 1 in 25). If carefully constructed, such social programs can actually facilitate market-based offers of weather index insurance for relatively more frequent (e.g., 1 in 10 year) catastrophic events.

problems with traditional crop insurance programs that focus on farm-level crop yields (Skees, Black, and Barnett, 1997; Martin, Barnett, and Coble, 2001). Such models were developed for higher income countries where many farmers specialize in specific crops and where data on crop yields and household income are abundant. Furthermore, the input packages used to grow crops in higher income countries are also significantly more homogenous than those used to grow crops in lower income countries.

Weather index insurance programs for specific crops are available in the United States and Canada. It is not surprising that much of the development of index insurance in lower income countries follows processes used in developed countries. However, this approach ignores data constraints in lower income countries and, to a large extent, is inconsistent with risk assessment findings regarding household livelihoods. Regarding data constraints, we return to our conceptual model for illustrative purposes. In our model, practitioners strive to identify the relationships between

Index ←→ cause of loss ←→ losses of the insured

and this process requires both an understanding of each of these distributions and of the relationship between distributions. An implicit assumption of the conceptual model is that the index around which the insurance is designed captures the loss exposure of the insured due to the cause of loss. When practitioners emphasize crop yields in lower income countries, they are using a (potentially poor) approximation for losses of the insured. Said differently, yield losses are simply one indicator of household well-being. While the relationship between yields of a specific crop and the well-being of the insured may be highly related for many farmers in developed countries where crop specialization has led to highly specialized farms, it is less clear that the yield of a specific crop is as important to households in developing countries (more on this below). Thus, these practitioners have added another step by focusing on yields.

Index ←→ cause of loss ←→ yield losses ←→ losses of the insured

As we discuss in Chapter 1, each distribution and each relationship that practitioners must estimate introduce additional basis risk. Thus, *the emphasis on yields can also increase basis risk*. The reader will also remember that quantitative loss data — for yields and for household losses more generally — often do not exist in lower income countries. Thus, practitioners emphasizing yields are put in the difficult position of working across several distributions for which they have no quantitative data. Compounding the problem: it is not clear that the relationship between crop yields and losses of the insured is necessarily linear, making it more difficult to estimate. In addition, the relationship will differ between individuals in the target market due to heterogeneity in their income strategies.

Risk assessments reveal the many concerns of households and firms. When these stakeholders talk about natural disasters, the most salient effects in the community are losses to well-being — households losing assets and depleting savings, increasing food prices, families starving, and loved ones dying. The second topic stakeholders identify includes the many diverse complications of the event. Floods in Peru are a good example. Household crops are destroyed; fertile topsoil is washed away; pest problems increase significantly; infectious diseases increase due to sedentary water; bridges and roads are destroyed and take months to repair; many households are isolated without access to food for extended periods, firms cannot transport goods nor receive supplies; etc. Clearly, losses caused by extreme weather events extend well beyond the impact on crop yields.

Broader research and risk assessments also tend to indicate that most smallholders in lower income countries do not rely on the yield of a single, specific crop (World Bank, 2007). Instead, they plant a variety of crops and often have livelihood portfolios that are diversified across labor activities besides farming. Berg and Schmitz (2008) demonstrate that weather index insurance for a specific crop is a much less effective risk management tool for households with a diversified portfolio than for those that specialize in a specific agricultural commodity. Therefore, insurance with a focus on a specific crop is likely of limited value for smallholders in the majority of the developing world.

Even within the limited realm of production agriculture, presenting index insurance only as a means of protecting crop yields can miss some of the more important production risks and the potential value of such products. For example, in the central highlands of Vietnam, smallholder coffee farmers are exposed to drought, but when drought occurs, these farmers often manage yield losses by increasing irrigation. However, when they extend the irrigation season, they also incur significantly higher costs as the water table is depleted and irrigation becomes more expensive. Some coffee plants also suffer from lower amounts of water resulting in coffee beans that are perhaps one-third the size of normal beans and prices that are less than one-half what they would be under normal weather conditions. In the worst conditions, coffee trees also die. In a pilot project supported by the Ford Foundation, the Vietnam insurance regulator has approved a drought business interruption insurance product designed to compensate for the consequential losses associated with severe drought conditions. A traditional crop insurance product would only pay for crop-yield losses and would not be interesting to these growers.

If practitioners conceptualize weather index insurance as a form of business interruption insurance to compensate for consequential losses associated with severe weather events, they could view data constraints and basis risk in a very different fashion, thus influencing their product design and marketing strategies. In many cases, weather index insurance will be the first form of weather-related insurance offered in rural areas of lower income countries. The views of the target market toward the insurance are likely to be significantly influenced by the marketing and education efforts of practitioners. If practitioners take the view that weather index insurance is a replacement for crop insurance, it may prevent the target market from recognizing the full value of the insurance. For example, if the drought insurance product in Vietnam would have been limited to yield losses only, farmers would likely have reported that they had means of managing yield losses associated with drought and so had no need for the insurance. By framing index insurance in a broader context, risk assessment interviews with Vietnam farmers demonstrated that they were interested in the product.

5.3 SKR Key Recommendations

As an extension of our analyses presented in this SKR and in response to the concerns described above, we present an alternative approach to the design and marketing of index insurance. How one frames a problem is critical to finding solutions. We believe that the approaches we recommend will expand the potential uses of weather index insurance, increase its potential application in data-constrained regions, and help overcome the cognitive failure problems described above. We propose the following three main premises, which emerge from our conceptual development presented above, our own experience with developing index insurance products, and the empirical findings presented in this document: 1) weather index insurance is best suited for consequential losses; 2) weather index insurance is best suited for catastrophes;

and 3) data constraints are lowest for risk aggregator products. These premises are interrelated and form the basis for much of our current work on index insurance products.

5.3.1 Index Insurance Is for Consequential Losses

Considering weather index insurance as an alternative to traditional crop insurance is an important first step in the evolution of the weather index insurance market. However, weather index insurance should now be recast in a much broader context. In short, the ways that extreme weather events retard economic growth in lower income countries extend far beyond crop-yield losses. Thus, index insurance products should be designed to protect against the variety of consequential losses that may occur during the critical period when an extreme weather event is most likely to occur rather than being designed only around the vulnerabilities of a specific crop. This product design should make the insurance more relevant to protecting the wealth positions and portfolio of activities of households and firms. Such an index insurance product would be designed in an encompassing fashion around the critical periods of the catastrophic event rather than the key vulnerabilities of a specific crop.

The discussion of consequential losses is quite relevant for risk aggregators as well. For example, for banks in northern Peru, El Niño is associated with borrowers having problems repaying their loans and depositors withdrawing savings. These difficulties create liquidity constraints, increase provisioning requirements, and cause higher administrative costs. These banks must optimize between the opportunity cost of maintaining poorly performing loans for months or years after their maturity date and taking large losses as a result of forgiving these debts. Risk aggregators, as well, may need assistance in capturing a vision for the variety of benefits associated with weather index insurance designed for consequential losses of a catastrophic weather risk.

Designing and marketing index insurance in terms of the consequential losses of an extreme weather event can have several benefits. First, insureds have the flexibility to purchase a sum insured to manage a variety of risks to which they are exposed. Insureds can use insurance payments for what they consider most important. The challenge of a heterogeneous target market is largely addressed by having such a flexible product.

Second, in some cases, such a flexible design may reduce basis risk when compared to weather index insurance contracts designed for crop yields. As discussed earlier, extreme weather events are context-specific; even the same weather peril in the same region may differ significantly from one event to the next, resulting in different types of loss across events. While the target market may expect to experience several types of loss associated with a catastrophic event, it is often unclear how the event will affect specific aspects of the wealth position and the portfolio of activities. Thus, a more general product designed to allow the insured to address a host of potential problems may more suitably match losses of the target market than a product based on one specific investment outcome such as crop yield.

Third, weather index insurance designed around consequential losses likely creates a much greater recognition of the value of index insurance among the target market than insurance designed around one crop. To determine the sum insured and potential uses of the insurance, practitioners marketing these products are likely to engage in a rich discussion with stakeholders in the target market regarding their exposure to the event and improving their strategies for managing consequences of the event.

Fourth, a focus on consequential losses also has important data implications. In particular, practitioners need not be so concerned about proving high in-sample correlations between the proposed index and yields for a specific crop. Instead, high-quality qualitative data obtained through carefully structured interactions with local experts are likely to be more useful for understanding the relationship between the index and the variety of consequential losses.

Finally, a design based on consequential losses would likely also increase demand for the insurance. We use a simplified version of the conceptual model from Chapter 1 to illustrate. For weather index insurance designed and marketed for a specific crop, individuals would be asked to determine the level of insurance they would like to purchase based on the weather risk for that crop

$$E(U)_{Crop} = \left(1 - \sum_{i=1}^{n} \pi_i\right) U\left(W_0 + R_g - pI\right) + \sum_{i=1}^{n} \pi_i U\left(W_0 + R_g - YL_i - pI + q_i I\right)$$

where π_i is the probability of a specific bad weather outcome, W_0 is initial wealth, R_g is the net return on investment in good years for all activities including, but not limited to, net returns for the specific crop, p is the premium rate, I is the sum insured, YL_i is the monetary value of yield losses associated with a specific bad weather outcome, and q_i is the insurance indemnity rate associated with a specific bad weather outcome. Notice that in this model individuals are considering yield losses as the only negative consequence of the bad weather outcome.

Alternatively, consider a weather index insurance contract designed and marketed based on the many consequences of an extreme weather event. Suppose that, beyond yield losses YL_i, the insured also experiences asset losses AL_i (e.g., losing property, home equipment, savings, etc.), health losses HL_i, reduced profits due to increased costs PL_i, losses in other labor opportunities LL_i (e.g., decreased employment in labor on other farms), and other losses OL_i associated with the disaster. In this case, individuals would make their insurance purchasing decisions based on

$$\begin{aligned} E(U)_{Disaster} &= \left(1 - \sum_{i=1}^{n} \pi_i\right) U\left(W_0 + R_g - pI\right) \\ &+ \sum_{i=1}^{n} \pi_i U\left(W_0 + R_g - YL_i - AL_i - HL_i - PL_i - LL_i - OL_i - pI + q_i I\right) \end{aligned}$$

The optimal level of insurance can be determined by taking the first derivative of the expected utility function with respect to I, setting this derivative equal to zero and solving for I. So for the crop-specific product

$$\frac{\partial E(U)_{Crop}}{\partial I} = (-p)\left(1 - \sum_{i=1}^{n} \pi_i\right) \frac{\partial U\left(W_0 + R_g - pI\right)}{\partial I} + (q_i - p) \sum_{i=1}^{n} \pi_i \frac{\partial U\left(W_0 + R_g - YL_i - pI + q_i I\right)}{\partial I} = 0$$

while for the consequential loss product

$$\frac{\partial E(U)_{Disaster}}{\partial I} = (-p)\left(1-\sum_{i=1}^{n}\pi_i\right)\frac{\partial U\left(W_0+R_g-pI\right)}{\partial I}$$
$$+(q_i-p)\sum_{i=1}^{n}\pi_i\frac{\partial U\left(W_0+R_g-YL_i-AL_i-HL_i-PL_i-LL_i-OL_i-pI+q_iI\right)}{\partial I}=0$$

The optimal sum insured depends on several factors (e.g., premium rate, magnitude and frequency of losses, basis risk, etc.), but the most relevant for this discussion is the recognition of consequential losses. To see this, note that, if the sum insured is the same, then equations for the crop-specific product and the consequential loss product are the same, except in bad years, ending wealth, and thus the utility, is lower for consequential loss product. The only way to increase wealth in the bad years for the consequential loss product is by increasing the sum insured. But increasing the sum insured also decreases wealth in the good years because the premium cost increases. However, by definition, risk-averse decision makers (and only risk-averse decision makers would purchase insurance) value additional wealth in the bad years (from receiving an indemnity) more than the loss of wealth in the good years (from paying the premium). More formally, for risk-averse decision makers, the marginal utility of wealth is higher in bad years than in good years (risk aversion implies concave utility functions or diminishing marginal utility of wealth). So considering not just yield loss but also other consequential losses should increase the sum insured.

Thus, *this model suggests that individuals would be expected to have higher demand for index insurance if the insurance were designed and marketed not just to protect against crop-yield losses but also for the various other consequential losses of an extreme weather event*. The model actually demonstrates a simple notion — if buyers can see that an insurance product can be used to protect against more types of losses, they will be willing to purchase more of the insurance.

5.3.2 Index Insurance Is for Catastrophes

There are at least two interesting data issues that relate to the question of whether index insurance should be used to protect against moderate loss events or only against catastrophic loss events. The first focuses on the variability of returns from a business's or household's portfolio of activities. We hypothesize that the pairwise covariances of returns among the various activities are not linear throughout all possible weather outcomes.[20] Specifically, we believe that the covariance of returns is greater for more extreme weather events. In other words, steps to diversify a portfolio by investing in several activities may be ineffective for extreme weather events. If so, this further supports our view that weather index insurance should focus primarily on addressing the range of consequential losses that result from catastrophic weather events.

Second, we hypothesize that the spatial covariance of some weather variables is not linear with respect to severity. A specific research agenda motivated by this SKR examines whether the

[20] Miranda develops a conceptual model to investigate these questions further. That model is presented in Appendix A. We plan to test this model with some empirical examples in the United States, where good data are widely available.

spatial correlation of drought or excess rainfall increases with severity. If spatial correlations do increase with severity, this suggests that the spatial specificity of data required for developing index insurance that protects against moderate loss events is greater than that required for developing index insurance that protects against catastrophic, extreme loss events.

A number of important research questions emerge from our concerns regarding the design of products that pay for frequent, moderate, losses. The dilemma remains that of the index insurance products developed thus far, most are designed to pay for losses that occur more frequently than 1 in 7 years. Users of these products express a strong preference to be paid frequently. While this may indicate they will gain more confidence in the product if they can see payments being made, we have raised questions about how long this confidence will last if they experience a large percentage of small payments and only a moderate payment when there is a catastrophe. Again, there is a clear tradeoff of premium. The tendency is to purchase a low sum insured given the higher price for a product that pays frequently. We tested this explicitly with livestock index insurance in Mongolia. Herders were given a choice between two policies; one that would pay when mortality rates exceeded 6 percent and one that would pay when mortality rates exceeded 10 percent. The premium rate was nearly 2 times higher for the 6 percent threshold than for the 10 percent threshold. Some educational effort was targeted at getting the herders to take the 10 percent catastrophic policy because with the same total premium they could purchase a higher sum insured given that the typical purchase was for only 30 percent of the value of animals. Yet, in the one year that this experiment was run, over 90 percent of the insured herders selected the 6 percent threshold policy. Similarly, many of those purchasing health insurance select low deductibles or co-payments rather than catastrophic health insurance. A counterexample seems to occur with life insurance. People in lower income countries are demonstrating a clear willingness to pay for market-based life insurance, which protects against a low-frequency, catastrophic event. Considerably more work is needed on the psychology of insurance purchase decisions, but it is interesting that life insurance protects against shocks to long-term wealth caused by loss of an important human asset.

The psychology and behavioral economics literatures clearly indicate that how people conceptualize an uncertain outcome (e.g., a loss, variable returns, a gamble, etc.) affects their strategies for managing it (Kahneman and Tversky, 1979). This research leads us to the conclusion that individuals likely evaluate risks to their long-term wealth positions (i.e., risks that affect the well-being and future opportunities of the household) differently from risks to current period returns (e.g., crop-yield risks). Many of the index insurance pilots have been marketed as a sort of pseudo crop insurance — that is, a way to reduce the variability in returns for a specific activity. Given the highly diversified portfolios of many households (and firms) in lower income countries, crop-yield losses in a single year may not significantly affect long-term wealth. If this is true, it is not surprising, given that insurance is a relatively expensive financial mechanism, that households would lose interest in insuring against weather risks designed to protect crop production.

In contrast, catastrophic insurance is about preserving the long-term wealth position of the household. By framing weather index insurance as protecting long-term wealth, it becomes much more akin to life insurance. As a starting point, we return to our expected utility theory model.[21] It seems consistent with the behavioral economics literature on framing that

[21] Psychology and behavioral economics researchers are often critical of an expected utility framework and provide many other suggestions for modeling the way individuals assess uncertain outcomes (e.g.,

individuals may use different utility functions based on how the problem is framed — specifically, individuals may have one utility function for wealth U_{wealth} and another for net investment returns $U_{returns}$. We postulate that individuals may be more risk averse to losses in their long-term wealth position, losses that may mean a lower quality of life for themselves or their family, than to losses in current period returns for a specific activity (i.e., $\frac{\partial^2 U_{wealth}}{\partial D^2} < \frac{\partial^2 U_{returns}}{\partial D^2} < 0$ where D is the total dollar value of the terms in the utility function). Thus, we can rewrite our expected utility model without insurance as

$$E(U)_{NI} = \left(1 - \sum_{i=1}^{n} \pi_i\right)\left[U_{wealth}(W_0) + U_{returns}(R_g)\right]$$
$$+ \sum_{i=1}^{n} \pi_i \left[U_{wealth}\left(W_0 - AL_i - HL_i - OL_i^w\right) + U_{returns}(R_g - YL_i - PL_i - LL_i - OL_i^r)\right]$$

where OL_i^w is all other losses in wealth experienced by the individual and OL_i^r is all other losses in current period returns experienced by the individual. Where does insurance fit into this model? We again suggest that it depends on how the insurance is designed and framed to the target market (Kahneman and Tversky, 1979). When individuals are more risk averse to losses in wealth than to losses in returns, they will purchase more insurance when they expect it to protect their wealth position than if they expect it to protect their net returns on investment. More directly, *individuals are likely to have a higher demand for weather index insurance if it is intended to protect against catastrophic risks that threaten the long-term well-being of the household than if it protects only against current period returns.*

Thus, designing and marketing index insurance for catastrophes may overcome the demand problems experienced by some weather index insurance pilots. Furthermore, it may be that if insureds believe that they are protected from such a disaster, they will be more likely to change their behaviors in the fashion predicted by insurance theory — engaging in higher-risk, higher-return activities that would contribute to increased growth in household wealth in the long term (Barnett, Barrett, and Skees, 2008).

5.3.3 Data Constraints are Lowest for Risk Aggregator Products

This report highlights the importance of risk aggregator products, especially in data-constrained regions.[22] To reiterate, products for these firms require assessment of a catastrophic weather event at a community or regional level, whereas household products require an assessment of the weather event at a specific geographic point. As a result, the risk aggregator product requires fewer data sources (e.g., fewer weather stations) in a region than products for households. As demonstrated by Ali, Lebel, and Amani (2005), combining the estimation of

Kahneman and Tversky, 1979; Quiggin, 1991). These researchers identify many biases that influence risk-taking behavior (e.g., judgment bias, hindsight bias, availability heuristics, gambler fallacy, etc.). We believe that future research on weather index insurance should more fully integrate these paradigms.

[22] This topic is directly addressed in Section 4.2.1. Data Constraints Are Less Binding for Risk Aggregator Products than Household Products, following the general review of weather station infrastructure in lower income countries.

several weather stations tends to lower estimation error of the weather event more than using a single weather station. Thus, basis risk associated with mismatches between the cause of loss and the index (**cause of loss ←→ index**) should be reduced. In sum, *risk aggregator products should require a less developed weather station infrastructure and increase opportunities for satellite-based products*.

Risk aggregator products are likely the only feasible mechanism for extending weather index insurance into many regions of the world. Because weather station infrastructure is so underdeveloped in many regions of the world (WMO, 2008) and satellite data are too coarse for many household risks (with a few notable exceptions), household products are simply inadvisable in many regions. Pursuing weather index insurance products for households despite inadequate data is likely to lead to 1) higher insurance prices by insurers and reinsurers due to uncertainty about the risk, and 2) products that poorly capture the risk of the target market and therefore contribute little to disaster risk management. As a result, risk aggregator products would seem to provide a better return on investment for economic development efforts in data constrained regions.

Some practitioners may question whether products designed for risk aggregators such as rural banks and agricultural value chain members substantially benefit the poor and would rather see insurance products that can be purchased by households. It is worth remembering that the highest poverty rates in lower income countries almost always occur in rural areas. While there are certainly risk aggregating firms that will only work with better-off households, many financial institutions, agricultural value chain members, etc., do work with poor populations. Moreover, one reason that some risk aggregator firms limit the services they provide to the poor is that they cannot manage the catastrophic weather risk associated with serving these clients. As we ourselves consider this question, we return to the risk management axiom: when losses occur, someone must pay for them. For example, households may pay banks higher interest rates because the bank is unable to efficiently manage the catastrophic risk exposure in the region. Also, agricultural input suppliers, commodity processors, and lenders alike may limit their presence in regions where households are vulnerable to catastrophic risk because these risk aggregators are unable to manage this correlated risk themselves. Weather index insurance products for risk aggregators that enhance the ability of these firms to manage catastrophic risk can increase household access to the services of these firms. Increased access to credit, inputs that increase crop productivity, and commodity export markets have all been shown to have important developmental outcomes (World Bank, 2007) and are the ultimate goal of many development projects.

An added benefit of working with risk aggregators is that these stakeholders, due to their professional experiences, could be expected to engage in a risk management discussion in a more sophisticated way than households. These firms seem prepared to understand weather index insurance more fully as they likely already use other financial contracts to manage risks. For example, banks coordinate bond holdings, interbank debt, loan maturity, certificates of deposit, etc., in asset-liability management. Also, commodity exporters often use forward (and sometimes futures) contracts. As a result, these risk aggregators may be more comfortable with evaluating weather index insurance contracts, and may also have more knowledge and experience with managing basis risk.[23] Since data requirements for risk aggregator products are relatively low and the market is relatively sophisticated, the potential for products using satellite data increases. Opportunities emerge for products that either: 1) use information from atypical sources (e.g., infrared sensors on orbiting satellites); or 2) integrate information from a variety of sources to create an index (combining data from weather station, satellite, SST, etc.). More sophisticated buyers should also be in a stronger position to understand more complex models that use combined data to capture the underlying catastrophic risk.

5.3.4 Summary of the Key Premises for Our Recommended Framework

In closing this discussion, we recommend a framework based on the three key premises: 1) weather index insurance is for consequential losses; 2) weather index insurance is for catastrophes; and 3) data constraints are lowest for risk aggregator products. We suggest that this framework be considered priorities that guide a *process* for developing sustainable insurance markets for disaster risks. Weather index insurance products addressing the consequential losses of catastrophic risks that improve the ability of risk aggregators to serve the poor may be a cost-effective entry point for new weather index insurance markets. Starting with risk aggregator products that cover consequential losses for disasters creates a foundation for future insurance products by building capacity among local insurers, the insurance regulator, and the target market. This foundation may also motivate data system investments as insurance awareness increases and local stakeholders develop a vision for extending products to other firms and to households.

We recognize that in some regions, conditions may motivate practitioners to deviate from the above recommendations — e.g., the presence of rich data sources, very specialized agricultural production, highly spatially correlated weather risk, etc. We are certain that there are many special cases where sound economic thinking and consideration of the data constraints *should* motivate a departure from these recommendations. Rather, the recommendations are general guidelines and are borne out of our evaluation of significant data constraints and of our goal of making weather index insurance more efficient, effective, sustainable, and scalable to more regions.

[23] Weather index insurance products for risk aggregators are typically designed with a small number of firms as the potential target market so it is much more feasible to address the specific needs for capacity building of each potential buyer.

Chapter 6 Important Research Areas and Vision for the Future

This SKR discusses data needs for weather index insurance and the challenges associated with offering index insurance in data sparse environments. We began this investigation with the recognition that, after several years and many pilot projects, practitioners, donors, and researchers are still searching for effective models to make weather index insurance scalable and sustainable. Index insurance has captured the attention of many new donors and new groups are venturing into the technically difficult process of index insurance product development. These stakeholders seem to recognize the several positive aspects of index insurance, but a number of unanswered questions exist that should rightly concern donors and multilateral agencies. For example:

1. Weather index insurance products have often been conducted as one-off demonstrations and, in some cases, have not clearly identified how these activities tie to a coherent framework of reducing poverty or developing sustainable markets. What implementation models should the development community prioritize for index insurance to enhance its effectiveness for increasing economic stability and long-term growth?
2. Throughout some regions of the world, including much of Africa, data constraints limit the development of weather index insurance products. These regions are often among those that could benefit the most from improved disaster risk management. The same data that can be used for index insurance will clearly also enhance disaster risk management efforts on a broader scale. Nonetheless, without investments in the establishing and maintaining data systems, it is not clear how index insurance can be developed and scaled in these regions.
3. Some of the index insurance products that currently appear to be gaining traction have been in regions where households engage in highly specialized means of production. For example, Mongolia relies so heavily on livestock production that it has an annual livestock mortality census. Many lessons can be learned from the Mongolia experience (e.g., designing markets and social programs so that they are complementary), but its data source is likely unique to Mongolia. As another example, index insurance in Malawi has been designed around specialty crops (e.g., tobacco) with vertically integrated value chains. Yet, we know that many of the world's poor do not engage in specialized agricultural production. How can these models be extended to reach more potential insureds?
4. Finance and economic theories suggest that index insurance is best suited for protecting against long-term wealth impacts of extreme natural disasters. Yet, based on their experiences, some practitioners have concluded that households will only remain interested in insurance if it covers moderate risks so that insureds receive frequent indemnities. How can index insurance markets be sustainable if theory posits that extreme risks should be prioritized while experience indicates households prefer insurance for moderate risks?

It is critical that these hard questions be addressed. Currently, index insurance is on the pathway of innovation. Proof of concept is only part of the innovation process. Next, it is imperative that products be designed that appeal to a wide range of stakeholders; otherwise, index insurance will not be scalable or sustainable. If donors believe the theoretical foundations and the models for implementation are sound, then further investments may lead to more effective and efficient models in the future. Yet, donor investments must be tempered by opportunity cost given the many unmet development needs. Thus, priority should be given to models of index insurance that seem likely to lead toward something that is efficient, effective, scalable, and sustainable.

Our initial intention in this document was to discuss how to assess and overcome data constraints to index insurance. Given the data constraints that we encountered as our investigation advanced and given our own experience with index insurance, we soon learned that we could not analyze and discuss data constraints in the absence of product design issues. Product design is predominantly affected by two important factors: 1) data constraints; and, 2) the practitioner's conceptual model for how index insurance contributes to economic growth. Implementation models with relatively high data requirements will not generally be scalable; models that seem to have a relatively minor effect on economic growth will not generally be sustainable. Our journey of review and analysis, which began long before writing this document, leads us to recommend three areas of focus for product design: consequential losses, catastrophes, and risk aggregator products (see conclusion of Chapter 5). We believe that these three areas of focus both *reduce data constraints and are generally the best means for index insurance to contribute to economic growth.* We recognize that our three recommendations run counter to most approaches that have been used in implementing index insurance to date. While these recommendations are grounded in theory and based on our experiences, we also recognize that they need to be tested more widely and rigorously.

In the next two years, our team will perform research with regard to some of the questions and research areas we put forth in Section 6.1. Our second SKR focuses on legal and regulatory component of developing weather index insurance, in particular, the legal and regulatory challenges of creating index insurance products designed to protect against consequential losses. We have experience designing consequential loss index products in two very different jurisdictions: Peru and Vietnam. The third SKR focuses on evaluating the scalability and sustainability of index insurance products. When our project is complete, we intend to integrate key analyses, recommendations, and themes from all SKRs into a single document to advance the ideas presented in the individual SKRs. While we hope to make our own contributions to the research needed to more fully understand the potential role of index insurance in helping lift rural people out of poverty, our efforts will not exhaust the researchable questions. Thus, we intend to widely distribute our work so as to motivate additional research on these important questions.

6.1 Research Questions Surrounding Demand for Index Insurance

Given our view that many important questions regarding the demand for index insurance are critical to understanding data needs, we develop the following important areas for research:

1. **Will framing index insurance around the many consequential losses of a weather event (not just yield losses) increase demand?** It seems rather obvious that decision makers are likely to purchase more insurance if the insurance product protects against more of their potential losses. But is it also possible that decision makers are more risk averse for some losses caused by extreme weather than they are for other losses? For example, decision makers may exhibit higher risk aversion with regard to loss events that destroy assets and thus affect the growth path of long-term wealth than they do to loss events that affect only some portion of current period net returns.

2. **Will framing index insurance as catastrophe insurance to protect against reductions in long-term wealth increase demand?** We have noted a common concern among practitioners that buyers will lose interest in weather index insurance unless they receive frequent indemnities. While there is certainly some research on the psychology of risk that would support this view, there is also widespread empirical evidence that people routinely purchase various types of insurance that will rarely pay an indemnity because the insurance protects against very low probability loss events (e.g., life insurance, flight insurance, earthquake insurance). Based, in part, on some of the psychology and behavioral economics literatures, we hypothesize that demand for weather index insurance depends in part on how the underlying risk problem is framed. There are many possible explanations for framing effects. As mentioned previously, decision makers may exhibit higher risk aversion with regard to loss events that destroy assets and thus affect the long-term trajectory of wealth accumulation than they do to loss events that affect only some portion of current period net returns. In addition, behavioral economists and psychologists have demonstrated that decision makers employ many biases and heuristics (e.g., judgment bias, hindsight bias, availability heuristics, and gambler's fallacy) in making risky decisions. Perhaps the nature of these biases and heuristics vary depending on whether the risky decision is defined in terms of wealth or annual returns (Quiggen, 1991; Quggin and Horowitz, 1995). This is another important area of research.

   A related research question emerges from our concerns regarding the long term sustainability of insurance products that pay more frequently. To what extent will experience with small and frequent payments for moderate losses and low payments for catastrophic events dampen the demand for these products over time?

3. **How will buyers react to index insurance products based on novel data sources?** As more complex models are developed using a combination of data sources, research will also be needed to assess buyer response to insurance products based on such novel data systems. This research should assess buyer response across different target markets (e.g., households, risk aggregators, etc.). We have argued that risk aggregators may be in a better position to understand sophisticated modeling efforts that lead to more complex indexes. Some analysis of potential buyer response has been conducted as part of the product development for an NDVI product in northern Kenya (Chantarat et al., 2009). To our knowledge, no other buyer response to satellite-based index insurance products has been conducted. Additionally, no fieldwork has been conducted comparing reactions to satellite-based index insurance products across different target markets.

6.2 Research Questions Surrounding Data

Because basis risk is such an important issue with index insurance, we first propose four different avenues of research pertaining to basis risk and product design.

6.2.1 Research Questions: What Product Designs Might Reduce Basis Risk?

1. **Does the spatial distribution of specific weather events change depending on the severity of the event?** This is a question of the physical presentation of weather. As we have suggested at various points in this SKR, it seems likely that for some weather events the spatial covariance is higher for more extreme weather occurrences. The

results of Bravar and Kavvas (1991) seem to support a hypothesis that extreme drought conditions are more spatially covariate than moderate drought conditions. This implies that an index insurance product designed to protect against drought would have lower basis risk for extreme droughts than for moderate droughts. This, in turn, has important implications for weather index insurance design, especially in data-sparse environments. Ultimately, this question requires good empirical research that employs advanced statistical procedures. We need to better understand the spatial presentation of weather variables and how that changes depending on the severity of a weather event. Appendix A describes an emerging research agenda designed to address these issues.

2. **Does the relationship between a specific weather event and realized losses change depending on the severity of the weather event?** Said differently, is the covariance between severe weather events and severe losses higher than the covariance between moderate weather events and moderate losses? This is a question of the physical processes underlying losses. The classic discussion of basis risk implicitly assumes that the covariance between a weather index and the realized loss is linear throughout the range of outcomes — that is, that the beta coefficient does not vary with the severity of the weather event. Again, the research findings on this question would have important implications for weather index insurance design.
3. **Does the covariance in investment returns change depending on the severity of the weather event?** We have hypothesized that the covariance of returns across the activities in a diversified portfolio often increases when severe weather events occur. Of course, this depends on the activities in the portfolio and how geographically concentrated they are. For example, if households have a family member working in another region and providing remittances, these returns would be unaffected by local weather events. Still, rural households invest in many activities exposed to the same extreme weather risks. For example, drought affects all rain-fed crops, reducing farm yields but also limiting opportunities for employment on other farms. At the community level, these aggregated losses hurt local businesses as well. If the covariance in returns is highest during extreme weather events, it is an indication that portfolio diversification is least effective under these conditions. Such an outcome would support using insurance to protect against extreme events and other risk management strategies, such as diversification, to protect against more moderate loss events.
4. **How does designing index insurance for consequential losses rather than for specific crops tend to affect basis risk?** This is a question of the physical presentation of the disaster. For example, is a specific level of excess rainfall a better estimator of one type of loss (low crop yields) or several types of losses (low crop yields, reduced small business revenues, property damage)? This is a difficult question, and the answer will depend in part on the risk and the context. Keeping in mind that most index insurance products are developed around a single index, the answer also depends on the correlation in the estimation error across losses. Finance theory on portfolio management will help demonstrate principles underlying this problem and the effects this correlation has on designing appropriate index insurance products.

6.2.2 Other Research Areas on Data Issues

Beyond the research questions targeted to learning more about how product design may reduce basis risks, there are several other research areas around data issues that merit attention.

1. **Improved risk assessment and qualitative data to enhance development of index insurance.** We have argued that qualitative data obtained from scientifically informed risk assessment procedures with key stakeholders may be equally if not more valuable than detailed historic data. In large part, this argument is based on both data constraints and the critical need to focus on infrequent catastrophic events. We have made it clear that even in situations where twenty or thirty years of loss data are available, these data may be inadequate for understanding the consequences of extreme and infrequent events. Such data also provide little or no information regarding how risk exposure may have changed (e.g., due to infrastructure investments, changes in production strategies, etc.). Some in the economic research community are quick to dismiss qualitative data as being "unscientific." Are there rigorous ways to examine using high quality but short time series of data to evaluate an extreme risk problem versus using qualitative data? How can protocols be improved to enhance the quality of risk assessment data for index insurance product development?

2. **Research that evaluates the value of emerging data systems needs to be supported.** While it is beyond our expertise, we place a high value on research that is being conducted to evaluate more efficient and accurate systems of weather data. In this document we raised questions about the sustainability and scalability of ground level instruments for measuring weather variables. Due to the current limited availability of weather stations, we have emphasized that novel data sources will likely be required to significantly increase the number of weather index insurance products offered in lower income countries. In particular, we have focused on satellite-based sources of vegetation and weather data. The underlying index for a weather index insurance product might be based exclusively on data from one of these satellite-based sources or might be based on a combination of satellite-based and ground-based (e.g., weather station) data sources. However, many questions remain about the novel data sources. Much research is needed to determine where, and for what variables, satellite-based data are likely to provide the most accurate measurements. For example, challenges exist with building consistent time series of data across different time periods due to changes in technology, satellite drift and aging of sensors. Higher spatial resolutions will likely be required for satellite-based data to support many types of index insurance offers. All of the cautions we raise about overfitting for simpler index insurance products could also apply to those developing more complex models that use a combination of data systems.

3. **Identifying regions where estimates of oceanic anomalies such as ENSO can improve weather risk transfer.** Many of the examples presented in this SKR have been based on our ongoing work developing ENSO-based insurance in Peru (see Appendix B for more description of this work). Oceanic anomalies can affect weather outcomes over large geographic areas. For example, ENSO affects weather in parts of South America, Central America, North America, and Australia but also affects weather in parts of Africa and Asia (Funk et al., 2003; McPhaden, 2003). Their widespread impact makes oceanic anomalies interesting candidates for weather index insurance but there is still much to learn. We have documented a strong relationship between sea surface temperatures off the coast of Peru and flooding in northern Peru due to El Niño. But how do ENSO measures affect other regions? What is the magnitude of the basis risk between ENSO measures and realized losses and (as discussed previously) does the magnitude of the basis risk depend on the severity of the oceanic anomaly? To what extent are extreme

weather events on the African continent associated with oceanic temperature anomalies measured in the north Atlantic or in the Gulf of Guinea? Are they strong enough to develop forecast insurance similar to what was developed in Peru?

4. **Identifying other potential indexes that could be used to make insurance payments before weather-induced loss events occur.** One of the interesting aspects of our ongoing work with El Niño Insurance in Peru is that extreme sea surface temperatures in November and December are indicators of flooding that generally does not occur until February or March. Thus, with an index insurance product based on sea surface temperature payments can be made *before* the extreme losses occur. We are calling this *forecast insurance* and, as we develop in Appendix B, this form of insurance can fit in the legal and regulatory environment as a form of consequential loss insurance. Policyholders can use the indemnities to fund loss mitigation efforts before flooding actually begins. There are significant economic efficiencies associated with policyholders being able to use indemnities to proactively prepare for a loss event rather than simply using the funds to recover after the loss occurs. Are there other regions of the world where oceanic anomalies or other measures could be used to make insurance payments before severe losses occur?

Chapter 7 Summary and Conclusions

Even in developed countries with sophisticated financial systems, financial market innovation is typically a long and slow process. The challenges are even greater in lower income countries where only the most basic financial services are available outside of major urban areas. Weather index insurance is an innovative financial instrument that holds great promise for helping decision makers in lower income countries manage their exposure to extreme weather events. Though the conceptual underpinnings of weather index insurance are rather simple and straightforward, the real-world application of those ideas can be extremely difficult. This SKR demonstrates that many of these difficulties are the result of data limitations.

Weather index insurance emerged from the need to develop agricultural insurance products that could be delivered to small-scale farmers in rural areas of lower income countries. After years of effort and failed experiments, development economists in the 1970s and 1980s essentially gave up trying to develop traditional crop insurance products for small-scale farmers in lower income countries. In the late 1990s a renewed interest in agricultural insurance for lower income countries stimulated research on weather index insurance as an alternative to traditional crop insurance. Since then, pilot programs have been instituted in several countries.

Despite about a dozen years of conceptualizing and working toward pilot programs and experimentation, market-based weather index insurance has still not been scaled to a significant level. While there are many reasons for this, data constraints are among the most important. A challenge for scaling up weather index insurance is that weather stations are generally quite sparse in rural areas of lower income countries. Furthermore, while there can be many developmental benefits to improving weather information systems within a country, the cost of installing and maintaining a sufficient density of weather stations specifically to support index insurance offers is likely prohibitive.

Alternative sources of weather data, generally collected from satellite platforms, have become available over recent decades. While these sources currently lack the spatial and/or temporal specificity required for many types of index insurance offers, the technologies are improving rapidly. At some point in the not too distant future, it may be feasible to base weather index insurance offers on weather measures collected from satellite-based platforms. More likely, weather measures will be available that integrate limited weather station data with more abundant data from alternative sources such as satellites. These advancements in sources of weather data will likely be spurred by developments in the market for weather index insurance. A comparison with the history of catastrophic modeling for earthquakes and hurricanes is insightful. Some thirty years ago insurers and reinsurers were challenged by the data and sophisticated modeling required to understand the underlying risk associated with insuring these major catastrophes. Academics from Stanford University were among the first to offer modeling services that met this emerging need. In 1988, the firm, Risk Management Solutions, emerged from these efforts. Others followed. There is evidence that a similar process is underway with weather index insurance modeling as part of the natural progression of market innovation.

While a number of techniques can be used in data-sparse environments to model relationships between weather variables and realized losses, we advise practitioners to exercise caution in utilizing these techniques. For example, when weather index insurance is targeted to losses for a specific crop, practitioners sometimes fit complex statistical relationships between the limited available weather data and crop-yield data. In so doing, they can show that the index "explains" a large part of the in-sample variability in crop yields. A concern with such approaches is that statistical relationships determined by overfitting the limited available in-sample data may not be supported when out of sample. Policyholders, who have been promised that the index is highly correlated with crop yield, are likely to end up being disappointed, frustrated, and, potentially, worse off for buying the insurance.

When few or no loss data are available, practitioners have sometimes used crop growth simulation models to determine relationships between weather variables and yield losses. A concern with this approach is that these models are parameterized for very specific crop varieties and regions. One cannot simply assume that the parameters contained in the models are generalizable to other crop varieties, regions, or farming practices. Also, while crop growth models are quite useful for estimating the effects of a change in a variable around the central tendency of the distribution, they are much less useful for predicting the effects of extreme weather events on yields — and it is just such extreme weather events that are the primary focus of weather index insurance.

We recommend a risk assessment process that utilizes both the limited available quantitative data and qualitative information collected from local sources. Those who have lived through previous extreme events tend to have a clearer picture of how households and businesses in the region were affected by those events. This risk assessment process operates under the recognition that weather risk and resulting losses occur in a larger system affected by many components: household livelihood strategies, geography, weather patterns, population dynamics, industry growth, cultural values, etc. As practitioners develop an understanding of risk in the local context, themes are likely to emerge that guide priorities in product development.

Our experience with developing index insurance products and our evaluations of other such efforts, have led us to three major recommendations regarding the role of index insurance in economic development efforts.

1) **Weather index insurance is for consequential losses.** Risk assessments indicate extreme weather events affect households and firms in many ways, reducing returns on investments (e.g., lower yields, reduced labor opportunities) and wealth positions (e.g., asset losses, household consumption demands). If weather index insurance is designed for a single aspect of the household portfolio of investments (e.g., crop risk), it may be of limited value to households. Because weather index insurance may be offered in regions where the target market has no previous experience with insurance, the onus is on practitioners to identify the needs of the target market through risk assessments and to design and market products with a vision for the ways in which extreme weather events are impeding growth for the target market.

2) **Weather index insurance is for catastrophic risk.[24]** Significant empirical evidence indicates that natural disasters can have extreme and long-term effects on poor households. The risk of low frequency, high severity, weather events can significantly hinder economic development. Insurance is a relatively expensive instrument so it is

[24] The focus of this SKR has been on developing market-based index insurance. Nevertheless, we acknowledge that an appropriate role for governments and donors may be needed for some catastrophic risk. GlobalAgRisk has made significant contributions to new thinking in this regard with our work in Mongolia.

best used to transfer extreme risks that cannot be managed using other methods. Other instruments such as savings and credit are more efficient mechanisms for managing moderate risks. Moreover, it seems likely that basis risk is higher for moderate weather risks than for extreme weather risks.

3) **Data constraints are lowest for risk aggregator products.** Risk aggregators such as rural banks and members of the agricultural value chain can use risk pooling to manage their exposure to idiosyncratic risks but not their exposure to correlated weather risks. Thus, they tend to limit the services they provide in rural areas that are highly exposed to correlated weather risks. Weather index insurance is only feasible for correlated risks, making it particularly well-suited for risk aggregators. The data systems required to support the offer of risk aggregator products also require less spatial specificity than those required for household insurance products. Therefore, risk aggregator products are particularly promising for data-constrained regions. Not only could these products potentially be supported by limited weather station infrastructure, but in some cases, satellite data may be sufficient.

We note that these recommendations address issues of index insurance product development and marketing, which may seem unusual in a document purportedly about data. However, as we have tried to emphasize throughout this document, data issues cannot be meaningfully separated from product development and marketing. Data limitations will inform what types of products are most feasible and for what target markets. Likewise, data requirements are always contextual and depend on the nature of the index insurance product being developed, its target market and application.

In closing, our analyses indicate that weather index insurance investments should be prioritized toward natural disaster risks that are likely to be impeding economic growth for poor households. These products are likely most effective if designed for the many consequential losses the target market experiences during a disaster — a much broader vision for weather index insurance than practitioners' historical focus on a specific crop, and, a considerable divergence from the practices of insuring moderate risks. Weather index insurance products that improve the ability of risk aggregators to serve the poor is consistent with this vision and may be a cost-effective entry point for new weather index insurance markets, especially in data-constrained regions. Starting with risk aggregator products that cover consequential losses for disasters creates a foundation for future insurance products by building capacity among local insurers, the insurance regulator, and the target market. This foundation may also motivate data system investments as insurance awareness increases and local stakeholders develop a vision for extending products to other firms and to households. In the long term, increased product offerings and target market specialization may create the way for a variety of new insurance products.

References

Ali A., T. Lebel, and A. Amani. "Rainfall Estimation in the Sahel. Part I: Error Function." *Journal of Applied Meteorology* 44(2005): 1691–1706.

Ali A., A. Amani, A. Diedhiou, and T. Lebel. "Rainfall Estimation in the Sahel. Part II: Evaluation of Rain Gauge Networks in the CILSS Countries and Objective Intercomparison of Rainfall Products." *Journal of Applied Meteorology* 44(2005): 1707–1722.

Alston, M. "Overview of the USDA Risk Management Agency Index-based Insurance Offering." 2nd Remote Sensing Index-Based Crop Insurance Working Group, San Francisco, CA, August 10–11, 2010.

Barnett, B. J., and O. Mahul. "Weather Index Insurance for Agriculture and Rural Areas in Lower-Income Countries." *American Journal of Agricultural Economics* 89(2007): 1241–1247.

Barnett, B. J., C. B. Barrett, and J. R. Skees. "Poverty Traps and Index-based Risk Transfer Products." *World Development* 36(2008): 1766–1785.

Bayarjargal, Y., A. Karnieli, M. Bayasgalan, S. Khudulmur, C. Gandush, and C. J. Tucker. "A Comparative Study of NOAA-AVHRR-derived Drought Indices Using Change Vector Analysis." *Remote Sensing of Environment* 105(2006): 9–22.

Berg, E., and B. Schmitz. "Weather-based Instruments in the Context of Whole-Farm Risk Management." *Agricultural Finance Review* 68(2008): 119–133.

Bohling, G. "Kriging." Course materials for C&PE940, Introduction to Geostatistics and Variogram Analysis, Kansas Geological Survey, October 19, 2005. Available at http://people.ku.edu/~gbohling/cpe940/Kriging.pdf
Accessed December 29, 2009.

Bravar, L., and M. L. Kavvas. "On the Physics of Droughts I. A Conceptual Framework." *Journal of Hydrology* 129(1991): 281–297.

Carter, M. R., F. Galarza, and S. Boucher. "Underwriting Area-based Yield Insurance to Crowd-in Credit Demand and Supply." Paper Prepared for the International Conference on Rural Finance Research: Moving Results into Policies and Practice, FAO Headquarters, Rome, Italy, March 17–19, 2007.

Carter, M. R., P. D. Little, T. Mogues, and W. Negatu. "Poverty Traps and Natural Disasters in Ethiopia and Honduras." *World Development* 35(2007): 835–856.

Chantarat, P. "Designing and Estimating Demand for Index-based Livestock Insurance in Northern Kenya: An Introduction." Presentation at the International Livestock Research Institute, Nairobi, Kenya, August 25, 2008.

Chantarat, S., A. G. Mude, C. B. Barrett, and M. R. Carter. "Designing Index Based Livestock Insurance for Managing Asset Risk in Northern Kenya." Working Paper, Cornell University, Ithaca, NY, July, 2009.

Chen, J., P. Jonsson, M. Tamura, Z. Gu, B. Matsushita, and L. Eklundh. "A Simple Method for Reconstructing a High-Quality NDVI Time-Series Data Set Based on the Savitzky-Golay Filter." *Remote Sensing of Environment* 91(2004): 332–344.

Collier B., A. L. Katchova, and J. R. Skees. "Loan Portfolio Performance and El Niño, an Intervention Analysis." Accepted for presentation at the Southern Agricultural Economics Association, February 8, 2010; *Agricultural Finance Review,* in review.

Collier, B., J. R. Skees, and B. J. Barnett. "Weather Index Insurance and Climate Change: Opportunities and Challenges in Lower Income Countries." *Geneva Papers on Risk and Insurance — Issues and Practice* 34(2009): 401–424.

CPC (Climate Prediction Center). "NOAA CPC Morphing Technique." Available at http://www.cpc.noaa.gov/products/janowiak/cmorph_description.html
Accessed November 25, 2009.

Dinku, T., S. Chidzambwa, P. Ceccato, S. Connor, and C. F. Ropelewski. "Validation of High-Resolution Satellite Rainfall Products over Complex Terrain." *International Journal of Remote Sensing* 29(2008):4097–4110.

ECMWF (European Centre for Medium Range Weather Forecasts. Web site. Available at http://www.ecmwf.int/
Accessed August 31, 2010.

Fensholt, R., K. Rasmussen, T. T. Nelson, and C. Mbow. "Evaluation of Earth Observation Based Long Term Vegetation Trends — Intercomparing *NDVI* Time Series Trend Analysis Consistency of Sahel from *AVHRR GIMMS*, Terra *MODIS* and SPOT VGT data." *Remote Sensing of Environment* 113(2009): 1886–1898.

Funk, C., J. Michaelsen, J. Verdin, G. Artan, G. Husak, G. Senay, H. Gadain, and T. Magadzire. "The Collaborative Historical African Rainfall Model: Description and Evaluation." *International Journal of Climatology* 23(2003): 47–66.

Giné, X., R. Townsend, and J. I. Vickery. "Patterns of Rainfall Insurance Participation in Rural India." *World Bank Economic Review* 22(2008): 539–566.

GlobalAgRisk. "Legal and Regulatory State of Knowledge Report." Report prepared for the Bill and Melinda Gates Foundation, Seattle, WA, February, 2010.

GlobalAgRisk. *Designing Agricultural Index Insurance in Developing Countries: A GlobalAgRisk Market Development Model Handbook for Policy and Decision Makers.* Lexington, KY: GlobalAgRisk, 2009. Available at http://www.globalagrisk.com/Pubs/2009_GlobalAgRisk_Market_Development_Model_Handbook.pdf
Accessed January 3, 2010

GlobalAgRisk. "Index Insurance for Weather Risk in Lower-Income Countries." Report prepared for the United States Agency for International Development, Washington, DC, November, 2006. Available at http://www.globalagrisk.com/Pubs/PRIMER_USAID.pdf
Accessed January 3, 2010

Grantham-McGregor, S., Y. B. Cheung, S. Cueto, P. Glewwe, L. Richter, and B. Strupp. "Development Potential in the First 5 Years for Children in Developing Countries." *Lancet* 369(2007): 60–70.

Hartell, J., and J. R. Skees. "Prefeasibility Analysis: Index-based Weather Risk Transfer in Mali." Report prepared for Save the Children Federation and the United States Agency for International Development, Washington, DC, February, 2009.

Hazell, P. B. R., C. Pomareda, and A. Valdes. *Crop Insurance for Agricultural Development: Issues and Experience*. Baltimore: The Johns Hopkins University Press, 1986.

Hellmuth, M. E., D. E. Osgood, U. Hess, A. Moorhead, and H. Bhojwani, eds. *Climate and Society Issue 2: Index Insurance and Climate Risk: Prospects for Development and Disaster Management.* New York, NY: International Research Institute for Climate and Society, Columbia University, 2009.

Huete, A. R., K. F. Huemmrich, T. Miura, X. Xiao, K. Didan, W. van Leeuwen, F. Hall, C. J. Tucker. "Vegetation Index Greenness Global Data Set." White Paper for NASA ESDR/CDR, April 2006. Available at http://landportal.gsfc.nasa.gov/Documents/ESDR/VI_Huete_whitepaper.pdf Accessed August 31, 2010.

Hipple, J. D. "Spatial, Temporal and Radiometric Variability issues that may affect the design of index insurance products from Remotely Sensed Data Sets". 2nd Remote Sensing Index-based Crop Insurance Working Group, San Francisco, CA, August 10–11, 2010.

Holecz, F. (Sarmap S.A.). Personal correspondence, November 9, 2009.

Joyce, R., J. Janowiak, P. Arkin, and P. Xie. "CMORPH: A Method that Produces Global Precipitation Estimates from Passive Microwave and Infrared Data at High Spatial and Temporal Resolution." *Journal of Hydrometeorology* 5(2004): 487–503.

Kahneman, D., and A. Tversky. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47(1979): 263–291.

Kalnay E., M. Kanamitsu, R. Kistler, W. Collins, D. Deaven, L. Gandin, M. Iredell, S. Saha, G. White, J. Woollen, Y. Zhu, M. Chelliah, W. Ebisuzaki, W. Higgins, J. Janowiak, K. C. Mo, C. F. Ropelewski, J. Wang, A. Leetmaa, R. Reynolds, R. Jenne, and D. Joseph. "The NCEP/NCAR 40-Year Reanalysis Project." *Bulletin of the American Meteorological Society* 77(1996): 437–471.

Khalil, A. F., H. H. Kwon, U. Lall, M. J. Miranda, and J. R. Skees. "El Niño Southern Oscillation-based Index Insurance for Floods: Statistical Risk Analyses and Application to Peru." *Water Resources Research* 43(2007); W10416, doi:10.1029/2006WR005281.

Kunreuther, H. "Mitigating Disaster Losses through Insurance." *Journal of Risk and Uncertainty* 12(1996): 171–187.

Kunreuther, H. "Limited Knowledge and Insurance Protection." *Public Policy* 24(1976): 227–261.

Kunreuther, H., and P. Slovic. "Economics, Psychology, and Protective Behavior." *American Economic Review* 68(1978): 64–69.

Laws, K. B., J. E. Janowiak, and G. J. Huffman. "Verification of Rainfall Estimates over Africa Using RFE, NASA MPA-RT, and CMORPH." Paper prepared for the American Meteorological Society 18th Conference on Hydrology, Seattle, WA, January 14, 2004. Available at http://www.cpc.ncep.noaa.gov/products/fews/AMS2004_Rainfall.pdf Accessed December 29, 2009.

Lebel, T., and A. Amani. “Rainfall Estimation in the Sahel: What is the Ground Truth?” *Journal of Applied Meteorology* 38(1999): 555–568.

Lotsch, A., W. Dick, and O. P. Manuamorn. *Assessment of Innovative Approaches for Flood Risk Management and Financing in Agriculture.* Washington DC: World Bank Press, 2009.

Mahul, O., and J. R. Skees. "Managing Agricultural Risk at the Country Level: The Case of Index-based Livestock Insurance in Mongolia." Policy Research Working Paper WPS4325, The World Bank, Washington, DC, August 1, 2007.

Martin, S. W., B. J. Barnett, and K. H. Coble. “Developing and Pricing Precipitation Insurance.” *Journal of Agricultural and Resource Economics* 26(2001): 261–274.

McPhaden, M. J. “El Niño and La Niña: Causes and Global Consequences.” *Encyclopedia of Global Environmental Change, Volume 1, The Earth System — Physical and Chemical Dimensions of Global Environmental Change.* Munn, T., M. M. MacCracken, and J.S. Perry, eds. Hoboken, NJ: John Wiley and Sons, 2003.

MDA (MacDonald, Dettwiler and Associates). “RADARSAT-2.” Available at
http://www.radarsat2.info
Accessed December 29, 2009.

Miranda, M. J. “Area-Yield Insurance Reconsidered.” *American Journal of Agricultural Economics* 73(1991): 233–242.

NASA (National Aeronautics and Space Administration). “MODIS Web.” Available at
http://modis.gsfc.nasa.gov/about/design.php
Accessed August 31, 2010(a).

NASA (National Aeronautics and Space Administration). “TRMM — Tropical Rainfall Measuring Mission.” Available at
http://trmm.gsfc.nasa.gov/data_dir/data.html
Accessed August 31, 2010(b).

NASA (National Aeronautics and Space Administration). “VIIRS — Visible Infrared Imaging Radiometer Suite.” Available at
http://jointmission.gsfc.nasa.gov/science/VIIRScontent.REM.html
Accessed August 31, 2010(c).

NASA (National Aeronautics and Space Administration). “Global Precipitation Measurement.” Available at
http://gpm.gsfc.nasa.gov/features/estimateofrainfall.html
Accessed August 31, 2010(d).

NASA (National Aeronautics and Space Administration). “Imaging Radar.” Available at
http://southport.jpl.nasa.gov/
Accessed August 31, 2010(e).

NationalAtlas.gov. “What is AVHRR?” Available at
http://rnp782.er.usgs.gov/atlas2/articles/mapping/a_avhrr.html
Accessed December 29, 2009.

Peters, A. J., E. A. Walter-Shea, L. Ji, A. Vina, M. Hayes, and M. D. Svoboda. “Drought Monitoring with NDVI-based Standardized Vegetation Index.” *Photogrammetric Engineering and Remote Sensing* 68(2002): 71–75.

Poccard, I., S. Janicot, and P. Chamberlin. “Comparison of Rainfall Structures between NCEP/NCAR Reanalyses and Observed Data over Tropical Africa. “ *Climate Dynamics* 16(2000): 897–915.

Quiggin, J. “Comparative Statics for Rank-Dependent Expected Utility Theory.” *Journal of Risk and Uncertainty* 4(1991): 339–350.

Quiggin, J., and J. Horowitz. “Time and Risk.” *Journal of Risk and Uncertainty* 10(1995): 37–55.

Sachs, J. D., and J. W. Arthur. “The Millennium Project: A Plan for Meeting the Millennium Development Goals.” *Lancet* 365(2005): 347–353.

Sandia National Laboratories. “Synthetic Aperture Radar.” Available at www.sandia.gov/radar/sar.html Accessed December 29, 2009.

Sapiano, M. R. P., and P. A. Arkin. “An Intercomparison and Validation of High-Resolution Satellite Precipitation Estimates with 3-Hourly Gauge Data.” *Journal of Hydrometeorology* 10(2009): 149–166.

Schiermeier, Q. “Radar Satellites: Mountains to Molehills.” *Nature* 446, pp. 22–23, March 1, 2007.

Schiermeier, Q. “US Tests Satellite Tool for Hurricane Monitoring.” *Nature* 437, p. 301, September 15, 2005.

Sharpe, W. F. “Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk.” *Journal of Finance* 19(1964): 425–442.

Shen, Y., A. Xiong, and P. Xie. “Performance of High-Resolution Satellite Precipitation Products over China.” Paper prepared for the American Meteorological Society 22nd Conference on Hydrology, New Orleans, LA, January 23, 2008. Available at http://ams.confex.com/ams/88Annual/techprogram/paper_131345.htm Accessed December 29, 2009.

Skees, J. R. “Challenges for Use of Index-based Weather Insurance in Lower Income Countries.” *Agricultural Finance Review* 68(2008): 197–217.

Skees, J. R., and B. J. Barnett. “Enhancing Microfinance Using Index-based Risk Transfer Products.” *Agricultural Finance Review* 66(2006): 235–250.

Skees, J. R., J. R. Black, and B. J. Barnett. “Designing and Rating an Area Yield Crop Insurance Contract.” *American Journal of Agricultural Economics* 79(1997): 430–438.

Skees, J. R., P. B. R. Hazell, and M. Miranda. “New Approaches to Public/Private Crop Yield Insurance in Developing Countries.” Environment and Production Technology Division (EPTD) Discussion Paper 55, International Food Policy Research Institute (IFPRI), Washington, DC, November, 1999.

Skees, J. R., A. G. Murphy, B. Collier, M. J. McCord, and J. Roth. “Scaling Up Index Insurance: What Is Needed for the Next Big Leap Forward?” Weather insurance review prepared by the MicroInsurance Centre and GlobalAgRisk for Kreditanstalt für Wiederaufbau (KfW), December, 2007.

Sylla, M. B., E. Coppola, L. Mariotti, F. Giorgi, P. M. Ruti, A. Dell'Aquila, and X. Bi. "Multiyear Simulation of the African Climate Using a Regional Climate Model (RegCM3) with the High Resolution ERA-Interim Reanalysis." *Climate Dynamics* 35(2009): 231–247.

Thenkabail, P. S., R. B. Smith, and E. DePauw. "Hyperspectral Vegetation Indices and Their Relationships with Agricultural Crop Characteristics." *Remote Sensing of Environment* 71(2000): 158–182.

Trenberth, K. E., R. Dole, Y. Xue, K. Onogi, D. Dee, M. Balmaseda, M. Bosilovich, S. Schubert, and W. Large. "Atmospheric Reanalyses: A Major Resource for Ocean Product Development and Modeling." *Proceedings of OceanObs'09: Sustained Ocean Observations and Information for Society (Vol. 2), Venice, Italy, September 21–25, 2009.* Hall, J., D. E. Harrison, and D. Stammer, eds. Paris: ESA Publications.

Tucker, C., J. Prinzon, M. Brown, S. Daniel, E. Pack, R. Mahoney, E. Vermote, and N. El Saleous. "An Extended AVHRR 8-Km NDVI Dataset Compatible with MODIS and SPOT Vegetation NDVI Data." *International Journal of Remote Sensing* 26(2005): 4485–4498.

Tversky, A., and D. Kahneman. "Availability: A Heuristic for Judging Frequency and Probability." *Cognitive Psychology* 5(1973): 207–232.

UCAR (University Corporation of Atmospheric Research). "NCEP/NCAR Global Reanalysis Products, 1948–Continuing." CISL Research Data Archives. Available at http://dss.ucar.edu/datasets/ds090.0/
Accessed August 31, 2010.

Uppala, S. M., P. W. Kållberg, A. J. Simmons, U. Andrae, V. da Costa Bechtold, M. Fiorino, J. K. Gibson, J. Haseler, A. Hernandez, G. A. Kelly, X. Li, K. Onogi, S. Saarinen, N. Sokka, R. P. Allan, E. Andersson, K. Arpe, M. A. Balmaseda, A. C. M. Beljaars, L. van de Berg, J. Bidlot, N. Bormann, S. Caires, F. Chevallier, A. Dethof, M. Dragosavac, M. Fisher, M. Fuentes, S. Hagemann, E. Hólm, B. J. Hoskins, L. Isaksen, P. A. E. M. Janssen, R. Jenne, A. P. McNally, J.-F. Mahfouf, J.-J. Morcrette, N. A. Rayner, R. W. Saunders, P. Simon, A. Sterl, K. E. Trenberth, A. Untch, D. Vasiljevic, P. Viterbo, and, J. Woollen. "The ERA-40 Re-analysis." *Quarterly Journal of the Royal Meteorological Society* 131(2005): 2961–3012.

Weler, J., and D. Heming. "Measuring Vegetation (NDVI & EVI)." NASA Earth Observatory, Featured Articles. Available at http://earthobservatory.nasa.gov/Features/MeasuringVegetation
Accessed January 2, 2010.

White, D. H., and J. J. Walcott. "The Role of Seasonal Indices in Monitoring and Assessing Agricultural and Other Droughts: A Review." *Crop and Pasture Science* 60(2009): 599–616.

WMO (World Meteorological Organization). *Capacity Assessment of National Meteorological and Hydrological Services in Support of Disaster Risk Reduction: Analysis of the 2006 WMO Disaster Risk Reduction Country-Level Survey.* Geneva: WMO Publications, 2008.

World Bank. *World Development Report 2008: Agriculture for Development*. Washington, DC: The World Bank, 2007.

Wright, B. D. "Why Government Crop Insurance?" *Income Stabilization in Agriculture: The Role of Public Policies*. Cafiero, C., and A. Cioffi, eds. Naples: Edizioni Scientifiche Italiane, 2006.

Xie, P., and P. A. Arkin. “Global Precipitation: a 17-Year Monthly Analysis Based on Gauge Observations, Satellite Estimates, and Numerical Model Outputs.” *Bulletin of the American Meteorological Society* 78(1997): 2539–2558.

Zeweldi, D. A., and M. Gebremichael. “Evaluation of CMORPH Precipitation Products at Fine Space–Time Scales.” *Journal of Hydrometeorology* 10(2009): 300–307.

Appendix A New Approaches for Index Insurance — El Niño Insurance in Peru

Jerry R. Skees[1]

[1] A version of this paper appears in Kloeppinger-Todd, R., and M. Sharma, eds. *Innovations in Rural Finance for Developing Countries, 2020 Vision Initiative Brief Series.* Washington, DC: International Food Policy Research Institute (IFPRI).

New Approaches for Index Insurance — El Niño Insurance in Peru

It is now possible for stakeholders in Peru to purchase a new form of insurance that pays in early January *before* catastrophic flooding created by an extreme El Niño begins in February through April in the northern regions of Peru. The El Niño Insurance product was introduced by a Peruvian insurance company in 2010. A major global reinsurer carries most of the risk. The Peruvian insurance regulator approved this insurance in May 2009. This new insurance product uses the U.S. National Oceanic and Atmospheric Administration (NOAA) measure of sea surface temperature known as ENSO (El Niño Southern Oscillation) as the event that triggers payments. This El Niño Insurance is the first regulated "forecast index insurance" product in the world. ENSO information from November and December is used to make payments in January. Opening the way for forecast insurance could enhance the overall progress associated with developing index-based insurance products for extreme weather events. Having early payments prior to an extreme weather event affords the opportunity to prepare and potentially adapt so as to lower the actual loss.

In recent years, there have been a growing number of pilot tests of index insurance for weather risk, motivated by an increased understanding of how natural disasters affect developing countries. Beyond immediate suffering (deaths, destroyed assets, lost income, etc.), the indirect effects are equally troublesome — economic growth can be disrupted, the poor are thrust into permanent poverty traps, and the mere presence of these risks constrain access to financial services and cause many decision makers to pursue low-return, low-risk strategies that impede economic progress.

Much of the development of index insurance focuses on agriculture, as activities associated with agriculture remain the primary livelihood strategies for the rural poor in developing countries. Index insurance uses an objective measure (an index) of a natural event known to cause losses (e.g., excess or shortfalls in rain, river levels, extreme sea surface temperatures, etc.). Using an index as a measure of the insured event dispenses with expensive loss assessments of individual policyholders. Furthermore, moral hazard and adverse selection, problems that plague traditional forms of insurance, are diminished. Given these advantages, index insurance may be well-suited to developing countries where data are sparse and delivery of financial services to smallholder households increases the per-unit cost of traditional insurance.

Despite the promise of index insurance, progress is slow at the micro level. Decision makers in smallholder households must still be educated about index insurance; demand can be low; cost-effective systems to sell to smallholder households must be developed; legal and regulatory systems must be developed, etc. Presently, index insurance may be better suited for risk aggregators — those lending to farmers, firms in the value chain, and farmer associations. Focusing first on risk aggregators may also accelerate the potential to build linkages and sustainable products that will directly serve smallholder households.

Index Insurance Is Suitable for Some Correlated Losses in Developing Countries

As a precondition of index insurance, losses created by the natural disaster to be insured must be strongly correlated. A clear measure of correlated losses is when a large number of individuals and risk aggregators suffer losses at the same time. Correlated losses from natural disasters constrain the development of credit markets for the rural poor, particularly for those involved in agriculture. Lenders cannot absorb the risk exposure of a large number of borrowers who may be unable to pay off loans after a major natural disaster.

Likewise, an insurer deciding to write any form of insurance against extreme weather events must have a means to transfer these risks — generally via a global reinsurer. Insurers in developing countries rarely have business practices that allow them to access global reinsurers. If the index being used is fully transparent, the global reinsurer understands the systems that are used to estimate the index. This is the certainly the case for ENSO measures, which have been developed for over 50 years by the U.S. agency, NOAA.

Extreme weather events such as drought and flooding can also have associated consequential losses that extend beyond what traditional crop insurance pays for losses of a specific crop. For example, in a number of African nations, where owning livestock is a form of savings, extreme droughts force large numbers of farmers to sell-off their livestock at the same time. These forced sales depress local prices compounding the losses. Floods and droughts also generally influence the quality of crops, not just the yields. And, strategies to diversify cropping enterprises to manage risk can quickly prove ineffective if the drought or floods negatively affect all of the crops at the same time. Processors, laborers, and any number of local businesses that depend on local crop production also suffer.

El Niño Insurance as a Form of Business Interruption Insurance for Consequential Losses

In Peru, where the El Niño Insurance is being tested, the consequential losses and problems associated with extreme rainfall (Figure A1) and catastrophic flooding are enormous — crops are lost, trees are killed, soils wash away, transportation systems break down, disease problems (e.g., malaria) increase, and markets are destroyed. When individuals and local markets suffer in this fashion, it is expected that firms in the value chain and the financial sector will also suffer.

Given the levels of rainfall in the region, it is full understandable why there were major disruptions in the northern region of Peru (Figure A1). The volume of water in the Piura River was also about 40 times normal in these two extreme El Niño years. In 1998, with a clear indication that El Niño was coming, farmers simply did not plant crops, resulting in a 27 percent drop in fertilizer sales in northern Peru. Agricultural lending was growing at a significant pace before the 1997–98 El Niño. The growth completely stopped after the event. Microfinance institutions had a significant increase in problem loans. The total agricultural portfolio for all MFIs in Piura had an increase of over 10 percentage points for loans that were 60 days late or more. Caja Piura, a leading MFI in Piura, restructured an estimated 3.8 percent of its total loan portfolio due to this event. Additionally, the major source of capital for the MFIs (member deposits and savings) suffered as people withdrew funds to cope with the problems created by the event.

Figure A1 Extreme El Niño Events of 1982/83 : 1997/98

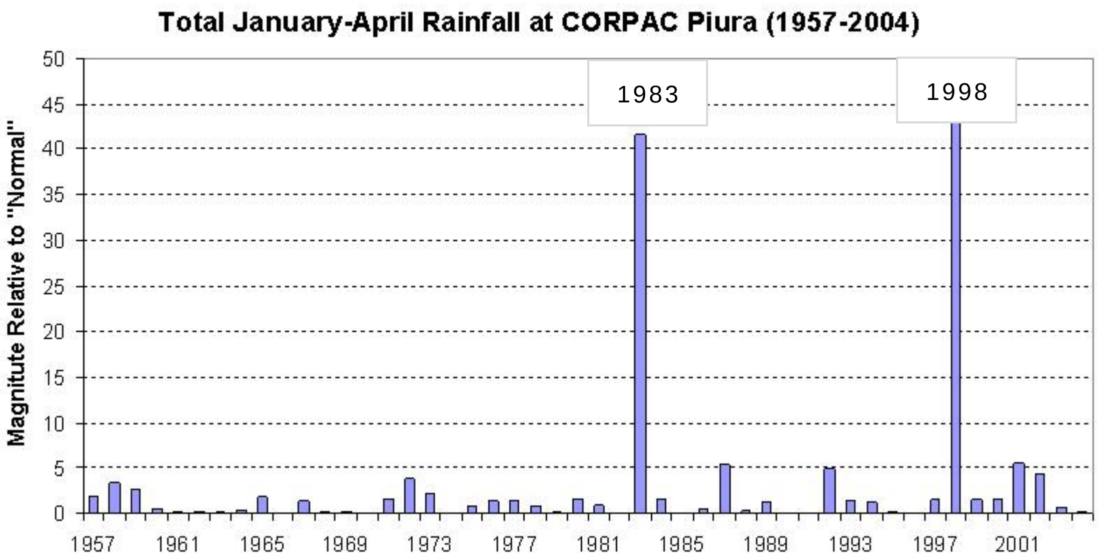

Based on an increased understanding of these types of associated problems, the El Niño Insurance in Peru was presented to the Peruvian regulator as a form of business interruption insurance designed to pay for consequential losses that are linked to extreme flooding that is highly correlated with ENSO. Furthermore, given that extreme ENSO measures in November and December are clear signals of an impending disaster, it was also accepted that stakeholders such as microfinance institutions would be incurring additional expenses even before the actual disastrous flooding begins in February through April. Assessments of the consequential losses, which are estimated using the ENSO measure, are done before the event. Assessing consequential losses including business interruptions is extremely difficult; therefore, the form of loss adjustment for the El Niño Insurance can be as acceptable to regulators as more traditional loss assessment processes of business interruptions (e.g., business revenue losses created by an event like a building fire that disrupts normal business). There is precedent for special forms of insurance referred to as “valued policies.” In the case of valued policies, there is a pre-agreed value and a pre-agreed event that will create losses. Experience in Peru demonstrates that index insurance can be presented in the same manner. These were important developments in properly positioning index insurance in the legal and regulatory environment.

The El Niño Insurance uses the monthly SST for ENSO Region 1.2 (0-10°South, 90°West-80°West), measured and reported by the NOAA Climate Prediction Center (CPC, 2010). The basis for payment is the average of two months — November and December. Payments begin when this measure exceeds 24.5 degrees Celsius, and payments reach a maximum when the measure reaches 27 degrees. The payout function is linear between these two temperatures. Thus, the payout rate[2] is calculated as:

[2] Using this calculation, the payout rate in 1998 would have been 71 percent of sum insured.

$$Payout\,Rate = \frac{(ENSOIndex - 24.5)}{(27 - 24.5)}$$

The insured selects the sum insured. Indemnity payments are made by multiplying the payout rate times the sum insured. The selection of sum insured should be based on a risk assessment that estimates the largest losses that may occur under the worst flooding event. The regulator could require documentation of these estimates to serve as the maximum value of insurance allowed. Prudent insureds will be more likely to select a value that is less than these estimates given the expense of this type of catastrophe insurance.

Since the El Niño Insurance pays before the catastrophe, educational efforts and workshops have been focused on helping the target markets understand how they might use the extra cash to mitigate the impending crisis to the extent possible. Farmer associations in remote regions of Piura, Peru, have expressed an interest in using the funds to clear drainage systems. Lenders are interested in using payments to ease the liquidity crisis as they work with problem loans at the same time that they see reductions in savings and deposits. Those in the value chain are interested in smoothing their losses. The El Niño Insurance is also being presented to local and regional governments to provide ready cash that may be able to mitigate some of the problems that are certain to emerge with catastrophic flooding.

At this stage, the El Niño Insurance is not being made available to smallholder households. However, the product can be tied to other financial services in a fashion that give smallholders greater access to these services at better prices. Targeting the El Niño Insurance to the risk aggregator first has proven a highly valuable exercise. The interest and involvement of the Peruvian financial regulators increases the potential that a sustainable index insurance product is being developed.

References

CPC (Climate Prediction Center). "Monthly Atmospheric and SST Indices." Webpage available at http://www.cpc.ncep.noaa.gov/data/indices/sstoi.indices Accessed August 31, 2010.

Skees, J. R. "Challenges for Use of Index-based Weather Insurance in Lower Income Countries." *Agricultural Finance Review* 68 (2008): 197–217.

Skees, J. R., and B. J. Barnett. "Enhancing Micro Finance Using Index-based Risk Transfer Products." *Agricultural Finance Review* 66(2006): 235–250.

Skees, J. R., J. Hartell, and A. G. Murphy. "Using Index-based Risk Transfer Products to Facilitate Micro Lending in Peru and Vietnam." *American Journal of Agricultural Economics* 89(2007): 1255–1261.

Appendix B Dependency Structure and Data Adequacy in Index Insurance Product Design and Market Development

Mario J. Miranda[1]

[1] Mario Miranda is the Andersons Professor of Finance and Risk Management in the department of agricultural, environmental, and development economics at The Ohio State University.

State of Knowledge Report: Dependency Structure and Data Adequacy in Index Insurance Product Design and Market Development

Mario J. Miranda

The Ohio State University
and
GlobalAgRisk, Inc.

December 23, 2009

Abstract

This technical report addresses various issues pertaining to the statistical methods used in index insurance product design and market development. Conventional actuarial methods based on univariate loss claim models are of limited applicability to index insurance because, unlike conventional insurance, index insurance indemnities are based, not on the verifiable losses, but rather a distinct random variable, the index. Statistical methods commonly used by economists to assess the potential demand for index insurance are also found wanting because of their reliance on linear correlation analysis, which is incapable of capturing nonlinear dependencies that may exist at the relevant extremes of the index and loss distributions. We propose novel approaches to the actuarial analysis of index insurance products based on copulas. Copulas are especially well-suited for capturing the complex dependencies that exist among extreme values of jointly distributed random variables, but remain little used by economists to analyze index insurance products.

1 Introduction

There are many issues that arise when assessing whether existing data and data collection practices are adequate to support index insurance product design and market development. Some issues are practical in nature. For example, to develop an index insurance product one must ask: Are data for candidate indices and insurable losses available, and, if so, are they readily accessible? Have the data been reliably collected and recorded in accordance with internationally recognized standards? Will data in the future be collected in a secure manner that will command the confidence of the insurer and the insured that index insurance indemnity payments will be settled in a transparent, fair, and timely manner?

However, there are also technical issues that arise when assessing the adequacy of data for index insurance product design and market development. From a purely actuarial perspective, available data are adequate only if they support accurate computation of actuarial statistics of interest to the insurer and the insured. For example, the expectation of the indemnities to be paid by an index insurance contract must be estimated as the first step in setting the premium. The standard error of the expected indemnity estimate must also be computed to allow the insurer to appropriately load the premium for parametric uncertainty. The probability distribution of indemnity payments for a portfolio of index insurance contracts must be derived in order to compute the maximum probable loss for an insurer's book of business. And how well indemnities track the losses of the insured must be reliably estimated to assess the potential demand for an index insurance product.

The precision with which these actuarial computations can be performed, and thus the adequacy of existing data, ultimately depend on the joint distributions of indices and losses and the statistical methods used to model them. Actuaries have developed a broad range of statistical techniques to model loss claim distributions based primarily on univariate statistical models that assume independence across claims. Univariate statistical methods, however, are of limited use in the analysis of index insurance contracts. With an index insurance contract, the variable used to determine indemnities, the index, is distinct from losses. Also, sets of indices, such as a rainfalls at different locations within a defined geographical area, usually exhibit correlation due to systemic weather effects, making the assumption of independence untenable for index insurance portfolio analysis and undesirable for efficient premium rate computation.

Economists have performed a substantial number statistical studies that have raised doubts about value of index insurance as a risk management tool.

The primary criticism of index insurance is that it is possible for the insured to suffer a loss yet not receive an indemnity. The potential severity of this problem, often referred to as “basis risk”, is typically measured empirically using the Pearson linear correlation coefficient, a statistical measure of the degree of linear dependence that exists between a pair of random variables. Linear correlation coefficients between losses and indices, however, are unreliable measures of basis risk for two reasons. First, the linear correlation coefficient between an index and insurable losses may be low, even though the index and the losses are strongly related, but in a nonlinear fashion. In these instances, basis risk can be substantially reduced through the use of an appropriate nonlinear indemnity schedule. Second, index insurance contracts provide indemnities only for extreme losses associated with extreme values of the index. The relationship between extreme values of losses and an index could be strong, yet be missed because an empirical linear correlation coefficient weighs all observations equally.

Actuarial assessments of index insurance products call for the use of multivariate statistical methods that can faithfully capture the distributional dependence that exists among indices and between specific indices and insurable losses, particularly in the extremes of the distributions. Copulas, which provide a flexible theoretical framework for capturing dependence among random variables, are well-suited for this task. Financial analysts began to take a strong interest in copulas in the wake of the financial crisis of 2007-9 ([1], [9], [11], [12], [18], [22], [28]; [6]). As a result of the crisis, financial analysts began to ask whether stock returns are more highly correlated during financial crises than in normal times, thus rendering stock portfolios riskier than predicted by conventional asset pricing models. The questions of interest to financial analysts are very similar to those that must be addressed in index insurance design: in both cases, one is concerned with the degree of dependence exhibited by multiple random variables at the extremes of their joint distribution. Copulas provide a formal framework for addressing questions such as these. However, copulas remain little used in index insurance applications.

In what follows we review the basic features of copulas and notions of dependence that are relevant to index insurance analysis. We survey common copula structures and related empirical methods and discuss how they might be adapted and extended to the actuarial analysis of index insurance products. We also propose a series of empirical applications that can illustrate and test the utility of copula methods in index insurance product design and market development.

2 Copulas

A copula is a function that describes how univariate marginal distributions are "coupled" together to form a multivariate distribution ([10]; [24]; [29]; [30]; [14]). Formally, an n-dimensional copula is a joint cumulative distribution function of n interdependent random variables, each of which, on the margin, is uniformly distributed on the unit interval [25]. Alternatively, an n-dimensional copula may be defined as a function C on the n-dimensional unit cube $[0,1]^n$, with values in the unit interval $[0,1]$, that satisfies the following conditions:

- $C(u) = 0$ whenever at least one component of u equals zero;
- $C(u) = u_i$ whenever all but the i^{ith} component of u equals 1;
- $C(u)$ assigns nonnegative probability to any n-dimensional cube in $[0,1]^n$.

The role that copulas play in capturing the interdependency among jointly distributed random variables is explained by Sklar's Theorem. Sklar's Theorem states that any continuous n-dimensional cumulative distribution function $F : \Re^n \mapsto [0,1]$ can be uniquely written

$$F(x_1, x_2, \ldots, x_n) = C(F_1(x_1), F_2(x_2), \ldots, F_n(x_n))$$

where C is an n-dimensional copula and F_i is the i^{th} marginal cumulative distribution function associated with F. Conversely, if C is an n-dimensional copula and $F_i : \Re \mapsto [0,1]$ are univariate cumulative distribution functions, then F as defined above is a cumulative distribution function on $\Re^n$ with marginal cumulative distributions F_i. The joint probability density function associated with a differentiable cumulative distribution function F can be recovered from its copula decomposition through the relation

$$f(x_1, x_2, \ldots, x_n) = c(F_1(x_1), F_2(x_2), \ldots, F_n(x_n))\Pi_{i=1}^n f_i(x_i)$$

where c is the joint probability density function associated with C and f_i is the univariate probability density function associated with F_i.

Sklar's Theorem asserts that any continuous multivariate distribution can be uniquely described by its marginal distributions and its copula. Copulas thus provide a general way to represent the dependency among jointly distributed random variables independently of their marginal behavior.

Copulas are invariant under strictly increasing transformations of the random variables. That is, if $\tilde{x}_1, \tilde{x}_2, \ldots, \tilde{x}_n$ are n jointly distributed random variables with copula C, and $g_1, g_2, \ldots, g_n$ are strictly increasing functions, then the random variables $g_1(\tilde{x}_1), g_2(\tilde{x}_2), \ldots, g_n(\tilde{x}_n)$ are also jointly distributed random variables with copula C. This implies that copulas capture dependency among random variables in a scale-free manner without regard to whether the dependency is linear.

Copulas can be useful in index insurance analysis because they provide a way to study the dependence among indices and between indices and losses, without regard to their marginal distributions. In particular, one is free to specify the forms of marginal distributions independently of one another and independently of the form of the copula function. For example, in building a bivariate model of an index and losses, it is possible to posit that one of the two random variables is log-normally distributed, the other is beta distributed, and the dependency between the two is captured by a Clayton copula. This flexibility is particularly useful in index insurance design, given that there is often no reason to suppose that the index and loss distributions belong to the same distributional family. The flexibility also offers the modeler the freedom to search among different copula functions to find the one that best explains the observed dependency between index and losses.

3 Parametric Families of Copulas

A number of parametric families of copulas are commonly used in statistical analysis of dependence. So-called "spherical" copulas include copulas generated by spherical multivariate distributions such as the normal and Student-t distributions. The Gaussian (e.g., normal) bivariate copula takes the form

$$C(u_1, u_2; \rho) = \Phi_\rho(\phi^{-1}(u_1), \phi^{-1}(u_2)), \qquad u_1, u_2 \in [0, 1]$$

where Φ_ρ is the cumulative distribution function of a bivariate standard normal distribution with correlation ρ and ϕ is the cumulative distribution for a univariate standard normal random variable. The Student-t bivariate copula takes the form

$$C(u_1, u_2; \rho, \nu) = \Phi_{\rho,\nu}(\phi_\nu^{-1}(u_1), \phi_\nu^{-1}(u_2)), \qquad u_1, u_2 \in [0, 1]$$

where $\Phi_{\rho,\nu}$ is the cumulative distribution function of a bivariate standard Student-t distribution with correlation ρ and ν degrees of freedom, and ϕ

Family	Parameter	$\psi_\theta(t)$
Clayton	$\theta \geq -1, \theta \neq 0$	$\frac{1}{\theta}(t^{-\theta} - 1)$
Frank	$\theta \neq 0$	$-\log\left(\frac{e^{-\theta t}-1}{e^{-\theta}-1}\right)$
Gumbel	$\theta \geq 1$	$(-\log t)^\theta$

Table 1: Archimedean Copula Generator Functions

is the cumulative distribution for a univariate standard Student-t random variable with ν degrees of freedom.

Another widely studied parametric family of copulas are the Archimedean copulas. A bivariate Archimedean copula takes the form

$$C(u_1, u_2; \theta) = \psi_\theta^{-1}(\psi_\theta(u_1) + \psi_\theta(u_2)), \qquad u_1, u_2 \in [0, 1]$$

where $\psi_\theta : [0, 1] \mapsto [0, \infty]$ is a continuous, strictly decreasing, convex function with $\psi_\theta(1) = 0$. The function ψ_θ is called the "generator function". Different generator functions give rise to different sub-families of copulas, of which the three most widely used are the Clayton, Frank, and Gumble copulas.[1] The generator functions for these sub-families of copulas are given in Table 1.

Another parametric family of copulas that have been used in financial analysis and may prove applicable to index insurance design are the "extreme-value" copulas. A bivariate extreme-value copula takes the form

$$C(u_1, u_2; A) = \exp\left(\log(u_1 u_2) A \left(\frac{\log(u_1)}{\log(u_1 u_2)}\right)\right)$$

where A, the called the "Pickands dependence function". The Pickands dependence function can be any real-valued convex function defined on the interval $[0, 1]$ such that $\max(t, 1-t) \leq A(t) \leq 1,\ t \in [0, 1]$. Different dependence functions give rise to different sub-families of copulas, of which the two most widely used are the Marshal-Olkin and Gumbel-Hougaard copulas. The dependence functions for these sub-families of copulas are given in Table 2.

One of our objectives in forthcoming research is to study copulas further in order to identify copulas that are well-suited for index insurance applications. Many copula structures in addition to those discussed here have found

[1]Nelsen [25] provides a list of 32 different Archimedean copulas

Family	Parameters	$A(t)$
Marshal-Olkin	$0 \le \alpha, \beta \le 1$	$1 - \min(\beta t, \alpha(1-t))$
Gumbel-Hougaard	$\theta > 1$	$\left(t^\theta + (1-t)^\theta\right)^{\frac{1}{\theta}}$

Table 2: Extreme-Value Copula Dependence Functions

application in the financial literature, and the list has been steadily growing in recent years. As we shall see, not all copulas are capable of capturing the complex dependence structures that may exist among proposed indices and between indices and loss distributions of interest. Which copula structures are best-suited for index insurance analysis will be an empirical question to be addressed by formal cross-specification tests based on goodness-of-fit statistics.

4 Dependence

Copulas allow an analyst to isolate, study, and model dependence among a random variables independently of their marginal distributions. But what does one mean by "dependence"? How does one measure it? And how does one best capture it for the purposes of index insurance product design and market development?

The most widely-used measure of dependence between random variables is the Pearson linear correlation coefficient. The linear correlation coefficient for two random variables $\tilde{x}$ and $\tilde{y}$ is defined by

$$\rho = \frac{\mathrm{Cov}(\tilde{x}, \tilde{y})}{\sqrt{\mathrm{Var}(\tilde{x})\mathrm{Var}(\tilde{y})}}$$

where Cov is the covariance operator and Var is the variance operator. The linear correlation coefficient is a *global* measure of *linear* dependence between two random variables. Its popularity stems from the fact that for the widely-used multivariate normal distribution, the linear correlation coefficients, together with the parameters of the univariate marginal distributions, fully characterize the joint distribution. This is true of all spherical multivariate distributions, including the Student-t distribution, but is not generally true of non-spherical multivariate distributions.

The linear correlation coefficient attempts to summarize in a single number the global linear dependence between two random variables. However, one cannot expect the linear correlation coefficient to adequately summarize complex dependencies. This is particularly true in the design and analysis of index insurance products. In index insurance design, one is primarily interested in whether an index and losses are monotonically associated, without regard to whether the association is linear. This so because indemnity schedules can be specified to be nonlinear, if necessary. Moreover, one is interested in whether an index and losses are strongly associated at the relevant extremes of the distributions, without regard to how they are associated throughout the remainder of their respective ranges. An index and losses may exhibit the desirable properties of being strongly monotonically associated at the relevant extremes, yet possess a very low linear correlation coefficient because the relationship is not strongly linear or global in scope. As such, linear correlation coefficients are inadequate for insurance insurance analysis.

More informative measures of association for the purposes of index insurance analysis are the Kendall's tau and Spearman's rho measures of association.[2] Kendall's tau and Spearman's rho both measure a form of dependence known as "concordance". Informally, a pair of random variables are concordant if "large" values of one tend to be associated with "large" values of the other and "small" values of one with "small" values of the other. Both measures are invariant under strictly increasing nonlinear transformations of the random variables, and assess how well an arbitrary monotonic function can describe the relationship between two variables without requiring the function to be linear.

Formally, Kendall's tau is defined for a pair of random variables $\tilde{x}$ and $\tilde{y}$ as

$$\tau = \mathrm{Prob}\{(\tilde{x}_1 - \tilde{x}_2)(\tilde{y}_1 - \tilde{y}_2) > 0\} - \mathrm{Prob}\{(\tilde{x}_1 - \tilde{x}_2)(\tilde{y}_1 - \tilde{y}_2) < 0\}$$

where $(\tilde{x}_1, \tilde{y}_1)$ and $(\tilde{x}_2, \tilde{y}_2)$ are independent and identically distributed as $(\tilde{x}, \tilde{y})$.

Spearman's rho for a pair of random variables $\tilde{x}$ and $\tilde{y}$ is simply the linear coefficient between their inverse cumulative distribution transforms:

$$\rho_s = \frac{\mathrm{Cov}(F_x^{-1}(\tilde{x}), F_y^{-1}(\tilde{y}))}{\sqrt{\mathrm{Var}(F_x^{-1}(\tilde{x}))\mathrm{Var}(F_y^{-1}(\tilde{y}))}}$$

where F_x^{-1} and F_y^{-1} are the inverses of the marginal cumulative probability distributions of $\tilde{x}$ and $\tilde{y}$, respectively. Kendall's tau and Spearman's rho

[2] We will adopt the convention suggested by Nelsen [25] and reserve the use of the term "correlation coefficient" to indicate the Pearson linear correlation coefficient.

Family	τ	ρ_s
Clayton	$\frac{\theta}{\theta+2}$	Complicated closed form
Frank*	$1-\frac{4}{\theta}\left(1-D_1(\theta)\right)$	$1-\frac{12}{\theta}\left(D_1(\theta)-D_2(\theta)\right)$
Gumbel	$\frac{\theta-1}{\theta}$	No closed form

$^*D_k(\theta)=\frac{k}{\theta^k}\int_0^\theta \frac{t^k}{e^t-1}dt$

Table 3: Kendall's Tau and Spearman's Rho for Selected Archimedean Copulas as a Function of the Generator Function Parameter θ.

possess certain features that are desirable for measures of association: Their values always lie between -1 and 1; they equal zero if the random variables are independent; they equal 1 if the two random variables are almost surely related by a strictly increasing function (which may, or may not be linear); and they equal -1 if the two random variables are almost surely related by a strictly decreasing function (which may, or may not be linear).

Given that Kendall's tau, Spearman's rho, and the copula of two jointly distributed random variables are invariant under arbitrary strictly increasing transformations of the underlying variables, it should not come as a surprise that both Kendall's tau and Spearman's rho are fully determined by the copula of the joint distribution. In general, if two random variables are related through a copula C, then it can be shown [25] that

$$\tau = 4\int_{[0,1]^2} C(u_1,u_2)dC(u_1,u_2)-1$$

and

$$\rho_s = 12\int_{[0,1]^2} C(u_1,u_2)du_1du_2-3.$$

It can also be shown that in the special case that the two random variables possess an Archimedean copula with generator function ϕ, then

$$\tau = 4\int_0^1 \frac{\phi(t)}{\phi'(t)}dt+1.$$

Table 3 presents Kendall's tau and Spearman's rho for the better-known Archimedean copulas, as functions of the underlying generator function parameter θ.

Given data, Kendall's tau and Spearman's rho can easily be estimated using their sample counterparts. Given a set of T joint observations (x_t, y_t), $t = 1, 2, \ldots, T$, there are $T(T-1)/2$ possible distinct pairings of these joint observations. A pair (x_t, y_t) and $(x_{t'}, y_{t'})$ are said to be concordant if $(x_t - x_{t'})(y_t - y_{t'}) > 0$ and discordant if $(x_t - x_{t'})(y_t - y_{t'}) < 0$. Kendall's tau for this sample is computed as

$$\tau = \frac{N_c - N_d}{0.5T(T-1)}$$

where N_c is the number of concordant pairs and N_d is the number of discordant pairs. Spearmans's rho for this sample is computed as

$$\hat{\rho}_s = \frac{\mathrm{Cov}(\mathrm{Rank}(x_t), \mathrm{Rank}(y_t))}{\sqrt{\mathrm{Var}(\mathrm{Rank}(x_t))\mathrm{Var}(\mathrm{Rank}(y_t))}}$$

where $\mathrm{Rank}(x_t)$ indicates the rank of the t^{th} observation of x among all T observations of x and $\mathrm{Rank}(y_t)$ indicates the rank of the t^{th} observation of y among all T observations of y.

Because index insurance indemnity schedules can be nonlinear, the ability to write an effective insurance contract for a particular loss distribution depends ultimately on whether the index and losses exhibit strong monotonic dependence, not linear dependence. As such, Kendall's tau and Spearman's rho provide measures of association that are superior to Pearson's linear correlation coefficient for assessing the viability of index insurance contracts. However, both Kendall's tau and Spearman's tau remain global measures of association that are unable to capture variations in the degrees of dependence throughout the range of the joint distribution, including asymmetries at the extremes of the distributions.

Efforts to explain asymmetries in dependence at extremes have lead to the introduction of the notion of *tail dependence*([25]; [19]; [5]). Tail dependence generally refers to the degree to which two random variables are related at the lower or higher extremes of their ranges. Two common measures of tail dependence are the asymptotic coefficients of lower and upper tail dependence. The coefficient of lower tail dependence for a pair of random variables λ_L is the limiting probability that one variable takes on a very low value, given that the other takes on a very low value. Similarly, the coefficient of upper tail dependence for a pair of random variables λ_U is the limiting probability that one variable takes on a very high value, given that the other takes on a very high value. Formally, for two random variables $\tilde{x}_1$ and $\tilde{x}_2$,

$$\lambda_L = \lim_{u \to 0} \Pr\{\tilde{x}_2 \leq F_2^{-1}(u) | \tilde{x}_1 \leq F_1^{-1}(u)\}$$

$$\lambda_U = \lim_{u \to 1} \Pr\{\tilde{x}_2 \geq F_2^{-1}(u) | \tilde{x}_1 \geq F_1^{-1}(u)\}$$

where F_1 and F_2 are the marginal cumulative distributions of the two random variables, respectively. The asymptotic coefficients of tail dependence for a joint distribution can be recovered from the copula of the joint distribution through

$$\lambda_L = \lim_{u \to 0} \frac{C(u,u)}{u}$$

$$\lambda_U = \lim_{u \to 1} \frac{2u - 1 + C(1-u, 1-u)}{u}.$$

Asymptotic coefficients of tail dependency must be carefully interpreted. A zero coefficient does not indicate that two random variables are independent over tails with positive probability. Rather, a zero coefficient indicates that the interdependence is weak at the tail and disappears in the limit as the tail probability goes to zero. For example, the upper and lower asymptotic coefficients of tail dependence for two jointly normally distributed random variables are zero, regardless of their Pearson linear correlation coefficient. The asymptotic coefficients of tail dependence for two jointly Student-t distributed random variables are positive, but symmetric, indicating that they exhibit the same dependence at the lower and upper tails.

The asymptotic coefficients of upper and lower tail dependence for standard Archimedean copulas as functions of the generator function parameter θ are presented in Table 4. The copulas differ markedly in their tail dependence. The Clayton copula exhibits lower tail dependence, but not upper tail dependence. The Gumbel copula exhibits upper tail dependence, but not lower tail dependence. The Frank copula exhibits neither upper nor lower tail dependence.

The nature of tail dependence exhibited by the Clayton and Gumbel copulas are illustrated in Figures 1 and 2. Figure 1 plots the contours of the joint probability density function of two marginally standard normal random variables related by a Clayton copula with generator function parameter $\theta = 3$ and a Gumbel copula with generator function parameter $\theta = 2.5$. As can be seen in the figure, with the Clayton copula, probability mass is more concentrated around the diagonal at the lower tail than at the upper tail, indicating lower tail dependence. Conversely, with the Gumbel copula, probability mass is more concentrated around the diagonal at the upper tail than at the lower tail, indicating lower upper dependence. Figure 2 plots simulated scatter diagrams generated by the two distributions.

Family	λ_L	λ_U
Clayton	$2^{-\frac{1}{\theta}}$	0
Frank	0	0
Gumbel	0	$2-2^{\frac{1}{\theta}}$

Table 4: Lower and Upper Asymptotic Tail Dependence for Archimedean Copulas

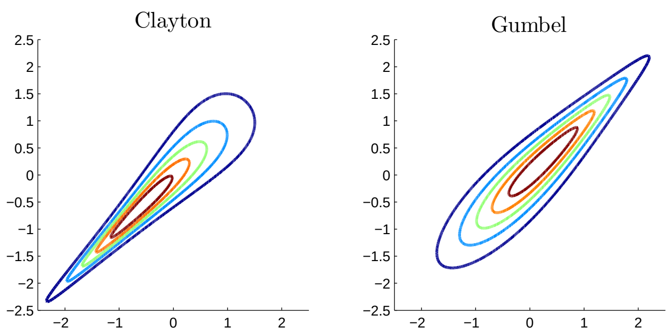

Figure 1: Probability Density Contour Plots of Two Archimedean Copulas

Whether a proposed index and losses of interest exhibit tail dependence is of fundamental interest in index insurance design. Suppose, for example, that one is interested in designing a rainfall index insurance product to address losses in income suffered by farmers due to drought. Such a contract would indemnify the insured for low values of rainfall. As such, the strength of the lower tail dependence exhibited by rainfall and income would be a key indicator of the viability of such a contract. If rainfall and income exhibit strong lower tail dependence, as would be the case if their dependence structure was similar to that of a Clayton copula, then it should be possible to design an index insurance contract, possibly with a nonlinear indemnity schedule, that carries acceptably low basis risk, even if the linear correlation coefficient between the two variables is low. Conversely, if rainfall and income exhibit no or low lower tail dependence, as would be the case if their dependence structure was similar to that of a Gumbel copula, then it would be difficult to structure an index insurance contract that would be useful to farmers as a risk management instrument.

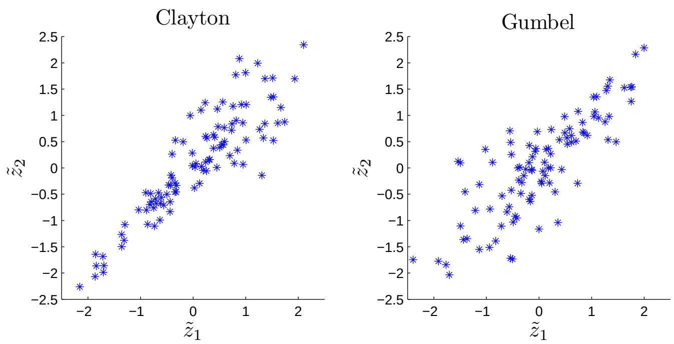

Figure 2: Simulated Scatter Plots of Two Archimedean Copulas

The nature of tail dependence among a series of indices is also of fundamental interest in rating index insurance products. Suppose, for example, that an insurer proposes to offer a variety of rainfall insurance contracts written for rainfalls at different locations within a defined geographical area. If these indices exhibit lower tail dependence, then they are more highly correlated for extreme low values than for mid-range or extreme high values. This would imply that the portfolio of the insurer's indemnity liabilities is risker, and the insurer's maximum probable losses would be higher, than would be indicated by conventional portfolio models based on Gaussian multivariate models.

5 Estimation of Copulas

We now discuss how to estimate parameters of copulas empirically. To this end, suppose we have a series of paired observations (x_{1t}, x_{2t}), $t = 1, 2, \ldots, T$, that we believe are independent realizations of a data generating process characterized by a joint cumulative probability function F that can be decomposed as

$$F(x_1, x_2; \theta, \beta_1, \beta_2) = C(F_1(x_1; \beta_1), F_2(x_2; \beta_2); \theta)$$

where β_i is a vector of parameters for the marginal distribution F_i of $\tilde{x}_i$, and θ is a dependence parameter of the copula C.

Various estimation procedures based on maximum likelihood principles are available to estimate copulas. First, if one is not sure about the form of the marginal distributions, it is possible to estimate the copula's dependence parameter using the method of pseudo-maximum likelihood suggested by Genest and Favre [15]:

$$\hat{\theta} = \text{argmax}_{\theta} \sum_{t=1}^{T} \log c(u_{1t}, u_{2t}; \theta)$$

where $u_{it} = \text{index}(x_{it})/(T+1)$ and c is the density function of the copula C.

If one is willing to specify the form of the marginal distributions, it is possible to estimate the dependence parameter of the copula and the parameters of the marginal distributions jointly via the method of full-information maximum likelihood:

$$(\hat{\theta}, \hat{\beta}_1, \hat{\beta}_2) = \text{argmax}_{\theta,\beta_1,\beta_2} \sum_{t=1}^{T} \log f(x_{1t}, x_{2t}; \theta, \beta_1, \beta_2)$$

where

$$f(x_1, x_2; \theta, \beta_1, \beta_2) = c(F_1(x_1; \beta_1), F_2(x_2; \beta_2); \theta) f_1(x_1; \beta_1) f_2(x_2; \beta_2)$$

Here, f is the marginal probability density function of the joint distribution and f_i is the probability density function associated with the cumulative distribution F_i of $\tilde{x}_i$. Full information maximum likelihood, however, can be computationally burdensome and is rarely used in practice.

The parameters of the copula and the parameters of the marginals can also be estimated via limited information maximum likelihood using a two step procedure. This procedure calls for the parameters of the marginals to be estimated independently first

$$\hat{\beta}_i = \text{argmax}_{\beta_i} \sum_{t=1}^{T} \log f_i(x_{it}; \beta_i)$$

then using the estimates thus derived to estimate the parameter of the copula via conditional maximum likelihood

$$\hat{\theta} = \text{argmax}_{\theta} \sum_{t=1}^{T} \log f(x_{1t}, x_{2t}; \theta, \hat{\beta}_1, \hat{\beta}_2)$$

or, equivalently,

$$\hat{\theta} = \text{argmax}_{\theta} \sum_{t=1}^{T} \log c(F_1(x_{1t}; \hat{\beta}_1), F_2(x_{2t}; \hat{\beta}_2); \theta)$$

This method is sometimes referred to in the literature as the "inference function for marginals" method ([19]; [21]).

Method-of-moments methods based on sample estimates of Kendall's tau are also available ([8]). To estimate the copula parameter, one could compute the sample Kendall's tau and invert the relation $\theta \mapsto \tau(\theta)$ given in Table 3 to recover an estimate of the copula parameter θ. Special methods for estimating copulas are further discussed in [13], [16], and [17].

6 Selection of Copulas

Another important task in working with copulas in empirical applications is to choose among various candidate copula types. Goodness-of-fit tests assess how well a statistical model fits a set of observations. The most widely used goodness-of-fit tests are the Pearson's chi-square and Kolmogorov-Smirnov tests.

The Pearson's chi-square goodness-of-fit test establishes whether a set of observations are realizations of a data generating process with cumulative distribution function F. To conduct the test, the range of the data generating process is divided into n "bins". One then computes

$$\chi^2 = \sum_{i=1}^{n} \frac{(O_i - E_i)^2}{E_i}$$

where O_i is the proportion of observations that lie in the i^{th} bin and E_i is the proportion of observations expected to lie in the i^{th} bin under the null hypothesis that the data are generated by a process with cumulative distribution function F. Under the null hypothesis, the statistic χ^2 is asymptotically distributed chi-square with $n-1$ degrees of freedom.

The Kolmogorov-Smirnov test is based on the statistic:

$$D_n = \sup_x |O(x) - F(x)|$$

where $O(x) = \frac{1}{n} 1(x_i \leq x)$ represents the empirical cumulative distribution function and $1(x_i \leq x)$ is the indicator function. Under the null hypothesis that the data are generated by a process with cumulative distribution function F, $\sqrt{n} D_n$ converges asymptotically to the Kolmogorov distribution.

The Akaike Information Criterion (AIC) may be used to chose among alternative model specifications. In general, the AIC is computed as:

$$\text{AIC} = 2k - 2\log(L)$$

where k is the number of parameters in the statistical model, and L is the maximized value of the likelihood function. Given a data set, several competing models may be compared according to their AIC. The one with the lowest AIC would be regarded as the best.

Another method for selecting among copulas is attributable to Li. Let

$$f(x,y) = \frac{1}{n}\sum_{i=1}^{n} K(\frac{x-x_i}{h_x})K(\frac{y-y_i}{h_y})$$

denote the kernel estimator based on a set of n paired observations (x_i, y_i) of two random variables. Also, let

$$g_j(x,y;\theta_j) = c_j(\hat{F}_X(x), \hat{F}_Y(y); \theta_j)\hat{f}_X(x)\hat{f}_Y(y)$$

denote fitted probability distributions based on distinct copulas c_1 and c_2. Then the integrated square difference for the two copulas

$$I_j = \int\int (f(x,y) - g_j(x,y;\theta_j))^2 dxdy.$$

may be used to discriminate among candidate copulas, with the one with the lowest difference providing the best fit.

Various other goodness-of-fit tests have been proposed. For example, Kolde, Koeidjk, and Verbek [23] suggests tests based on variants of the Kolmogorov-Smirnov and Anderson-Darling statistics.

7 Spatial Contagion Among Weather Indices: The Case of Iowa Rainfall

One of the questions I was asked to address for this State of Knowledge Report is whether spatially separated weather variables commonly used in index insurance design, such as rainfall at different weather stations within a defined geographical area, are “more highly correlated at the tails”. There are a number of papers that have appeared in the finance literature that have addressed this question in the context of financial asset markets, where the phenomenon of high correlation at the tails is refereed to as “spatial contagion” ([2], [3], [4], [7]). Spatial contagion in the context of index insurance, however, has not been, to our knowledge, addressed in the agricultural economics literature.

Spatial contagion is an important question in the design of index insurance products for two reasons. First, suppose an insurer offers a range of index insurance contracts written on weather variables, say, rainfalls, at different locations in a defined geographical area. The insurer will be interested in assessing the distribution of payouts of his entire portfolio of index insurance contracts in order to calculate the maximum probable loss associated with his entire book of business. If the underlying weather variables exhibit spatial contagion, then standard portfolio risk assessment methods based explicitly or implicitly on normal distribution theory could result in serious underestimates of the riskiness of the portfolio, leaving the insurer exposed to greater business risk than he realizes.

Second, an important task in index insurance design is to compute the expected indemnity associated with a given indemnity schedule. Indemnities, however, are paid only when the index falls below a certain threshold, an event that occurs only infrequently. As such, the data available to support the calculation of this critical statistic is usually very limited. One way to address the paucity of data is to estimate the expected indemnities of multiple contracts jointly. This should lead to gains in efficiency that will depend primarily on the degree of dependence exhibited at the critical extremes of the underlying index distributions. In other words, in the presence of spatial contagion, it may be possible to achieve substantial gains in efficiency by jointly estimating the expected indemnities of various contracts, provided the tail dependencies are faithfully captured.

For our first case study, we propose assess the degree of spatial contagion exhibited by Iowa June and July county-level rainfalls, employing data for all 99 Iowa counties from 1954-2008 obtained from National Climatic Data Center (NCDC). For our second case study, we propose assess the degree of spatial contagion exhibited by growing seasons rainfalls in Northwest Peru. We will employ a variety of methods, including a test proposed by Nelsen [25], structural econometric estimation, and copula function estimation.

7.1 Kendall's Tau Test

Suppose we have a series of paired observations (x_{1t}, x_{2t}) on two random variables $\tilde{x}_1$ and $\tilde{x}_2$. Whether the two random variables are more highly correlated at the lower tails of their distributions than at the upper tail of their distributions may be tested using a method developed by Nelsen [25] that compares Kedall's correlation coefficient computed for subsamples of lower and upper tail observations:

- sort the paired observations (x_{1t}, x_{2t}) according to the values x_{1t};
- delete the $n/4$ observations in the middle, partitioning the observations into two sets of equal size, one containing the lower ranked observations

and one containing the higher ranked observations;

- compute the Kendall correlation coefficients for each subset of observations, τ_L and τ_U, and set $D = \tau_U - \tau_L$.

For large samples, D is normally distributed, allowing us to test the hypothesis $H_0 : \tau_U = \tau_L$ against the alternative $H_A : \tau_U < \tau_L$ using a one-tailed test. Due to the limited amount of data, the standard error of the estimate will be computed using bootstrapping methods.

We will test the hypothesis of asymmetric lower and upper tail correlations among Iowa rainfalls following one of two procedures. Either we will perform this test for all 4851 possible pairs of counties in Iowa, or we will limit the test to counties and their contiguous neighbors. Regardless of the method chosen, I will report the number of tests that result in rejection of the null hypothesis.

7.2 Structural Spatial Econometric Analysis

A second approach to testing for spacial contagion will employ more conventional structural spatial econometric estimation methods. The approach would begin by positing that rainfall y_{it} in county i in year t are generated as

$$y_{it} = \alpha_i + \beta_i z_t + \epsilon_{it}$$

where the ϵ_{it} are serially independent zero-mean normal variates and z_t is a systemic factor that affects rainfalls in all counties, for which Iowa state-wide average rainfall will be used as a proxy.

We will posit that the variance-covariance structure of the idiosyncratic error terms ϵ_{it}, conditional on the the systemic factor z_t, are given by

$$Var(\epsilon_{it}) = \sigma_i^2 z_t$$

and

$$Corr(\epsilon_{it}, \epsilon_{jt}) = \rho(d(i,j), z_t)$$

where $d(i,j)$ is the distance between county i and county j. Higher correlation at the lower tail of the distribution would be indicated if ρ is decreasing in z_t.

7.3 Copula Analysis

A third approach would employ copulas. Specifically, I will search among various candidate copulas and, using goodness of fitness tests, attempt to identify the copula structures that best explain the nature of dependence between adjacent county rainfall series. If tail symmetry can be rejected in favor of lower tail dependence in most county pairs, spatial contagion will be said to exist among Iowa county-level rainfalls. Copulas to be tested include the Gaussian, Student-t, Clayton, Gumbel, Frank, extreme-value, and possibly others. Vine copulas, pair coputals, and hierarchial Archimedean copulas will also be examined ([20]; [27]; [26]).

We will also attempt to fit some candidate mixture distributions. For example, the Clayton copula exhibits strong lower tail dependence and the Gumbel copula exhibits strong upper tail dependence. Neither is, on its own, exhibits the flexibility to capture varying relative degrees of upper and lower tail dependence. Thus, we propose to estimate a mixture of the Clayton and Gumbel copulas.

$$C(u_1, u_2; \mu, \theta_1, \theta_2) = \mu C_1(u_1, u_2; \theta_1) + (1 - \mu)C_2(u_1, u_2; \theta_1)$$

where C_1 indicates the Clayton copula with dependence parameter θ_1 and C_2 indicates the Gumbel copula with dependence parameter θ_2.

8 Copulas in Index Insurance Product Design

Our primary objective will be to explore effective uses of copula methods for designing and analyzing index insurance products. In particular, we are interested in developing a protocol for designing and analyzing index insurance products that is scalable and adaptable to a wide variety of index insurance design settings. We intend to produce suitable computer code, written in Matlab, that will perform the necessary computations for arbitrary index and loss data series. We will illustrate the use of these methods in two settings: rainfall insurance for Henan Province, China and rainfall insurance for the Department of Piura, Peru.

Essential actuarial computations required in the design of index insurance products include estimating the expected indemnity (i.e., fair premium rate) associated with a hypothetical index insurance contract, computing an approximate standard error surrounding the expected indemnity estimate, and measuring basis risk.

We propose the following general procedure. Formally, let $\tilde{x}$ denote a given index, call it “rainfall”, and let $\tilde{y}$ denote an indicator of income, call it "yields". We are given a hypothetical indemnity schedule $f(x;\alpha)$ that specifies the indemnity to be paid in terms of the observed index x and a vector of contract parameters α. Given T observations on rainfall x_t and yields y_t, we construct a model of the joint distribution of $\tilde{x}$ and $\tilde{y}$ using a selected parametric copula family. Given the fitted distribution, we compute an estimate of the expected indemnity

$$\pi(\alpha) = Ef(\tilde{x};\alpha)$$

and then compute an estimate of its standard error using bootstrapping principles. We the posit a utility of wealth function u, and compute willingness to pay $\omega(\alpha)$ by numerically solving the nonlinear inequality

$$Eu(\tilde{y}) = Eu(\tilde{y} + f(\tilde{x},\alpha) - \omega).$$

Given a procedure to estimate the fair premium and the willingness to pay, we search for an optimal design, that is, for an optimal value of the contract parameters α, by maximizing willingness to pay $\omega(\alpha)$ subject to constraint $\pi(\alpha) = \pi^*$ where π^* is a target fair premium level, say 10%.

The approach above can be repeated under different assumptions regarding the underlying dependency between rainfall and yields. In particular, one can perform the procedure using conventional assumptions of symmetric tail dependence and then allowing for varying degrees of tail dependence. A hypothesis to be tested is that conventional approaches lead to significant underestimates of the value of index insurance due to the incidence of substantial lower tail dependency.

References

[1] Andrew Ang and Joseph Chen. Asymmetric correlations of equity portfolios. *Journal of Financial Economics*, 63:443–494, 2002.

[2] Brendan O. Bradley and Murad S. Taqqu. Framework for analyzing spatial contagion between financial markets. *Finance Letters*, 2(6):8–15, 2004.

[3] Brendan O. Bradley and Murad S. Taqqu. Empirical evidence on spatial contagion between financial markets. *Finance Letters*, 3(1):77–86, 2005.

[4] Brendan O. Bradley and Murad S. Taqqu. How to estimate spatial contagion between financial markets. *Finance Letters*, 3(1):64–76, 2005.

[5] Ana C. Cebrian, Michel Denuit, and Philippe Lambert. Analysis of bivariate tail dependence using extreme value copulas: An application to the soa medical large claims database. *Belgian Actuarial Bulletin*, 3(1):33–41, 2003.

[6] Kevin Dowd. Copulas in macroeconomics. *Journal of International and Global Economic Studies*, 1(1):1–26, 2008.

[7] Fabrizio Durante and Piotr Jaworski. Spatial contagion between financial markets: A copula-based approach. *Applied Stochastic Models in Business and Industry*, 2010 (forthcoming).

[8] John H.J. Einmahl, Andrea Krajina, and Johan Segers. A method of moments estimator of tail dependence. *Bernoulli*, 14(4):1003–1026, 2008.

[9] Paul Embrechts. Extremes in economics and the economics of extremes. In Barbel Finkenstadt and Holger Rootzen, editors, *Extreme Values in Finance, Telecommunications, and the Environment*. Chapman and Hall, 2001.

[10] Paul Embrechts, Filip Lindskog, and Alexander McNeil. Modelling dependence with copulas and applications to risk management, 2001.

[11] Paul Embrechts, Alexander McNeil, and Daniel Straumann. Correlation and dependence in risk management: Properties and pitfalls, 1999.

[12] Dean Fantazzini. Dynamic copula modelling for value at risk. *Frontiers in Finance and Economics*, 5(2):72–108, 2008.

[13] Gabriel Frahm, Markus Junker, and Rafael Schmidt. Estimating the tail-dependence coefficient: Properties and pitfalls. *Insurance: Mathematics and Economics*, 37(1):80–100, 2005.

[14] Edward W. Frees and Emiliano A. Valdez. Understanding relationships using copulas. *North American Actuarial Journal*, 2(1):1–25, 1998.

[15] Christian Genest and Anne-Catherine Favre. Everything you always wanted to know about copula modeling but were afraid to ask. *Journal of Hydrolic Engineering*, 12(4):437–368, 2007.

[16] Christian Genest and Johan Segers. Rank-based inference for bivariate extreme value copulas. *The Annals of Statistics*, 37(5B):2990–3022, 2009.

[17] Gordon Gudendorf and Johan Segers. Nonparametric estimation of an extreme-value copula in arbitrary dimensions, 2009.

[18] Luiz Koodi Hotta, Edimilson C. Lucas, and Helder P. Palaro. Estimation of var using copula and extreme value theory. *Multinational Finance Journal*, 12(3/4):205–218, 2008.

[19] Harry Joe. *Multivariate models and dependence concepts*. Chapman and Hall, London, U.K., 1997.

[20] Harry Joe, Haijun Li, and Aristidis K Nikoloulopoulos. Tail dependence functions and vine copulas. Technical Report 2008-3, Department of Mathematics Washington State University, 2008.

[21] Harry Joe and James J. Xu. The estimation method of inference functions for margins for multivariate models. Technical Report 166, Department of Statistics, University of British Columbia, 1996.

[22] J.-F. Jouanin, G. Rapuch, G. Riboulet, and T. Roncalliy. Modelling dependence for credit derivatives with copulas, 2001.

[23] Erik Kole, Kees Koedijk, and Marno Verbeek. Selecting copulas for risk management. *Journal of Banking and Finance*, 31:2405–2423, 2007.

[24] Nikolai Kolev, Ulisses dos Anjos, and Beatriz Vaz de M. Mendes. Copulas: A review and recent developments. *Stochastic Models*, 22(4):617–660, 2006.

[25] Roger B. Nelsen. *An introduction to copulas*. Springer-Verlag, New York, NY, 2nd edition, 2006.

[26] Ostap Okhrin, Yarema Okhrin, and Wolfgang Schmid. Properties of hierarchical archimedean copulas. Technical report, Humboldt-Universitt, Berlin, Germany, 2009.

[27] Doris Schirmacher and Ernesto Schirmacher. Multivariate dependence modeling using pair-copulas. Technical report, Society of Actuaries, 2008.

[28] Denitsa Stefanova. Dependence modeling of joint extremes via copulas: A dynamic portfolio allocation perspective, 2007.

[29] Pravin K. Trivedi and David M. Zimmer. Copula modeling: An introduction for practitioners. *Foundations and Trends in Econometrics*, 1(1):1–111, 2007.

[30] Jun Yan. Enjoy the joy of copulas: With a package copula. *Journal of Statistical Software*, 21(4):1–21, 2007.

Appendix C Climate Indices and Global Teleconnections: Review of State of Knowledge and Potential Utility for Index Insurance

Upmanu Lall[1]

[1] Upmanu Lall is the Alan and Carol Silberstein Professor of Engineering in the department of earth and environmental engineering at Columbia University.

Climate Indices and Global Teleconnections: Review of State of Knowledge and Potential Utility for Index Insurance

Abstract

Index insurance is a potentially useful risk transfer mechanism, at least at a regional aggregation level where basis risk associated with individual exposure is less of an issue. In this setting, climate extremes are often determined by extremes in atmospheric circulation attributes that lead to the advection or blockage of moisture to the region of interest, or to persistent conditions that allow extreme hot or cold temperatures to develop. Even extreme wind anomalies are often related to changes in large scale atmospheric circulation. Atmospheric circulation anomalies are in turn driven by identifiable spatio-temporal patterns of surface temperature. Given their relatively slower evolution, Sea Surface temperatures (SSTs) are often considered the "carrier" of the climate information. Thus, an entire set of regional weather (e.g., precipitation, temperature, wind) extremes could in concept be related to specific SST patterns. Consequently, some attention has been focused on understanding how seasonal and inter-annual climate evolution may be related to specific SST or atmospheric circulation indices. The current state of knowledge of these teleconnections at a global scale is reviewed here with the perspective of the potential use of some of the well known climate indices as a concurrent or as a predictive surrogate index for regional weather. The ability to assess these indices, their predictability, and the predictability of the associated weather extreme in a region, using publicly available information that cannot be easily manipulated by the intended beneficiaries are considered. Given that climate data quality and availability are highly variable across the globe, the readily available analyses are either specific to certain regions (e.g., the USA) where the investigators have access to high resolution and long data records, or are directed at commonly available global climate data parameters that are often established by spatial interpolation across a sparse data set. Here, we provide only an overview of analyses that have been conducted globally, since a comprehensive review across local /regional literature requires considerable investment of time. Even though many of the regional analyses have been published in peer reviewed journals, there is a need to evaluate the quality of the analysis as well as the specific assessment. Given the context provided by this review, a perspective on analyses that can be conducted at a regional level to develop suitable climate indices is presented.

Introduction and overview

This note explores whether and how a routinely available climate index can be used for hemispheric to global or even as a regional index for the purposes of weather index insurance. The basic idea is that there are several well established, routinely updated and reported indices of selected climate parameters that may inform weather extremes over large land areas, at least in terms of seasonal probabilities of exceedance, and perhaps even with respect to event attributes. The index could be a) concurrent, or for the season in which the weather extremes are a concern, or b) predictive, i.e., computed prior to the beginning of the season of concern. Ideally, the index would be well correlated with the variable (e.g., precipitation or temperature, as represented by a drought or flood, or cold or heat waves) of regional interest. Since the index may be based on a remote teleconnections, prior analyses would establish how the index is to be quantitatively mapped to specific outcomes in different regions, and the attendant uncertainty, for each type of weather extreme or impact. We recognize that the use of an index that applies to many regions provides unique opportunities for the design of index insurance products and risk reduction for the insurance provider. However, broad based indices may expose the purchaser of index insurance to high basis risk, since their outcomes may or may not map well to the large scale index. Consequently, we consider regional banks or loan providers or governments or other aggregators as the target beneficiaries of the proposed insurance, rather than individual farmers. With this context in mind, this section briefly discusses the potential indices and summarizes some key recommendations. A more detailed discussion of the correlation of the indices with potential regional outcomes, and of related index design issues is provided in the next section.

At approximately the beginning of the 20^{th} century, it was already recognized that global teleconnections or spatial relationships in atmospheric pressure existed. Often these were persistent and anomalous patterns that led to drought or floods in specific regions, and recurred with almost the same frequency. The modern discovery of the El Nino Southern Oscillation (ENSO) attributed to Sir Gilbert Walker, a British meteorologist in India is a prime example of such phenomena. During the El Nino phase of the oscillation, the central to eastern equatorial Pacific Ocean is anomalously warm, with a coincident shift in the location of the low pressure center, and of tropical convection. Sir Walker established that such conditions translated into drought over India, while the opposite side of the oscillation, La Nina corresponding to anomalously warm conditions in the Western Pacific, led to floods in India. His work was also able to establish some broader global connections with the ENSO phenomena. Today, the understanding of this phenomenon has been advanced significantly: the way in which these conditions evolve and recur, as well as the nature of its global teleconnections to precipitation and temperature variations has been established to a degree. Even so, over all land areas this phenomena explains only about 20% of the inter-annual variations in rainfall and temperature if a linear relationship between key indices and the observed precipitation is used (Mason and Goddard, 2001). However, most of these “climate shifts” are associated with extreme conditions at the locations impacted. Thus, climate indices associated with ENSO may be potentially useful for index insurance. The understanding of ENSO impacts and their prediction continues to be an open research area. This is due in part to the complexity of global climate dynamics, which is manifest through the nonlinear interaction of a variety of phenomena at multiple time scales. The understanding of the physics of these phenomena and their interactions continues to be incomplete. Thus, when we hear that no two ENSO events are alike, either in their duration, intensity or predictability or in their subsequent impact, implicitly we are told that the interactions with other climate processes may lead to differences in the global manifestation of what is characterized as an El Nino or La Nina event.

Some of the commonly identified climate indices that have at least hemispheric teleconnections, and their characteristic recurrence intervals or frequency of occurrence, predictability and a brief description are presented in Table 1 and Figure 1. A general review article on climate teleconnections is presented by Nigam (2003). Given the relatively short climate records that are available, by and large only linear analyses of the individual connections of these indices with global precipitation and temperature data fields have been pursued. The ENSO phenomenon is the most studied, and its modulation by the PDO has been discussed. Typically, the longer the average recurrence interval of the phenomenon, the more likely that it modulates the effect or occurrence of a phenomenon that occurs more often. Also, the longer the time scale, the longer the climate shift and the associated change in the nature of the climate impacts. However, there is some evidence that these shifts may translate into an increased frequency of specific climate impacts (drought vs flood, hot vs cold spells), but may have a limited contribution to the direct, inter-annual variance explained for precipitation and temperature in a given region.

In the context of index insurance for precipitation, temperature or wind, i.e., weather variables, that have a time scale of evolution of hours to days to months, given their time scale of variation, it is not clear that these indices will connect directly to event intensity and duration attributes. The situation is further compounded by the fact that the temporal variation of these indices and the associated spatial expression of their teleconnections are typically modulated by indices that are characteristically lower frequency. This implies that the independent effects of each index on the weather extremes of interest are even harder to isolate, and that even the combined effect will likely be nonlinear. As a matter of fact, it is unlikely that the direct linear correlation of any of these indices with regional weather extremes will be particularly high. Most of the analyses presented in the literature consider a correlation of seasonal precipitation and temperature with the index, or consider the composite "average" of the regional conditions for an extreme end of the oscillation (e.g., El Nino or La Nina event), or consider probability of exceedance of a specific percentile of regional rainfall or temperature conditional on the exceedance of some percentile of the climate index. Examples of these analyses are presented in the next section. Unfortunately, these are only indicative of how well a climate index may relate to regional weather, but are not directly informative as to how well an insurance index using large scale predictors may be designed for the region.

Table 1. Commonly used Climate Indices that at least have hemispherical structure or impact

Phenomenon	Indices	Recurrence Interval	Predictability	Description /Notes
El Nino Southern Oscillation (ENSO)	NINO1, NINO1.2,NINO2, NINO3, NINO3.4, MEI, BEST, CTI, SOI. All NINOx.x indices and the CTI are defined using average SST values over specific boxes in the equatorial Pacific Ocean, the SOI is an atmospheric pressure index, the MEI and BEST are derived from both SST and sea level pressure data in the region	3 to 7 years	Statistical models offer some predictability at lead times of 1 to 24 months, while General Circulation Models (coupled Ocean-Atmosphere) offer predictability up to 1 year or more. A spring barrier (April) to prediction is cited by many. Forecasts made before April often deteriorate for the period following April. Predictability may increase again for forecasts made after April. Generally predictability decreases with lead time. Teleconnections in many regions are predictable.	See saw of SST and atmospheric pressure between western and central/eastern Pacific. El Nino and La Nina events persist typically for a year or so. La Nina typically follows El Nino. Impacts vary by season and location. Changes storm track in the sub-tropics and mid-latitudes. Changes in location of convection in equatorial and tropical regions, impact rainfall and temperature in those regions. Hurricane birth landfall probabilities may be influenced
North Atlantic Oscillation (NAO)	NAO (2 indices) computed as the difference in atmospheric pressure over the North West Atlantic (e.g., Iceland) and East Central Atlantic (e.g., Azores or Gibraltar)	8 to 12 years	No demonstrated predictability of the index, despite some claims. Identified predictability of impacts especially in Northern Europe, Central Asia and the Middle East, and North Eastern Americas.	See saw of atmospheric pressure between South Central /Eastern Atlantic and North Western Atlantic, especially prominent in Dec-Feb. Significant winter and spring/fall impacts identified in some regions. Possible interactions with ENSO and influence on hurricane probabilities.
Pacific Decadal Oscillation (PDO/IPO)	PDO computed using a Principal Component Analysis of atmospheric pressure data over the Northern Pacific, and removing the effects of ENSO. The Pacific North American Index (PNA) is a real time index that is related to the PDO in terms of the spatial hemispheric expression	16 to 22 years	No demonstrated predictability of the index, despite some claims. Identified predictability of impacts especially in North Asia and North America. Evidence of modulation of ENSO impacts. Predictability of impacts high particularly when modeled in conjunction with ENSO.	This pattern is derived from statistical analysis of a spatial field, rather than directly from a simple combination of observed values. Some have argued that the PDO and the NAO are part of a common annular mode of N. Hemispheric circulation. The PDO is defined in the N. Hemisphere and a similar mode called the Interdecadal Pacific Oscillation (IPO) is also defined. Both typically influence the mid to high latitude climate.
Atlantic Multidecadal Oscillation (AMO)	AMO index defined by detrending N. Atlantic SSTs and then averaging the SST over 0 to 66 N in the Atlantic	60-80 years	No known predictability, though some statistical models to detect the shift in the AMO phase have been postulated and tested. Demonstrated changes in N hemisphere precipitation and temperature, especially in the region ranging from N E Brazil to corresponding regions northwards in Africa, Eurasia and the Americas	There is some controversy as to whether or not this is an oscillation or random switching of climate regimes. Modulation of other modes is also noted, and there is also a bi-decadal rhythm which may be related to NAO/PDO

Figure 1. Time series of indices (using data from http://climexp.knmi.nl/). Left pane =monthly values. Right pane moving averages (2 years for ENSO, 5 years for all others), 95% confidence interval for the mean. Note the different characteristic time scales of variation.

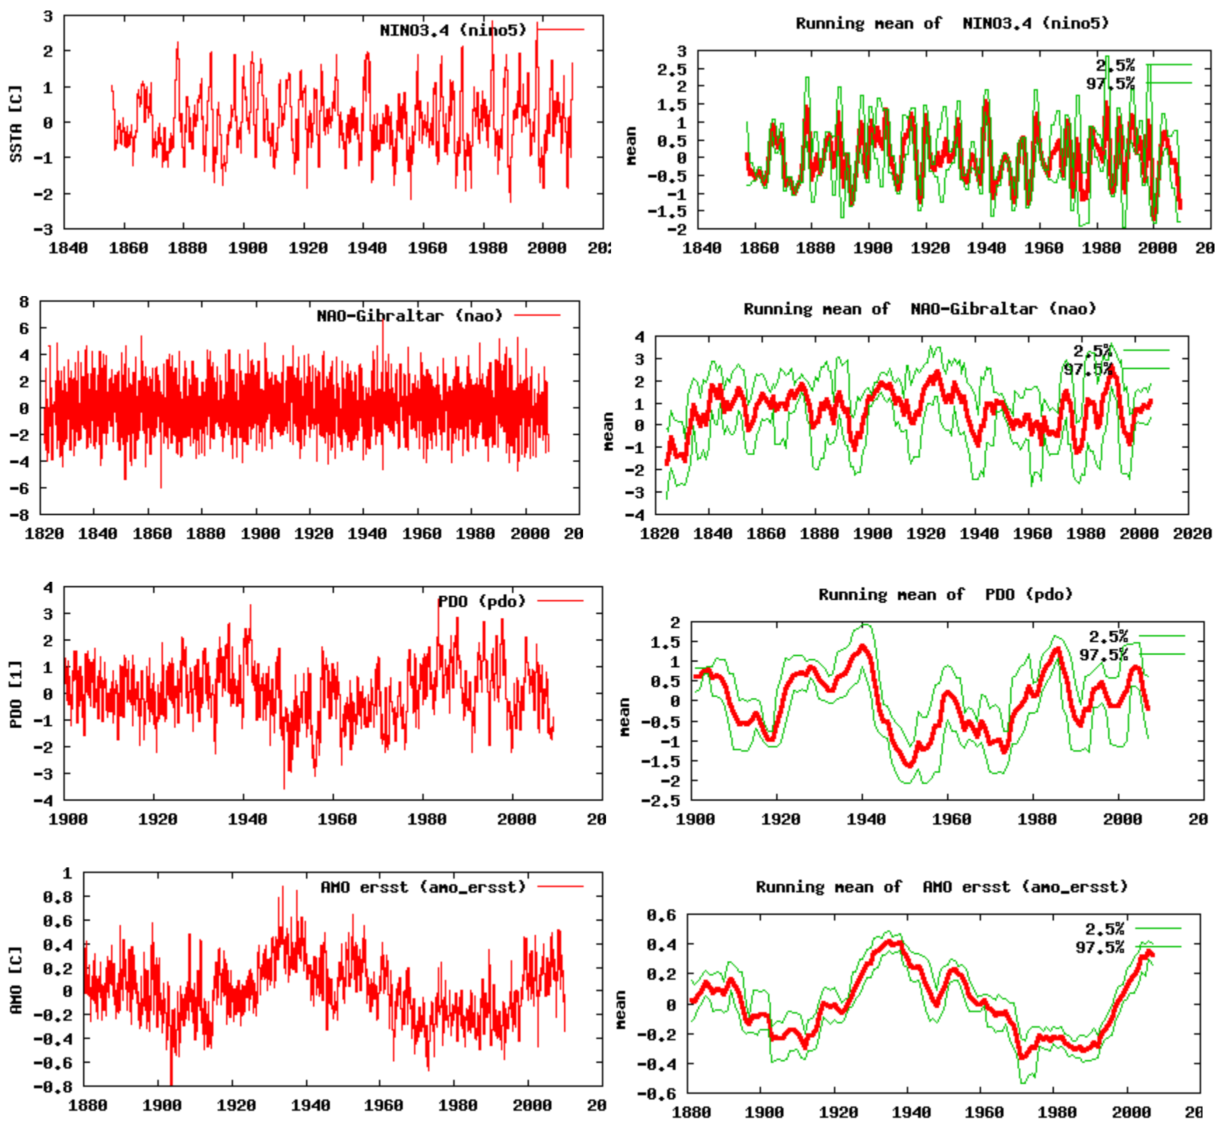

For index insurance applications the following design issues are highlighted by the properties of these indices:

1. The probability distribution of regional weather extremes for a given season is likely to change systematically and slowly over time. For a region influenced by ENSO/NAO/PDO/AMO, depending on the phase of the oscillation, conditions may be abnormally wet/dry or warm/cold, for periods that on average are about 3/4/10/30 years long respectively, if the independent effects of each mechanism on regional weather are considered. The situation is more complex if the interactive effects are considered. The implication for index insurance are that :

a. Long records are essential for assessing the sensitivity of regional weather to these long period oscillations. Further, the longer the period of the oscillation that influences regional weather, the higher the probability that there will be runs of years over which either there are no payouts from an insurance contract, or of years when there are significant payouts.
b. Some investigators may have success modeling the regional weather extremes and clustered payout processes using long memory models or heavy tailed probability distributions. Practically this means that in a system influenced by such climate phenomenon, there is a much higher than expected potential for extreme, persistent weather events that could be catastrophic. These can be thought to result from the potential for the superposition of like effects that may result from the individual climate modes. It is important that such factors be assessed and modeled where possible and used to design an appropriate risk management strategy.

2. For all indices, there are regions that are typically positively correlated with the climate index, and others that are negatively correlated. For the index insurance provider, this offers the possibility of reducing the aggregate risk of payouts by offering an appropriate balance of insurance products across these regions. The resulting reduction in payout risk faced by the provider could be used to reduce the cost of offering insurance products.
3. There is at least some capacity for ENSO prediction at least a few months ahead. For regions that have a strong ENSO teleconnection, this has the following implications for index insurance:
   a. Sales of an insurance contract that is based on an ENSO sensitive index defined for the season of regional weather impact need to be closed prior to the longest possible period of predictability of the index. The prediction lead times and skill continue to improve, so the design of the index insurance policy needs to keep this in mind, and re-evaluate the offerings periodically.
   b. At least in regions where ENSO teleconnections are strong, there is a potential to define the index insurance contract using an ENSO index for the *pre-season* conditions. This has the potential benefit that regional agencies who may need funds to prepare for upcoming contingencies
4. For the climate oscillations with the longer periods, there is the possibility of probabilistically detecting a change in state, and thus predicting a change in the probabilities of weather extremes in a region with strong teleconnections to that index. The implications for index insurance are:
   a. There is an opportunity to use insurance rates to communicate the changing risk over time as identified through these teleconnections. The communication of the changing risk could potentially lead to changes in lending and hence in user behavior that may provide an additional level of adaptive risk management to that offered by the index product directly. The advantage of this communication has to be assessed relative to the

increase in potential transaction costs and the actual predictable shift in the associated probability distributions.

b. Since, change detection typically requires at least some data indicating that the change has taken place, either the index insurance provider and/or the purchaser will face adverse conditions during the transition period, or the high variance (uncertainty) in the associated payout structure will need to be accounted for in the premiums. It is worth emphasizing that this, as well as other aspects of the implications of structured climate variability apply not just to index insurance but also to any other risk management programs that rely on an estimate of the probability distribution of events. The time varying nature of inter-annual and longer climate risk exposure has been ignored in past analyses, where stationarity (past represents the future) has been assumed as a matter of convenience. These factors continue to be relatively ignored in the anthropogenic climate change debate, where only potential changes in the mean and variance of regional weather are explored with high uncertainty as part of a carbon emission scenario, and the quasi-periodic nature of climate and how it may change in the future may change is not being discussed. While the long term mean and variance of precipitation and temperature have a bearing on the climate risk exposure, as discussed earlier, the tail probabilities that mark the risk of extremes are perhaps determined to a great degree by the potential clustering of events due to multiple time scale quasi-periodic phenomena. The assessment of how these tail probabilities correspond to specific climate regimes and the prediction of the upcoming probability of a regime could potentially be used to better parse and manage climate risk exposure today, rather than focus on a 2080-2100 climate change scenario.

5. Given the potential for the interaction across the climate modes represented by these indices in any given region, it is likely that a more effective index would be derived by a statistical combination of the effects of these indices. The potential nonlinearity adds a layer of complexity in the derivation of such an index. It is critical that the derivation of this index be relatively transparent and easy to compute, so that the resulting product and its properties are easily understood. This suggests that complex, black box statistical models such as artificial neural networks or even nonparametric regression methods would not be appropriate for developing such an index. A product that uses categories that represent a combination of threshold exceedances of 2 or more indices (e.g. NINO3.4 greater than 1.5 and PDO greater than 1.0) to define a payoff structure could be useful. A parametric regression model with an easily interpreted equation may also be useful.

Many other indices (e.g., Pacific North American or PNA, or Madden-Julian Oscillation or MJO, or All India Monsoon rainfall) have also been developed using daily or monthly data on sea level pressure, outgoing long wave radiation, rainfall or other parameters. These may potentially be much better related to weather extremes in some regions than the indices introduced here. However, these indices are in turn modulated by the low frequency climate indices defined earlier. Hence, it is important to preface their potential utility with a discussion of the low frequency climate indicators. Often, these high frequency climate indices are identified on the basis of explaining a reasonably large amount of spatial variance in a global or hemispheric analysis, i.e., they have large scale applicability. This leads to two potential outcomes. First, areas that have publicly available, high quality, long run data sets, such as N. America, Europe, Australia and India, tend to be unduly weighted in the analysis and this is reflected in the indices developed. Second, the index explains some variance almost everywhere, but not as much of the variance in the regional or local time series as some other index that could in concept be developed for regional applications.

In summary, we offer the following recommendations for index insurance development:

1. For regional applications, it may be best to develop a targeted index that is appropriate for that region, for either concurrent or predictive index insurance. This index may be developed based on
    a. a statistical relationship between specific local weather extreme data (e.g., duration of longest dry spell in the season, or heating degree days) and a suite of global climate indices, or
    b. An exploration of the climate factors that influence local weather extremes more directly and then a relationship of the statistics of local weather extremes to those variables in the same way.

   In either case, Monte Carlos simulations of the temporal variability and uncertainty associated with the contracts will be needed to assess the effectiveness of the proposed index. The spatial and temporal distribution of weather extremes within a region would also need to be simulated so that basis risk issues are better understood. One could use local rainfall or temperature directly as an index, but this exposes both the user and the provider to high basis risk, particularly where the index contracts are tied to short records of local data, since these records will invariably not represent the inter-annual and longer climate variability that has been discussed in this paper. Hence, an approach that ties local /regional climate records to regional and hemispheric long climate indices is to be preferred, particularly if the recommendations as to assessment of the associated uncertainties are followed. The work in Khalil et al (2007) provides an example of how a regional index can be developed, and its predictability and long term variability evaluated via simulation.

2. An effort should be pursued to develop regional indices and to assess their long term variability and predictability in a systematic way across the major regions where offering index insurance is of interest. By and large, at least 30 to 50 years of regional weather data is available in many places today. It may be possible to supplement this with longer proxy records in many places. A strategy for record extension using MERRA, NCAR and ERA40 re-analysis records back to 1948 may be possible globally. Unfortunately, at the regional scale these models do not always provide very representative results for weather extreme interpolation. However, no systematic effort has been made so far to develop appropriate global and regional indices of climate

variation that inform regional weather extremes. Investment of effort in this direction will be very valuable in assessing what can be done considering both regional and global sources of information (both data and climate models). As seen in the next section, a fair amount of knowledge exists with respect to regional correlation of seasonal weather averages with the leading climate indices. Much less has been done to assess the connection weather extremes. The specific analyses that should relate to suitability of index insurance applications have been pursued in very few places indeed, and a ready suite of results for specific statistics of weather extremes that would best inform the targets of index insurance is currently not available.

6. Analyses at a nearly global scale may be useful primarily for portfolio management for an index insurance provider. The global or hemispheric climate indices discussed here vary in their strength of impact by region. Typically, the magnitude of the global variance explained in precipitation is much less than for precipitation. Thus, the direct use of these indices for regional index insurance is not likely to be too useful. As indicated in the first recommendation, some statistical combination of these indices may indeed be useful for regional analyses. However, this combination would be regionally specific, and hence a single index could not be directly applied globally. The most significant utility of the global indices reviewed here is that they represent rather different time scales of natural climate variability and hence provide a context for how regional climate and weather extremes may vary over time. The relatively long records available for these indices permit an exploration of how regional climate in relatively data sparse regions could possibly change. For instance, in the Sahel region of Africa, precipitation has undergone nearly a bid-decadal switching regime over the last century, punctuated to an extent by higher frequency changes in ENSO and SST conditions in the Gulf of Guinea. The use of the global indices as well as regional climate influences jointly could help diagnose how index insurance could be applied and monitored. The recommendation is to pursue detailed analyses of the regional to global teleconnections with the intention of tagging the index insurance and other risk management strategies to some prediction or possibility of shifts in the operative regimes, either retrospectively or looking forward. This is important since in many places (e.g., Western Australia) a relatively short wet spell of a few years starting in 1998 was viewed as "breaking the drought", and conditions have since reverted to those persisting prior to 1998. Tailoring the index insurance pricing to either the higher variance role in global weather systems, much of the analyses relate to mid-latitude jet stream related dynamics, or if in the tropics to ENSO related activity, and more regional influences that interact with the larger signal but provide more precision to regional outcomes are missed.

Seasonal correlations of global climate indices to regional precipitation and temperature

Much of the published literature provides either an analysis of the linear correlations of the indices with regional rainfall and temperature, or an analysis of the regional outcomes that correspond to the extreme phases of the oscillation represented by each index individually. Sometimes, especially for an

ENSO index, the probability of exceedance of selected thresholds of rainfall or temperature is also presented conditional on a threshold of an ENSO index. Unfortunately, the data sets used by the authors vary and so do the conclusions as to the some of the specific areas and seasons where the teleconnections are strong. By and large, most of the literature is ENSO related and focused on either the Americas or S. Asia and Australia, with the NAO connections being described largely for parts of N. Africa, the middle-east and Europe. Here, we first present a direct analysis of correlation of the indices discussed in the previous section with seasonal rainfall and temperature using the global gridded 1° by 1°, data set from the Climate Research Unit (CRU) that is usually a benchmark used for many of the IPCC studies for climate change trend analysis. Several efforts at data quality control and interpolation have gone into creating the monthly data series at CRU and hence it is useful to examine the results across all indices and across all seasons for this data set. However, we note, that sparsely sampled areas in Africa, N. Asia, and S. America will still not be well represented since the gridded data is interpolated from very limited observations in these locations and filled in over the full century. The presentation of the results of this analysis is followed by a discussion and a summary of the key features or teleconnections noted by other authors, including, in the case of ENSO a discussion of some of the conditional probability analyses. A note to the reader is that when we have investigated specific regional teleconnections, the best relationships of rainfall and temperature with these indices are typically nonlinear, with an amplified response as the index anomaly is more extreme and a dampened response as the index anomaly is closer to its average value. Consequently, the linear correlation analyses presented here likely understate the potential relationship with the index in such places.

The correlations of each index with seasonally averaged (not extremes) daily maximum and minimum temperature and rainfall for 4 seasons per year are shown in Figures 2 through 5. For several seasons and regions, statistically significant (against a null hypothesis with $p=0.1$ that does not consider memory or persistence in either series) but relatively modest correlations are indicated. Typically, a physical connection to a region is meaningful only if a generally large, contiguous region shows a correlation of the same sign as the index.

ENSO correlations and conditional probabilities:

There are several ENSO indices available. A commonly used one is the NINO3.4 index. This is the one that was used here for illustrations in Figure 2. For the daily maximum temperature, the correlations are broadly similar for the Jan-Mar (JFM), Apr-Jun (AMJ) and Oct-Dec (OND) seasons. Typically, the index is positively correlated with daily maximum temperature in the Northern part of N. America, and across the equatorial belt around the world in these seasons. It is negatively correlated with daily maximum temperature in the Southern parts of N. America for these seasons. In the July-Sep (JAS) season, the daily maximum temperature and NINO3.4 correlations are positive for central and N. Africa and S. Asia (excluding the Himalayan and northern region) and negative in the Philippines and the Eastern Indonesian islands. The spatial correlation patterns of the daily minimum temperature with NINO3.4 are similar to those for daily maximum temperatures, but weaker.

For seasonally averaged daily precipitation, the NINO3.4 correlations for JFM are negative in North Western N. America, the Great Lakes, equatorial parts of S. America, S. Africa, Western Australia, and Indonesia and positive in S. W. N. America and Central Asia. The spatial correlation patterns for AMJ and OND are similar but with some spatial shifts. In OND, central Eastern S. America is positively correlated. In JAS, the monsoonal regions of Asia and the Americas, essentially the equatorial and tropical band are negatively correlated.

The next set of figures from the IRI website (http://iri.columbia.edu) are from Mason and Goddard (2001). Their analysis was based on the 45 year record of 0.5 by 0.5 degree data set compiled from 12000 stations around world by New et al (1999, 2000). They consider the rainfall data for a given season to be partitioned into 3 terciles – the upper tercile is termed "above normal", the middle "normal" and the lowest one "below normal". The long term of climatological probability associated with each tercile are 1/3 by construction. The authors consider 10 years which had positive values of the NINO3.4 index representing El Nino events of different magnitudes and compute the probability associated with the above normal and the below normal categories for each season's rainfall for this set of years, for each season. Even before reviewing the spatial patterns that result, it is worth noting that such a plot can be misleading since it is based on a relatively small set of years, and thus has high uncertainty. However, it is a conditional probability plot relating extremes of seasonal precipitation to a one sided extreme of the index, and hence is more directly informative as to potential impacts to be treated by index insurance than the linear correlations presented earlier.

Comparing the shifts in the probabilities for the above and below normal cases, the JFM, AMJ and OND results are by and large consistent with those one would expect from the linear correlation analysis. The exception is central Asia, where the below normal category does not undergo a significant change in probability but the above normal does, suggesting that the "normal" category has a reduced probability and the changes are all in the normal and above normal category, i.e., a more subtle shift during El Nino conditions in that region. For JAS, the results are largely consistent with those seen with the linear correlation analysis. Mason and Goddard (2001) and the IRI web site also present the results for their analysis for a composite of La Nina events. The patterns are similar in general but the strength of the shifts is not usually the same suggesting that there are nonlinearities in the response of seasonal precipitation to ENSO. The tropical Pacific islands in general demonstrate a stronger response to La Nina than to El Nino events, but in general La Nina conditions are not as far from normal as El Nino conditions for most other places. Mason and Goddard (2001) note that only about 20%–30% of land areas experience significantly increased probabilities of above- or below-normal seasonal precipitation during at least some part of the year during El Nino events, and that since different areas are affected at different times of the year, the fraction of global land affected in any particular season is only about 15%–25%. In their paper they review a large number of past studies on ENSO teleconnections to seasonal precipitation and find them consistent in general with the patterns reported by them and in the linear correlation analysis presented here.

The statistical predictability of ENSO indices is reviewed in Latif et al (1998), Mason and Mimack (2001), Khalil et al (2007) and most recently by Lima et al (2009). The highest forecasting skills under cross validation are established by Lima et al. They find that the variance explained in the prediction of the NINO3.4 and NINO1.2 indices decreases from about 0.8 to 0.9 at 1 month lead time to about 0.5 or lower at 4 to 6 months lead time by most methods, and may then be less than 0.2 for longer leads except for their method where they are able to preserve predictive skill for up to 18 months explaining between 0.2 to 0.5 of the fractional variance of the index. However, their testing period is relatively short, so these results should be viewed as indicative trends for the potential predictability rather than as measures of actual predictive skill. It is worth noting that while explaining 0.5 of the fractional variance in an ENSO index through a forecast may be comparatively impressive, it may not mean much for practical predictions of a regional insurance index, since the teleconnections of ENSO explain only a fraction of the variance associated with regional rainfall or temperature. Goddard and Dilley (2005) compare the relative predictability of seasonal precipitation during ENSO extremes (El Nino and La Nina years) compared to “neutral” years using a ranked probability skill score. This figure is reproduced as the last visual in Figure 2. Rather modest positive skills are to be noted in most seasons in most locations. Some regions do exhibit strong skills. These locations by and large coincide with those identified earlier as having statistically significant correlations. Some locations with weak correlations exhibit negative skill, as may be expected by chance given the short records used.

Where multiple indices are used or nonlinear models are used, it is possible to demonstrate considerably higher skills for precipitation or streamflow extremes with lead times of a season to a year. Examples include those in Khalil et al (2007) and Souza et al (2004) and other regionally rather than globally focused analyses.

Figure 2: Linear correlation map for each index with 1 latitude and longitude gridded, seasonally averaged weather data, for each of 4 seasons. First with daily max temperature, then daily minimum temperature, and then precipitation. In all cases the monthly CRU data for weather and for the climate index as available at http://climexp.knmi.nl/ were used. For ENSO probability shifts in precipitation are also provided.

ENSO correlations: Daily maximum temperature

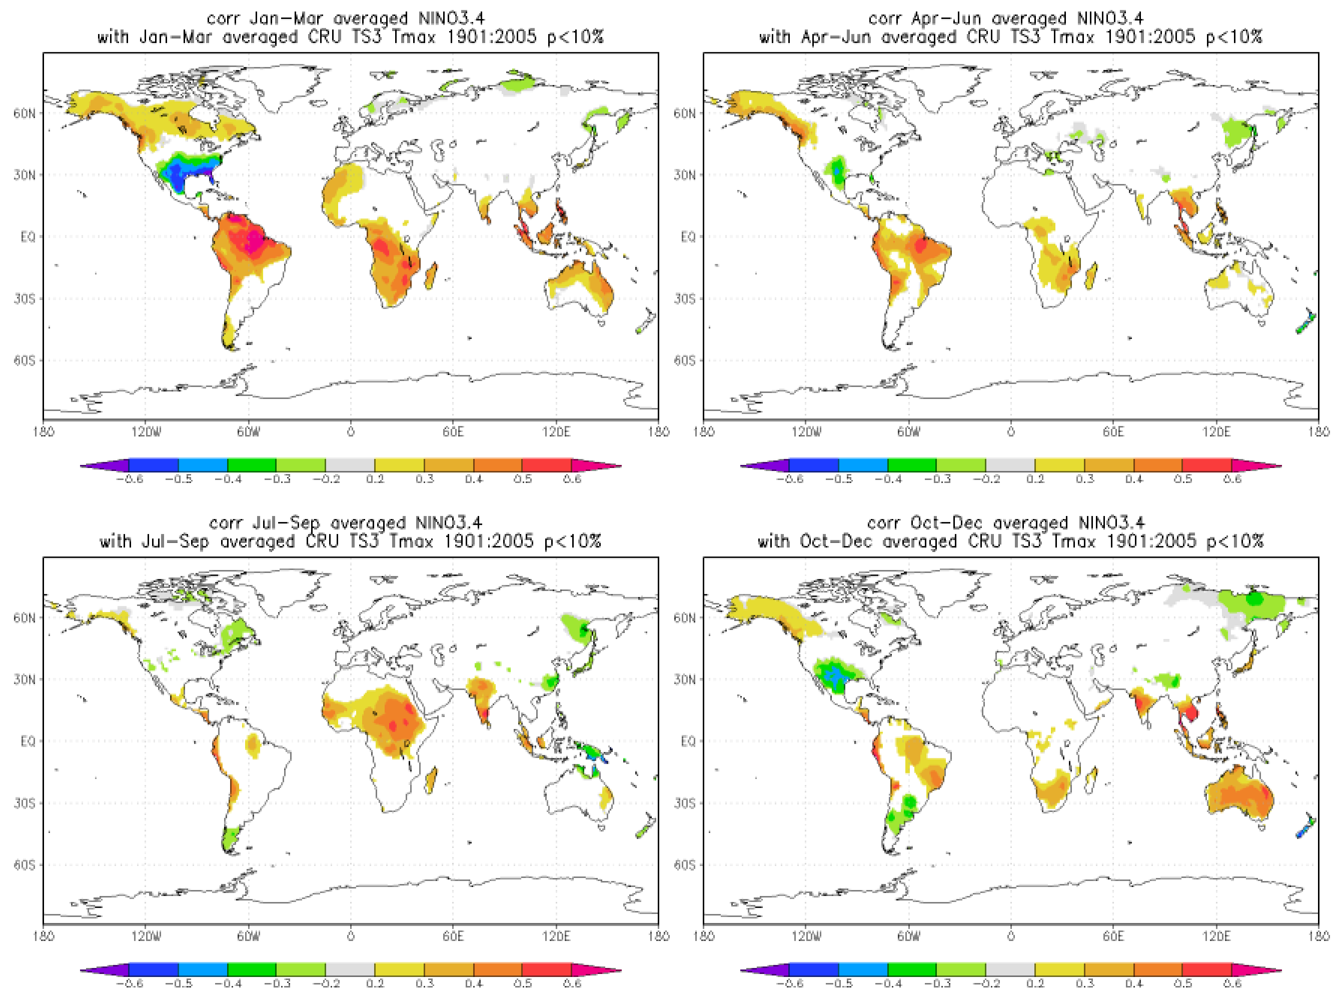

ENSO correlations: Daily minimum temperature

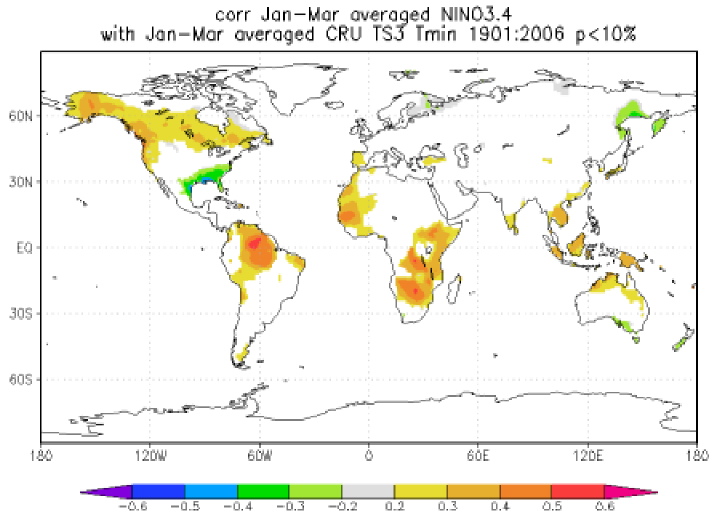

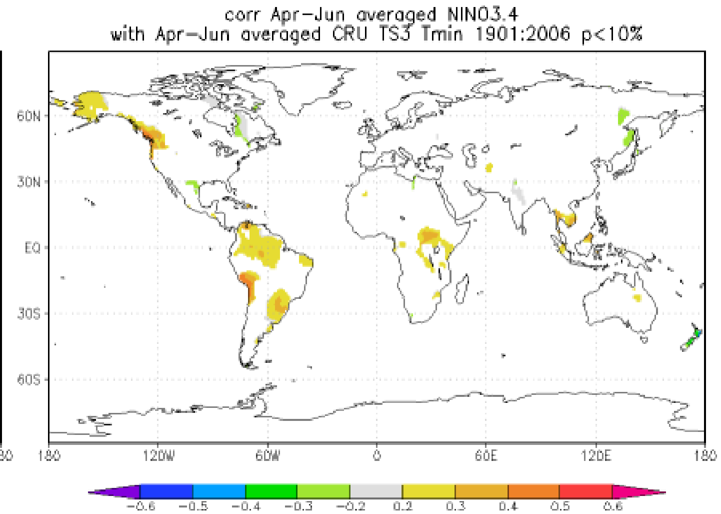

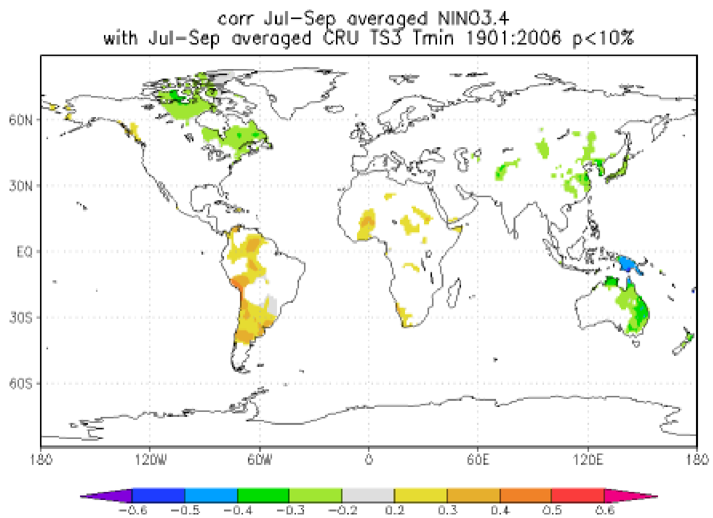

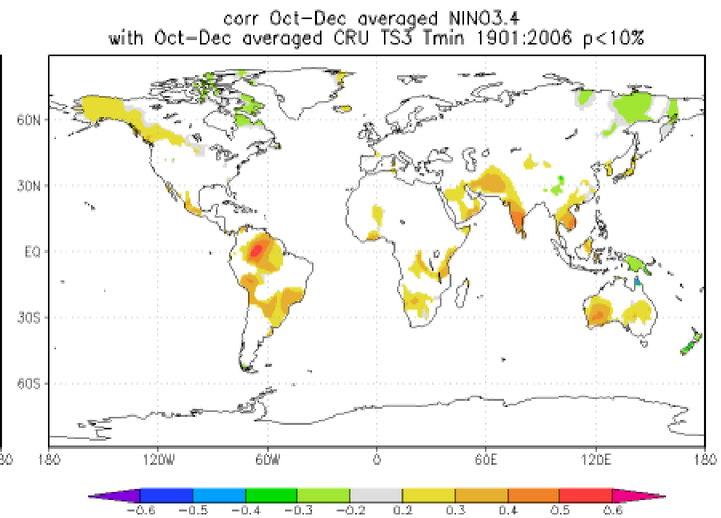

ENSO correlations: Daily precipitation

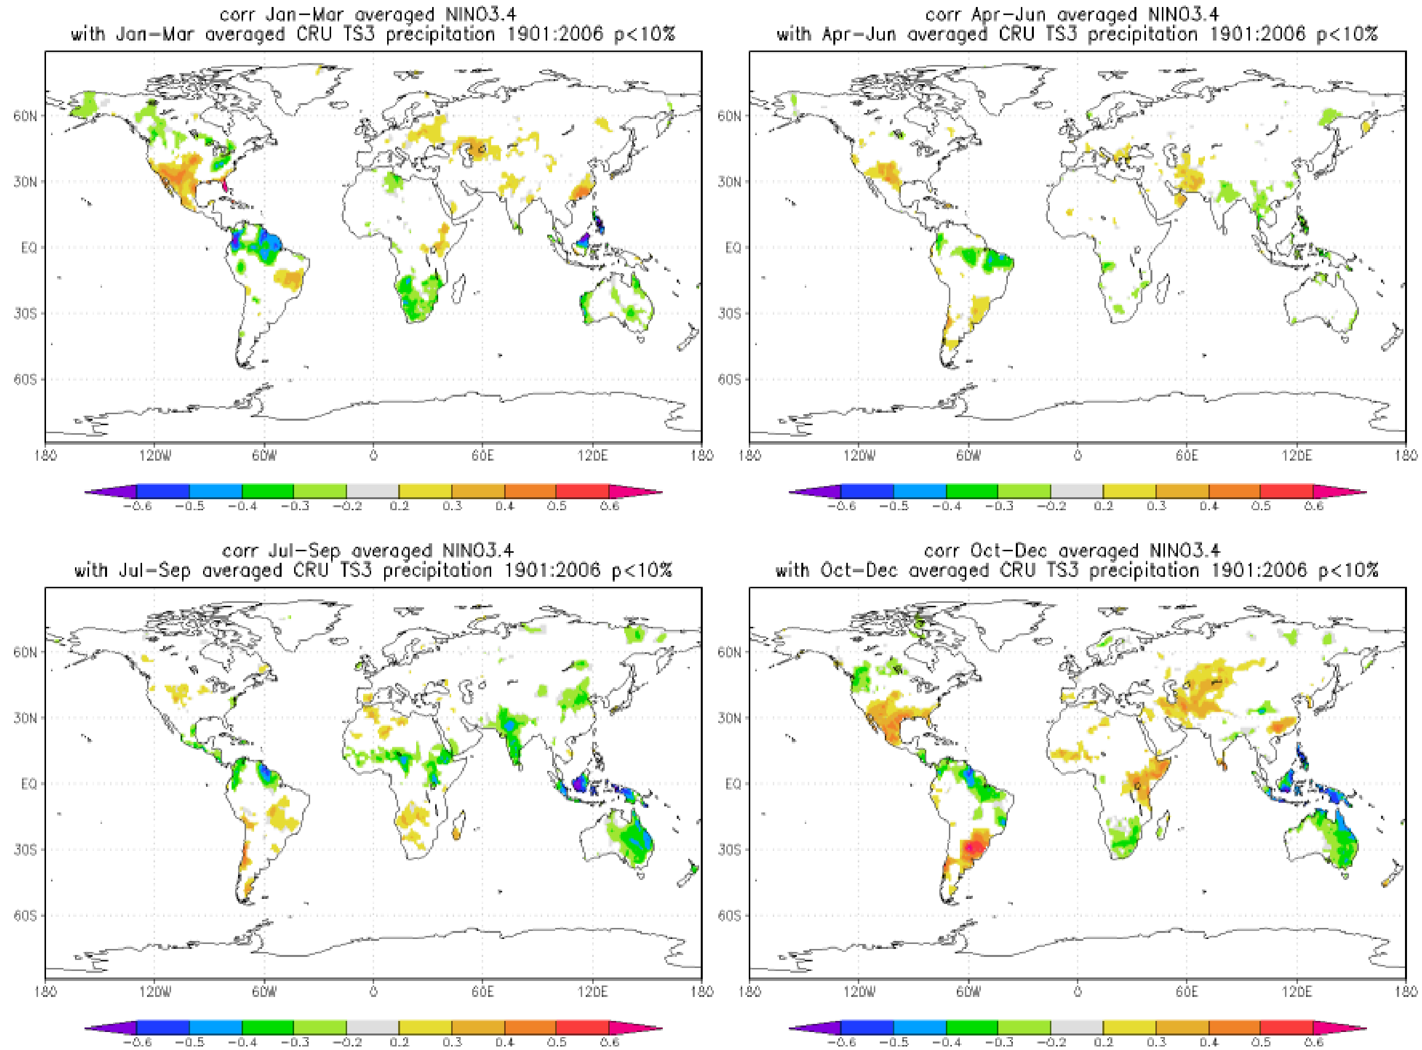

ENSO probability shifts: seasonal precipitation (from IRI forecast products library, http://iri.columbia.edu, also Mason and Goddard 2001)

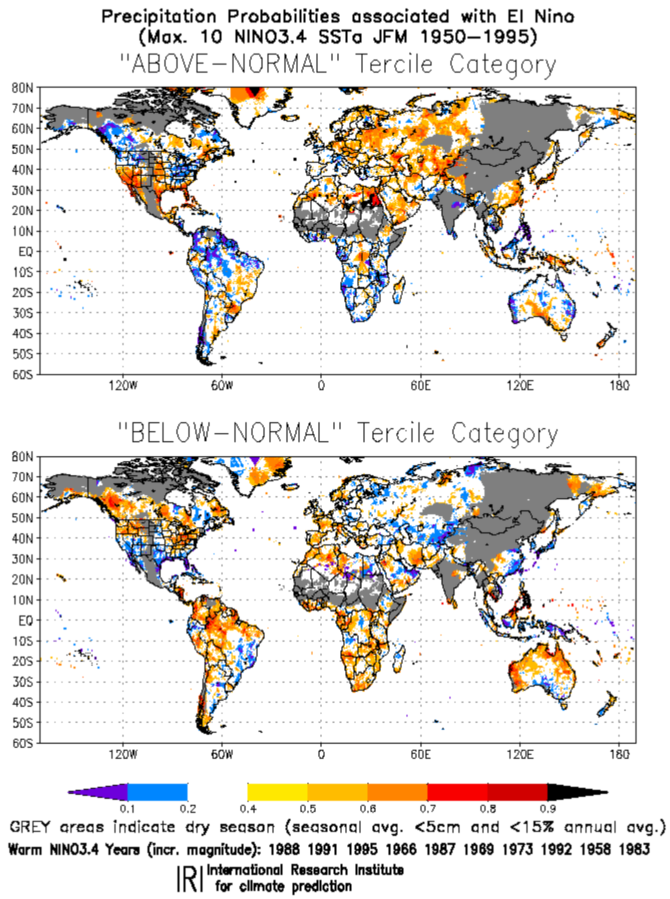

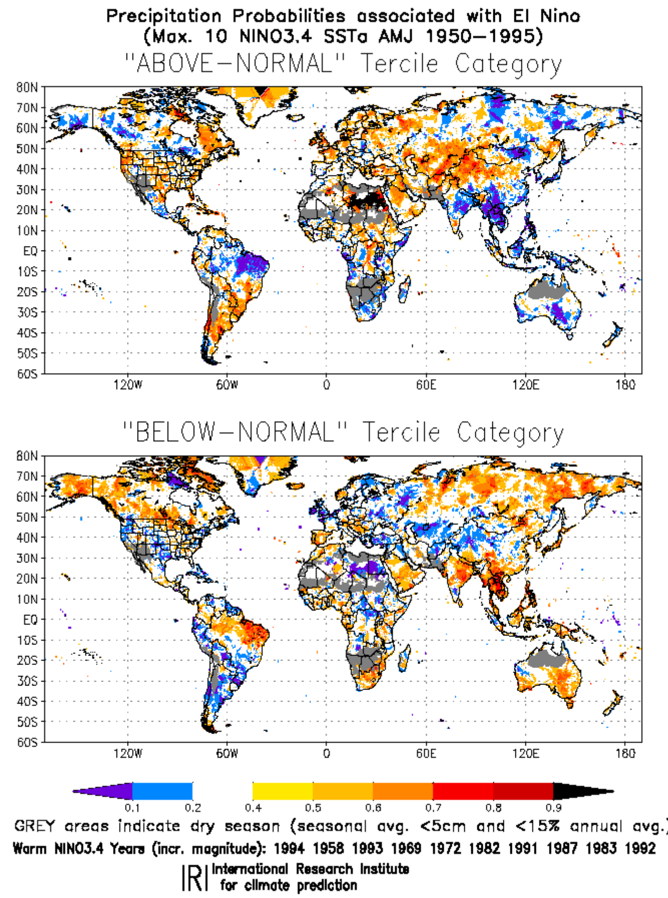

ENSO probability shifts: seasonal precipitation

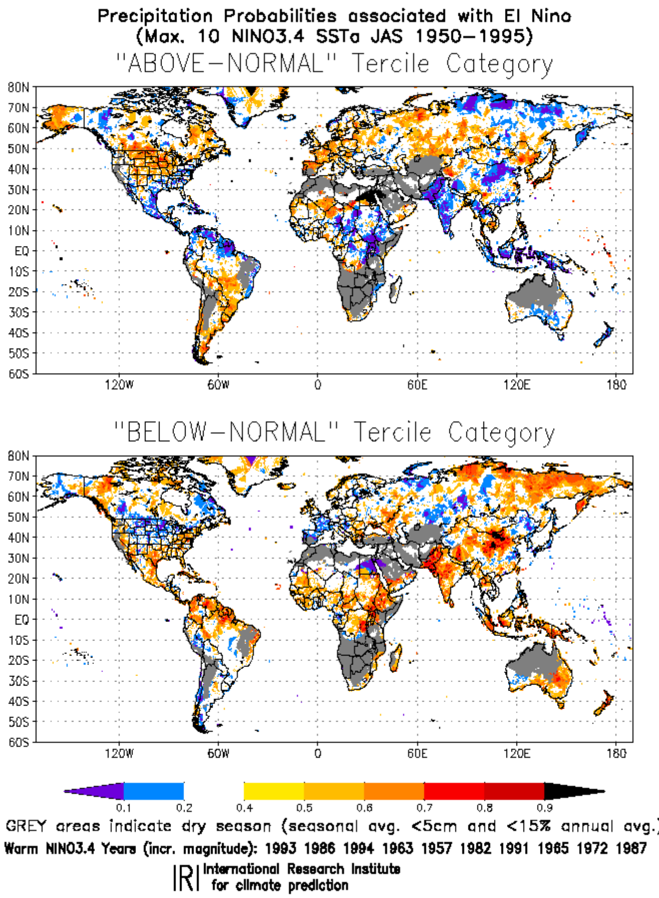

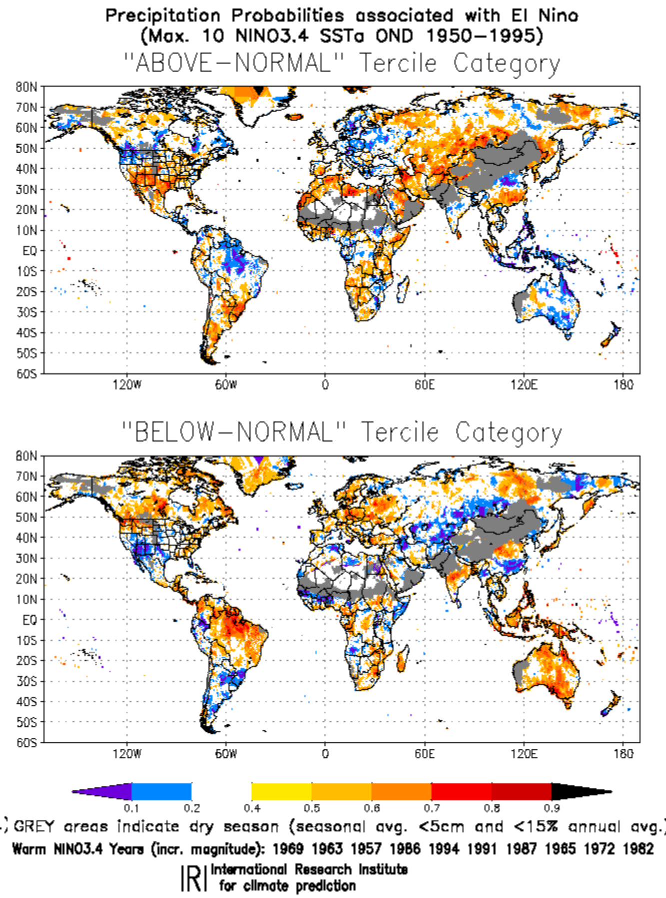

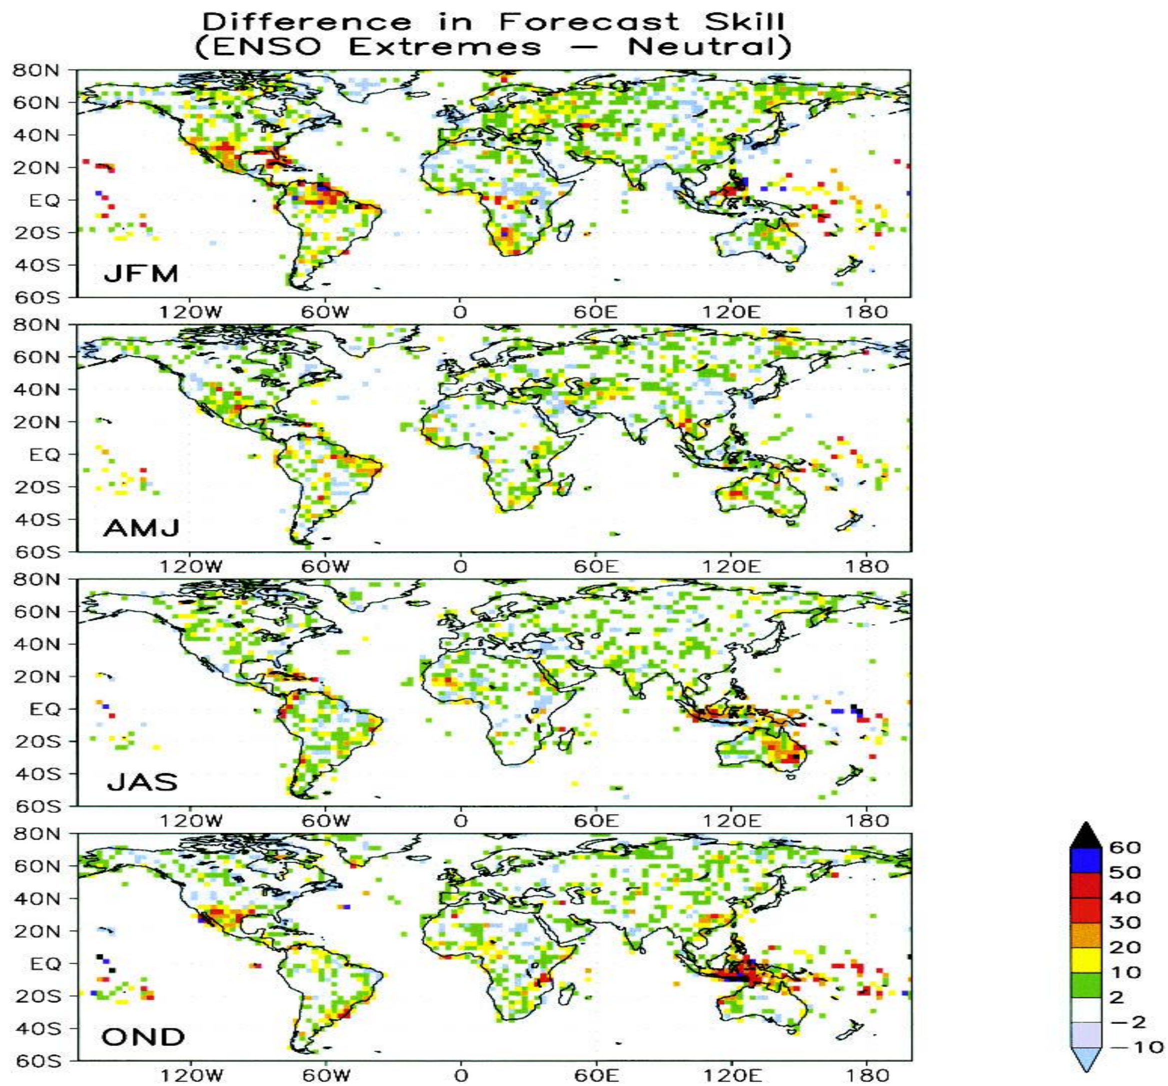

Differences in skill (Ranked Probability Skill Score) for three categories of seasonal rainfall forecasts between ENSO extremes and neutral conditions for the 1950–95 period. Positive values indicate higher skill during ENSO extremes. From Goddard and Dilley (2005)

NAO correlations

The NAO, also discovered by Sir Gilbert Walker in the 1920s is defined over the N. Atlantic (Hurrell et al, 2001) and is most prominently associated with the winter season (DJF) and spring (MAM) jet stream dynamics over the N. hemisphere. It is not quite a global index in the sense of an ENSO index. Kushnir et al (2006) present a recent review of the literature on the NAO and its predictability, including the potential for interactions with ENSO that may then manifest them in the teleconnections of ENSO or NAO to regional variability. Kushnir et al indicate that in DJF and MAM there are large areas where 16%–36% (and above) of the surface temperature and rainfall variance (correlation of 0.4 to 0.6) can be reconstructed. Typically, whereas ENSO influences the tropics, the NAO influences the mid to high latitudes. An interesting aspect of the NAO index is its rapid decline in predictability across the season of the year (it is most persistent in the winter-spring), but its persistence or re-emergence year after year. Based on the discussion in Kushnir et al, one can understand the NAO as a phenomenon whose center of activity shifts spatially across the seasonal transition, and hence the usual NAO index as it is computed best measures the strength of the phenomenon in the DJF and MAM seasons. The SST pattern that emerges in the other seasons appears to be related to the strength of the DJF/MAM patterns and subsequently seems to determine the DJF and MAM patterns for the following year. Thus, there is some hope for NAO prediction over longer time scales, and an explanation of the re-emergence of the phenomenon across years. This observation also explains the DJF/MAM NAO and ENSO connection to N. and Central American hurricane landfall incidence noted among others by Elsner et al (2001).

The graphics in Figure 3 for the linear correlations of the seasonal NAO index to the seasonal averages of daily maximum and minimum temperature, and daily rainfall are quite interesting with large spatially coherent patterns in the N. Hemisphere. The most prominent season of impact is JFM with OND and AMJ exhibiting similar but considerably weaker patterns. The NAO index in JFM is negatively correlated with daily maximum temperature over Iceland, Greenland, N. Africa, the Middle East, W. Asia and N. India. It is positively correlated with N. E. USA, Europe and N. Asia. The pattern for daily minimum temperature is similar but weaker. In AMJ, the pattern shrinks to a negative correlation over Greenland and the S. Mediterranean with a modest positive correlation along the coastal N. Mediterranean. In JAS, the Mediterranean part of the pattern is all that remains. For precipitation, statistically significant and spatially coherent relationships are considerably weaker and are manifest only in JFM with positive correlations over the Scandinavian countries, and the S. Mediterranean, and negative correlations over N. Mediterranean and the equatorial, central Africa.

A number of authors (e.g, Dugam et al 1997; Chang et al, 2001, Gong and Ho, 2003) have argued for links between the winter NAO and the subsequent summer Indian and East Asian monsoons, particularly through a modulation of the ENSO impact. However, others (e.g., DelSole and Shukla, 2002) report that linear models using ENSO and NAO indices are unable to successfully forecast the Indian monsoon any better than other predictors in use. On the other hand Wu et al (2009) demonstrate success in prediction the East Asian summer monsoon using NAO and other predictors.

Figure 3: NAO correlations

NAO Correlations: Daily maximum Temperature

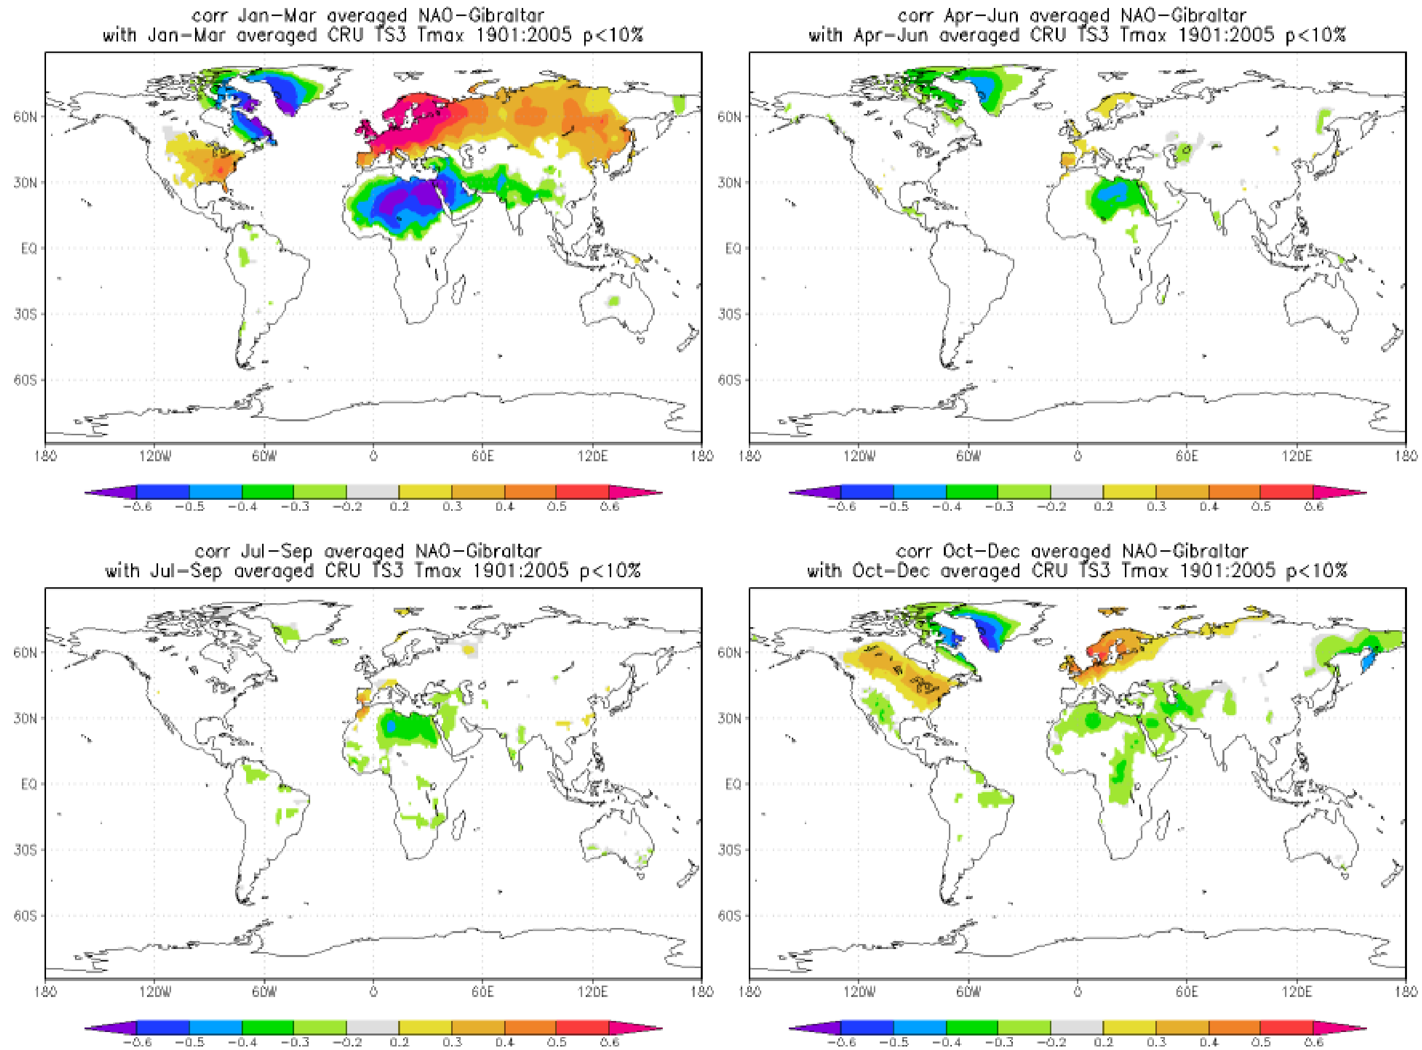

NAO Correlations: Daily minimum Temperature

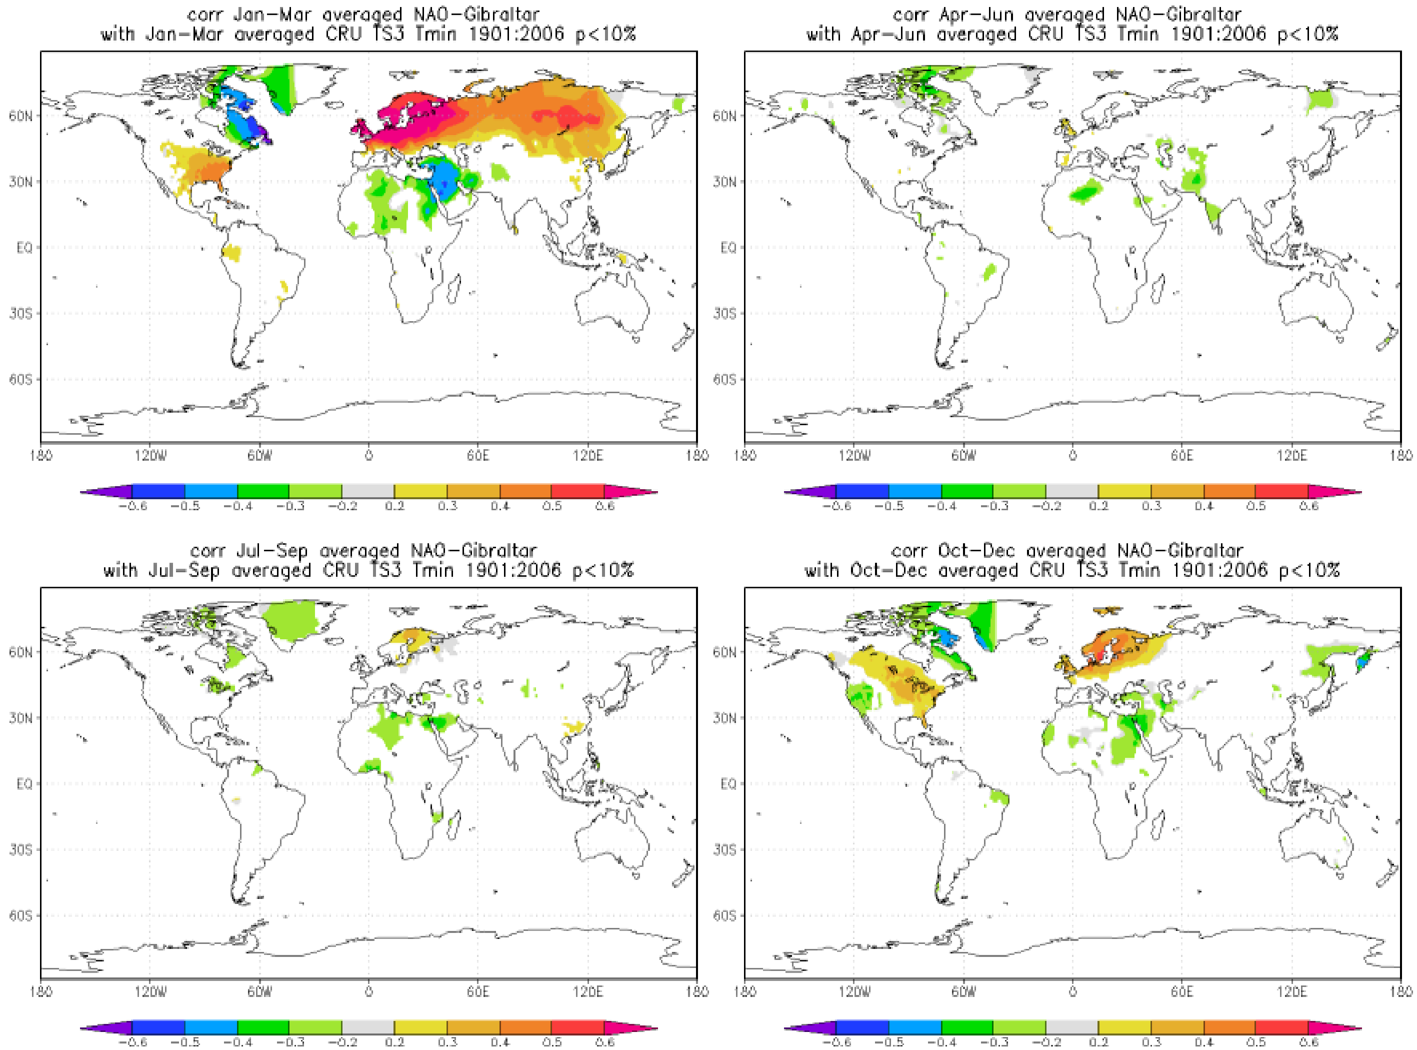

NAO Correlations: Daily precipitation

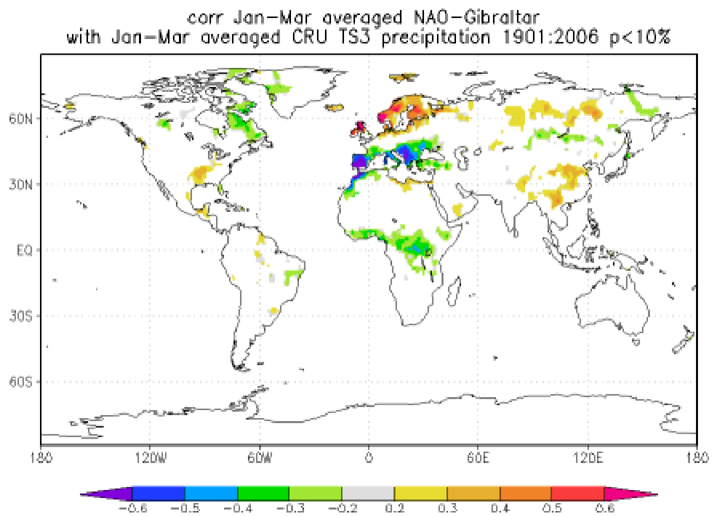

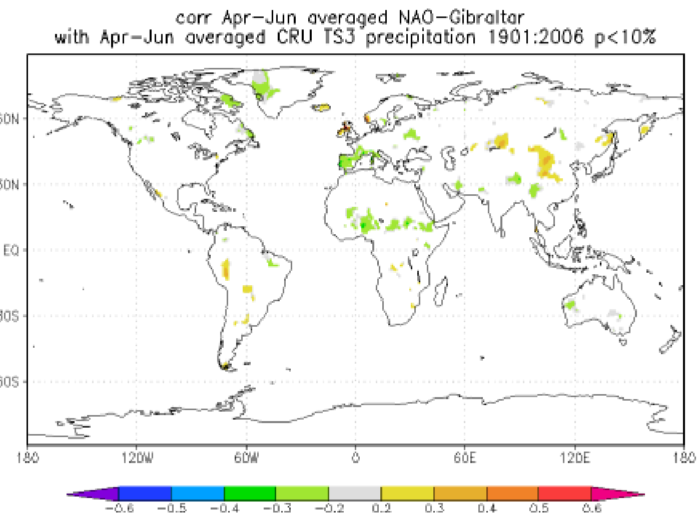

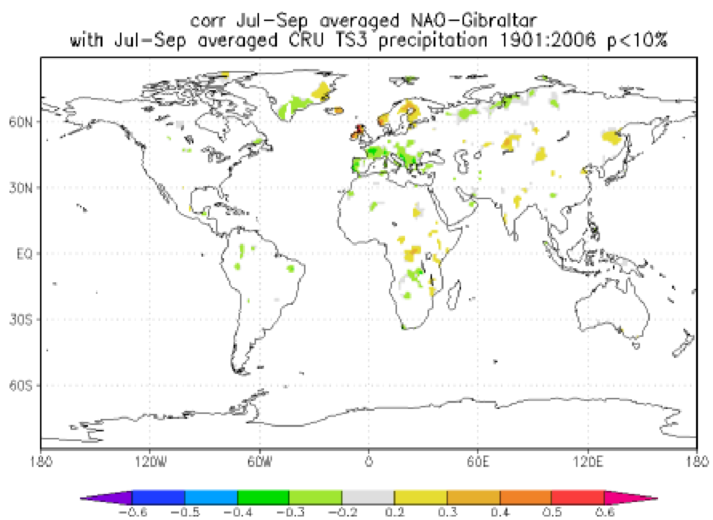

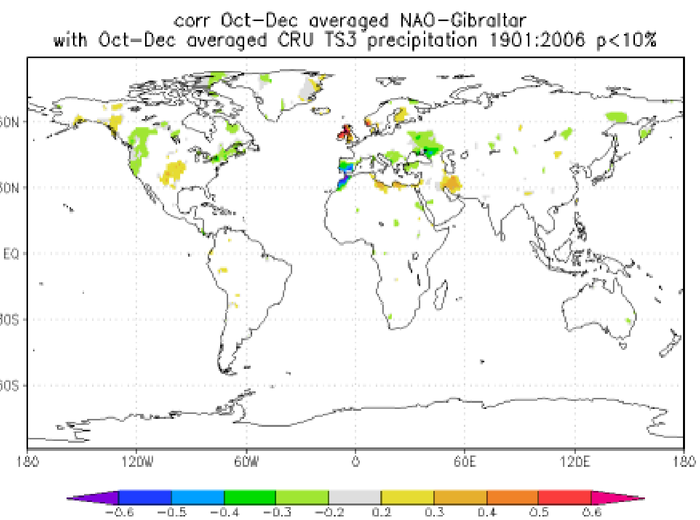

PDO correlations

The PDO refers to a see saw pattern in the N. Pacific Ocean SSTs, that was discovered and named by Steven Hare in 1997 (Mantua, 1997). A closely related term is the Interdecadal Pacifical Oscillation or IPO which establishes similar variability in the N. and S. Pacific. Given the relatively recent discovery of this climate mode, its long period, and the paucity of long records, not as much is understood about the PDO, its teleconnections and their predictability. Often (e.g., Gershunov, et al, 1998) it is discussed as an inter-decadal mode that modulates ENSO teleconnections, and most of the research on teleconnections has been confined to the Pacific Rim region and the Western United States (Mccabe et al, 2004).

Given the long memory in the time series, computing the significance of correlations using the usual tests is suspect. With that in mind, we can explore how this index correlates with the global min/max temperature and precipitation as illustrated in Figure 4. The correlations with max/min temperature are generally very similar and are marked by similar patterns in JFM, AMJ and OND that vary a bit spatially. The PDO index in these seasons is positively correlated with the North Western quadrant of N. America, the North Eastern section of S. America, North Eastern Australia, Indonesia and the Phillipines, Equatorial to Southern Africa and S. India. It is negative correlated with Eastern North America and Scandinavia. In JAS, the index is positively correlated with central equatorial Africa and S. Asia, with negative correlations with eastern central Asia.

For precipitation, the linear relationships are generally much weaker, and are positive for South Western N. America and the Hudson Bay region in JFM and negative for Northern N. America in JFM and for the equatorial belt including C. Africa and the Sahel in JAS.

As was noted earlier, as the period of the oscillation increases, its impact on weather extremes likely comes through a modulation of the impacts of the higher frequency components of the climate system rather than directly. Nevertheless, large spatially coherent regions for temperature teleconnections are identifiable from this analysis. Note that all data were linearly detrended prior to computing correlations to remove the potential impacts of anthropogenic global warming. However, since the trend in global temperature over the 20th century is probably more exponential than linear, it is possible that much of the temperature correlations seen in this analysis and the analysis with the NAO actually still relate to the global warming signal shared in both data sets rather than a real response. The issue is more pronounced with PDO since it is an SST based index, while the NAO is an atmospheric pressure based index and is less closely related to temperature trends.

Figure 4: PDO correlations

PDO correlations: Daily Maximum Temperature

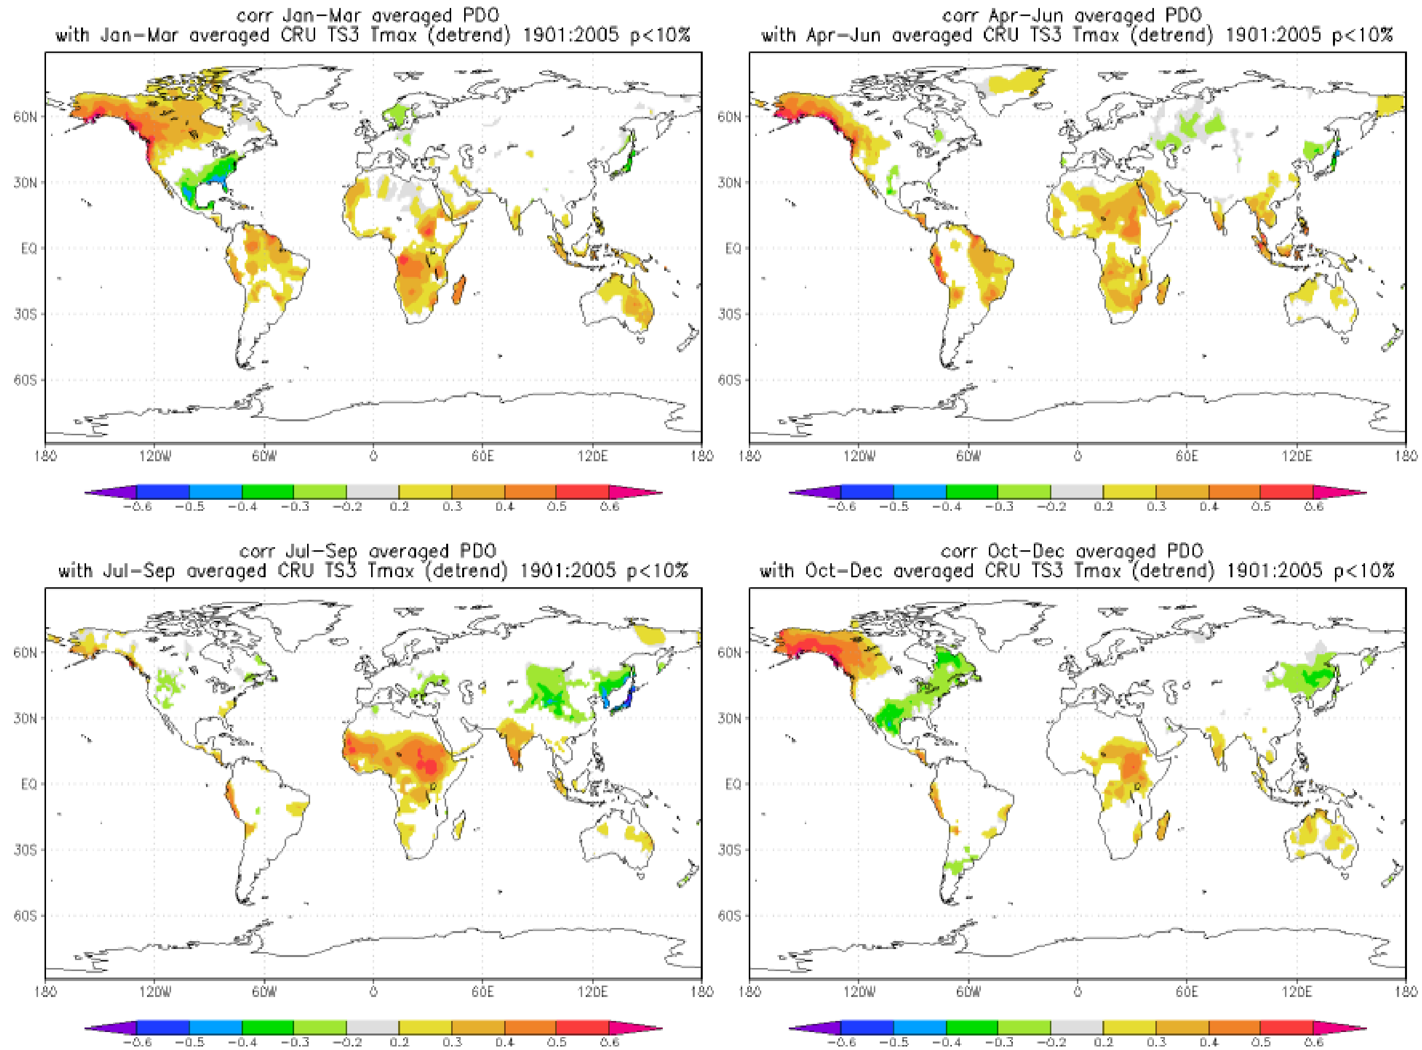

PDO correlations: Daily Minimum Temperature

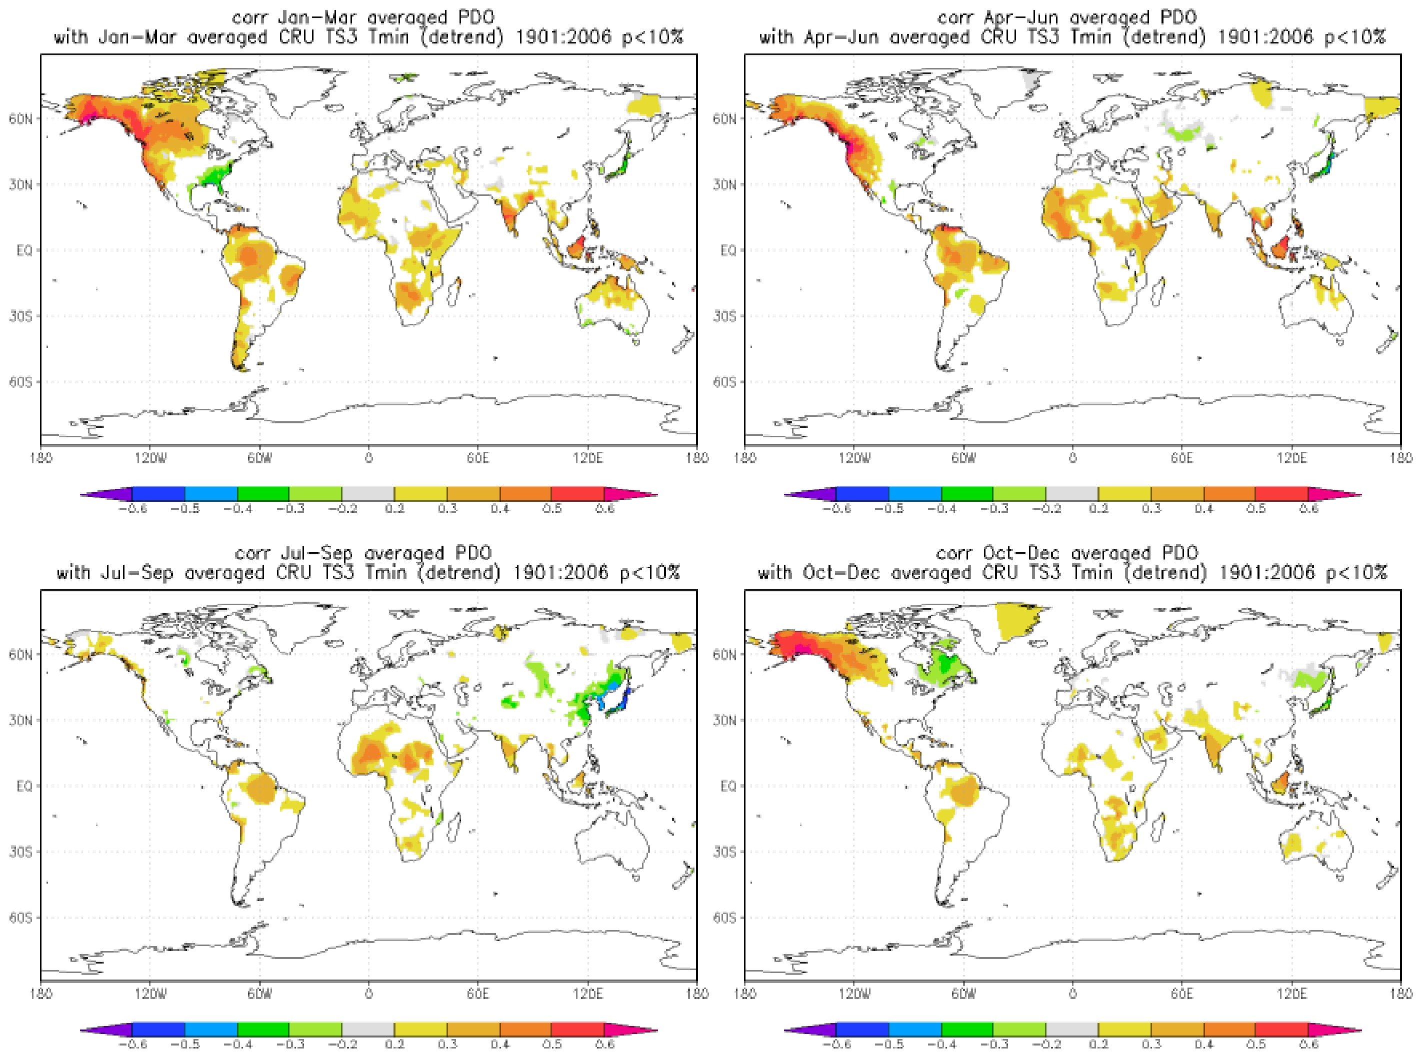

PDO correlations: Daily Precipitation

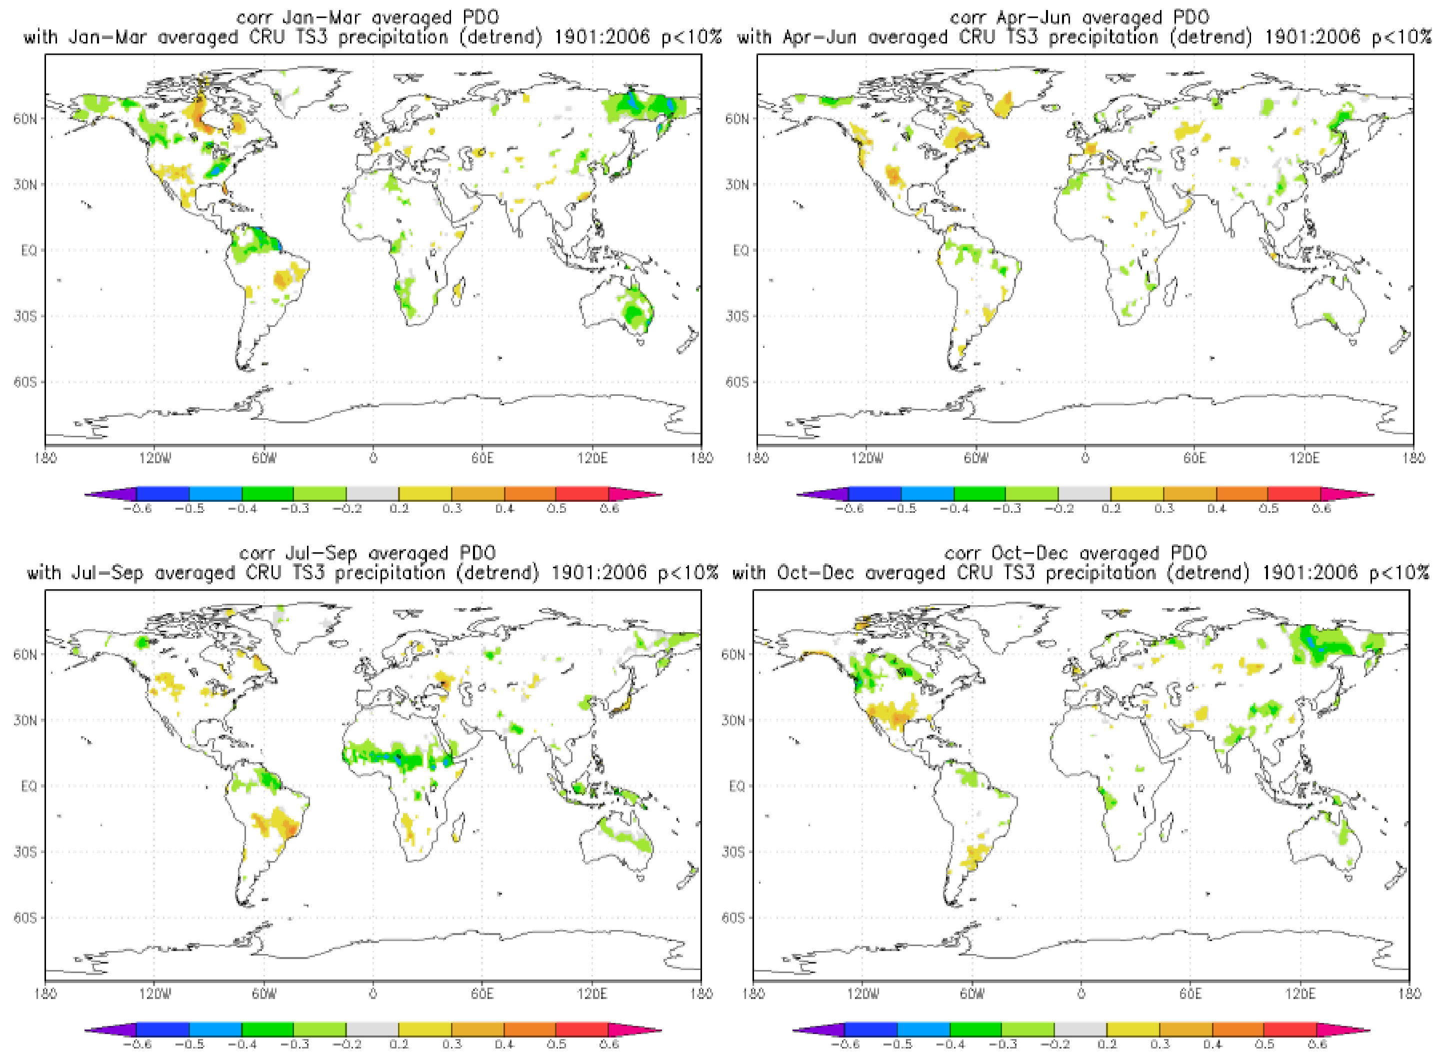

AMO correlations

The AMO is defined (Kerr, 2000; Enfield et al 2001) using the SST field in the N. Atlantic, and is characterized by a 60-80 year period. Given that the instrumental records are only about a century long, the existence of the AMO is somewhat controversial. However, it has been reconstructed from long proxy records (Gray et al, 2004; Hettzinger et al, 2008) and hence has a following. Its teleconnections to precipitation and to streamflow have been studied (Knight et al , 2006; Mccabe et al, 2004) and its role in moderating hurricane frequency in the Atlantic basin has also been assessed (Goldenberg, 2001; Hettzinger et al, 2008; Elsner et al 2001).

With the same caveats for correlation analyses as were indicated for the PDO we can examine the correlation of the AMO with the weather extreme variables in Figure 5. We note similar patterns for daily maximum and minimum temperatures with weaker correlations by far for the minimum temperature, consistent with previous analyses. The JFM and OND correlation patterns are generally similar but somewhat weaker than the AMJ pattern. In JFM/OND the N.E. United States and Greenland, Central to N. Africa, and Central S. America are positively correlated with the AMO index, while parts of C. Asia are negatively correlated. For AMJ, Greenland, parts of the Middle East and Central China are positively correlated while parts of Central S. America, S. Australia and S. India are negatively correlated. For JAS, only Canada and Greenland show significant positive correlations. For precipitation there is very little direct correlation. Perhaps, the only region to mention is Central Africa/Sahel which is positively correlated in JAS. However, note that several authors point to strong determination of seasonal precipitation by a combination of AMO, NAO, PDO and ENSO in places such as S. Florida. Such a feature really does not emerge using a detrended correlation analysis using only AMO.

Figure 5 AMO Correlations

AMO correlations: Daily Maximum Temperature

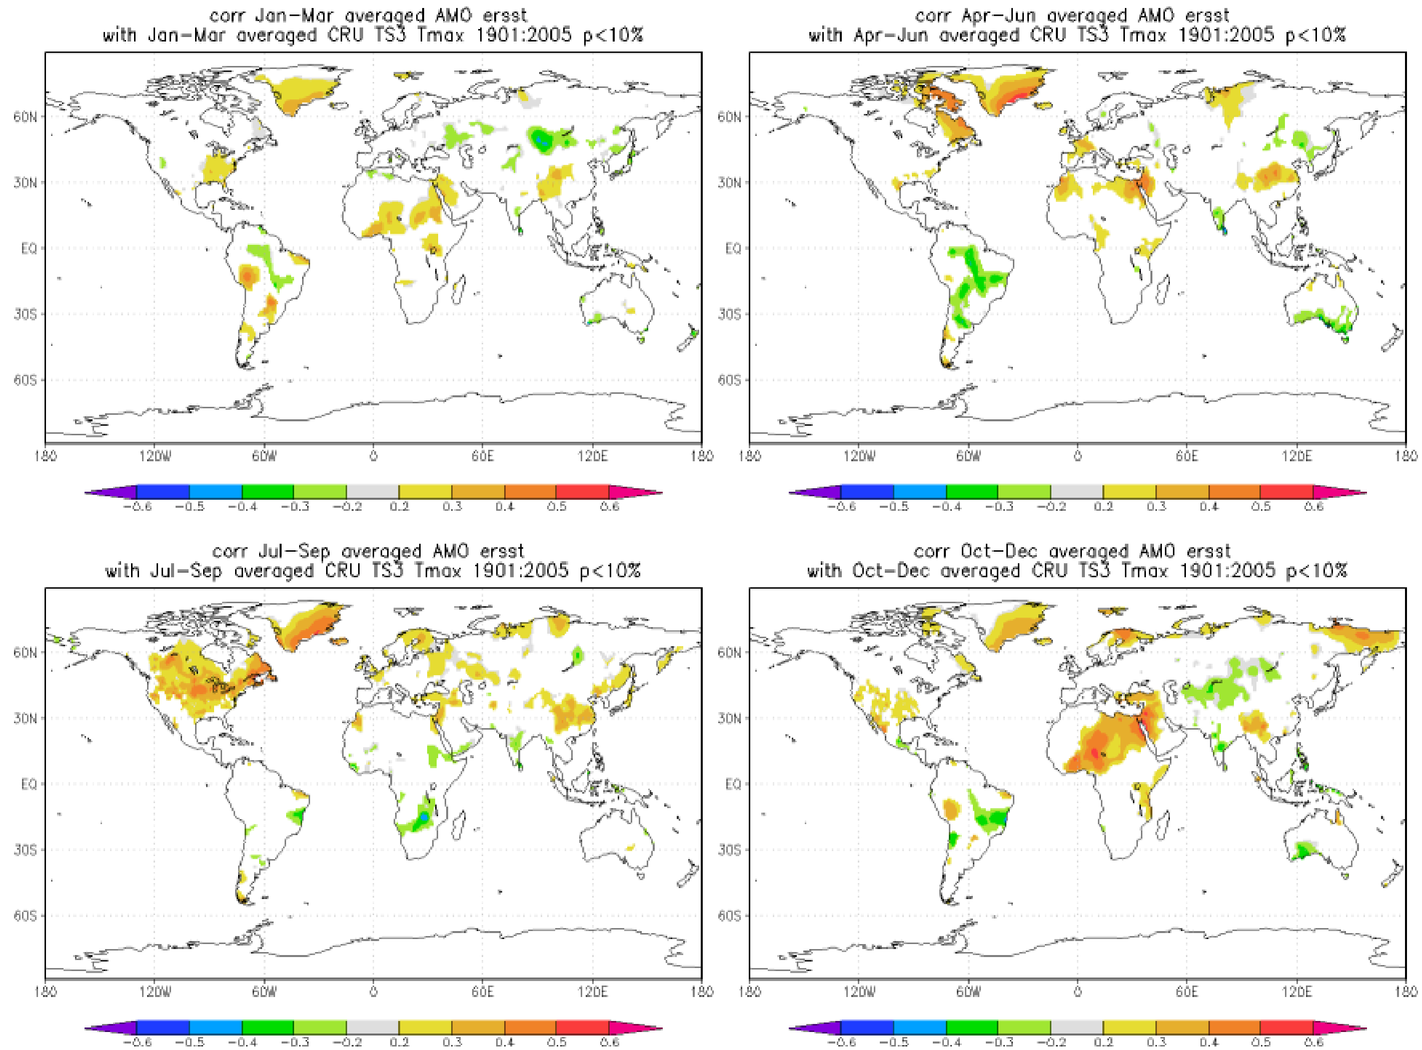

AMO correlations: Daily Minimum Temperature

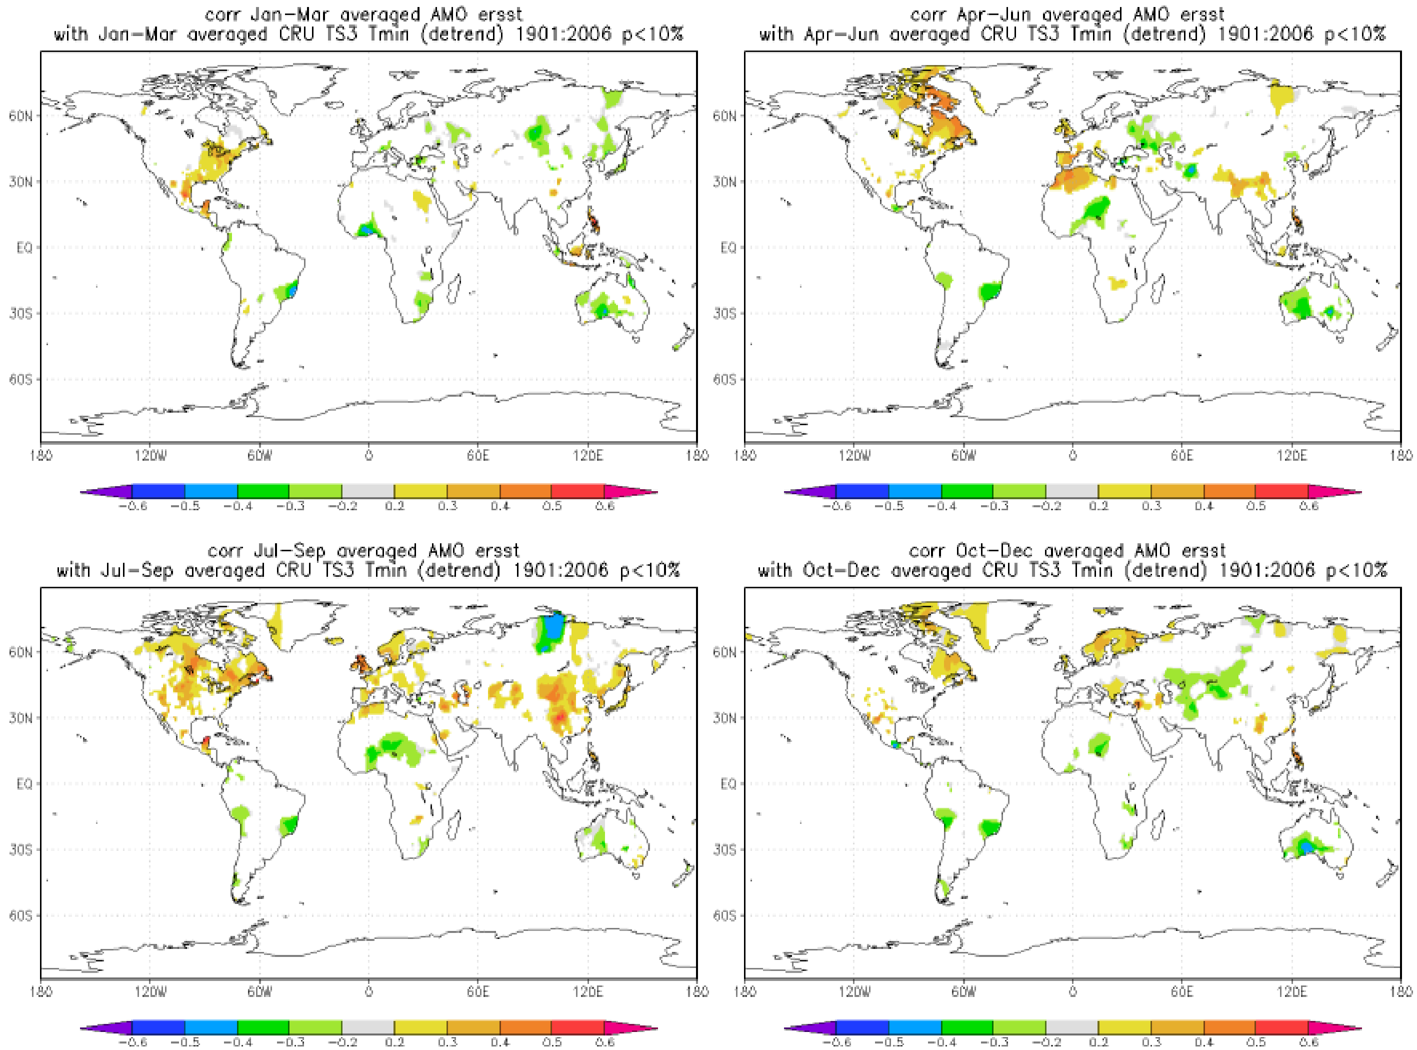

AMO correlations: Daily Precipitation

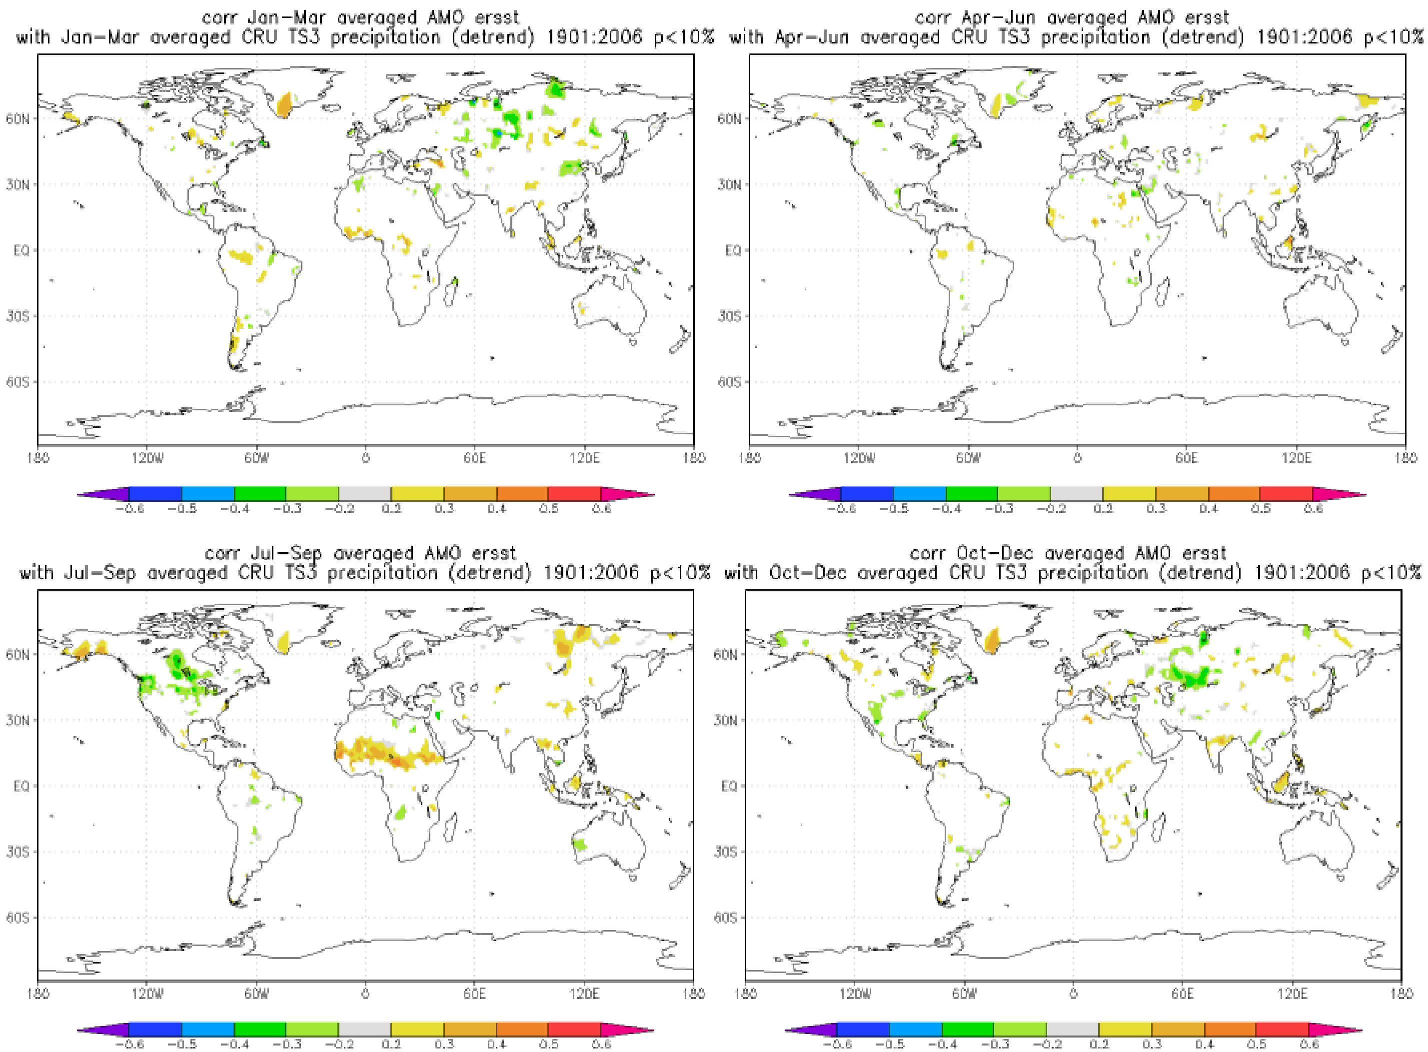

References

Chang, C. P., P. Harr, and J. Ju (2001), Possible roles of Atlantic circulations on the weakening Indian Monsoon Rainfall–ENSO relationship, J. Clim., 14, 2376– 2380.

DelSole, T., and J. Shukla, 2002: Linear prediction of Indian monsoon rainfall. *J. Climate,* **15,** 3645–3658.

Dugam, S.S., S.B. Kakade and R.K. Verma 1997. Interannual and long-term variability in the North Atlantic Oscillation and Indian summer monsoon rainfall. *Theor. Appl. Climatol.* **58**, 21-9.

Elsner, J. B., Bossak, B. H. & Niu, X.-F. 2001 Secular changes to the ENSO–U.S. hurricane relationship. Geophys. Res. Lett. 28, 4123–4126. (doi:10.1029/2001GL013669)

Enfield, D. B.; Mestas-Nunez, A. M.; Trimble, P. J. (2001). "The Atlantic Multidecadal Oscillation and its relationship to rainfall and river flows in the continental U.S.". Geophys. Res. Lett. 28: 2077–2080. doi:10.1029/2000GL012745

Gershunov, A., and T. P. Barnett, 1998: Interdecadal modulation of ENSO teleconnections. *Bull. Amer. Meteor. Soc.,* **79,** 2715–2725.

Goddard, L. and M. Dilley. 2005. El Nino: Catastrophe or opportunity. J. Climate, 18: 651-665.

Goldenberg, S. B.; et al. (2001). "The recent increase in Atlantic hurricane activity: Causes and implications". Science 293: 474–479. doi:10.1126/science.1060040.

Gong, D. Y., and C. H. Ho (2003), Arctic Oscillation signal in the East Asian summer monsoon, J. Geophys. Res., 108(D2), 4066, doi:10.1029/ 2002JD002193.

Gray, S. T.; et al. (2004). "A tree-ring based reconstruction of the Atlantic Multidecadal Oscillation since 1567 A.D.". Geophys. Res. Lett. 31: L12205. doi:10.1029/2004GL019932.

Hetzinger, Steffen; et al. (2008). "Caribbean coral tracks Atlantic Multidecadal Oscillation and past hurricane activity". Geology 36 (1): 11–14. doi:10.1130/G24321A.1.

Hurrell, J. W., Y. Kushnir, and M. Visbeck (2001), The North Atlantic Oscillation, Science, 291, 603– 605.

Kerr, R. A. (2000). "A North Atlantic climate pacemaker for the centuries". Science 288 (5473): 1984–1986. doi:10.1126/science.288.5473.1984.

Khalil, A. F., H.-H. Kwon, U.Lall, M. J. Miranda, and J. Skees, 2007: El Niño – Southern Oscillation – based index insurance for floods: Statistical risk analyses and application to Peru. Water Resources Res., 43, doi:10.1029/2006WR005281.

Knight, J. R.; C. K. Folland, and A. A. Scaife (2006). "Climate impacts of the Atlantic Multidecadal Oscillation". Geophys. Res. Lett.: L17706. doi:10.1029/2006GL026242.

Kushnir Y, RobinsonWA, Chang P, Robertson AW. 2006. The physical basis for predicting Atlantic sector seasonal-to-interannual climate variability. J. Climate 19: 5949–5970.

Latif, M., D. Anderson, T. Barnett, M. Cane, R. Kleeman, A. Leetmaa, J. O'Brien, A. Rosati and E. Schneider (1998): TOGA Review Paper: "Predictability and Prediction". J. Geophys. Res., 103, C7, 14,375-14,393.

Lima, C. H. R., U. Lall, T. Jabara, and A. G. Barnston. Statistical Prediction of ENSO from Subsurface Sea Temperature Using a Nonlinear Dimensionality Reduction. Journal of Climate 22:4501-4519, 2009.

Mason, S. J., and G. M. Mimmack, 2002: Comparison of some statistical methods of probabilistic forecasting of ENSO. Journal of Climate, 15, 8-29.

Mason, S. J., and L. Goddard, 2001: Probabilistic precipitation anomalies associated with ENSO. Bulletin of the American Meteorological Society, 82, 619-638.

Mantua, Nathan J.; *et al.* (1997). "A Pacific interdecadal climate oscillation with impacts on salmon production". *Bulletin of the American Meteorological Society* **78**: 1069–1079. doi:10.1175/1520-0477(1997)078<1069:APICOW>2.0.CO;2

McCabe, G. J., M.A.Palecki, and J.L. Betancourt (2004), Pacific and Atlantic Ocean influences on multidecadal drought frequency in the United States, *Proceedings of the National Academy ofSciences,* 101, 4136-4141

New, M., M. Hulme, and P. D. Jones, 1999: Representing twentieth century space–time climate variability. Part I: Development of a 1961–90 mean monthly terrestrial climatology. J. Climate, 12, 829–856.

New, M., M. Hulme, and P. D. Jones, 2000: Representing twentieth century space–time climate variability. Part II: Development of 1901–96 monthly grids of terrestrial surface climate. J. Climate, 13, 2217–2238.

Nigam, S., 2003:Teleconnections. Encyclopedia of Atmospheric Sciences; Editors: J.R. Holton, J.A. Pyle & J.A. Curry; Academic Press, Elsevier

Science, London, 2243-2269

Souza Filho, F. A., and U. Lall (2003), Seasonal to interannual ensemble streamflow forecasts for Ceara, Brazil: Applications of a multivariate, semiparametric algorithm, Water Resour. Res., 39(11), 1307, doi:10.1029/2002WR001373.

Wu, Z., B. Wang, J. Li, F. -F. Jin, 2009b: An empirical seasonal prediction model of the East Asian summer monsoon using ENSO and NAO. J. Geophys. Res., 114, doi:10.1029/2009JD011733