Climate and Agricultural Risk: Measuring the Effect of ENSO on U.S. Crop Insurance*

Jesse B. Tack†‡ David Ubilava§

Abstract

Predictive models of climatic phenomena can aid in insurance program-design and decision-making. Extreme weather outcomes have been linked to the El Niño Southern Oscillation (ENSO), which globally impacts agricultural production. This study demonstrates that extreme ENSO events alter cotton yield distributions in the Southeastern U.S. These impacts translate into economically meaningful effects on crop insurance premium rates. Commercial insurers can use publicly available information to determine if government-set premium rates are mis-priced, and in turn extract economic rents via the federally mandated Standard Reinsurance Agreement. By ceding underpriced policies in El Niño and La Niña years, we find that private insurance companies can reduce paid indemnities by 10-15 percent on average.

JEL Codes: G22; Q11; Q18; Q54

Keywords: Climate; Cotton; Insurance; El Niño Southern Oscillation; ENSO

*This is the peer reviewed version of the following article: Tack, J.B. and D. Ubilava, (2015), Climate and Agricultural Risk: Measuring the Effect of ENSO on U.S. Crop Insurance, *Agricultural Economics, 46*(2), pp. 245-257, which has been published in final form at http://onlinelibrary.wiley.com/doi/10.1111/agec.12154/abstract.

†E-Mail: tack@agecon.msstate.edu; Phone: +1 662 325 7999; Fax: +1 662 325 8777.

‡Assistant Professor in the Department of Agricultural Economics at Mississippi State University. P.O. Box 5187, Mississippi State, MS 39762, USA.

§Lecturer in the School of Economics at the University of Sydney. Room 370, Merewether Building, The University of Sydney, NSW 2006, Australia.

1 Introduction

Climate impacts agricultural production in a number of different ways, including droughts, excess rainfall, and more hazardous manifestations such as wildfires and hurricanes. Recent U.S. droughts have served as a reminder that extreme weather events do occur. Notably, these drought conditions materialized following a La Niña episode of the El Niño Southern Oscillation (ENSO) cycle – a linkage that has been extensively documented (Handler, 1990; Solow et al., 1998; Adams et al., 1999; Chen et al., 2002). Several studies suggest that extreme weather events are likely to occur more frequently in the future (e.g. Meehl et al., 2000; Rosenzweig et al., 2001). Others have conjectured that the frequency and intensity of ENSO anomalies will increase as a result of climatic change (Timmermann et al., 1999; Chen et al., 2008).[1] In the wake of improved predictability of the ENSO phenomenon, and a better understanding of its effects on agriculture, we ask in this article whether additional information due to climate anomalies can aid decision making in the context of crop insurance.

A better understanding of the intricacies surrounding ENSO dynamics have resulted in improved forecasts of this phenomenon (Kirtman and Schopf, 1998; Hall et al., 2001; An and Jin, 2004; Deng and Tang, 2009; Ubilava and Helmers, 2013). These efforts have substantially improved the predictive ability of ENSO forecasts, which in turn has created opportunities for crop producers to mitigate adverse climatic conditions and/or take advantage of more favorable ones (e.g. Meza et al., 2008). The value of these improved forecasts has been well recognized by researchers (e.g. Mjelde et al., 1996; Marshall et al., 1996; Hansen, 2002; Cabrera et al., 2007). For example, Solow et al. (1998) use simulation methods to analyze economic gains of improved ENSO predictions within the U.S. agricultural sector, and estimate the expected value of ENSO prediction to range from US$240 – US$323 million. Similarly, Jones et al. (2000) focus on the ENSO – crop yield linkage and estimate the economic value of ENSO forecasts in Georgia (U.S) as well as the Pampas region of Argentina. They find the per hectare value of optimal maize management using ENSO forecasts to be between

[1] Whether or not the more frequent extreme events are to become the new norm is the subject of an ongoing debate, and the literature on this topic is still evolving. For example, Rajagopalan et al. (1997) questions the statistical significance of the possible change in ENSO activities during the past several decades. The current consensus is that none of the competing hypotheses can be rejected at this time (Fedorov and Philander, 2000).

US$3 – US$5 in Georgia and between US$11 – US$35 in Pampas. In addition, several other studies report positive impacts of improved ENSO prediction in different regions of the world (e.g. Phillips et al., 1998; Messina et al., 1999; Adams et al., 2003; Stige et al., 2006).

While impacts of climate forecasts on production decision-making has been well studied, more work needs to be done to address the problem from the insurance standpoint (Carriquiry and Osgood, 2011). To date, there have been several attempts to tackle the issue. Ker and McGowan (2000) investigate adverse selection among commercial crop insurers. They hypothesize that a private insurer's information set nests the U.S. Risk Management Agency's (RMA) information set, due to the very nature of the timing of the decision making in relation to the signals associated with ENSO events. In addition, they note that while an individual farmer's information set is comparable to the private insurer's, the former may still behave sub-optimally due to bounded rationality and higher risk-aversion. Nadolnyak et al. (2008) examine the value of information from ENSO forecasts and consider cases of symmetric and asymmetric information by private insurers and farmers in selecting optimal coverage rates for insurance products. They find that crop yield distributions vary considerably across El Niño and La Niña regimes, and discuss implications regarding the actuarial soundness of the current insurance program.

In the U.S. a number of federal programs, coupled with commercial insurers' products, create a "safety net" for crop producers, particularly against extreme weather events. Crop insurance programs mitigate producers' exposure to downside risk by providing indemnity payments when yield or revenue outcomes fall bellow a contractually specified guarantee level. The largest provider of crop insurance in the U.S. is the Federal Crop Insurance Corporation (FCIC). This heavily subsidized program provided coverage on over 250 million acres and US$109 billion in liabilities in 2012, with the major program crops being corn (81 million acres), soybeans (65 million acres), wheat (46 million acres), and cotton (11 million acres).

The insurance products offered by the FCIC can be grouped into farm-based policies and area-based policies, and our focus here will be on the area-based Group Risk Plan (GRP). The main distinction between these groupings is that indemnities under farm-based policies are triggered by farm-level outcomes, while area-based policies offer insurance against more wide-spread losses as indemnities are triggered at the area level. While farm-based policies such as Yield and Revenue Protection have attracted much of farmers' interest historically, from a policy perspective there has been considerable traction towards area-based programs. Currently, expansion of area-based products is playing a key role in shaping the future of federal farm policy as the U.S. Congress has enacted legislation that includes new area-based insurance products, Stacked Income Protection (STAX) and Supplemental Coverage Option (SCO). Both designs are similar to current area products such as GRP in that they have the potential to mitigate adverse selection and moral hazard problems, as well as rate inaccuracies, that are associated with farm-based insurance products (e.g Barnett et al., 2005; Deng et al., 2007; Harri et al., 2011).

The GRP, also known as area-yield insurance, is based on average yields at the county level. It was introduced in 1993 and is available for several U.S. crops. Under GRP, indemnity payments are made when the county yield falls below a county-yield guarantee, which is independent of a particular producer's actual yield. The underlying assumptions of area-based crop insurance programs, and GRP in particular, are that yield distributions are "properly" determined (Goodwin and Ker, 1998), and that agents have no superior information about the determinants of this distribution (Harri et al., 2011). Under these assumptions the program is considered actuarially fair; however, when and if the assumptions are violated the rates may no longer be actuarially fair and possibilities of over- or under-charging emerge (Harri et al., 2011).

What factors can sabotage actuarially fair GRP rates, and how might ENSO events be linked to these deficiencies? Moreover, what opportunities could these effects create for crop insurers? Using the insight from Goodwin and Ker (1998) that area-specific idiosyncrasies such as floods and/or extreme hail likely affect localized yield distributions, ENSO's established linkage with hazardous weather conditions in turn suggests a linkage with the overall shape of these distributions.[2] Furthermore, previous research has shown that the spatial correlation of crop yields strengthens during extreme weather events (Goodwin, 2001), which implies that ENSO's distributional impacts are likely spatially correlated as well. Thus, premium rates that ignore ENSO-driven impacts may

[2]ENSO has been shown to induce a wide range of extreme weather outcomes including damaging storms, drought, and flooding (Handler, 1990; Bove et al., 1998; Saunders et al., 2000); as well as wildfires (Swetnam and Betancourt, 1990; Brenner, 1991; Legler et al., 1999)

not reflect “true” yield distributions, and the spatial dependence of these rating inefficiencies suggests an overall program vulnerability similar in spirit to the “formula that felled Wall St” criticism that has been (at least) partially linked to the 2008-2009 global financial crisis (Jones, 2009).

The chronology of events surrounding crop insurance rate-setting and producer purchase-decisions creates an interesting economic dynamic. First, RMA is legally required to set premium rates well in advance of the planting period, a requirement that coincidentally also requires these rates to be set prior to the establishment of a particular ENSO regime (e.g. El Niño or La Niña).[3] Next, ENSO conditions establish well in advance of the deadline for crop insurance enrollments, and perhaps more importantly, before the deadline for private crop insurance companies to decide whether to retain the policies they have sold or cede some of them back to the FCIC.[4] Thus, it is quite reasonable to presume that private crop insurers can exploit this ENSO-driven information asymmetry to identify pricing inaccuracies in the RMA premium rates, and capitalize on this advantage in an economically meaningful way.[5]

Given these insights, the focus of this research is twofold. First, we utilize a historical dataset of yield and weather outcomes to evaluate the extent to which ENSO events perturb the “true” distribution of crop yields. Second, we analyze whether these perturbations provide an opportunity for private insurance companies to accrue economic rents through their reinsurance decisions. This study uses cotton as a crop of interest and the southeastern tier of the U.S. as a region of interest to illustrate that economic impacts of publicly available ENSO information on the crop insurance industry are indeed significant. The reasons for this focus is that cotton is one of the major crops produced in the southeastern tier of the U.S., a region that has been an attractive venue for analyzing ENSO impacts on agriculture due to the numerous hazardous weather conditions it may cause (Hansen et al., 1998; Jones, 1999; Solís and Letson, 2013).

[3]ENSO conditions start establishing during the northern hemisphere Summer – Fall, and become most apparent during the month of December.

[4]A nice overview of the FCIC Reinsurance Agreement can be found at http://www.rma.usda.gov/pubs/ra/, which notes that the Standard Reinsurance Agreement (SRA) and the Livestock Price Reinsurance Agreement (LPRA) are cooperative financial assistance agreements between the Federal Crop Insurance Corporation (FCIC) and an insurance company. Both agreements establish the terms under which FCIC provides reinsurance and subsidies on eligible crop insurance contracts sold by the insurance company.

[5]See also Ker and McGowan (2000) for the conceptual representations in the context of insurance, reinsurance and adverse selection.

2 Empirical Framework

In order to evaluate the effects of ENSO on both yields and crop insurance rates, we require a method for linking historical weather outcomes to the entire distribution of cotton yields. Goodwin (2008) cautions against the use of nonparametric density estimators in the current context where their lack of structure can be problematic given the limited historical data available for measuring ENSO events. Thus, we forgo the nonparametric approaches of Ker and McGowan (2000) and Nadolnyak et al. (2008), and instead follow recent research evaluating the effects of climate change on cotton yields that proposes using a moment based maximum entropy (MBME) model (Tack et al., 2012). This approach has also been used by Tack and Ubilava (2013) to condition corn yield distributions on ENSO events. Since this is a relatively new approach, we provide a brief description here that follows Tack and Ubilava (2013).

The MBME model has two components, the first of which is a data-based regression framework used to predict moments of the yield distribution under three ENSO regimes: La Niña, La Nada[6], and El Niño (we discuss how we distinguish between these three regimes in the Data section below). The second component utilizes these predicted moments within a maximum entropy framework to identify how the El Niño and La Niña regimes perturb the distribution of yields relative to the La Nada regime. Importantly, conditioning these distributions on the ENSO regimes allows use to measure their effect on actuarially fair crop insurance premium rates.

As in Tack et al. (2012), we focus on the first three raw moments of the cotton yield distribution as mean, variance, and skewness properties have been the focus of much of the yield modeling literature (e.g. Day, 1965; Gallagher, 1987; Nelson and Preckel, 1989; Moss and Shonkwiler, 1993; Goodwin and Ker, 1998; Ker and Coble, 2003; Ramirez et al., 2003; Sherrick et al., 2004; Hennessy, 2009a,b). The moments model for the $j = 1, 2, 3$ raw moments of the cotton yield distribution is:

$$\begin{aligned} y_{it}^{j} &= \alpha_{ij} + \beta_{j1} low_{ijt} + \beta_{j2} med_{ijt} + \beta_{j3} high_{ijt} + \beta_{j4} prec_{it} + \beta_{j5} prec_{it}^{2} \\ &+ \beta_{j6} nino_{t} + \beta_{j7} nina_{t} + \beta_{j8} trend_{t} + \varepsilon_{ijt} \end{aligned} \quad (1)$$

[6] La Nada is a relatively new jargon in the climate literature, as compared to the terms El Niño and La Niña, and is used to denote *normal* or *neutral* conditions.

where the dependent variable y_{it}^{j} is the j^{th} power of the yield variable for county i in period t.

We include a county-by-equation fixed effect α_{ij}, a quadratic effect for precipitation, and two dummy variables, $nino_t$ and $nina_t$, that respectively take a value of one when an El Niño or La Niña event is experienced, and zero otherwise.[7] The low, medium, and high temperature variables are measured as degree day intervals.[8] For example, low might capture degree days between, say, 0 and $10°C$, which is just the number of degree days above $0°C$ minus the number of degree days above $10°C$. Using these degree day intervals is equivalent to the piecewise linear specification of Schlenker and Roberts (2009), which compared favorably against other nonlinear representations of temperature effects. Generalizing the approaches of Tack et al. (2012) and Tack and Ubilava (2013), we allow the cut-points that define these intervals to vary across equations. This is an important generalization as it allows different temperature thresholds to exist for each of the included moments, and is discussed more fully in the Data section below. Finally, we include an equation specific linear trend to control for technological change over time. Under the assumption $E(\varepsilon_{ijt}) = 0$, equation (1) provides a convenient framework for conditioning yield moments in the spirit of Antle (1983) and Just and Pope (1978), and the parameters can be consistently estimated using Ordinary Least Squares.

This moments model provides a mechanism for predicting moments, however the entire density of the yield distribution is needed to calculate premium rates. While the moments alone cannot determine this density (Shohat and Tamarkin, 1943), previous work in the yield modeling literature has demonstrated how this problem can be ameliorated using the concept of maximum entropy (Stohs and LaFrance, 2004; Tack et al., 2012). Define by $\mu_{ij} \equiv E\left[Y_i^j\right]$ the j^{th} raw moment of the yield variable for county i. Also define by $\mathbf{X}_i$ a county specific random vector for the right-hand side variables of equation (1), and define the outcomes $\mathbf{x}_i^{nino}$, $\mathbf{x}_i^{nina}$, and $\mathbf{x}_i^{nada}$ as values for the right-hand side variables that represent the El Niño, La Niña, and La Nada regimes. The conditional

[7] While Deschenes and Greenstone (2007) suggest amending this fixed effects approach to include state-by-year effects, Fisher et al. (2012) caution against such an approach. They note that after removal of both county and state-by-year fixed effects, remaining weather variance pertains only to yearly within-state deviations from county means. Thus, we follow the approach of Schlenker and Roberts (2009) and include county fixed effects.

[8] As discussed in the appendix to Schlenker and Roberts (2009), degree days are a special case of time-separable growth, typically defined as the sum of truncated degrees between two bounds. Using their example with bounds of $8°C$ and $32°C$, a day of $9°C$ contributes 1 degree day, a day of $10°C$ contributes 2 degree days, up to a temperature of $32°C$, which contributes 24 degree days. All temperatures above $32°C$ also contribute 24 degree days.

raw moments are defined as $\mu_{ij}^{r} \equiv E\left[Y_i^j \mid \mathbf{X}_i = \mathbf{x}_i^r\right]$ for each regime $r \in \{nino, nina, nada\}$.

For an arbitrary county i and regime r, the maximum entropy distribution is defined by:

$$f_{ir}^{*} = \arg\max_{f} - \int f(y) \ln f(y)\, dy \tag{2}$$

subject to the moment constraints

$$\int f(y)\, dy = 1 \text{ and } \int y^j f(y)\, dy = \mu_{ij}^{r}, j = 1, \ldots, 3. \tag{3}$$

The associated Lagrangian for this maximization problem is:

$$\mathcal{L} = - \int f(y) \ln f(y)\, dy - \gamma_0 \left[\int f(y)\, dy - 1\right] - \sum\nolimits_{j=1}^{3} \gamma_j \left[\int y^j f(y)\, dy - \mu_{ij}^{r}\right] \tag{4}$$

and the implied solution is the maximum entropy density:

$$f_{ir}^{*}(y) = \frac{1}{\psi(\gamma_{ir}^{*})} \exp\left[- \sum\nolimits_{j=1}^{3} \gamma_{ijr}^{*} y^j\right] \tag{5}$$

where

$$\psi(\gamma_{ir}^{*}) = \int \exp\left[- \sum\nolimits_{j=1}^{3} \gamma_{ijr}^{*} y^j\right] dy.$$

The parameter vector γ_{ir}^{*} represents the solution to the maximization problem and $\psi(\gamma_{ir}^{*})$ is the normalizing factor that insures the density integrates to unity. A key benefit of using maximum entropy in this context is that the moments of the estimated densities are guaranteed to match the predicted moments obtained from equation (1). Furthermore, a straightforward transformation of the parameters γ_{ir}^{*} demonstrates that using the $j = 1, ..., 3$ moments as constraints generates densities of the form:

$$f(y) = \exp\left[\delta_0 + \sum_{j=1}^{3} \delta_j y^j\right]. \tag{6}$$

This will be the exact density for y if it is a member of the exponential family and the sample counterparts to the moments μ_{ij}^{r} are sufficient statistics. Since we do not know the true density in practice, we are more interested in whether the exponential function in equation (6) closely approximates the true (unknown) distribution. Taking the natural log of both sides we obtain

$$\ln\left[f\left(y\right)\right] = \delta_0 + \sum_{j=1}^{3} \delta_j y^j \tag{7}$$

which shows that the maximum entropy distribution can be thought of as a 3^{rd} order approximation of the log density.[9]

For a more general case in which an arbitrary J moment constraints are included, the corresponding maximum entropy distribution corresponds to a J^{th} order approximation of the true, unknown density function. Therefore, the flexibility of this modeling approach increases with the number of moment constraints. We focus on the first three moments ($J = 3$) as their importance has been repeatedly emphasized in the yield modeling literature.[10] We follow Tack et al. (2012) and utilize the sequential updating method developed in Wu (2003) to estimate the maximum entropy distribution for each county-regime combination.[11] The next section presents the data used to estimate these distributions, and then we present our empirical findings.

3 Data

We use the same weather and sea surface temperature (SST) anomaly data as Tack and Ubilava (2013), and augment it with cotton yield data from the National Agricultural Statistics Service (NASS) to construct a panel that spans 1968-2005. The time dimension is determined by the following limitations: NASS county-level cotton yield data for Texas (the largest producing state) does not extend beyond 1968, and the weather data ends in 2005. We measure yield as production divided by harvested acres in 10 lb units per acre, and focus on the major cotton producing states: Alabama, Arkansas, Georgia, Louisiana, Mississippi, North Carolina, Tennessee, and Texas. We include all 224 counties with a complete 38 year yield history.

[9]These log density approximations have a long history in the statistics literature. If the number of included moment conditions increases with sample size at an appropriate rate, then a consistent nonparametric estimator of the underlying density can be obtained. See Wu (2010) and references therein for further discussion.

[10]The appropriate J at which one obtains sufficient flexibility for approximating the true density is an open question, and thus a ripe area for future research in the crop insurance literature.

[11]Matlab code for this approach is available on Wu's web page at http://agecon2.tamu.edu/people/faculty/wu-ximing/

We use a monthly time series of the ENSO anomaly, *Niño 3.4*, derived from the index tabulated by the Climate Prediction Center at the National Oceanic and Atmospheric Administration. This index is based on SST anomalies in the area of the Pacific Ocean between $5°N - 5°S$ and $170°W - 120°W$, and is a strong indicator of ENSO occurrence. The anomaly in a particular month is defined as the deviation from the average historic measure (during the 1981 – 2010 base period) for that month in the *Niño 3.4* region.

To measure the occurrence of an El Niño and La Niña event, we construct dummy variables based on the monthly *Niño 3.4* anomaly data. The construction of these variables is not straightforward as the length of time for teleconnections to reach the U.S. is not well understood.[12] As noted in Ker and McGowan (2000), the climatology literature has varying definitions of what constitutes El Niño and La Niña events. Thus, we follow Tack and Ubilava (2013) and consider four well recognized methods for determining El Niño and La Niña occurrences. The first method uses the current period growing season (April–October) average anomaly as in Cane et al. (1994). The second method uses the preceding May-December average as in Hsiang et al. (2011), while the third is a slight variant that uses the November-December-January average as in Phillips et al. (1998). The final method considered is based on the Japan Meteorological Agency approach that utilizes a five month rolling average across the preceding May to the current period April. We chose the construction that yields the highest R^2 for the yield regressions, which is the preceding May-December average. As a result, there are ten growing seasons that fall into the El Niño regime (1970, 1973, 1983, 1987, 1988, 1992, 1995, 1998, 2003, 2005), another ten that fall into the La Niña regime (1971, 1972, 1974, 1975, 1976, 1985, 1989, 1999, 2000, and 2001), and the remaining eighteen years are considered the La Nada Regime.

This study uses the weather data of Schlenker and Roberts (2009), which spans 1950-2005 and is based on the rectangular grid system underlying PRISM covering the contiguous United States.[13] The authors construct daily temperature distributions from observed minimum and maximum

[12] Findings from Ropelewski and Halpert (1986) and Chiang and Sobel (2002) suggest very weak, if any, teleconnections between SST in the tropical Pacific and local weather in the U.S., however what has been found has generally occurred after a December ENSO peak.

[13] PRISM stands for the Parameter-elevation Relationships on Independent Slopes Model. The data from Schlenker and Roberts (2009) has been updated to more recent years, but is not yet publicly available.

temperatures, which is used to estimate time-spent in all one-degree temperature intervals from $-5°C$ and $50°C$. Daily time intervals are aggregated across PRISM grid cells up to the county level using acreage-weights, and then summed over the seven month (April–October) growing season.[14]

As mentioned in the Empirical Framework section above, we allow the low, medium and high temperature intervals to vary across equations. The low temperature measure is constructed as the number of degree days above $0°C$ minus the number of degree days above some cut point x_1, the medium measure is constructed in the same way but with bounds x_1 and $x_2 > x_1$, and the high measure is the number of degree days above x_2. For each of the regression equations $j = 1, 2, 3$, we consider the set of all cut-points $\{(x_1, x_2) : x_1 \in [5, 20], x_2 \in [x_1 + 10, 30]\}$ and select the optimal values (x_1^*, x_2^*) as those with the highest R^2. Interestingly, the optimal cut-points $(7, 29)$, $(7, 28)$, and $(7, 27)$ are nearly invariant across equations.

Precipitation is measured in centimeters, and is aggregated across the growing season in the same way as the temperature variables. Descriptive statistics for the data are reported in the online Supplementary Table S1. The first set of statistics correspond to the entire data, and the remaining sections correspond to the La Nada, El Niño, and La Niña regimes. The sample means suggest that El Niño is mean enhancing relative to the La Nada regime, while La Niña is mean reducing. Given that previous studies linking this weather data to yield outcomes (Tack et al., 2012; Schlenker and Roberts, 2006, 2009) have found that high temperatures have a strong negative influence on yields, this suggests that El Niño (La Niña) should be associated with a reduction (increase) in high temperatures; however, this relationship is not supported by the sample data as La Niña is associated with a nearly 10 percent reduction in high temperature exposure. In addition, a change in precipitation is likely not a major driver of the La Niña/La Nada discrepancy as it is nearly identical across these two regimes. Thus, our inclusion of the ENSO event indicators seems particularly important in this analysis as there are likely other drivers beyond temperature and precipitation involved.

[14] A more detailed description of the weather data can be found in the appendix to Schlenker and Roberts (2009), which is publicly available at: http://www.pnas.org/content/106/37/15594/suppl/DCSupplemental.

4 Results

The first subsection presents and discusses the results for the regression based estimation of the moments given by equation (1) and discusses the mean effects. The second subsection presents the GRP premium rate changes under alternative regimes, and the third subsection discusses the simulation exercise we conduct to demonstrate the economic significance of these rate changes.

4.1 Estimation of Moments and Mean Effects

An important concern for our empirical framework is if the weather, ENSO, and trend parameters are spatially invariant across states. We formally test this by estimating a version of equation (1) that allows for state specific parameters:

$$\begin{aligned} y^j_{ist} &= \alpha_{ij} + \beta_{js1} low_{ijt} + \beta_{js2} med_{ijt} + \beta_{js3} high_{ijt} + \beta_{js4} prec_{it} + \beta_{js5} prec^2_{it} \\ &+ \beta_{js6} nino_t + \beta_{js7} nina_t + \beta_{js8} trend_t + \varepsilon_{ijt} \end{aligned} \tag{8}$$

where the dependent variable is now the j^{th} power of yield for county i in state s during period t. For each equation $j = 1, 2, 3$, we conduct an *F-test* of the null hypothesis that the parameter vectors defined by $\beta_{js} = [\beta_{js1}, ..., \beta_{js8}]'$ are common across states. The *p-values* for all three equations are less than 0.001, suggesting that the model parameters should not be held constant across states. Thus, we estimate the parameters of the moment based model for each state separately as in equation (8).

The r-squared values for the three moment equations average 0.62, 0.59, and 0.54 across states, and thus suggest that our model provides a reasonable level of fit for the data. One might also consider including additional weather variables to increase the explanatory power of the model, such as those used in Vedenov and Barnett (2004). The authors note that adding additional - and perhaps more refined - weather measures could improve the explanatory power of the yield-weather model, however this comes at the cost of a more complicated model. As our r-squared values for the mean equation are similar to those for cotton in Vedenov and Barnett (2004), we proceed with the current specification.

We also evaluate whether alternative approaches in the spirit of Ker and McGowan (2000) and Nadolnyak et al. (2008) should be considered. The former argue in favor of using the underlying SST data over more coarsely defined "phase-based" dummy variables, while the latter stratify the sample data by ENSO events and estimate separate models for each. To investigate whether these alternatives should be employed here, we consider six different alternative models and use an out-of-sample forecasting exercise to compare them to our preferred specification, equation (8). The first four alternatives omit the weather and phased-based dummy variables and replace them with functions of the underlying May-December average SST anomaly. We consider linear, quadratic, cubic, and quartic functions of the SST anomaly (AM1-AM4) to best capture potential nonlinearities. The final two alternatives stratify the sample by the El Niño and La Niña years defined in the Data section above, and estimate separate models for each of the three regimes (El Niño, La Niña, and Neutral). We include two specifications based on equation (8) for this stratified approach, the first restricts the model to include only county fixed effects and a trend (AM5) while the second uses the specification of equation (8) exactly (AM6).

Supplementary Table S2 reports the root mean squared errors (RMSE) from the out-of-sample forecasting exercise. The forecast errors are obtained as follows: we omit an entire year from the data, estimate the parameters of the model, and then predict yields for the omitted year. The difference between the predicted yields and the actual yield realization in the omitted year generates the forecast error. We iterate over all years in the data and report the square root of the average squared forecast error for our model as well as the six alternatives. Although Ker and McGowan (2000) cite evidence suggesting the use of SST anomaly data over phase-based approaches, our approach - which combines a phased-based approach with actual weather outcomes - is preferred in the current context. In addition, our results caution against a stratified-sample approach similar to Nadolnyak et al. (2008). Noting that AM6 is the same specification as our preferred model but with the parameters allowed to vary across ENSO phases, we see that this additional flexibility introduces substantial prediction error. This outcome is consistent with the general understanding in the forecasting literature that the more parsimonious models can outperform the more flexible counterparts in an out-of-sample setting, because of exposure to parameter uncertainty in the latter models (e.g. Stock, 2001). Overall, we find that the RMSE is lowest for our preferred model in every state, thus lending support for our specification compared to previous approaches.

Supplementary Figures S1-S3 report 95 percent confidence intervals for the parameter estimates of the moment equations for the eight states using OLS with robust standard errors clustered by year. As in Schlenker and Roberts (2009), exposure to low and medium temperatures have relatively minor effects on yields relative high temperatures. As in Tack et al. (2012), these high temperature effects are statistically significant for all states and moment equations. Although not reported in these figures, the quadratic effect of precipitation is an important determinant of yields as the *p-values* associated with *F-tests* of the null hypotheses $H_o : \beta_{js4} = \beta_{js5} = 0$ are less than 0.05 for nearly all of the state-equation combinations.

As discussed in the data section above, the inclusion of the El Niño and La Niña dummy variables are important control variables for ENSO effects beyond temperature and precipitation in some states. The *p-values* associated with *F-tests* of the null hypotheses $H_o : \beta_{js6} = \beta_{js7} = 0$ are less than 0.05 for nearly half of the state-equation combinations, and the parameter estimates for the La Niña coefficients are statistically significant at the 5 percent level for over half of the state-equation combinations. The absolute value of the La Niña coefficient estimates are four times larger than those for El Niño on average, implying that controlling for additional impact vectors is especially important for La Niña.

The impact of El Niño (La Niña) on mean yields is measured as the percentage change in the mean of the El Niño (La Niña) regime relative to the mean of the La Nada regime. We use the estimated coefficients from the first moment equation and calculate mean impacts for each county i according to:

$$impact_i^r = 100\frac{\left(\hat{\alpha}_{i1} + \hat{\beta}_{1s}\bar{\mathbf{x}}_i^r\right) - \left(\hat{\alpha}_{i1} + \hat{\beta}_{1s}\bar{\mathbf{x}}_i^{nada}\right)}{\hat{\alpha}_{i1} + \hat{\beta}_{1s}\bar{\mathbf{x}}_i^{nada}}, \quad r \in \{nino, nina\} \tag{9}$$

where the $\bar{\mathbf{x}}_i^r$ are county-specific predictors under each regime. For the El Niño regime, we fix the county-level temperature and precipitation variables at their sample average for the El Niño years, and the El Niño and La Niña dummy variables are set to one and zero respectively. We do the same for the La Niña regime using the La Niña years, but fix the El Niño and La Niña variables to zero and one respectively. For the La Nada regime, we use the La Nada years to construct county-level sample averages and fix both dummy variables to zero.

Figure 1 reports kernel density plots of the county-level El Niño and La Niña impacts for all counties in the data (Cotton Belt), counties in Texas, Louisiana, Arkansas, and Mississippi (Western Cotton Belt), and counties in Alabama, Tennessee, Georgia, and North Carolina. The acreage-weighted average impacts for the Cotton Belt are mean enhancing for El Niño at 2.01% and mean reducing for La Niña at −5.87%, however these averages mask a considerable amount of heterogeneity across counties as the density plots indicate. It is interesting to note that this heterogeneity is primarily driven by the counties in the Eastern region, as the El Niño (La Niña) impacts are positive (negative) for 87% (96%) of the Western Corn Belt counties. The acreage-weighted average impacts in the West are 2.02% for El Niño and −7.09% for La Niña, while these impacts are 1.93% and −0.118% respectively for the East. The similarity of the Western impacts to overall averages is driven by the relatively large concentration of the sample acreage in the Western region, roughly 75%.

4.2 GRP Insurance Rate Effects

Following the procedure outlined in the Empirical Framework section, we construct three maximum entropy distributions for each county. The parameter estimates from equation (8) are used to predict moments as:

$$\hat{\mu}^r_{isj} = \hat{\alpha}_{ij} + \hat{\beta}_{js}\bar{\mathbf{x}}^r_{ij}, \quad i = 1, \ldots, 224, \quad s = 1, .., 8, \quad j = 1, 2, 3, \quad r \in \{nino, nina, nada\} \tag{10}$$

where $\bar{\mathbf{x}}^r_{ij}$ are county- and equation-specific predictors under each regime as discussed above. These moments are then used to estimate the maximum entropy density functions $\widehat{f^r_i(y)}$.

In this empirical application, we were able to obtain densities for 199 of the 224 counties in the sample. This is because the predicted moments for some of the regime-county combinations implied non-positive variance, i.e. $\hat{\mu}^r_{is2} - [\hat{\mu}^r_{is1}]^2 \leq 0$, in which case the maximum entropy density cannot be estimated. If this occurred for any county-regime combination, we removed the county from the sample. These were mostly smaller-acreage counties as the total acreage represented in the sample was only reduced by 8.7 percent. More generally, this suggests that the MBME model as formulated in Tack et al. (2012) could be prone to the this type of empirical irregularity.

There are far too many estimated densities to report here, so we include the regime-specific densities for the largest producing counties within each state in Supplementary Figures S4 and S5. Given estimated densities for all county-regime pairs, expected indemnities across alternative coverage levels c are calculated as:

$$E(indem_{ic}^{r}) = \int_{0}^{\widetilde{y}_r} \frac{1}{c} \left(\widetilde{y}_r - y\right) \widehat{f_i^r\left(y\right)} dy \tag{11}$$

where $\widetilde{y}_r \equiv \hat{\mu}_{is1}^{r} \times c$ is the yield guarantee and $1/c$ is the "disappearing deductible" factor that is built into GRP contracts (Barnett et al., 2005). The associated GRP premium rate is then just the expected indemnity over the liability:

$$rate_{ic}^{r} = \frac{E(indem_{ic}^{r})}{\widetilde{y}_r}. \tag{12}$$

The acreage-weighted county-averaged GRP rates are reported in Figure 2. Within each column the top row graphs the GRP rates across coverage levels ranging from 50 to 90 percent, while the lower panel reports corresponding pair plots constructed as the ratio $rate^{nino}/rate^{nada}$ and $rate^{nina}/rate^{nada}$. For the entire sample (first column), El Niño generates a reduction in the GRP rate while La Niña generates an increase across all coverage levels. The pair plot ratio for El Niño ranges from 0.92 to 0.96, which correspond to rate reductions between 4 and 8 percent respectively. The rate impacts for La Niña are much larger as the ratios range from 1.4 to 1.1, which correspond to rate increases between 10 and 40 percent respectively. As with the mean effects, these averages mask a considerable amount of heterogeneity across regions as indicated by the findings for the Eastern region (third column). As opposed to the Western region, the El Niño and La Niña effects are nearly identical as the pair plots range from 0.69 to 0.89 for El Niño and from 0.70 to 0.88 for La Niña. These values correspond to rate decreases between 12 and 31 percent in the Eastern region.

In general it is difficult to determine the drivers of the GRP rate differences across regions as the impact vectors can vary widely across relatively small regions within the Southern U.S. However, we can attempt to evaluate the relative importance of the weather and temperature variables included in this analysis using an analysis of variance. First, for each county-coverage level combination, ic, we construct regime-differenced variables according to:

$$\Delta y_{ic}^{r} \equiv rate_{ic}^{r} - rate_{ic}^{nada} \text{ and } \Delta \mathbf{x}_{i}^{r} \equiv \bar{\mathbf{x}}_{i}^{r} - \bar{\mathbf{x}}_{i}^{nada}, \quad r \in \{nino, nina\} \tag{13}$$

Next, we pool these observations together and fit a regression model of the form:

$$\Delta Y_i = \alpha_c + \Delta \mathbf{X}_i' \beta_s + \varepsilon_i \tag{14}$$

where α_c are coverage level fixed effects and the β_s vary by state as in equation (8). The baseline model with β_s set to zero generates an R^2 of 0.001, so the coverage level fixed effects alone do not explain much of the variation in rate differences. Interestingly, we next add just the high temperature differenced variable and the R^2 improves substantially to 0.51, implying that differences in high temperatures alone account for about 50 percent of the variation in rate differences. The inclusion of the other temperature variables improves the R^2 to 0.67, and subsequently adding in precipitation improves it to 0.71. Finally, including all weather variables as well as a dummy for El Niño generates an R^2 of 0.76. Based on this analysis, we conclude that changes in exposure to extreme heat is an important driver of GRP rate differences across ENSO regimes, but does not explain all variation.

4.3 Economic Significance of Rate Effects

The previous two subsections suggest that El Niño and La Niña alter the mean of the cotton yield distribution as well as the implied actuarially fair GRP premium rates. In this subsection we demonstrate that these effects are economically significant by conducting a repeated game of insurance selection similar to Harri et al. (2011). This game is especially important in this context because of the timing of the extreme ENSO events. Recall from the Data section above that the method for determining an El Niño or La Niña year is based on information from the preceding

May-December. This has two important implications for the crop insurance industry. First, the RMA *cannot* make use of this information because they are required by law to set rates *before* the previous December. Second, private crop insurance companies *can* make use of this information as the deadline for reinsurance decisions occurs *after* the previous December. Essentially, there exists an information asymmetry between the RMA and the insurance companies, and these companies can exploit this asymmetry to extract economic rents when and if they believe contracts are mis-priced.

Due to the infrequency of extreme ENSO events and the relatively small time dimension of our data, we cannot follow Harri et al. (2011) explicitly in that we cannot simulate outcomes for a decision rule based on out-of-sample forecasting. Thus, we develop an in-sample counterpart to their repeated insurance game, which essentially measures the historical rents that could have been extracted if the insurance companies based their reinsurance decision on the asymmetric ENSO event information.

The repeated game is characterized as follows. The entire sample is used to derive actuarially fair premium rates for each county while ignoring the influence of both El Niño and La Niña. In practice, this is accomplished by fixing the temperature and precipitation variables at their overall sample average and setting the El Niño and La Niña dummy variables to zero in equation (10). This can be thought of as the "ignorant" historical average GRP rate for each county, and is denoted $rate_{ic}^{ig}$. Next, we assume the role of a private insurance company and compare these rates to their El Niño and La Niña counterparts from the previous section, $rate_{ic}^{nino}$ and $rate_{ic}^{nina}$, and abide by the following time-invariant decision rule: if $rate_{ic}^{ig}$ is lower than the El Niño (La Niña) rate then we cede the policy in every El Niño (La Niña) year as we believe them to be under-priced. Conversely, we retain policies where the ignorant rate is higher as we believe they are overpriced.

To evaluate the performance of the insurance company's decision rule, we calculate actual indemnities for all county-year observations in the data using the realized yield observations y_{it}. That is, for each of the $7,562$ county-year combinations it we calculate:

$$indem_{itc} = \max\left\{\frac{\widetilde{y}_{it} - y_{it}}{c}, 0\right\}, \quad t = 1, ..., 38, \quad c \in \{50, 60, 70, 80, 90\} \tag{15}$$

where the yield guarantee $\widetilde{y}_{it} \equiv \widehat{y}_{it} \times c$ is constructed using the fitted value $\widehat{y}_{it}$ from the trend regression:

$$y_{it} = \alpha_i + \beta trend_{it} + \varepsilon_{it} \tag{16}$$

and the trend parameter, β, is allowed to vary across states as in equation (8).

The results of the insurance selection game are summarized in Table 1. Each cell reports the acreage-weighted average indemnity across all county-year pairs for the various coverage levels. The All Policies column represents the status quo as it reports the average indemnity across all policies, essentially assuming that all policies are retained. Conversely, the Retained Policies column allows insurance companies to cede policies based on the above decision rule, and reports average indemnities of the policies they retain. Looking at the first two columns, it is clear that insurance companies are extracting large economic rents as indemnities for the retained policies are roughly 10-15 percent lower than their all-policy counterparts. These rents accrue relatively equally across El Niño and La Niña years as the indemnities are roughly 10-65 percent lower for each set of years. Interestingly, the largest indemnity reductions occur at low coverage levels, suggesting that the proposed decision rule does a relatively better job at shielding private insurers against indemnities triggered by deep versus shallow losses.

It is worth noting that 24.6 percent of the counties do not cede any policies, and these counties are included in the findings summarized in Table 1. Focusing only on the counties that cede, Figure 3 reports acreage-weighted average indemnities for ceding counties in El Niño (top row) and La Niña (bottom row) years. The bar charts in the first column report indemnities for both the ceded and retained policies, and the pair-plots in the second column show that retained indemnities are 30-40 percent lower than their ceded counterparts for La Niña, and 10-30 percent lower for El Niño. Supplementary Figures S6 and S7 distinguish between the Western and Eastern regions, and again we see that the full sample results mask a considerable amount of heterogeneity across regions. As opposed to the Western region, the retained to ceded ratios in the Eastern region are much smaller at the lowest coverage level under both El Niño and La Niña, thereby implying that policy ceding offers relatively more protection against large-scale indemnities in the East compared to the West.

5 Conclusion

While ENSO impacts on agricultural production and commodity prices has been well documented (e.g. Handler, 1990; Brunner, 2002; Ubilava and Holt, 2013), research focusing on ENSO linkages with the entire distribution of crop yields, and subsequently agricultural risk and crop insurance, is relatively thin. This is the main focus of this study as we analyze the effect of ENSO on crop insurance premium rates using a panel of U.S. county-level cotton yield data spanning 1968 to 2005. Our findings provide insights for commercial insurers, policy makers, and crop producers, and are particularly relevant given that U.S. agricultural policy has become increasingly focused on crop insurance.

This study confirms previous findings that ENSO events impact not only mean cotton yields, but the distribution as well. We find that ENSO events impact actuarially fair premium rates, with rates among Eastern Cotton Belt counties decreasing under both El Niño or La Niña episodes. This is in contrast to the findings of Nadolnyak et al. (2008) who found that expected losses for cotton are generally highest during El Niño. This difference could be due to alternative empirical models and/or definitions of ENSO events, neither of which have been established in the literature. We also find that the pattern of premium rate differentials is distinct across the Western vs. Eastern Cotton Belt, where the former exhibits increasing rates under La Niña and decreasing rates under El Niño.

Our simulation results indicate that private insurance companies can use publicly available information to determine if government-set premium rates are mis-priced, and in turn extract economic rents via the federally mandated Standard Reinsurance Agreement. This finding complements Ker and McGowan (2000) as they focused on Texas winter wheat. They note that in response to the current information asymmetry between the RMA and private insurers, the RMA could alter the parameters of the Standard Reinsurance Agreement. However, this insight was made back when the prewarning time for ENSO events was limited to roughly six months ahead, in which case the RMA could not effectively incorporate forecasts into premium rates to offset the informational advantage held by private insurers. Interestingly, recent empirical evidence suggests that a significant extension of this prewarning time is feasible (e.g. Ludescher et al., 2013). Such modeling advances could in turn erode the current informational advantage. Thus, future work that proposes methods for RMA premium rates based on these longer range ENSO forecasts would be of interest.

It is important to note that the overall importance of this research contribution, as well as those of Ker and McGowan (2000) and Nadolnyak et al. (2008), will largely reveal itself through future work that considers a wider range of crop-policy combinations. Of particular interest would be extensions to farm-based insurance products, however the relatively short time-dimension of available farm-level data could prove difficult to overcome. We focus here on area insurance because there exists a long enough time series of data to demonstrate the ENSO-premium rate linkage in a compelling fashion. It is our opinion that these results are likely applicable to farm-based products for cotton since it is subject to the same weather relationships that apply to the area insurance product studied here. In addition, previous work linking ENSO dynamics to U.S. corn yields (e.g. Tack and Ubilava, 2013) further suggests that these findings will likely apply to other crops as well. Thus, demonstrating that the linkage we find here applies to farm-based products as well appears challenging but well worth future efforts.

This research also contributes to the burgeoning scientific literature linking climatic phenomena to insurance program-design and decision-making in a developing country context. Much of the research in the latter has focused on linking ENSO events to insurance-related outcomes (e.g. Khalil et al., 2007; Skees et al., 2007; Carriquiry and Osgood, 2011; Collier et al., 2011; Miranda and Farrin, 2012), however some interesting linkages with other climatic systems have also been found (e.g. Chang et al., 2011). Research focusing on ENSO linkages is typically conducted in a developing county context where severe weather events can have widespread effects beyond agricultural production. For example, Khalil et al. (2007) note that climate-related hazards can impact not only crop production and the associated banking/microfinance institutions, but also unrelated basic infrastructure and microenterprises. While our focus on the U.S. crop insurance program is somewhat of a departure compared to these studies, it is likely that these findings can be useful in guiding insurance design and implementation in developing countries.

References

Adams, R., C. Chen, B. McCarl, and R. Weiher (1999). The Economic Consequences of ENSO Events for Agriculture. *Climate Research 13*(3), 165–172.

Adams, R. M., L. L. Houston, B. A. McCarl, M. Tiscareño L, J. Matus G, and R. F. Weiher (2003). The Benefits to Mexican Agriculture of an El Niño–Southern Oscillation (ENSO) Early Warning System. *Agricultural and Forest Meteorology 115*(3), 183–194.

An, S. and F. Jin (2004). Nonlinearity and Asymmetry of ENSO. *Journal of Climate 17*(12), 2399–2412.

Antle, J. M. (1983). Testing the Stochastic Structure of Production: A Flexible Moment-based Approach. *Journal of Business & Economic Statistics 1*(3), 192–201.

Barnett, B., J. Black, Y. Hu, and J. Skees (2005). Is Area Yield Insurance Competitive with Farm Yield Insurance? *Journal of Agricultural and Resource Economics 30*(2), 285–301.

Bove, M. C., J. J. O'Brien, J. B. Eisner, C. W. Landsea, and X. Niu (1998). Effect of El Niño on U.S. Landfalling Hurricanes, Revisited. *Bulletin of the American Meteorological Society 79*(11), 2477–2482.

Brenner, J. (1991). Southern Oscillation Anomalies and their Relationship to Wildfire Activity in Florida. *International Journal of Wildland Fire 1*(1), 73–78.

Brunner, A. (2002). El Nino and World Primary Commodity Prices: Warm Water or Hot Air? *Review of Economics and Statistics 84*(1), 176–183.

Cabrera, V. E., D. Letson, and G. Podestá (2007). The Value of Climate Information when Farm Programs Matter. *Agricultural Systems 93*(1), 25–42.

Cane, M., G. Eshel, and R. Buckland (1994). Forecasting Zimbabwean Maize Yield using Eastern Equatorial Pacific Sea Surface Temperature. *Nature 370*(6486), 204–205.

Carriquiry, M. and D. Osgood (2011). Index Insurance, Probabilistic Climate Forecasts, and Production. *Journal of Risk and Insurance 79*(1), 287–300.

Chang, C. C., S. K. Lin, and M. T. Yu (2011). Valuation of Catastrophe Equity Puts With Markov-Modulated Poisson Processes. *Journal of Risk and Insurance 78*(2), 447–473.

Chen, C., B. McCarl, and C. Chang (2008). Strong El Niño – Southern Oscillation Events and the Economics of the International Rice Market. *Climate Research 36*(2), 113.

Chen, C., B. McCarl, and H. Hill (2002). Agricultural Value of ENSO Information under Alternative Phase Definition. *Climatic Change 54*(3), 305–325.

Chiang, J. and A. Sobel (2002). Tropical Tropospheric Temperature Variations Caused by ENSO and Their Influence on the Remote Tropical Climate. *Journal of Climate 15*(18), 2616–2631.

Collier, B., A. L. Katchova, and J. R. Skees (2011). Loan Portfolio Performance and El Niño, an Intervention Analysis. *Agricultural Finance Review 71*(1), 98–119.

Day, R. (1965). Probability Distributions of Field Crop Yields. *Journal of Farm Economics 47*(3), 713.

Deng, X., B. Barnett, and D. Vedenov (2007). Is There a Viable Market for Area-Based Crop Insurance? *American Journal of Agricultural Economics 89*(2), 508–519.

Deng, Z. and Y. Tang (2009). The Retrospective Prediction of ENSO from 1881 to 2000 by a Hybrid Coupled Model: (II) Interdecadal and Decadal Variations in Predictability. *Climate Dynamics 32*(2), 415–428.

Deschenes, O. and M. Greenstone (2007). The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather. *The American Economic Review 97*(1), 354–385.

Fedorov, A. and S. Philander (2000). Is El Niño Changing? *Science 288*(5473), 1997–2002.

Fisher, A., W. Hanemann, M. Roberts, and W. Schlenker (2012). The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment. *American Economic Review 102*(7), 3749–3760.

Gallagher, P. (1987). US Soybean Yields: Estimation and Forecasting with Nonsymmetric Disturbances. *American Journal of Agricultural Economics 69*(4), 796.

Goodwin, B. and A. Ker (1998). Nonparametric Estimation of Crop Yield Distributions: Implications for Rating Group-Risk Crop Insurance Contracts. *American Journal of Agricultural Economics 80*(1), 139.

Goodwin, B. K. (2001). Problems with Market Insurance in Agriculture. *American Journal of Agricultural Economics 83*(3), 643–649.

Goodwin, B. K. (2008). Climate Variability Implications for Agricultural Crop Production and Risk Management: Discussion. *American Journal of Agricultural Economics 90*(5), 1263–1264.

Hall, A., J. Skalin, and T. Teräsvirta (2001). A Nonlinear Time Series Model of El Niño. *Environmental Modelling & Software 16*(2), 139–146.

Handler, P. (1990). USA Corn Yields, The El Niño and Agricultural Drought: 1867–1988. *International Journal of Climatology 10*(8), 819–828.

Hansen, J., A. Hodges, and J. Jones (1998). ENSO Influences on Agriculture in the Southeastern United States. *Journal of Climate 11*(3), 404–411.

Hansen, J. W. (2002). Realizing the Potential Benefits of Climate Prediction to Agriculture: Issues, Approaches, Challenges. *Agricultural Systems 74*(3), 309–330.

Harri, A., K. Coble, A. Ker, and B. Goodwin (2011). Relaxing Heteroscedasticity Assumptions in Area-Yield Crop Insurance Rating. *American Journal of Agricultural Economics 93*(3), 707–717.

Hennessy, D. (2009a). Crop Yield Skewness and the Normal Distribution. *Journal of Agricultural and Resource Economics 34*(1), 34–52.

Hennessy, D. (2009b). Crop Yield Skewness Under Law of the Minimum Technology. *American Journal of Agricultural Economics 91*(1), 197.

Hsiang, S., K. Meng, and M. Cane (2011). Civil Conflicts are Associated with the Global Climate. *Nature 476*, 438–441.

Jones, J. (1999). Atlantic and Pacific Sea Surface Temperatures and Corn Yields in the Southeastern USA: Lagged Relationships and Forecast Model Development. *International Journal of Climatology 31*(4), 592–Ű604.

Jones, J., J. Hansen, F. Royce, and C. Messina (2000). Potential Benefits of Climate Forecasting to Agriculture. *Agriculture, Ecosystems & Environment 82*(1-3), 169–184.

Jones, S. (2009, April 24). The Formula that Felled Wall St. Financial Times.

Just, R. E. and R. D. Pope (1978). Stochastic Specification of Production Functions and Economic Implications. *Journal of Econometrics 7*(1), 67–86.

Ker, A. and K. Coble (2003). Modeling Conditional Yield Densities. *American Journal of Agricultural Economics 85*(2), 291.

Ker, A. P. and P. McGowan (2000). Weather-Based Adverse Selection and the U.S. Crop Insurance Program: The Private Insurance Company Perspective. *Journal of Agricultural and Resource Economics 25*(2), 386–410.

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 Research 43*(10), W10416.

Kirtman, B. and P. Schopf (1998). Decadal Variability in ENSO Predictability and Prediction. *Journal of Climate 11*(11), 2804–2822.

Legler, D., K. Bryant, and J. O'Brien (1999). Impact of ENSO-related Climate Anomalies on Crop Yields in the U.S. *Climatic Change 42*(2), 351–375.

Ludescher, J., A. Gozolchiani, M. I. Bogachev, A. Bunde, S. Havlin, and H. J. Schellnhuber (2013). Improved el niño forecasting by cooperativity detection. *Proceedings of the National Academy of Sciences 110*(29), 11742–11745.

Marshall, G. R., K. A. Parton, and G. Hammer (1996). Risk Attitude, Planting Conditions And The Value Of Seasonal Forecasts To A Dryland Wheat Grower. *Australian Journal of Agricultural and Resource Economics 40*(3), 211–233.

Meehl, G. A., F. Zwiers, J. Evans, T. Knutson, L. Mearns, and P. Whetton (2000). Trends in Extreme Weather and Climate Events: Issues Related to Modeling Extremes in Projections of Future Climate Change. *Bulletin of the American Meteorological Society 81*(3), 427–436.

Messina, C., J. Hansen, and A. Hall (1999). Land Allocation Conditioned on El Niño-Southern Oscillation Phases in the Pampas of Argentina. *Agricultural Systems 60*(3), 197–212.

Meza, F. J., J. W. Hansen, and D. Osgood (2008). Economic Value of Seasonal Climate Forecasts for Agriculture: Review of Ex-Ante Assessments and Recommendations for Future Research. *Journal of Applied Meteorology and Climatology 47*(5), 1269–1286.

Miranda, M. J. and K. Farrin (2012). Index Insurance for Developing Countries. *Applied Economic Perspectives and Policy 34*(3), 391–427.

Mjelde, J. W., T. N. Thompson, and C. J. Nixon (1996). Government Institutional Effects on the Value of Seasonal Climate Forecasts. *American journal of agricultural economics 78*(1), 175–188.

Moss, C. and J. Shonkwiler (1993). Estimating Yield Distributions with a Stochastic Trend and Nonnormal Errors. *American Journal of Agricultural Economics 75*(4), 1056.

Nadolnyak, D., D. Vedenov, and J. Novak (2008). Information Value of Climate-Based Yield Forecasts in Selecting Optimal Crop Insurance Coverage. *American Journal of Agricultural Economics 90*(5), 1248–1255.

Nelson, C. and P. Preckel (1989). The Conditional Beta Distribution as a Stochastic Production Function. *American Journal of Agricultural Economics 71*(2), 370.

Phillips, J., M. Cane, and C. Rosenzweig (1998). ENSO, Seasonal Rainfall Patterns and Simulated Maize Yield Variability in Zimbabwe. *Agricultural and Forest Meteorology 90*(1-2), 39–50.

Rajagopalan, B., U. Lall, and M. Cane (1997). Anomalous ENSO Occurrences: An Alternate View. *Journal of Climate 10*(9), 2351–2357.

Ramirez, O., S. Misra, and J. Field (2003). Crop-Yield Distributions Revisited. *American Journal of Agricultural Economics 85*(1), 108.

Ropelewski, C. and M. Halpert (1986). North American Precipitation and Temperature Patterns Associated with the El Niño/Southern Oscillation (ENSO). *Monthly Weather Review 114*(12), 2352–2362.

Rosenzweig, C., A. Iglesias, X. Yang, P. R. Epstein, and E. Chivian (2001). Climate Change and Extreme Weather Events: Implications for Food Production, Plant Diseases, and Pests. *Global Change & Human Health 2*(2), 90–104.

Saunders, M., R. Chandler, C. Merchant, and F. Roberts (2000). Atlantic Hurricanes and NW Pacific Typhoons: ENSO Spatial Impacts on Occurrence and Landfall. *Geophysical research letters 27*(8), 1147–1150.

Schlenker, W. and M. Roberts (2006). Nonlinear Effects of Weather on Corn Yields. *Review of Agricultural Economics 28*, 391–398.

Schlenker, W. and M. Roberts (2009). Nonlinear Temperature Effects Indicate Severe Damages to US Crop Yields under Climate Change. *Proceedings of the National Academy of Sciences 106*(37), 15594.

Sherrick, B., F. Zanini, G. Schnitkey, and S. Irwin (2004). Crop Insurance Valuation under Alternative Yield Distributions. *American Journal of Agricultural Economics 86*(2), 406.

Shohat, J. and J. Tamarkin (1943). *The Problem of Moments*, Volume 1. American Mathematical Society.

Skees, J. R., J. Hartell, and A. G. Murphy (2007). Using index-based risk transfer products to facilitate micro lending in peru and vietnam. *American Journal of Agricultural Economics 89*(5), 1255–1261.

Solís, D. and D. Letson (2013). Assessing the Value of Climate Information and Forecasts for the Agricultural Sector in the Southeastern United States: Multi-Output Stochastic Frontier Approach. *Regional Environmental Change 13*(1), 5–14. doi: 10.1007/s10113-012-0354-x.

Solow, A., R. Adams, K. Bryant, D. Legler, J. O'Brien, B. McCarl, W. Nayda, and R. Weiher (1998). The Value of Improved ENSO Prediction to US Agriculture. *Climatic Change 39*(1), 47–60.

Stige, L. C., J. Stave, K.-S. Chan, L. Ciannelli, N. Pettorelli, M. Glantz, H. R. Herren, and N. C. Stenseth (2006). The Effect of Climate Variation on Agro-Pastoral Production in Africa. *Proceedings of the National Academy of Sciences of the United States of America 103*(9), 3049–3053.

Stock, J. H. (2001). Forecasting economic time series. In B. H. Baltagi (Ed.), *A Companion to Theoretical Econometrics*, pp. 562–584. Blackwell Publishing.

Stohs, S. M. and J. T. LaFrance (2004). A Learning Rule for Inferring Local Distributions over Space and Time. *Advances in Econometrics 18*, 295–331.

Swetnam, T. and J. Betancourt (1990). Fire–Southern Oscillation Relations in the Southwestern United States. *Science 249*(4972), 1017.

Tack, J., A. Harri, and K. Coble (2012). More than Mean Effects: Modeling the Effect of Climate on the Higher Order Moments of Crop Yields. *American Journal of Agricultural Economics 94*(5), 1037–1054.

Tack, J. B. and D. Ubilava (2013). The Effect of the El Niño Southern Oscillation on U.S. CornProduction and Downside Risk. *Climatic Change 121*(4), 689–700.

Timmermann, A., J. Oberhuber, A. Bacher, M. Esch, M. Latif, and E. Roeckner (1999). Increased El Niño Frequency in a Climate Model Forced by Future Greenhouse Warming. *Nature 398*, 694–697.

Ubilava, D. and C. Helmers (2013). Forecasting ENSO with a Smooth Transition Autoregressive Model. *Environmental Modelling & Software 40*(1), 181–190.

Ubilava, D. and M. Holt (2013). El Niño Southern Oscillation and its Effects on World Vegetable Oil Prices: Assessing Asymmetries using Smooth Transition Models. *Australian Journal of Agricultural and Resource Economics 57*(2), 273–297.

Vedenov, D. V. and B. J. Barnett (2004). Efficiency of Weather Derivatives as Primary Crop Insurance Instruments. *Journal of Agricultural and Resource Economics 29*(3), 387–403.

Wu, X. (2003). Calculation of Maximum Entropy Densities with Application to Income Distribution. *Journal of Econometrics 115*(2), 347–354.

Wu, X. (2010). Exponential Series Estimator of Multivariate Densities. *Journal of Econometrics 156*(2), 354–366.

Figures

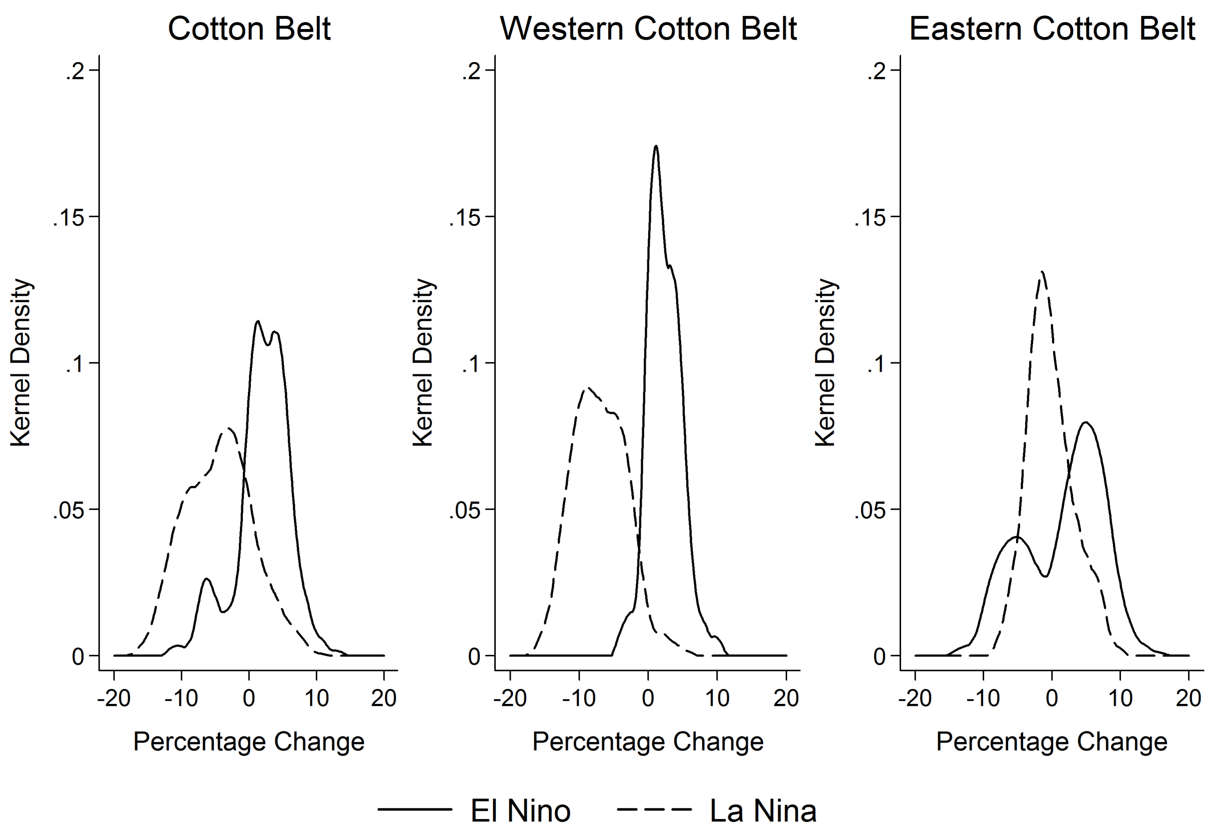

Figure 1: ENSO impacts on mean yield

Notes: Each graph contains kernel densities of the county-level mean-yield impacts for El Niño and La Niña relative to La Nada.

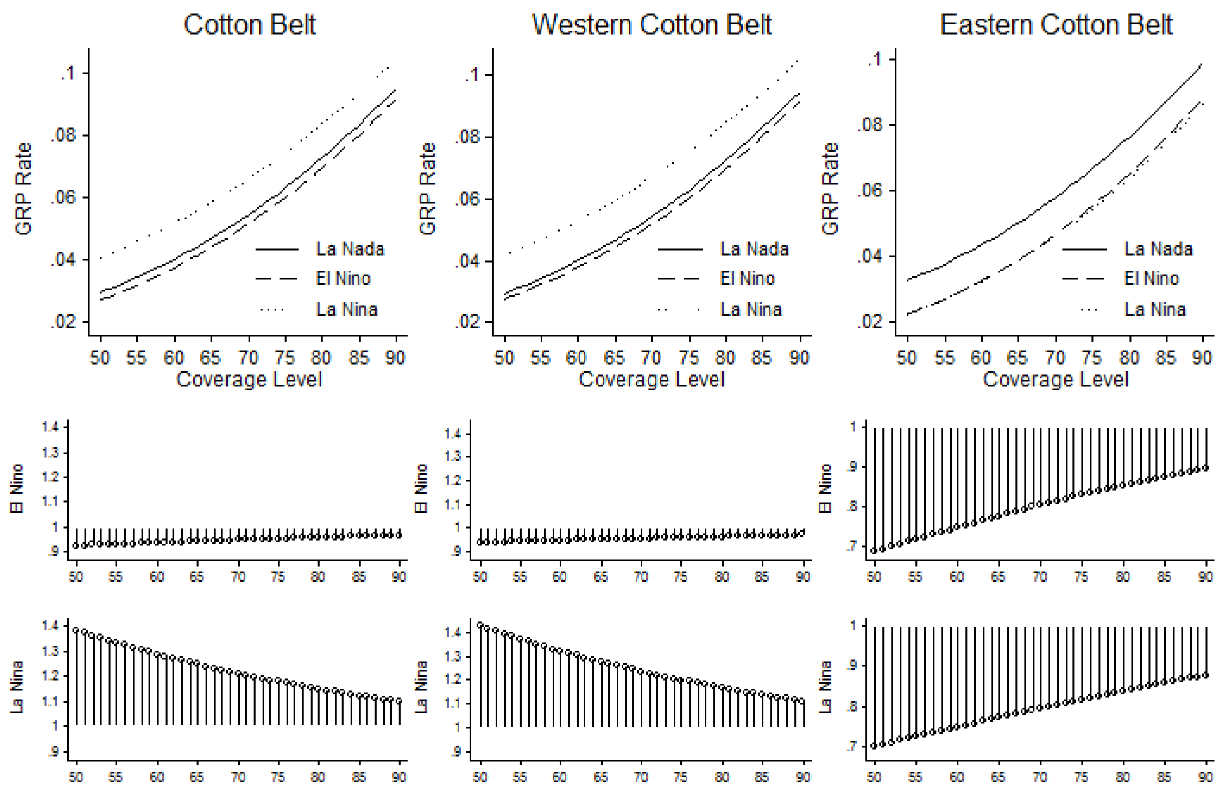

Figure 2: GRP premium rates across ENSO regimes

Notes: The first row presents acreage-weighted rates across counties for various coverage levels. The second row presents pair-plots of the rates for El Niño and La Niña relative to La Nada. Each pair-plot is a simple ratio of the rates with the La Nada rate in the denominator.

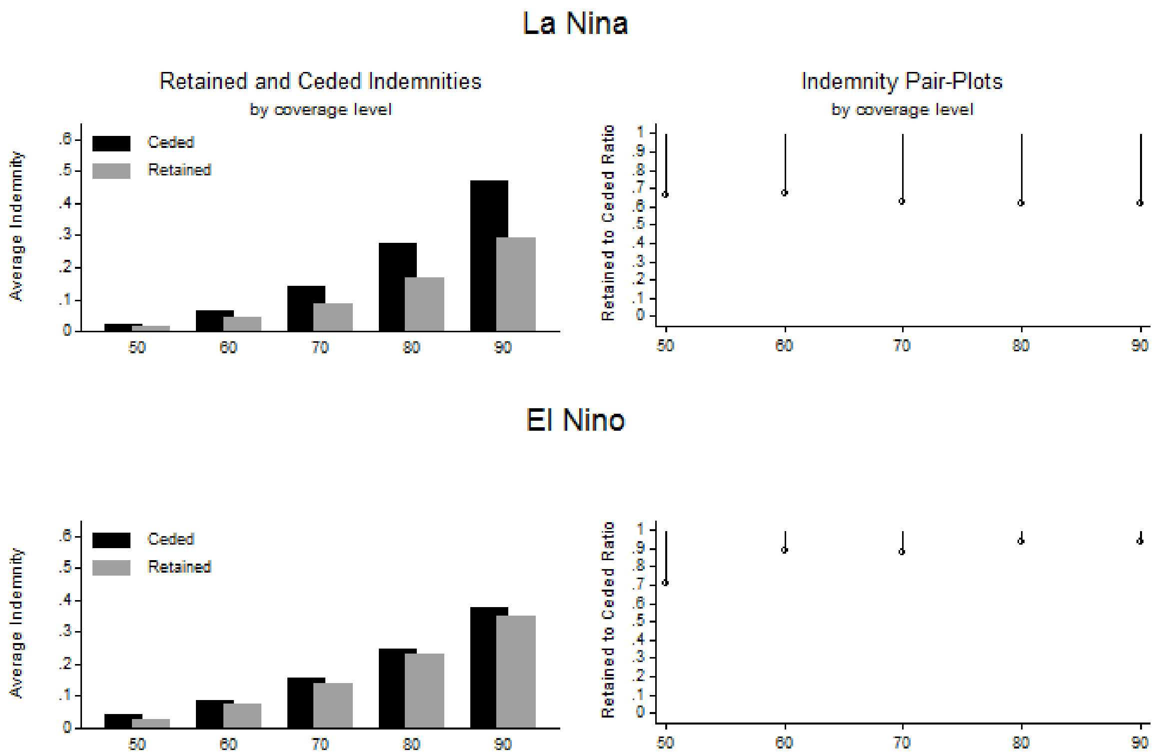

Figure 3: Indemnities for ceders: Cotton Belt

Notes: The first column reports acreage-weighted indemnities across counties for the repeated game of insurance selection. The second column presents pair-plots of these indemnities for ceded and retained policies. Each pair-plot is a simple ratio of the rate for retained polices over the rate for ceded policies.

Tables

Table 1: Average County-Level Indemnities for Repeated Insurance Game

	All Years		El Nino Years		La Nina Years	
Coverage Level	All Policies	Ret Policies	All Policies	Ret Policies	All Policies	Ret Policies
50	0.013	0.011	0.014	0.005	0.011	0.004
60	0.039	0.034	0.028	0.013	0.041	0.017
70	0.090	0.078	0.056	0.032	0.108	0.070
80	0.185	0.162	0.126	0.095	0.245	0.194
90	0.334	0.298	0.250	0.220	0.455	0.398

Notes: Entries in the table report the acreage-weighted average indemnity across all counties. The first two columns include all years 1968-2005, the second (last) two columns just include the El Niño (La Niña) years. The *All Policies* columns includes both the ceded and retained policies, while the *Ret Policies* columns include only the retained policies.