2
El Niño-Southern Oscillation (ENSO)

This chapter presents my analysis of the El Niño-Southern Oscillation (ENSO) climate phenomenon and its economic impacts. This work provides the first line of evidence in support of traded markets in ENSO risk. The chapter includes:

- a description of the ENSO as a climate phenomenon;
- a brief discussion of what current climate science tells us about ENSO's impacts;
- an introduction to indexes of ENSO (see chapter 3 for additional information on ENSO-related SST dataset); and
- statistical analysis of the correspondence between the ENSO index and disaster damages (estimated in chapter 1) around the world.

Introduction to El Niño-Southern Oscillation

ENSO refers to a coupled oceanic/atmospheric cycle, its occasional break-down called El Niño, and supercharging called La Niña. In normal years, the ENSO cycle refers to currents and winds (each reinforcing the other) that bring water along the surface of the Pacific ocean from South America (from the eastern side of the Pacific) to Indonesian and the South Pacific (the western side of the Pacific). As that water travels along the ocean surface, it warms, thanks to the intense sunlight in the tropics. This results in water piling up on the Pacific's western side[1] actually making sea levels measurably higher in Indonesian than in Peru. As this mass of warm water accumulates, much of it sinks deeper into the ocean, where it naturally flows back east, across the Pacific, toward South America. By the time that the subsurface mass or water has reached the South American coast it is cold, allowing it to store more of the nutrients that serve as the basis of a vibrant aquatic ecosystem. So, as it springs up to replace the water moving west, it enriches the fisheries off Peru and Chile.

[1] Confusingly, the Western Pacific bumps up against South*east* Asia.

During an El Niño anomaly, this cycle weakens. (See figure 2.1 modified from Rosenzweig and Hillel [2008] below.) As less water reaches the western end of the Pacific, sea-surface temperatures rise. Over the course of the year, a plume of warmer-than-normal water creeps eastward across the Pacific. When that plume of warm water reaches Peru, it parks a moisture laden air mass off the coast. When that mass meets cold air coming east to west over the Andes mountains, Peru suffers catastrophic downpours and flooding.

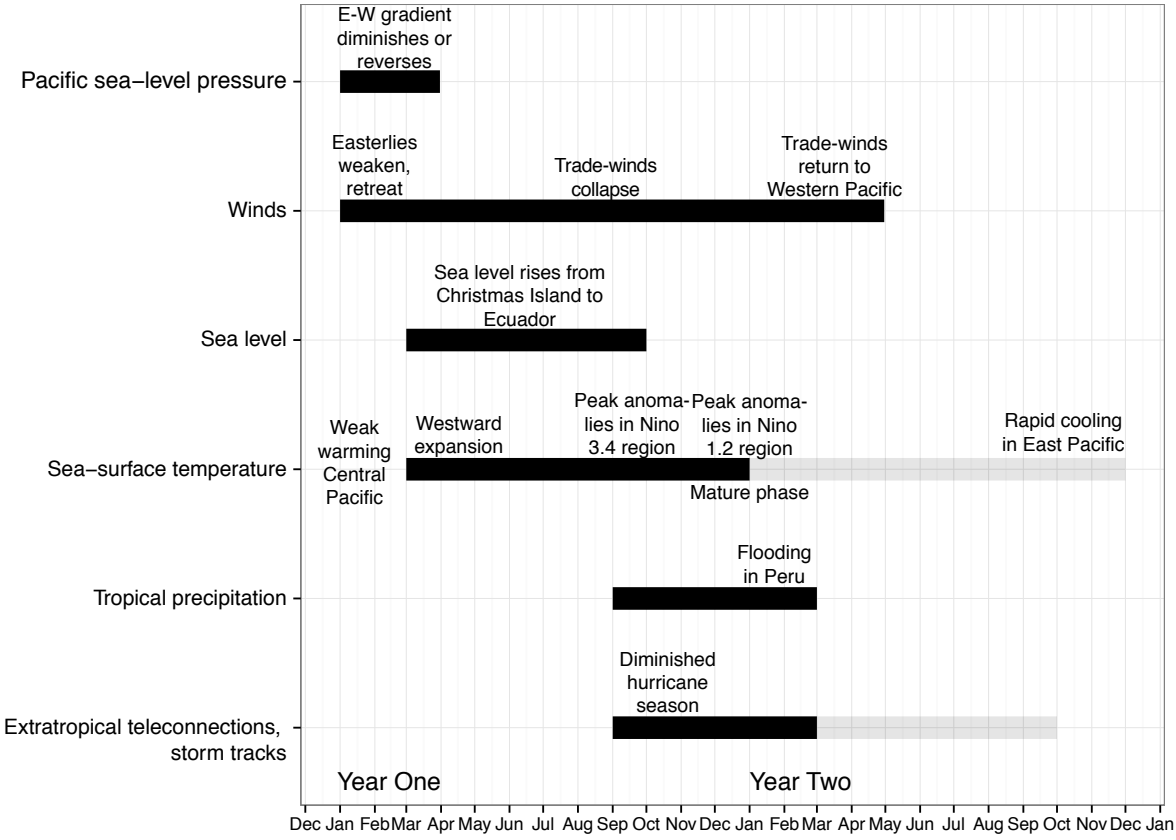

Figure 2.1: Calendar of average El Niño event, modified from Rosenzweig and Hillel [2008]. Gray indicates impacts contingent on the strength of the event. Note that the calendar for any one ENSO events can vary greatly.

By contrast, during a La Niña anomaly the normal cycle enhances. More water gets pushed from the South American coast, raising sea-surface temperatures in Australia and Indonesia above normal. That leaves Southeast Asia and Oceania with the same problem as Peru during El Niño. A warm air mass sits in the region waiting for the opportunity to cause extreme rains and floods.

The ENSO cycle drives weather patterns well beyond Australia, Indonesia, and Peru. Figures 2.2 and 2.3 summarize global precipitation and temperature impacts for El Niño. La Niña shows opposing impacts, with a similar geographic footprint, but not necessarily of the

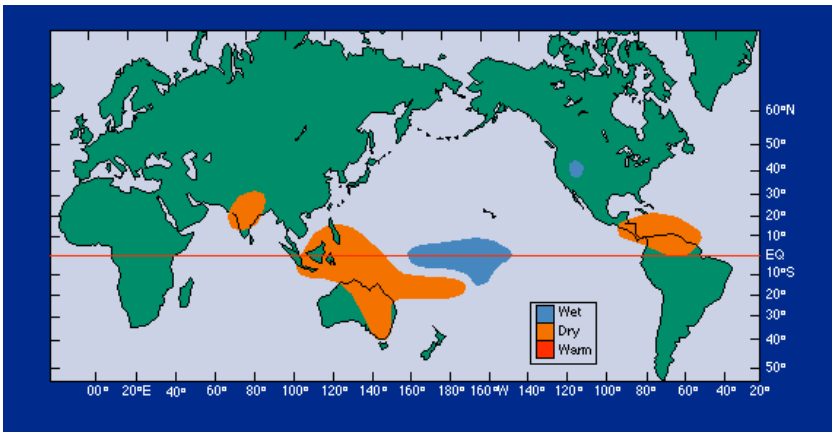

Figure 2.2: El Niño global impacts during the Northern Hemisphere summer

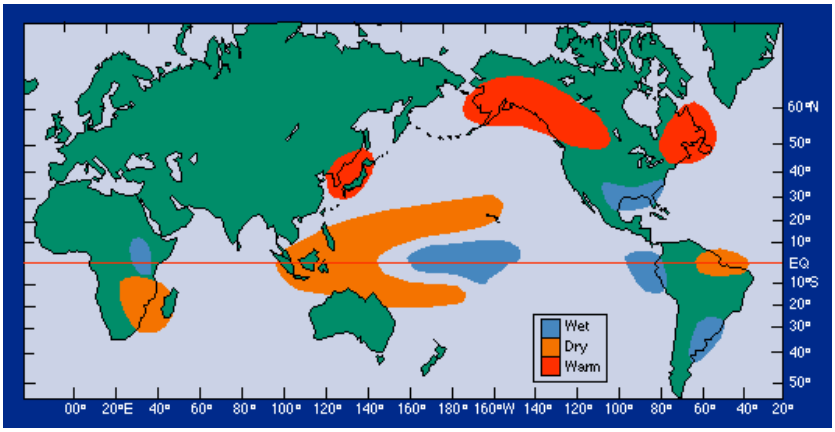

Figure 2.3: El Niño global impacts during the Northern Hemisphere winter

same magnitude.[2]

Below are indicative publications covering regional or peril-specific ENSO impacts in greater depth:

Global impacts

- Worldwide precipitation patterns

  Ropelewski and Halpert [1987] and Ropelewski and Halpert [1989] are seminal papers looking at the footprints of ENSO anomalies between 1875 and 1983. The basis for figures 2.2 and 2.3, they identify regions (19 for El Niño and 15 for La Niña) where precipitation has a statistically significant link to the ENSO cycle.

  Mason and Goddard [2001] provides a more recent probabilistic estimates of ENSO's influence on precipitation across the globe.

Global studies provide an excellent starting point for understanding ENSO's importance to catastrophic weather. But they rely on global datasets with uneven coverage in the developing world. Judging by Mason and Goddard [2001], alone you might conclude that El Niño has a stronger influence on precipitation in the southwestern United States than in southern Peru and southern Ecuador. That is an artifact of the data that regional studies can address. Below are indicative citations that illustrate hedging opportunities that may be obscured in global ENSO research:

Regional impacts

- Flooding in the tropical Andean countries during El Niño

  Khalil et al. [2007] was prepared in association with GlobalAgRisk's Gates Foundation-supported work on El Niño insurance for northern Peru. It looks at the link between different ENSO indexes and extreme rainfalls in the Department of Piura, the local basis risk on those indexes, and the influence of climate change on regional flooding. It also addresses the trade-off between basis risk and advanced payments using earlier months' index values for insurance.

- La Niña/El Niño flooding/drought in Australia

  Chiew et al. [1998] provides an overview of the relationship between ENSO and rainfall, drought and streamflow in Australia. The analysis shows that ENSO is a statistically significant predictor of hydrological conditions across Australia. In particular, dry conditions in Australia tend to be associated with El Niño. The authors suggest that ENSO is, on its own, a useful forecasting tool for spring rainfall in eastern Australia and summer rainfall in

[2] Rosenzweig and Hillel [2008] provides an excellent non-technical overview of research related to the economic impacts ENSO around the globe.

north-east Australia. It is also helpful in predicting spring runoff in south-east Australia and summer runoff in the north-east and east coasts of Australia. However, autocorrelations diminish ENSO's value as a stand-alone predictor of Australian streamflows.

- Suppressed Atlantic hurricane activity during El Niño

  Klotzbach [2011] finds ENSO is the the primary interannual driver of variability in Caribbean hurricane activity, boosting hurricane activity in La Niña years and suppressing it in El Niño years. The article also examines interaction effects between ENSO and the Atlantic multidecadal oscillation.

- Flood and drought in the Southern Cone during El Niño and La Niña respectively

  Grimm et al. [2000] analyzes precipitation and circulation across South America's Southern Cone. It finds significant links to the ENSO cycle, both across the region and in eight distinct subregions. The strongest subregional association links above-average rainfall in Southern Brazil to El Niño.

- Drought in Northeastern Brazil during El Niño

  Hastenrath [2006] looks at ENSO's influence on the short rainy season (covering just March and April) in the Nordeste region of Brazil. Like northern Peru, the Brazilian Nordeste has a particularly high incidence of poverty and a history of dramatic precipitation events (drought in Brazil) coincident with extreme El Niño. Hastenrath [2006] examines the climate drivers behind the region's recurrent Secas (droughts) with a focus on ENSO.

- La Niña/El Niño flooding/drought in Southeast Asia

  Murty et al. [2000] looks at the acute airborne pollution in Malaysia likely sparked by the 1997/1998 El Niño. That season brought Indonesia's worst drought in 50 years which in turn sparked a forest fire on the island of Borneo that engulfed over one million acres. This article summarizes the climatic roots of that disaster and looks at how they interacted with land management decisions to export the catastrophic consequences of an ENSO anomaly beyond its core region.

- Suppressed Indian monsoon activity during El Niño

  Kumar et al. [2006] suggests that over the last 132 year El Niño events have been a necessary, but not sufficient, prerequisite to shortfalls in the Indian monsoon. The article suggests that Central-Pacific (Modoki) El Niños have a stronger link to Indian drought than classical Eastern Pacific El Niños. The article explains these

differing El Niño signatures using an atmospheric general circulation model.

- Drought in the West African Sahel during El Niño

  Janicot et al. [2001] explores the unstable relationship between Sahel rainfall and ENSO in the northern summer. Looking at 20 year running correlations between Sahel rainfall index and ENSO SST between 1945 and 1993, the article suggests that the correlation between El Niño and drought has changed over time. While it was not signficant in the 1960s, it strengthened and has been significant since 1976. The article proposes interactions with multi-decadal oscillations as a cause of that change.

- Flooding in East Africa during El Niño

  Indeje et al. [2000] investigates above-average rainfall linked to El Niño through in the data of 136 weather stations across Kenya, Uganda and Tanzania between 1961 and 1990. Using both an empirical orthogonal function (EOF) and basic correlations, the article identifies eight subregions with distinct rainfall patterns. The article agrees with previous studies suggesting a modest tendency toward above-average rainfall in El Niño years followed by below-average rainfall the next year.

- Drought in Southern Africa during El Niño

  Camberlin et al. [2001] looks at the connection between ENSO and precipitation anomalies across Africa. The article confirms previous findings that El Niño is linked to drought in East Africa (shortfalls in the rainy season betweens July and September in Ethiopia and between October and December in the east equatorial countries) and in Southern Africa, especially during the second part of its rainy season. Southern African rains also show a link to teleconnections based in the Indian Ocean, which may account for droughts in South Africa not associated with the ENSO cycle.

Index construction

ENSO anomalies are multifaceted phenomena involving feedback loops from many climate systems. However, most major NMS define El Niño/La Niña just by looking at one simple index, the temperature of the sea-surface, relative to its seasonal average in specific regions across the Pacific. Generally, NMS prefer to average their SST measurements across a month or months, but they also issue more frequent measurements. Hence, in its most basic form, the index tracking ENSO anomalies is directly interpretable.[3]

[3] The indexes of of some other regional climate anomalies like the AO (discussed in chapter 10) require graduate-level mathematics to calculate and are not denominated in simple units like degrees.

The index-based insurance purchased by Caja Nuestra Gente in 2012[4] used as its sole payment trigger November and December measurements of the NOAA-defined region know as Niño 1.2, which lies directly off the Peruvian coast. (See figure 2.4 for a map of NOAA's Niño regions.) If the average of NOAA's November and December 2013 SST readings for the Niño 1.2 region is 24°C or above, then Caja Nuestra Gente will receive an insurance payment for the occurrence of a severe ENSO anomaly.

[4] 2012 marks the second year in a row that the bank has purchased the coverage designed by GlobalAgRisk.

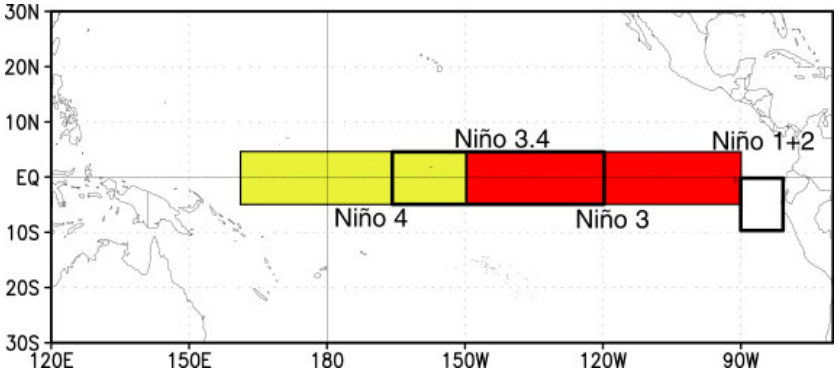

Figure 2.4: NOAA's Niño SST regions from http://www.cpc.ncep.noaa.gov

Niño 1.2 is the best predictor of catastrophic flooding in Peru and Ecuador, El Niño's flagship impact. However, NMS generally mark ENSO anomalies using the Niño 3.4 region[5] (roughly, from 5°N to 5°S and from 120° to 170°W), which stretches across the central Pacific[6] [7]. Both regions, Niño 1.2 and the Niño 3.4, have a very high correlation during extreme anomalies. But Niño 3.4 is generally considered a better proxy for the worldwide teleconnections associated with ENSO. In particular, it does a better job capturing ENSO anomalies with different geographic signatures. During the 1972/1973 El Niño, for example, most of the sea-surface temperature warming occurred in the central Pacific, closer to Niño 3.4. El Niño events focused on the Central Pacific are also called *Modoki* Niños and can have large global impacts[8].

While month-by-month sea-surface temperatures alone provide a functional benchmark for extreme ENSO anomalies, NOAA's default index for ENSO anomalies, the Oceanic Niño Index (ONI), attempts to correct for two important statistical dynamics related to ENSO. First, the teleconnections associated with ENSO, correspond best to high sea-surface temperatures sustained across a few months. Consequently, ONI uses a 3-month mean SST anomaly (i.e. each month is reported as degrees above its average temperature) averaged over the Niño 3.4 region. Second, average sea-surface temperatures in the Niño 3.4 region have demonstrated a slight upward bias in recent decades. You can clearly see the bias in figure 2.5, where monthly averages over

[5] Niño 3.4, straddles two separate regions, Niño 3 and Niño 4.

[6] A.F. Khalil, H. Kwon, U. Lall, M.J. Miranda, and J. Skees. El Niño-Southern Oscillation-based index insurance for floods: Statistical risk analyses and application to Peru. *Water Resources Research*, 43(10): 10416, 2007

[7] A.G. Barnston, M. Chelliah, and S.B. Goldenberg. Documentation of a highly ENSO-related SST region in the equatorial Pacific. *Atmosphere Ocean*, 35(3):367, 1997

[8] Karumuri Ashok, Swadhin K Behera, Suryachandra A Rao, Hengyi Weng, and Toshio Yamagata. El Niño Modoki and its possible teleconnection. *Journal of Geophysical Research: Oceans (1978–2012)*, 112 (C11), 2007

successive 30 year periods have been creeping upward. This raises the possibility that "El Niño and La Niña episodes that are [normalized to] a single fixed 30-year base period (e.g. 1971-2000) are increasingly incorporating longer-term trends that do not reflect inter-annual ENSO variability."[9] To correct for this, the ONI index takes each 3-month mean sea-surface temperature from the Niño 3.4 regions and divides it by a corresponding average for a rolling base period. For example, the March 1950 ONI value is equal to the average of Niño 3.4 temperatures for January, February, and March, divided by the January, February, and March average between 1936 and 1965. For recent data, NOAA uses the 1981-2010 base period. This means that recent values are subject to revision. NOAA currently changes the base period for readings every decade, but as of 2016 will begin updating the base period every 5 years.

[9] Rebecca Lindsey. In watching for El Niõ and La Niña, NOAA adapts to global warming. *ClimateWatch Magazine*, February 5 2013. URL http://www.climatewatch.noaa.gov/

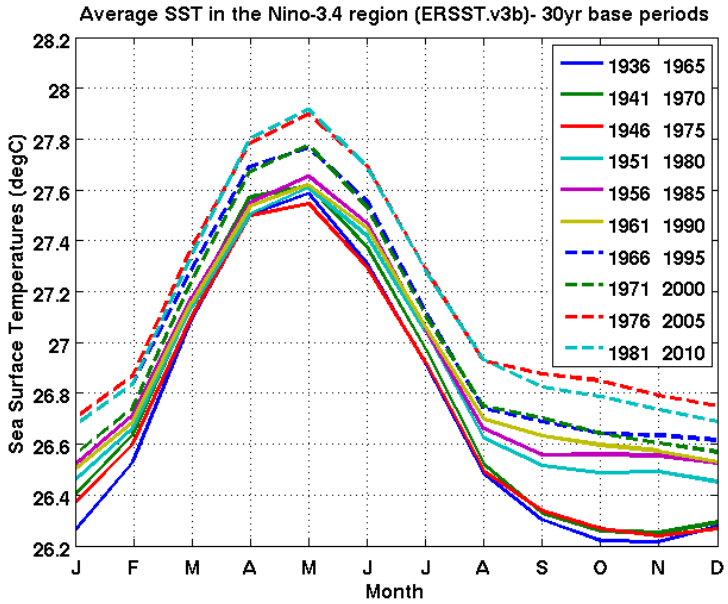

Figure 2.5: Long-term warming trend in Niño 3.4 region from NOAA http://www.cpc.ncep.noaa.gov/products/analysismonitoring/ensostuff/ONIchange.shtml

The ONI index is more difficult to interpret than simple monthly sea-surface averages. I suspect that this makes it less suitable as the basis of an exchange traded risk management contracts. However, I believe the smoothed index provides a solid foundation for this initial statistical analysis.

Statistical analysis of EM-DAT disasters

Researchers have used the EM-DAT database to estimate ENSO's global impacts. But there are clear opportunities to enhance that literature. Bouma et al. [1997], for example, identified a strong link

between the ENSO cycle and the number of people affected by disasters globally. But curiously, part of the uptick is linked to increased volcanic activity in years following El Niño events.

Goddard and Dilley [2005], by contrast, found that the overall frequency of hydrological disasters in the EM-DAT database was not significantly higher during El Niño or La Niña events than during ENSO neutral periods. That analysis also found weak evidence of trends in aggregate precipitation over land areas associated with ENSO extremes.

Goddard and Dilley [2005]'s finding are not surprising. First, ENSO represents shifts in burden of disaster across the globe and changes in the magnitude of disaster impacts. Indeed without some zero-sum-like shift in disaster burden, there would be little advantage to managing ENSO risk on an exchange, with relative winners from any ENSO state trading risk with relative losers in that state. Second, as discussed in chapter 1, the EM-DAT database has some evident short-comings as a proxy for ENSO impacts, which I have worked to correct. Even after my augmentation of missing data from the EM-DAT database, flooding in northern Peru between January and April, the largest and most dramatic impact of El Niño, hardly appears in the database. For that reason, I believe that accurate disaster impact statistics require Bayesian analysis which allows me to reference outside assessments of regional ENSO impacts as I extrapolate from the disaster statistics I compiled in the last chapter.

Figure 2.6, shows disaster burden data aggregated across the world next to the historic time series of the ONI index, with extreme El Niño events marked in red and La Niña events marked in blue. As Goddard and Dilley [2005] noted, it is difficult to identify clear tends in any of the disaster types.

To identify groups of ENSO hedgers, my analysis segregates country-disasters into groups that likely have similar hedging interest. In my ENSO analysis I use four large groups:

- Flood and epidemics on South America's Pacific Coast - Countries on the Pacific coast of South America tend to face flood and epidemic risk associated with El Niño. Some countries, such as Peru have experienced both flood and drought in extreme La Niña years, but the physical and statistical link with regional drought is less strong than for El Niño.

- Drought across the Southern Atlantic and Indian Ocean Basin - Historically many countries have experienced drought in ENSO years. The strongest links are with Pacific Asia and Oceania and Atlantic South America. There are also important potential links between ENSO and droughts in Southeast Asia, and Eastern/Southern

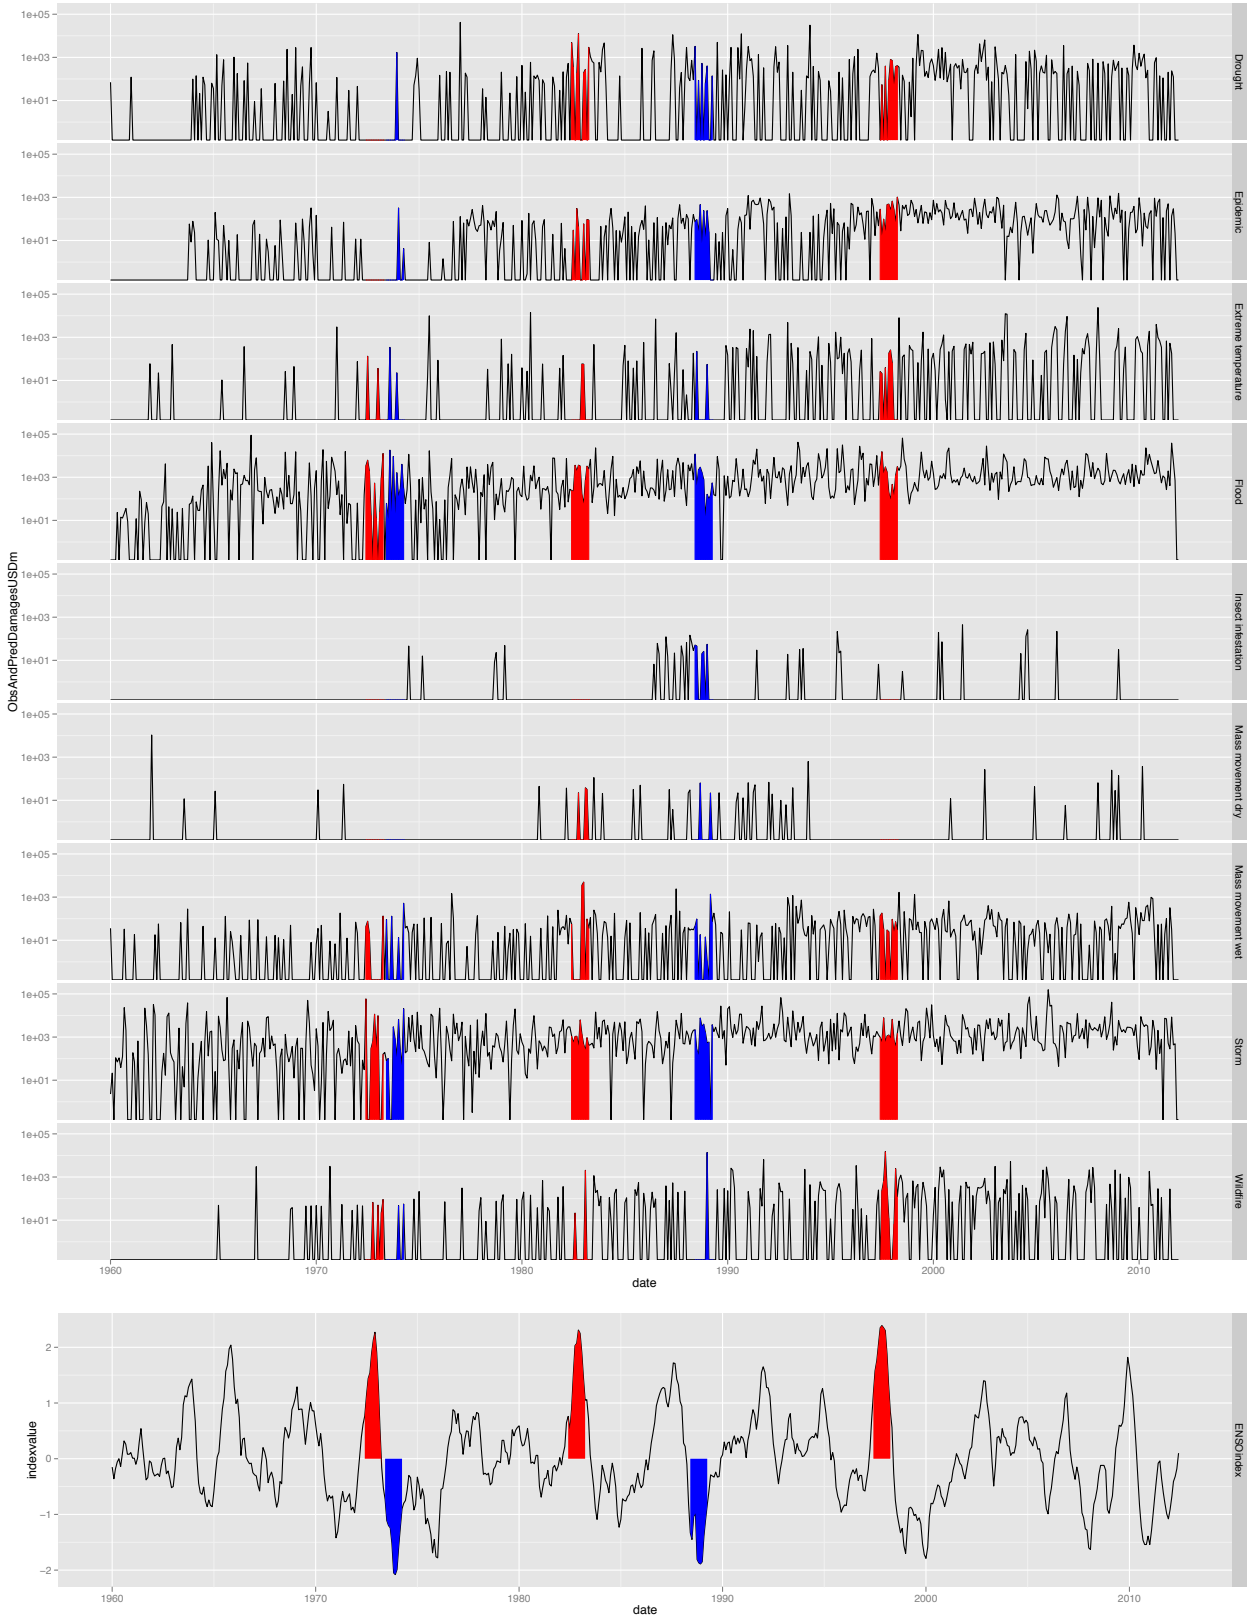

Figure 2.6: Worldwide disaster damage estimates by disaster type compared to ENSO (ONI) index

Africa. Given the link between drought and wildfire, wildfire incidence is also included in this grouping.

- Storms in North America and the Caribbean - Perhaps the most economically important offset for an ENSO market stems from the an inverse correlation between ENSO and storm activity in the Western Atlantic.

- Flooding in Pacific Asia and Oceania - This impact is generally associated with La Niña.

Undoubtedly, there are other groups with important exposure to the ENSO index and I could achieve a more accurate estimate of ENSO damage by further distinguishing subgroups. However, I believe that this grouping should be large enough to avoid spurious correlations in the data but small enough that they will not mask regional exposures to specific disaster-types.

South America's Pacific Coast - Flooding, landslides and epidemics from El Niño

The economic impacts of El Niño are well known in Peru and Ecuador. Despite the clear link between El Niño and disaster in the region, there are relatively few extreme El Niño events in recent historical record (1972/73, 1982/83, and 1997/98) and the region covers relatively few countries so statistical inference about the economic burden of El Niño in the region must come from relatively few disaster events.

In the typical extreme El Niño, sea surface temperatures off the coast of Peru hit anomaly levels in the last months of a given year and flooding begins in the first months of the following year. Based on that pattern, I aggregated damage due to flood, landslide, and epidemic in the first six months of each year between 1961 and 2010 and divided by the median annual damage over the period of study (roughly 259 m in 2010-USD). This sample included 192 separate disaster events.

To measure the influence of ENSO on, for example damages from January through June 2010, I averaged the ONI index[10] from October 2009 through January 2010. Using this technique, it is easy to distinguish the three extreme El Niño events in the recent series.

[10] Note that this index is based on a running average of monthly Niño 3.4 data.

Using these time series (for the index and damages as a percent of the seasonal median), I performed the Augmented Dickey-Fuller Test and the Phillips-Perron Unit Root Test. Both tests favored the alternative hypothesis of stationarity with greater than 95 percent confidence. Neither, the index nor the damage time series showed significant autocorrelation using a standard autocorrelation function, indicating that there is only weak interaction between the values of

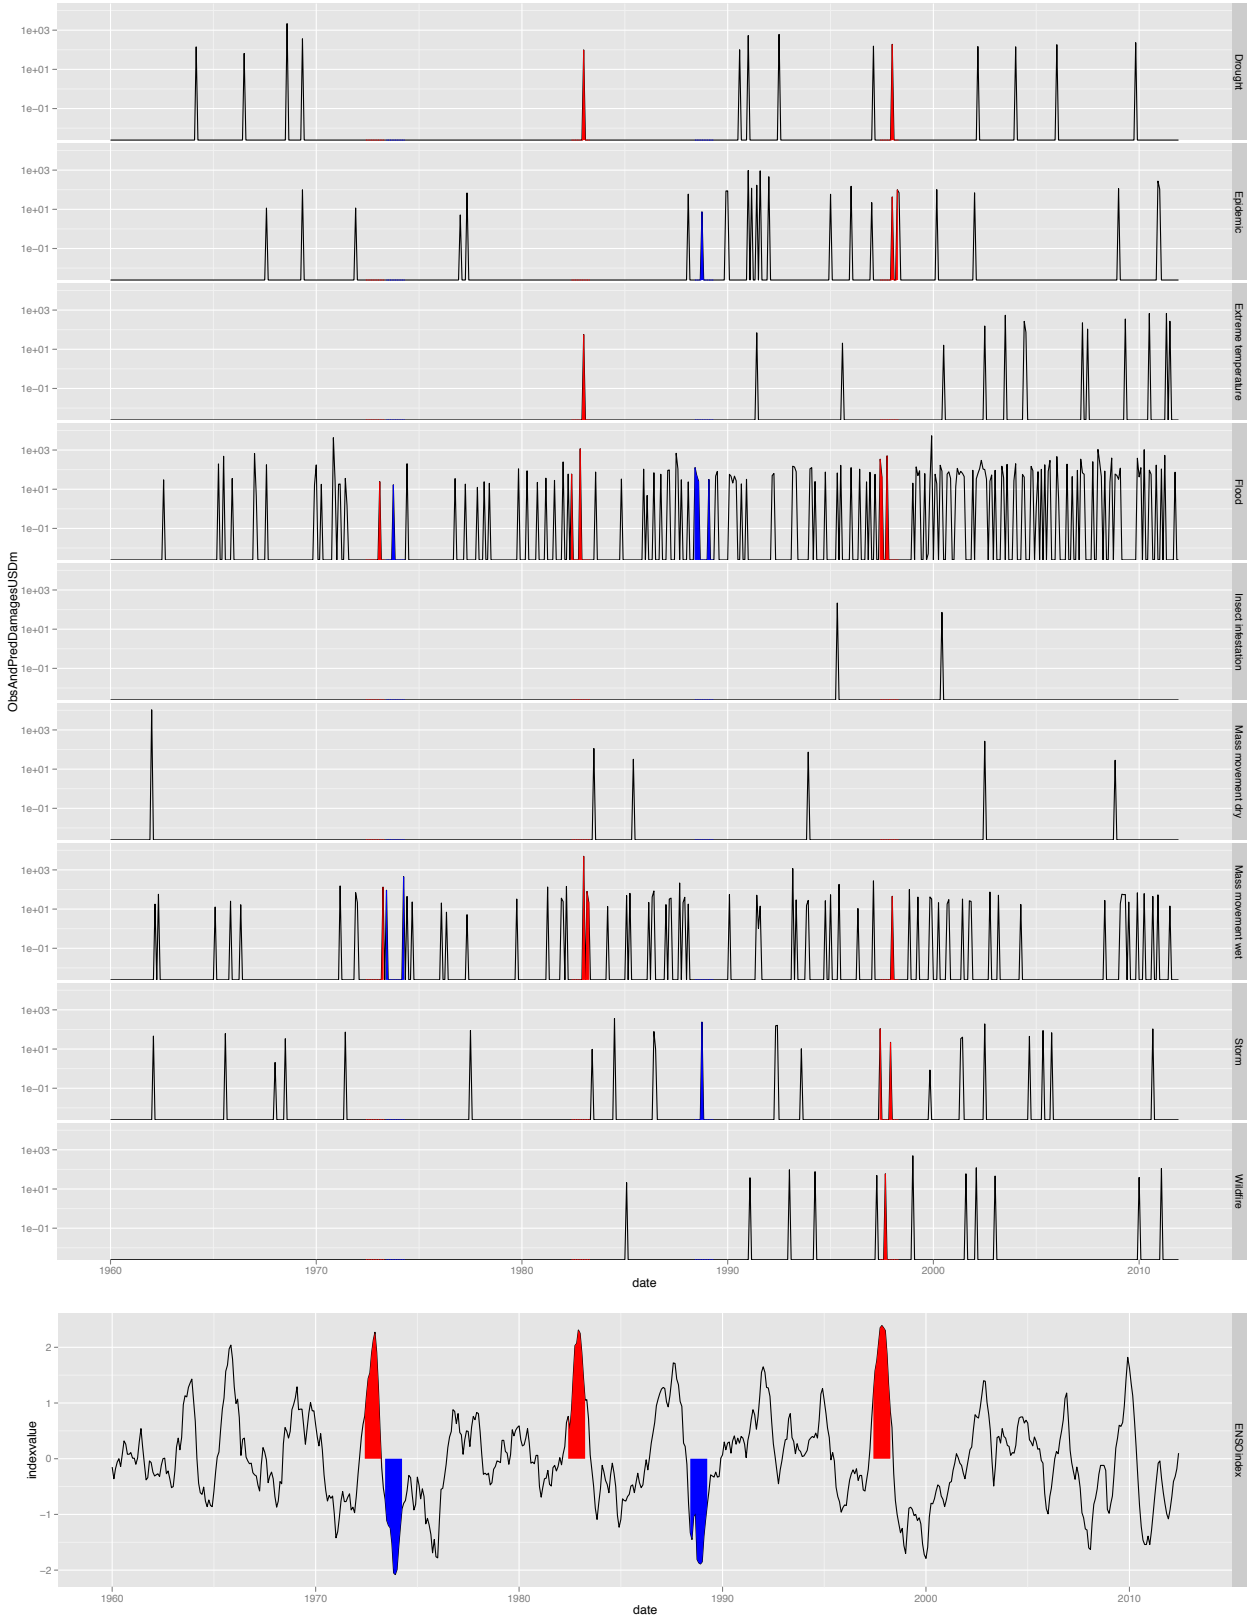

Figure 2.7: Disaster damage estimates by disaster type for countries along South America's Pacific coast compared to ENSO (ONI) index

one season and the next. As we have discussed earlier, many of the most prominent ENSO measurement indexes have show some upward bias in recent decades. However, the ONI index corrects for that bias.

I segregated the dataset to run separate regressions on:

1. El Niño seasons - those with an ONI average for October through January above 1[11]; and
2. Normal or La Niña seasons - those with a seasonal index below 1.

[11] The ONI index is normalized such that a value of 1 indicates one degree deviation above the average value for the corresponding historical window (see index construction section for more details). The standard deviation of the dataset is 0.82, so a value of 1 is slightly greater than a one standard deviation anomaly.

I chose to run separate regressions because I believe that the underlying process producing flooding and related disasters is distinct during moderate to strong El Niño conditions. This modeling decision likely reduces the power of inference, but allows for opposing slopes during each phase. I separate Bayesian regressions on each subset as indicated in equation 2.1.

I selected a diffuse prior for both the slope and the intercept for normal and La Niña years. The intercept's diffuse prior was centered on 1 and the slope's on 0, simply to account for the fact that the regression was stated in terms of median damages so most years in the sample will have a value of one and show no trend related to ENSO.

I selected an informative prior for El Niño based on damages estimates from Ecuador and Peru from the 1982/83 El Niño compiled in Rosenzweig and Hillel [2008] using data from Glantz et al. [1991], as well as Peruvian and UN reports. Those estimates placed the damages of that disaster at roughly USD 10.5 b. (This estimate is presumably in 1982/83 USD, so it would be larger if adjusted to present dollars.) That disaster corresponded to an ONI index value slightly above 2. I also assume that an ONI index of 1 (the point that distinguishes between normal and El Niño conditions) results in median damages. This gives me two points on which I can base my prior beliefs about the slope and intercept of the line describing El Niño damages. The move from an ONI of 1 to an ONI index of 2.165 resulted in approximately 4045 percent greater damage than the median year. That is equivalent to a slope parameter of 33.93.

I set a prior of $b_{\text{El Niño}} \sim \mathcal{N}(33.93, 5.5^2)$. Based on Gelman and Hill [2007] this is equivalent to providing one direct observation of the slope parameter with a weight that is slightly less than one single data point (because the standard deviation of observed damages is 5.37 less than the standard deviation of the prior.) In other words, this prior is slightly less influential to the final estimation than any single data point in the regression.

Implicit in my belief about the slope of the line, is a similar prior about the intercept parameter in the regression. If the effect of El Niño is negligible, then the intercept of the El Niño regression is 1, indicating losses that are 100 percent of the median. If the effect of a

one point rise in my ONI index is to raise disaster damages to a level to approximately 35 times the median (the prior I assigned above), then then the intercept will be -33 (i.e. the line has a slope of 34 and runs through point $(1,1)$). This range (an intercept between 1 and -33) is summarized in the prior $a_{\text{El Niño}} \sim \mathcal{N}(-16.0, 17^2)$. This is a relatively diffuse prior, and has a weight of considerably less than one data point.

$$\begin{aligned}
\log \text{Jan-Jun damage as percent of median}_{\text{year t}} &\sim \mathcal{N}(\hat{y}_i, \sigma_y^2) \\
\hat{y}_i &= a_{\text{Niño phase}} \\
&\quad + b_{\text{Niño phase}} * \\
&\quad \text{mean Oct-Jan ONI index}_{\text{year t-1 through t}} \\
a_{\text{La Niña to normal}} &\sim \mathcal{N}(1, 1000) \\
a_{\text{El Niño}} &\sim \mathcal{N}(-16.0, 17^2) \\
b_{\text{La Niña to normal}} &\sim \mathcal{N}(0, 1000) \\
b_{\text{El Niño}} &\sim \mathcal{N}(33.93, 5.5^2) \\
\sigma_y^2 &\sim \mathcal{U}(0, 100)
\end{aligned} \tag{2.1}$$

The output from those regressions in table 2.1, indicates that with 95 percent probability, the slope on the El Niño regression is positive indicating that more extreme ONI index values are indeed associated with increased disaster damages. The mean slope for La Niña and normal seasons is close to 0, and 0 is with the 95 percent probability interval. That indicates a weak or non-existent relationship between disaster damages and ONI index values outside the El Niño range.

The slopes of the two regressions show no overlap in their 95 percent probability intervals. The regression indicates that an average ONI index reading of 2 for the months of October through January (historically strong El Niño conditions, which happened three times since 1970) was associated on average with a 1326 percent increase in economic damages due to flooding, mudslides, and epidemics - equivalent to roughly USD 3.4 b in absolute damages higher than during normal or La Niña conditions. (See figure 2.8 for more details.) Using the limits of the 95 percent probability intervals for the slope and intercept parameters, which imply the strongest and weakest link to ONI respectively, the regression suggests that the credible range for this figure is between USD 2.2 and 4.7 b.

Assuming that the probability of an extreme El Niño is roughly $\frac{3}{40}$ (with three large El Niño events since 1970), then the mean damages estimate suggests that the region will, at any given time, be interested in roughly USD 250 m of risk coverage against the ONI index.

The difference in slopes is clear in figure 2.8 which shows a scatter plot of damage data alongside the mean regression line and the 95 percent probability interval for that regression.

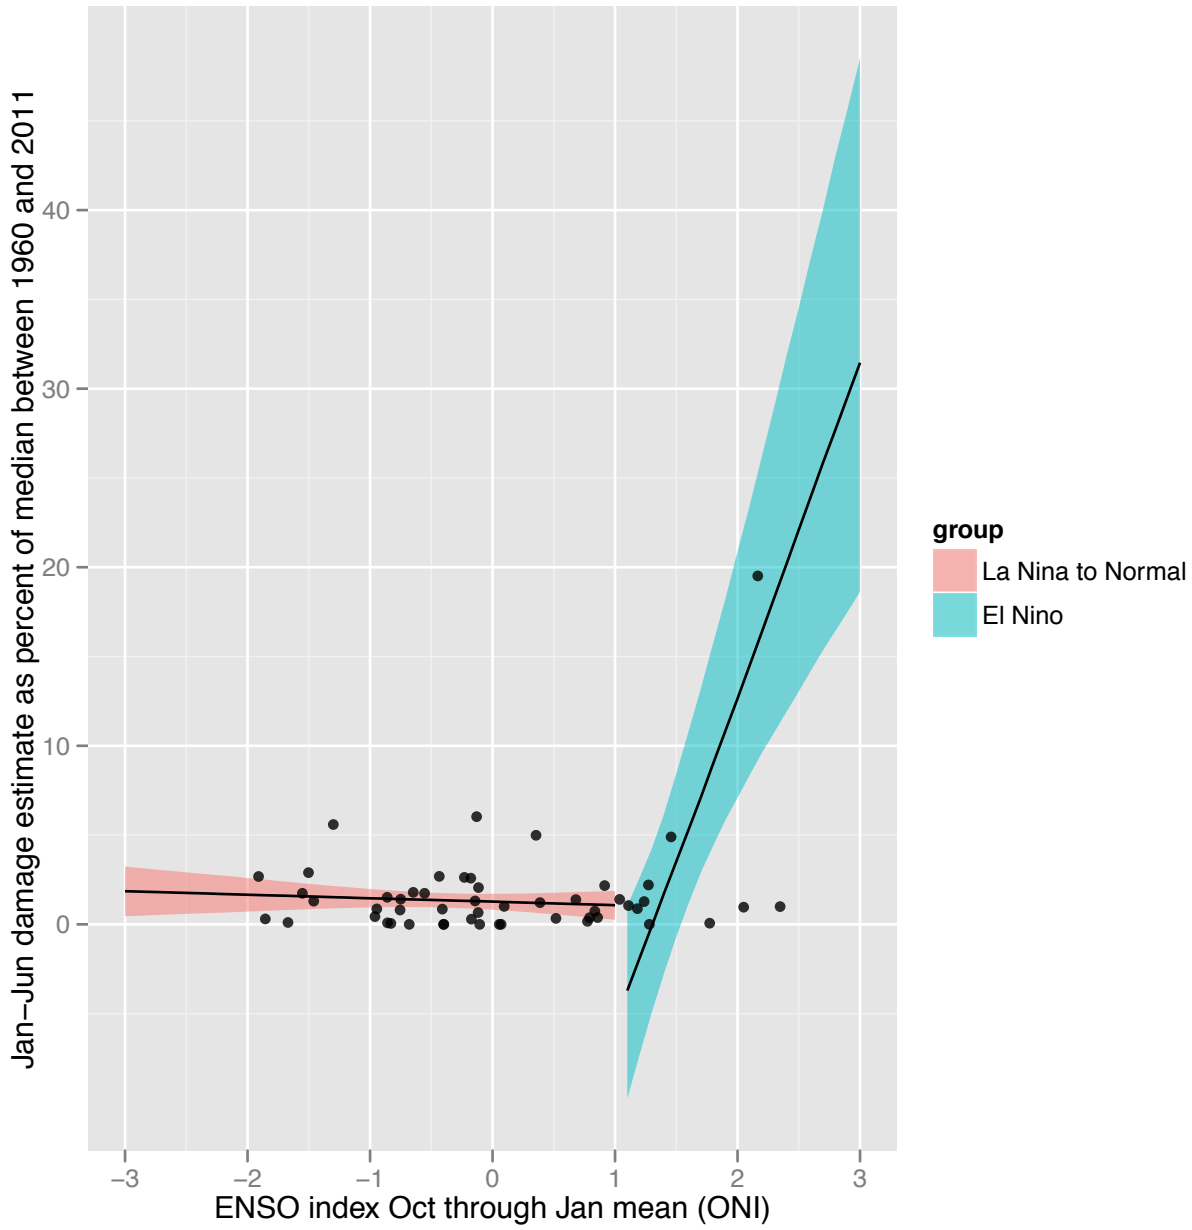

Figure 2.8: Bayesian regression of flood, landslide, and epidemic damages estimates from South America's Pacific coast, 1960-2011 predicted by ONI index

El Niño	Observed seasons	12							
	mean	sd	2.50%	25.00%	50.00%	75.00%	97.50%	$\hat{R}$	n.eff
a	-25.292	8.098	-42.289	-30.567	-24.921	-19.542	-10.655	1.0009	11000
b	19.278	5.371	9.637	15.445	19.046	22.855	30.392	1.0009	11000
sigma.y	9.910	3.464	5.059	7.442	9.282	11.665	18.451	1.0010	11000
La Nina to Normal	Observed seasons	41							
	mean	sd	2.50%	25.00%	50.00%	75.00%	97.50%	$\hat{R}$	n.eff
a	1.274	0.276	0.735	1.091	1.273	1.457	1.817	1.0013	4200
b	-0.194	0.322	-0.822	-0.406	-0.195	0.018	0.450	1.0009	11000
sigma.y	1.555	0.183	1.247	1.427	1.538	1.667	1.953	1.0010	11000

Table 2.1: Diagnostics for Bayesian regression of economic damages along South America's Pacific coast from January to June on ONI October to January average

East Pacific Asia and Oceania - Flooding from La Niña

Stable, liquid markets in teleconnection index risk will require balanced populations of hedger with opposing risks. Assume that flooding in South America creates a group of hedgers that want to receive payment in El Niño years. What region is the natural counter-party for this hedge? There are two ways to identify likely counter-parties:

1. We could look for regions which could be considered winners from an El Niño. For example a region that is often in water deficit and receives above-average rainfall without suffering floods during El Niño. However, these gains are likely to be modest relative to the sudden and catastrophic losses caused by extreme El Niño. To balance the market, you would need participation from many of these counter-parties.

2. Alternatively, we could look for regions and industries that face opposing losses across time. These counter-parties are perhaps less desirable than El Niño winners. Setting up trades between La Niña/normal phase losers and El Niño losers requires long-term commitments from both parties. In interest rate markets hedgers swap fixed and adjustable rates on loans based on indexed contingencies. Similarly, El Niño hedgers could receive a lower interest rate on an outstanding loan (in the case of El Niño) and visa versa for La Niña hedgers. That type of hedge, while normally accomplished through OTC swaps markets today, could be mediated by futures and options on futures.

La Niña related flooding in Pacific Asia and Oceania could be the major driver of hedging activity in that second category. In this section, I analyze the EM-DAT database for trends related to La Niña.

Using the average ONI index between October and January, the largest La Niña events in recent history occurred in 1973/74, right on the heels of the Modoki El Niño of 1972/73, and in 1988/89[12]. See figure 2.9 for details.

[12] By some index measures the 2010 and 2011 La Niñas were also among the strongest on record.

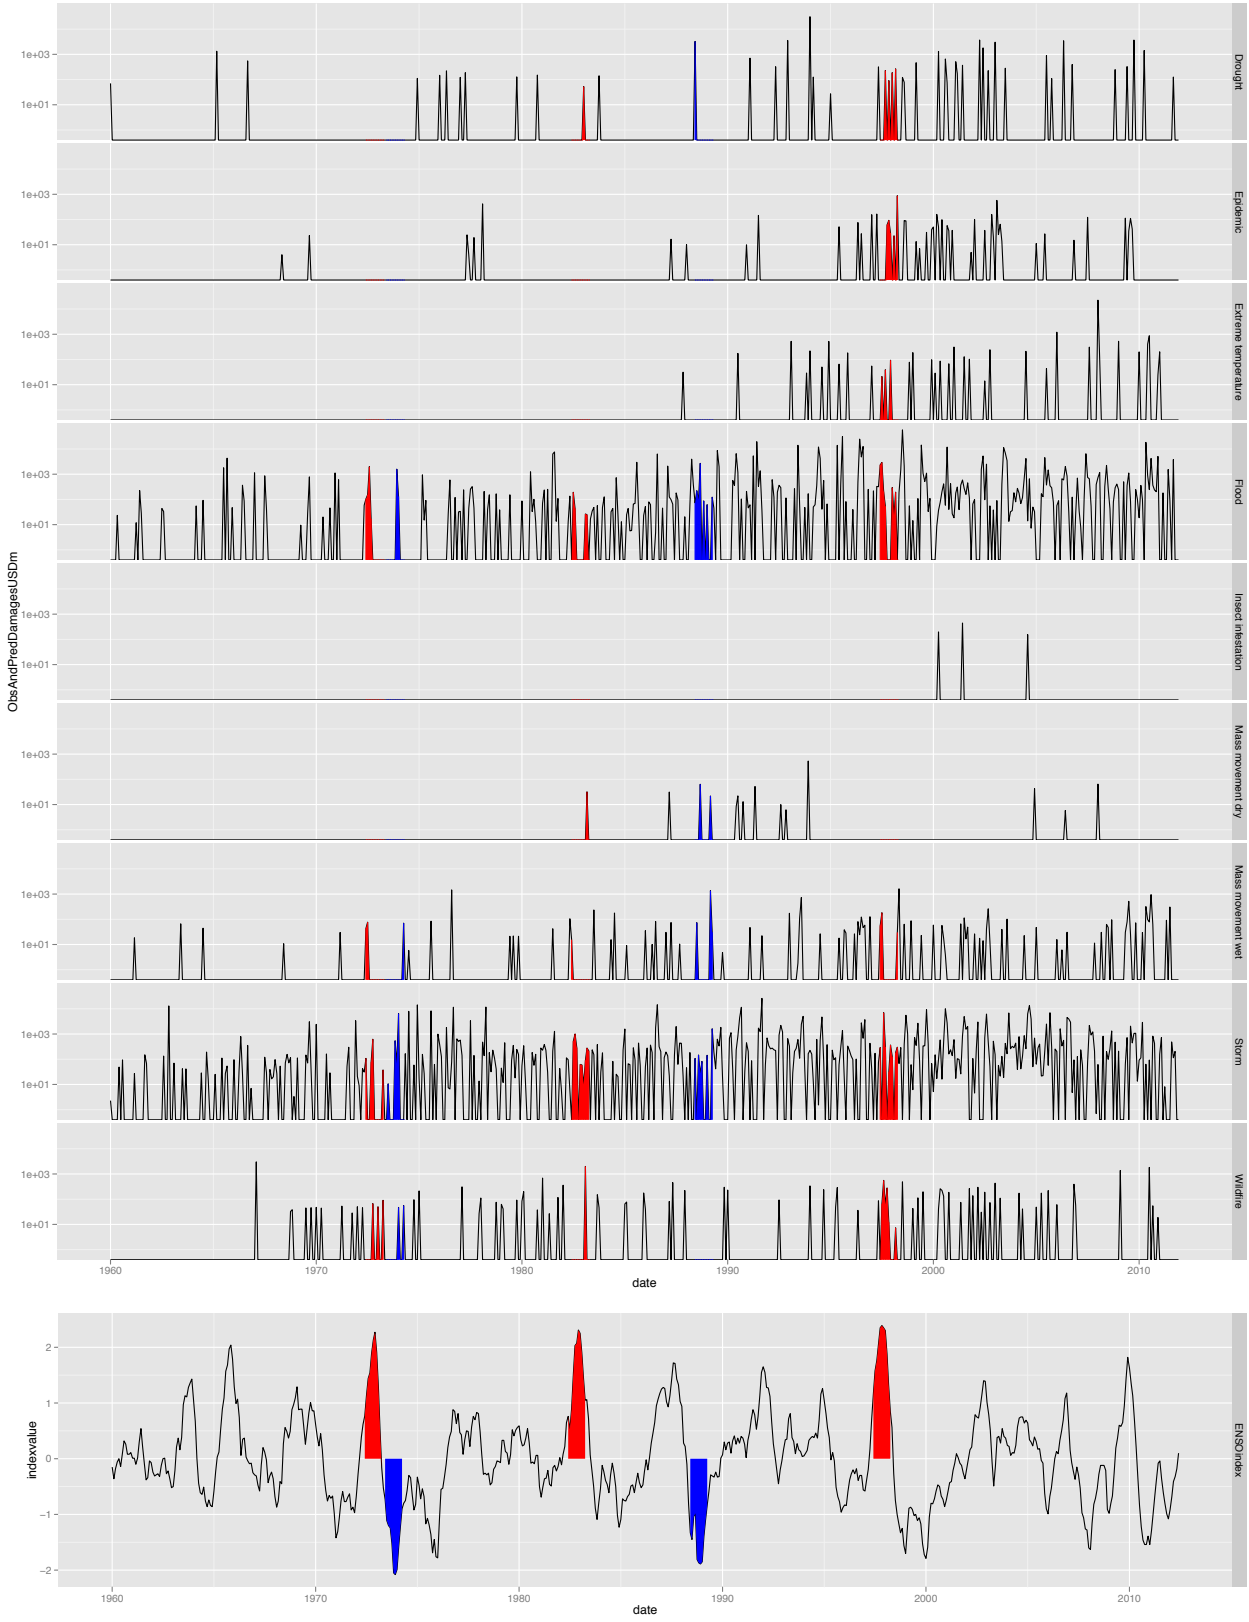

Figure 2.9: Disaster damage estimates by disaster type for countries in Pacific Asia and Oceania compared to ENSO (ONI) index

ENSO events generally begin in the Central Pacific with a slowdown of the atmospheric/oceanic cycle that brings water upwelling off the South American coast toward Indonesia. As that cycle slows, often beginning as early as January (i.e. January 1997 for the 1997/98 El Niño), by April of that year those changes are visible in the Eastern and Central Pacific sea surface temperatures. Roughly by September, still in advance of the impacts felt on South America's Pacific coast (roughly in the first six months of the following year, 1998 for the 1997/98 El Niño), persistent sea surface temperate anomalies result in changes in patterns of precipitation in Pacific Asia and Oceania. (See figure 2.1 for more detail on the calendar of events.) For that reason, I analyze aggregate flood damage in the region from September of year t though August of year $t+1$ for its connection to the average ONI index between October of year t and January of year $t+1$ (the same index used in the South America section). This division means that in 1972/73, when El Niño conditions quickly changed to La Niña conditions in 1973/74, the flooding that occurred in September of 1973 in Pacific Asia and Oceania is matched with the 1973/74 ONI readings.

Rather than analyze raw disaster data, I again set the seasonal disaster impact as a percentage of the median through the period of study. The median estimated flood damages in Pacific Asia and Oceania between 1960 and 2010 was USD 3.07 b between September of year t and August of year $t+1$. Those damages covered 523 disasters in the EM-DAT database.

I performed the Augmented Dickey-Fuller Test and the Phillips-Perron Unit Root Test on the damage data. While the Phillips-Perron test favored the alternative hypothesis of stationarity with greater than 95 percent confidence, the Augmented Dickey-Fuller Test failed to reject the null hypothesis of non-stationarity. This indicates that there may be long term trends in flood damage in the region which could produce spurious correlations on OLS regressions. Stationarity is not strictly required for Bayesian analysis, because the underlying parameters of the regression are considered stochastic.

The damage time series did not, however, show significant autocorrelation using a standard autocorrelation function, indicating that there is only weak interaction between the values of one season and the next.

I selected an informative prior for La Niña damages by referencing my inference for the damages of El Niño on South America's coast. I believe that the influence of ENSO on flooding across this large region (Pacific Asia and Oceania) is more subtle than El Niño's effects on the Pacific Coast of South America. In the latter case, my analysis indicated that on average, a move from an ONI value of 1 to 2 provoked a twenty-fold increase in flood damages across the region relative to

normal or La Niña conditions. Hence, I assume extreme La Niña (an ONI value of -2) will result in flooding in Pacific Asia and Oceania somewhere between the median for region and five times above the median.

When I combine this belief with the belief that normal conditions will result in median losses across the region, I can also make inferences about the intercept parameter in my regression. If the effect of La Niña is negligible, then the intercept of the La Niña regression is 1, indicating losses that are 100% of the median. If the effect of La Niña is equal to the effect of extreme El Niño in South America, then the intercept will be -4. This range (an intercept between 1 and -4 is summarized in the prior $a_{\text{La Niña}} \sim \mathcal{N}(-1.5, 2.5^2)$.

Given the tendency of some Pacific Islands to suffer from catastrophic flooding during El Niño, despite the regions' tendency toward drought, I broke the regression into three parts, rather than two [13]. This resulted in the regression equation listed in equation 2.2.

[13] UCAR. El Niño and climate prediction: Reports to the nation on our changing planet. Technical Report 3, University Corporation for Atmospheric Research (UCAR), 1994

$$
\begin{aligned}
\log \text{Jan-Jun damage as percent of median}_{\text{year t}} &\sim \mathcal{N}(\hat{y}_i, \sigma_y^2) \\
\hat{y}_i &= a_{\text{Niño phase}} \\
&\quad + b_{\text{Niño phase}} * \\
&\quad \text{mean Oct-Jan ONI index}_{\text{year t-1 through t}} \\
a_{\text{La Niña}} &\sim \mathcal{N}(-1.5, 2.5^2) \\
a_{\text{normal}} &\sim \mathcal{N}(1, 1000) \\
a_{\text{El Niño}} &\sim \mathcal{N}(-1.5, 2.5^2) \\
b_{\text{La Niña}} &\sim \mathcal{N}(2.5, 2.5^2) \\
b_{\text{normal}} &\sim \mathcal{N}(0, 1000) \\
b_{\text{El Niño}} &\sim \mathcal{N}(-2.5, 2.5^2) \\
\sigma_y^2 &\sim \mathcal{U}(0, 100)
\end{aligned}
\tag{2.2}
$$

The output from those regressions in table 2.2, indicate that:

- With 90 percent probability, the slope on the La Niña regression is negative. Hence, more extreme ONI index value are associated with increased flood damage in the region;
- With 95 percent probability, the slope on the El Niño regression is positive. So, extreme positive ONI index values are also associated with increased flood damage in the region;
- A slope of 0 for normal conditions is within the 95 percent confidence interval, but that interval is biased toward negative values;
- The 50 percent probability intervals of each the regressions' slopes are distinct, but there is some overlap between the 95 percent probability interval of all three slope parameters.

This indicates that while extreme values of the ONI index likely influence flood damage in the region, the connection would be clearer if I included information from climate research through Bayesian priors (as I do below for Atlantic storm damage) and changed the scale of analysis to the sub-regional level.

Based on the regression, the expected impact of a La Niña event of the same magnitude as that of 1988, an ONI index of -1.85 which was reached twice since 1970, was an 261 percent increase in regional flood damages relative to the median of slightly more than USD 8 b in absolute damages. (See figure 2.10 for more details.) While the impact of La Niña on flooding across the region is less pronounced than that of El Niño in South America, the expected damages are large in aggregate. In fact they are large enough to fully offset the hedging interest generated by El Niño along South America's Pacific coast, even after accounting for the fact that there have been only two major La Niñas since 1970 versus three major El Niños over the same period (i.e. adjusting for the probability of the extreme event in question by $\frac{2}{40}$ rather than $\frac{3}{40}$.)

My analysis indicates that futures and options on futures for ENSO index risk would enjoy large balanced hedging interest. However, market professionals will need to find clever ways to link natural counterparties in the market across time such that Asian hedgers are willing to insure the losses of South American hedgers during El Niño years and visa versa for La Niña years.

La Niña	Observed seasons	9							
	mean	sd	2.50%	25.00%	50.00%	75.00%	97.50%	$\hat{R}$	n.eff
a	-1.045	1.756	-4.511	-2.213	-1.032	0.142	2.373	1.0010	11000
b	-1.979	1.208	-4.376	-2.778	-1.973	-1.174	0.393	1.0010	11000
sigma.y	2.205	0.713	1.286	1.721	2.060	2.517	3.913	1.0012	5800
Normal	Observed seasons	33							
	mean	sd	2.50%	25.00%	50.00%	75.00%	97.50%	$\hat{R}$	n.eff
a	2.012	0.559	0.927	1.642	2.002	2.379	3.128	1.0011	10000
b	-0.534	0.922	-2.352	-1.143	-0.529	0.063	1.286	1.0010	11000
sigma.y	3.097	0.420	2.397	2.802	3.051	3.343	4.031	1.0009	11000
El Niño	Observed seasons	12							
	mean	sd	2.50%	25.00%	50.00%	75.00%	97.50%	$\hat{R}$	n.eff
a	-0.714	2.086	-4.826	-2.118	-0.708	0.694	3.347	1.0010	11000
b	3.354	1.546	0.298	2.316	3.347	4.389	6.427	1.0012	5600
sigma.y	7.334	1.792	4.764	6.077	7.032	8.260	11.760	1.0012	6000

Table 2.2: Diagnostics for Bayesian regression of economic damages in Pacific Asia and Oceania from September to August on ONI October to January average

North America and Caribbean - Storms from El Niño

So far, I have discussed catastrophes that are direct results of ENSO anomalies. What distinguishes ENSO as a teleconnection index is its

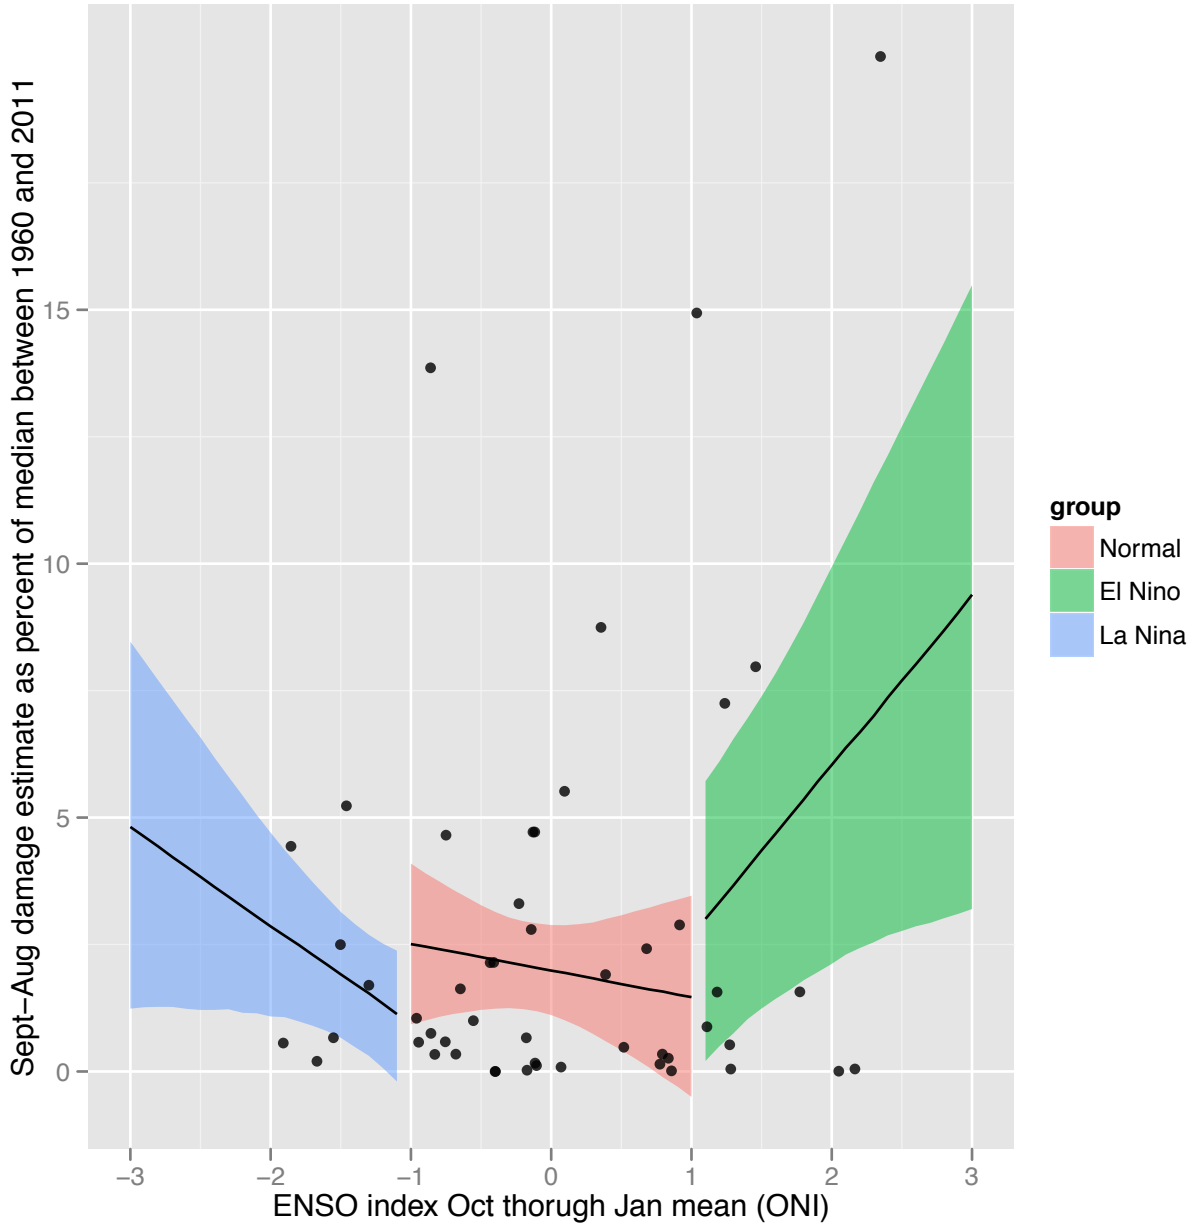

Figure 2.10: Bayesian regression analysis of flood damage in the East Asian Pacific and Oceania predicted by ENSO index

ability to reshape weather patterns across the globe, affecting regions and weather phenomenon with no obvious immediate connection to the ENSO event itself. This presents an opportunity. The hedging activity generated by otherwise disparate weather events could, in part, be driven to one central ENSO market, providing the liquidity, competition, and collective information that will drive ENSO protection prices steadily downward.

ENSO's link to Atlantic hurricanes is likely the most important of the true teleconnections (as opposed to the more direct linkages between ENSO and precipitation in the equatorial Pacific) that would drive liquidity on an ENSO market. Hedging interest related to hurricane damage might naturally come from the global reinsurance industry, the large insurers of insurers that specialize in spatially correlated risks. (See chapter 7 for more details on hedging interest from reinsurers.)

US hurricanes represent the largest single non-life risk in the portfolios of large reinsurers, hence any cost-effective hedging instrument that could help them share Atlantic hurricane risk, or bring non-hurricane risk into their portfolios to offset that hurricane risk, should be very valuable to reinsurers[14]. Figure 2.11 shows my own measure of CAT bond issuance by peril type between 1996 and 2011 (2012 data runs through March). CAT bonds are often considered a substitute for reinsurance. They also reflect the risk that reinsurers want to transfer out of their portfolios. The thick pink line in figure 2.11 is my estimate of cat bond issuance specific to US hurricane risk. As you can see, Atlantic hurricane risk has dominated catastrophe bond issuance virtually every year since the market's inception in 1996. (Figure 6.6 provides another estimate from a reinsurance brokerage with more aggregation of issuance.)

[14] IAIS. Reinsurance and financial stability. policy paper, International Association of Insurance Supervisors (IAIS), November 2011

Given the concentration of hurricane risk in reinsurance portfolios, economic theory suggests that reinsurers should be excited by the opportunity to hold positions that are not highly correlated to hurricanes, particularly in a market where they feel they have expertise.

But ENSO wouldn't be just an uncorrelated market. It would be a negatively correlated market, actually offsetting hurricane risk. Climatological studies suggest that landfalls of major hurricanes along the east coast of the United States and the Caribbean are historically less likely during El Niño years than during normal ENSO phase or La Niña conditions[15] [16] [17] [18]. This result, corroborated by repeated studies of different data sets spanning three decades, means that reinsurers selling El Niño protection will be paying out on contracts in years where their hurricane losses are light and receiving payments in years where the rest of their portfolios are suffering.

[15] William M Gray. Atlantic seasonal hurricane frequency. part 1: El Niño and 30 mb Quasi-biennial Oscillation influences. *Monthly Weather Review*, 112(9):1649–1688, 1984a

[16] William M Gray. Atlantic seasonal hurricane frequency. part 2: Forecasting its variability. *Monthly Weather Review*, 112:1669, 1984b

[17] R.M. Wilson. Statistical aspects of major (intense) hurricanes in the Atlantic basin during the past 49 hurricane seasons (1950-1998): Implications for the current season. *Geophysical Research Letters*, 26(19): 2957–2960, 1999

[18] P.J. Klotzbach. El Niño-Southern Oscillation's impact on Atlantic basin

The EM-DAT database includes 616 separate catastrophic storm

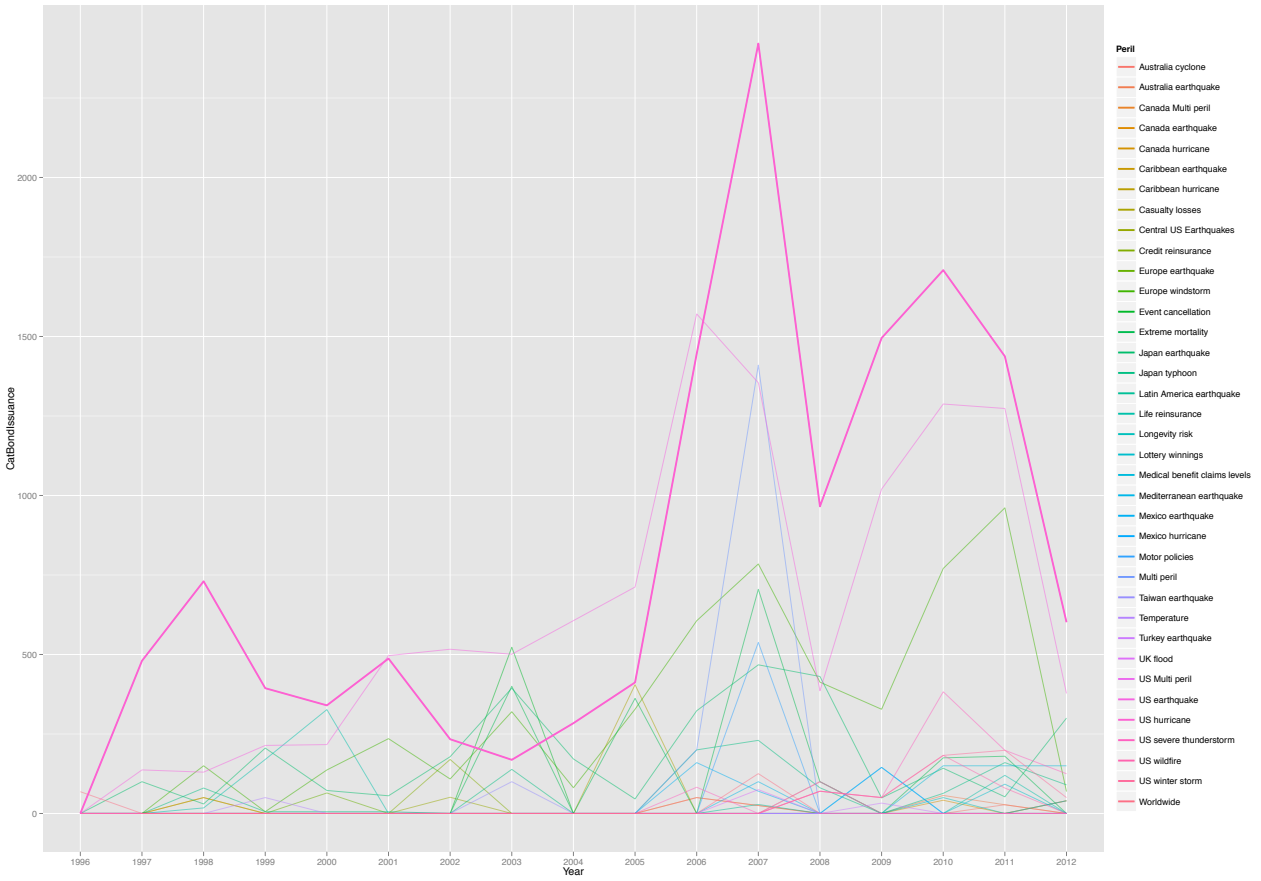

Figure 2.11: Catastrophe bond issuance based on Artemis.bm's Catastrophe Bond Deal Directory with analysis by author

events between 1960 and 2010 impacting North America and the Caribbean between the months of June and November (the traditional hurricane season). The median damage across the region for each hurricane season is roughly USD 10.5 b, many times larger than for floods in Southeast Asia and Oceania, or Pacific South America. Economic damage estimates for all weather-linked disaster types across the region is displayed alongside the ONI index in figure 2.12. No clear trends are visible in the raw data, apart from a rise in the average damages due to flooding and storms.

The Phillips-Perron Unit Root test favored stationarity for both time series (for the index and damages as a percent of the seasonal median) with 95 percent confidence. The Augmented Dickey Fuller Test found a p-value of 0.05 for the damage series. Neither time series showed significant autocorrelation using a standard autocorrelation function.

Based on Klotzbach [2011] I created an ENSO index for hurricanes by averaging ONI index values between August and October for any given year (e.g. the 2010 hurricane season spanning June through November 2010 is matched with the average ONI index values for August through October 2010.) I then regressed hurricane damages on the seasonal ONI index average as in equation 2.3. Distinct from the other risk-regions analyzed in this chapter, I estimated regression coefficients both for the individual ENSO phases (e.g. $a_{\text{La Niña}}$, a_{normal}, $a_{\text{El Niño}}$) and for all ENSO phases pooled together (e.g. a_{pooled}). A pooled regression only makes sense in this case because previous literature suggested that there may be a straight-forward inverse relationship between ENSO and hurricanes, with high ENSO values producing low levels of hurricane damage and visa-versa.

The pooled and normal coefficients were given uninformative priors while the El Niño and La Niña coefficients were given priors based on Klotzbach [2011], Bove et al. [1998], and Pielke Jr and Landsea [1999]. Klotzbach [2011] found that average number of major hurricanes per year for El Niño years was 1.5, compared to 2.1 for years when ENSO was in a normal phase. Bove et al. [1998] suggest that the ratio of probabilities of a major US landfall in El Niño was 23 percent versus 58 percent in normal years. Finally, Pielke Jr and Landsea [1999] suggested that the mean damage in El Niño seasons was USD 1997 2.0 b compared to 6.9 b for normal seasons.[19] So given that three studies agreed that the burden of major hurricanes was roughly a third of its normal value in El Niño years[20], I constructed a prior that would put an average El Niño event (roughly an ONI index of 1.2) at a level of damage that was one-third the median across all seasons. I further constrained that prior so as to avoid making the prior so steep as to suggest no damage for an ONI index of 3 (larger than any on record

[19] Note that Pielke Jr and Landsea [1999]'s median damage estimate, made in 1997, is well below the current USD 10.5 b suggested by EM-DAT.

[20] Note that the studies used slightly different definitions of El Niño.

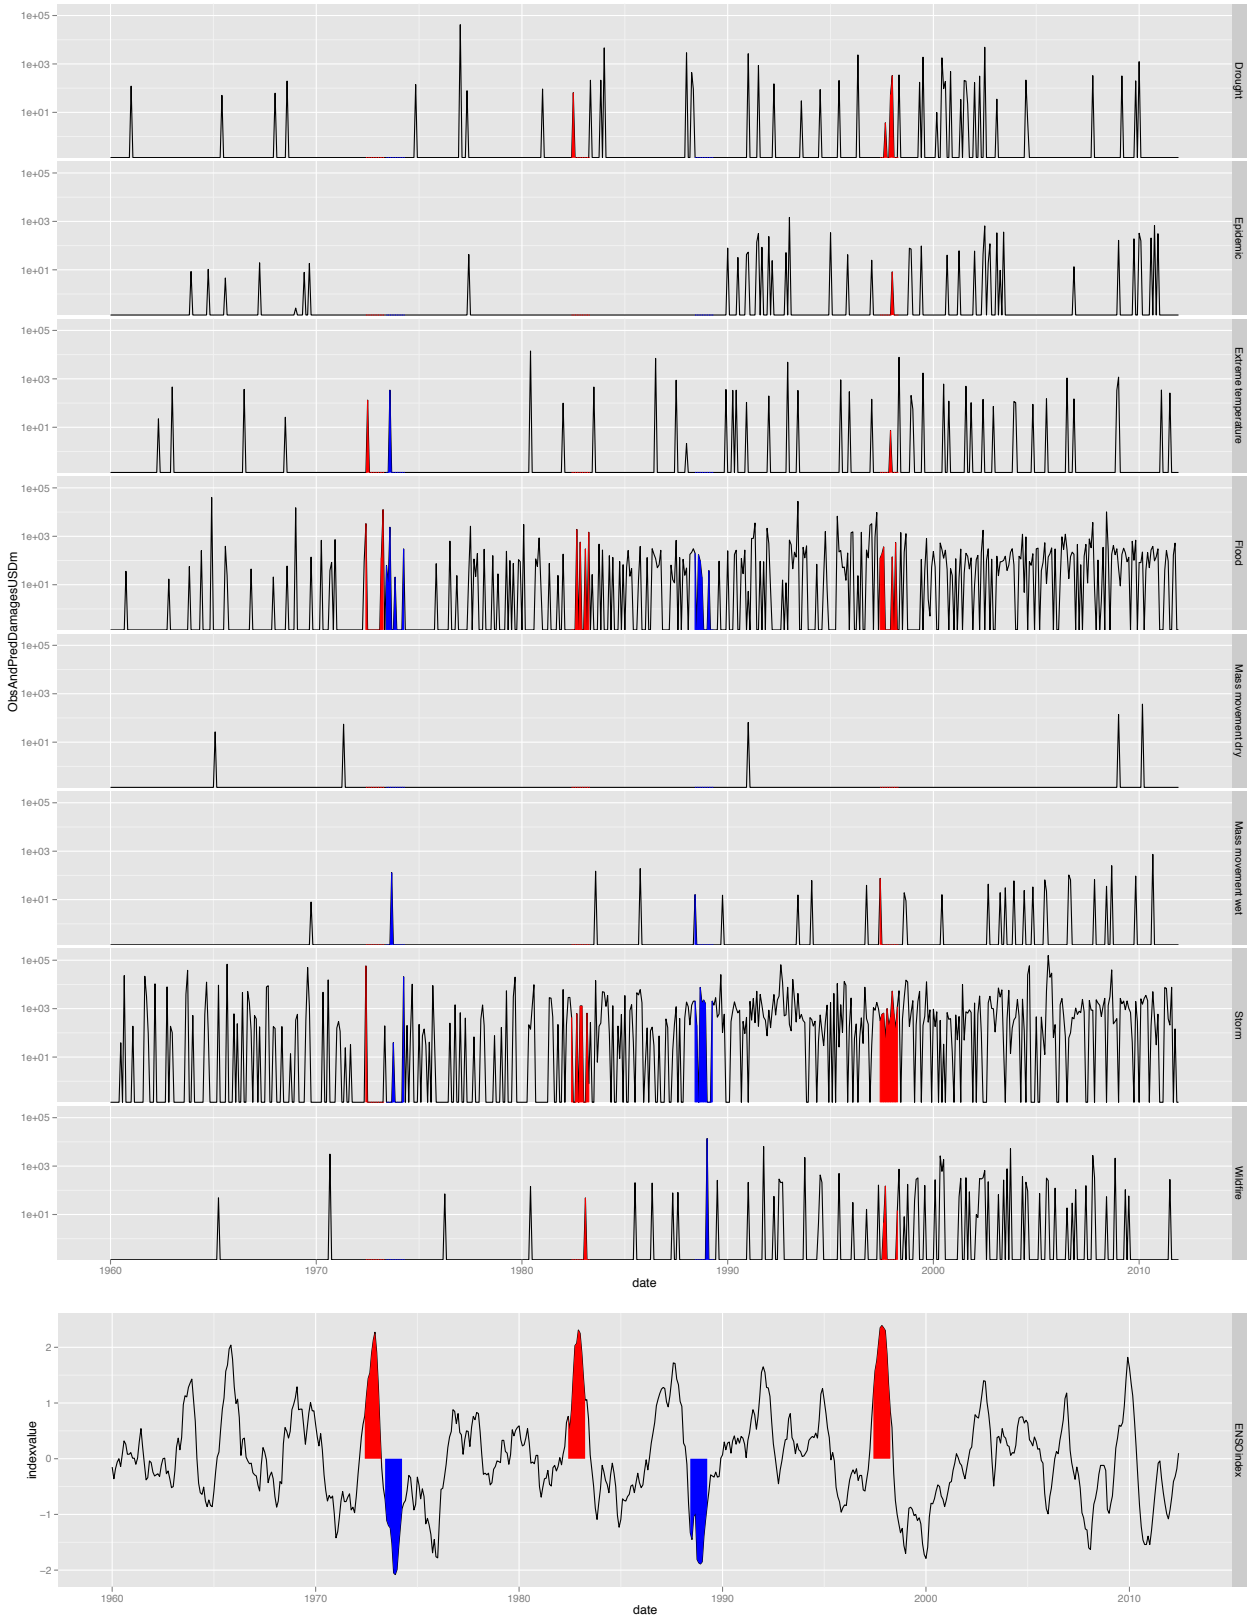

Figure 2.12: Disaster damage estimates by disaster type for countries in North America and the Caribbean compared to ENSO (ONI) index

but well within the realm of possibility).

The three studies showed similar agreement for La Niña. Klotzbach [2011] suggested that the average La Niña season has a hurricane burden two-thirds that of a normal season. Bove et al. [1998] and Pielke Jr and Landsea [1999] found that the average La Niña season has hurricane burden slightly lower than that of a normal phase season. Based on those studies, I constructed a prior suggesting that the average La Niña season (an ONI value of roughly -1.2) had damages that were between one-third higher than the median and one-half of the median. Both these sets of priors are highly informative. I believe that informative priors are justified by the concurrent findings of previous climate research on separate databases.

$$
\begin{aligned}
\log \text{Jun-Nov damage as percent of median}_{yeart} &\sim \mathcal{N}(\hat{y}_i, \sigma_y^2) \\
\hat{y}_i &= a_{\text{Niño phase}} \\
&\quad + b_{\text{Niño phase}} * \\
&\quad \text{mean Aug-Oct ONI index}_{\text{year t}} \\
a_{\text{pooled}} &\sim \mathcal{N}(1, 1000) \\
a_{\text{La Niña}} &\sim \mathcal{N}(1, 1.75^2) \\
a_{\text{normal}} &\sim \mathcal{N}(1, 1000) \\
a_{\text{El Niño}} &\sim \mathcal{N}(0, 0.5^2) \\
a_{\text{pooled}} &\sim \mathcal{N}(0, 1000) \\
b_{\text{La Niña}} &\sim \mathcal{N}(-0.66, 1.25^2) \\
b_{\text{normal}} &\sim \mathcal{N}(0, 1000) \\
b_{\text{El Niño}} &\sim \mathcal{N}(-0.13, 0.4^2) \\
\sigma_y^2 &\sim \mathcal{U}(0, 100)
\end{aligned}
\tag{2.3}
$$

Table 2.3 provides the output of the regressions in equation 2.3. The grouped regression suggests the opposite of what I expected to find - a slightly positive slope coefficient. For the phase-specific regressions, a slope coefficient of 0 is well within the 95 percent probability interval for all three phases, suggesting no clear relationship between changes in the ONI index and storm damages in EM-DAT. However, the intercept parameter for El Niño is likely below that for normal ENSO phase seasons. The El Niño and normal phase intercept parameters begin to overlap at the 94 percent and 6 percent quantiles respectively, suggesting that they are distinct with 88 percent probability.

Figures 2.13 and 2.14 provide a more complete picture of the regression findings. Even with strong guidance from previous studies, the regional damage data in EM-DAT remains noisy and involves too few observations from La Niña and El Niño years to provide a strong inference about hedging interest on an ENSO exchange. However, damages

grouped	observed seasons	51							
	mean	sd	2.50%	25.00%	50.00%	75.00%	97.50%	$\hat{R}$	n.eff
a	1.931	0.472	1.009	1.615	1.931	2.245	2.876	1.0012	6300
b	0.454	0.561	-0.657	0.084	0.457	0.825	1.552	1.0009	11000
sigma.y	3.357	0.348	2.769	3.112	3.326	3.569	4.130	1.0012	5300
La Niña	observed seasons	6							
	mean	sd	2.50%	25.00%	50.00%	75.00%	97.50%	$\hat{R}$	n.eff
a	0.984	1.214	-1.381	0.176	0.975	1.800	3.341	1.0012	6800
b	-0.194	0.941	-2.040	-0.829	-0.186	0.440	1.652	1.0010	11000
sigma.y	1.374	0.659	0.662	0.958	1.211	1.592	3.032	1.0011	10000
Normal	observed seasons	39							
	mean	sd	2.50%	25.00%	50.00%	75.00%	97.50%	$\hat{R}$	n.eff
a	1.940	0.593	0.776	1.552	1.934	2.328	3.112	1.0009	11000
b	0.310	1.185	-2.036	-0.478	0.312	1.073	2.637	1.0011	9000
sigma.y	3.706	0.449	2.951	3.388	3.662	3.982	4.692	1.0011	9100
El Niño	observed seasons	7							
	mean	sd	2.50%	25.00%	50.00%	75.00%	97.50%	$\hat{R}$	n.eff
a	0.259	0.493	-0.702	-0.070	0.268	0.590	1.219	1.0011	9300
b	0.131	0.390	-0.661	-0.125	0.142	0.394	0.886	1.0014	3900
sigma.y	4.142	1.656	2.174	3.071	3.785	4.768	8.262	1.0010	11000

Table 2.3: Diagnostics for Bayesian regression of economic damages in North America and the Caribbean from June to November on ONI August to October average

are likely lower in El Niño years than during the normal ENSO phase.

On average, the estimated losses for a modest El Niño season (August through October ONI index average) are roughly 40 percent of the median across all seasons (USD 10.5 billion). Given the magnitude of median losses, this is a difference that, while not statistically significant with 95 percent probability, is of great economic importance. Given a 40 percent drop in hurricane damage during El Niño season, the reinsurance industry should gladly act as counter-party for any firm looking to purchase El Niño protection on a futures or options market. Their windfall due to the drop in hurricane damages should be enough to cover the full range of estimated impacts on South America's Pacific costs (USD 2.2 to 4.7 billion) generated by an extreme El Niño.

El Niño drought regions - Southern Atlantic, Indian Ocean basin, East Asia, Oceania

El Niño is associated with drought and wildfire across large swaths of the globe. In this analysis I grouped disaster data from all the regions strongly suspected of suffering from El Niño related drought. That includes most of of the Indian Ocean Basin, as well as the region most associated with La Niña flooding, Pacific East Asia and Oceania. It also includes the Brazil and the countries of the Sahel. See the damage time series for this region is displayed along side the ONI index in figure 2.15 for details.

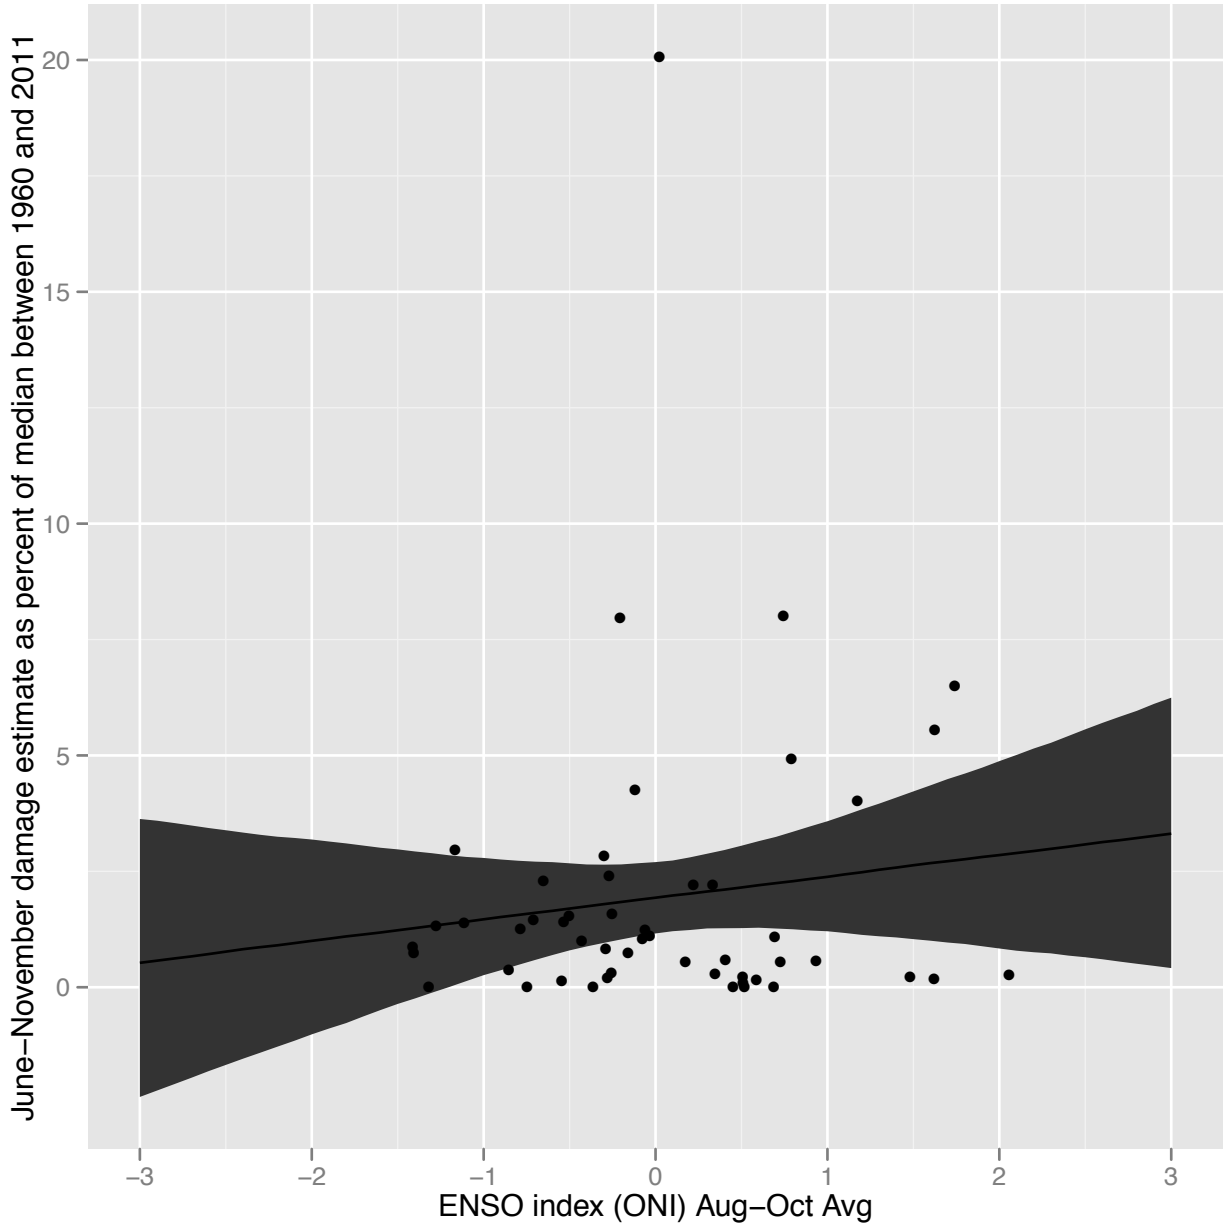

Figure 2.13: Bayesian regression analysis of damage estimates from storms and flooding in North America and the Caribbean predicted by ENSO index, pooled across ENSO phases

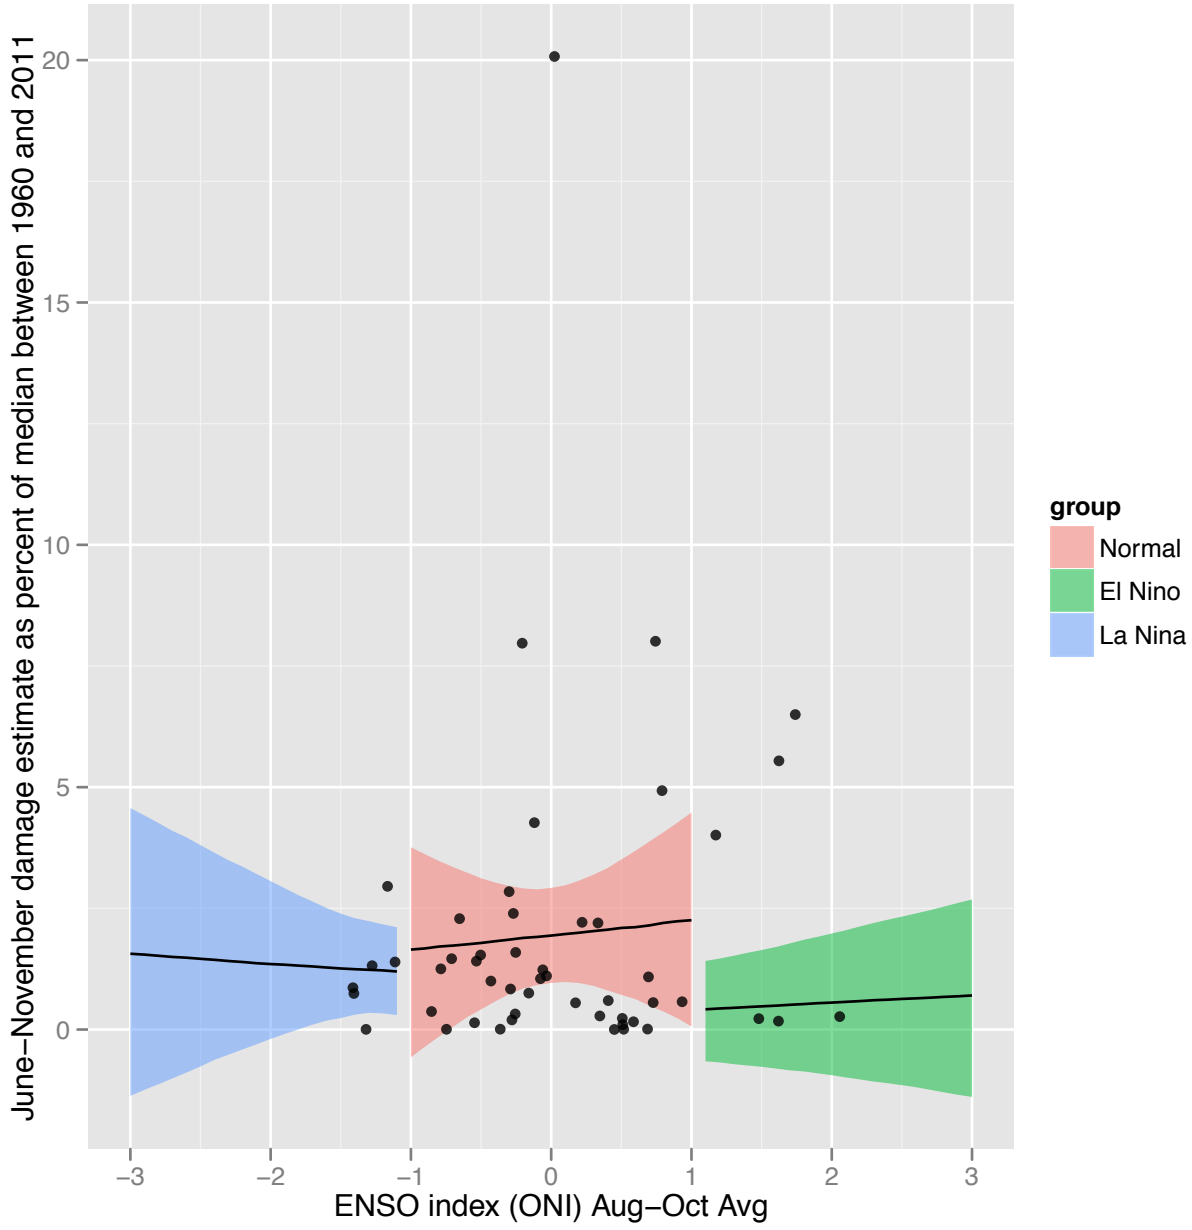

Figure 2.14: Bayesian regression analysis of damage estimates from storms and flooding in North America and the Caribbean predicted by ENSO index, separate regressions for each ENSO phase

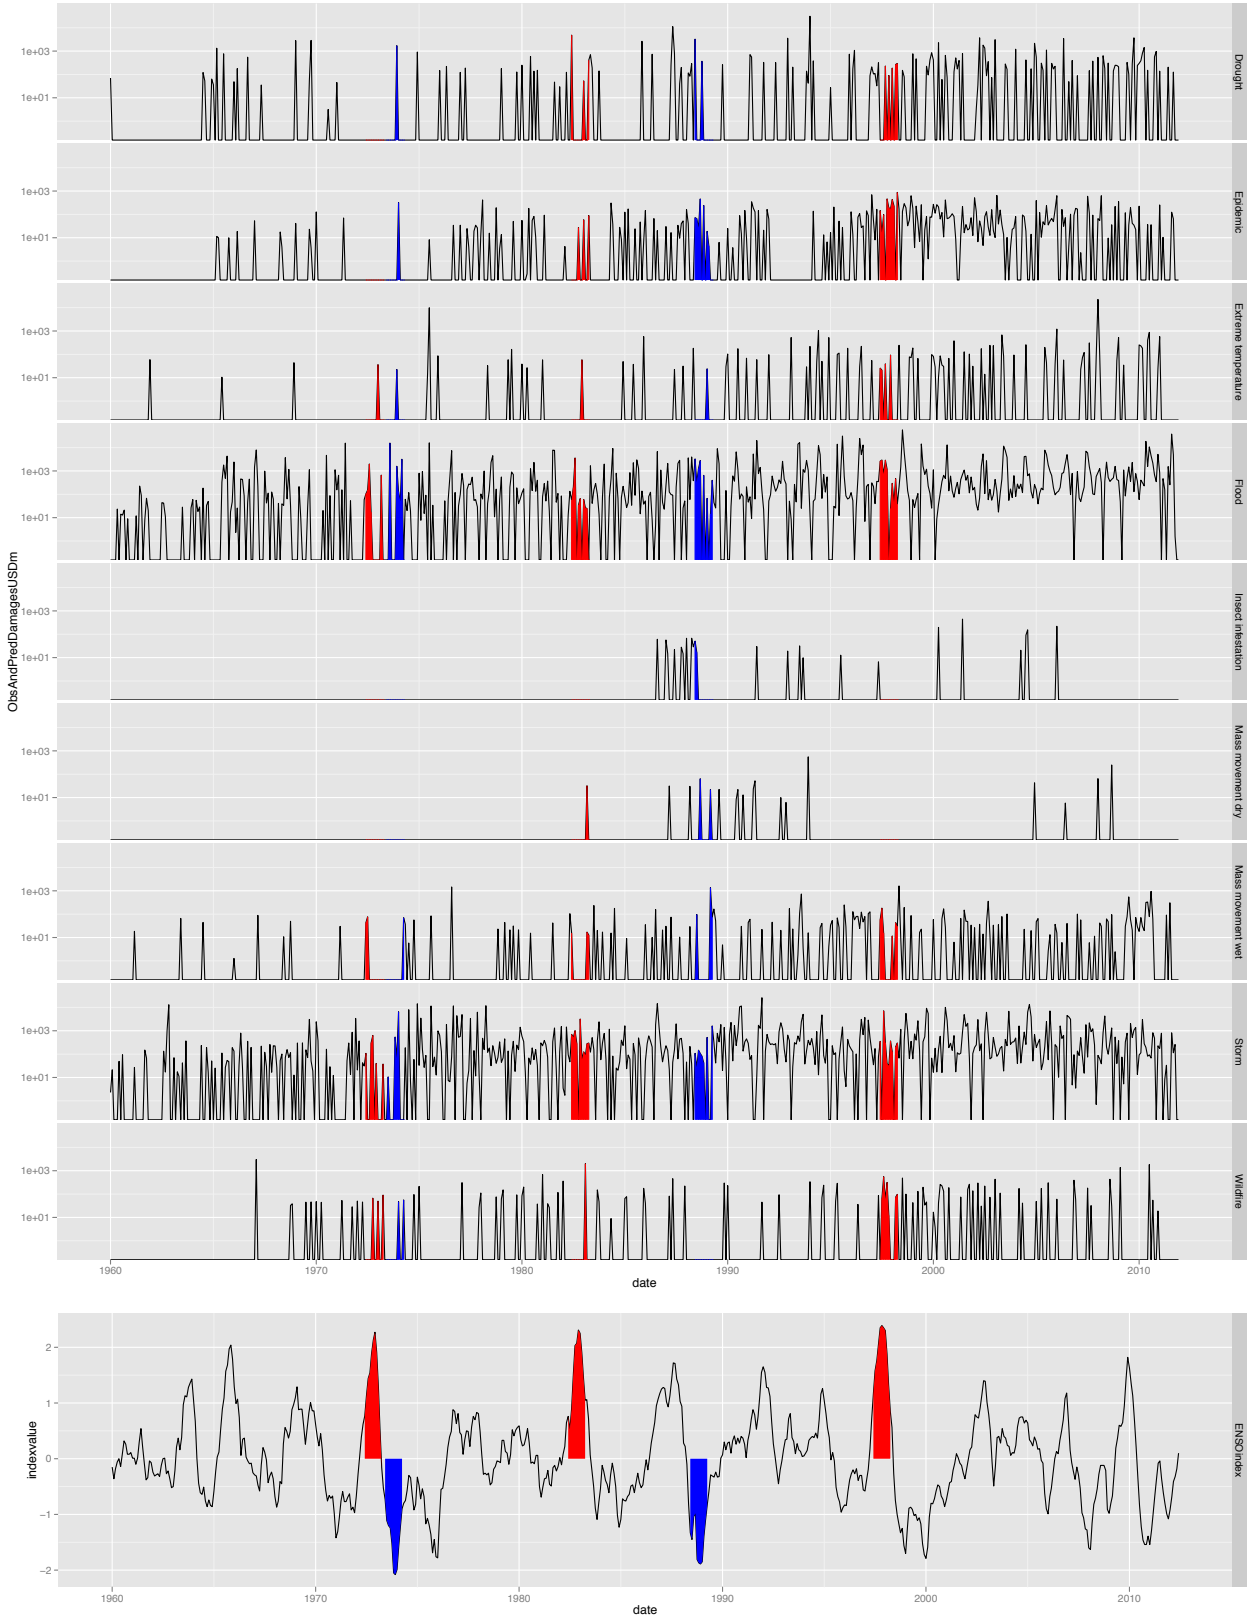

Figure 2.15: Disaster damage estimates by disaster type for countries in regions that are suspected to experience drought during El Niño events compared to ENSO (ONI) index

There are 299 droughts and wildfires in the EM-DAT database corresponding to this group of countries, with a median September through August (the same benchmark months used to measure flooding in Pacific Asia and Oceania) damage due to drought and wildfire of USD 1.7b. Each season's damage corresponded to average ONI index values measured between October and January of that damage season.

I performed the Augmented Dickey-Fuller Test and the Phillips-Perron Unit Root Test on the damage time series. (See the Pacific South America section for results from the corresponding index time series.) Both tests favored the alternative hypothesis of stationarity with greater than 95 percent confidence. The time series showed no significant autocorrelation using a standard autocorrelation function.

Dai et al. [1998] used a linear regression to describe the relationship between ENSO and drought across many regions of the world, without segregating the dataset into its constituent ENSO phases. I follow that convention here, presenting a single grouped regression, as in equation 2.4, rather than a series of regressions. To the extent that I believe (based on Dai et al. [1998]) that the relationship between ENSO and drought in these countries can be represented by a single line, then I prefer a single regression because that would maximized the number of observations in the sample.

The priors I chose for the regression are based on figure 2.16 which suggests that between 1979 and 1995, an El Niño with a severity that is two standard deviations away from the average, caused the burden of drought across the world to be approximately 1.75 standard deviations above its average, over the full period of study. Between 1900 and 1978 that same magnitude El Niño was associated with drought burden approximately 1.5 standard deviations above the sample average. Assuming that the drought index used in the study is a reliable proxy for drought in my sample, I used the standard deviation from the sample to translate these observations from Dai et al. [1998], along with observations for the intercept of each regression line, into a likely range for the parameter values in my regression. I doubled the standard deviation of the prior relative to the standard deviation suggested by the range in figure 2.16 to allow additional flexibility in the regres-

sion. The resulting priors are presented in equation 2.4

$$
\begin{aligned}
\log \text{Sep-Aug damage as percent of median}_{yearthought+1} &\sim \mathcal{N}(\hat{y}_i, \sigma_y^2) \\
\hat{y}_i &= a \\
&\quad + b * \\
&\quad \text{mean Oct-Jan ONI index}_{\text{year t-1 through t}} \\
a &\sim \mathcal{N}(2.625, 3.25^2) \\
b &\sim \mathcal{N}(1.8, 2.5^2) \\
\sigma_y^2 &\sim \mathcal{U}(0, 100)
\end{aligned}
\tag{2.4}
$$

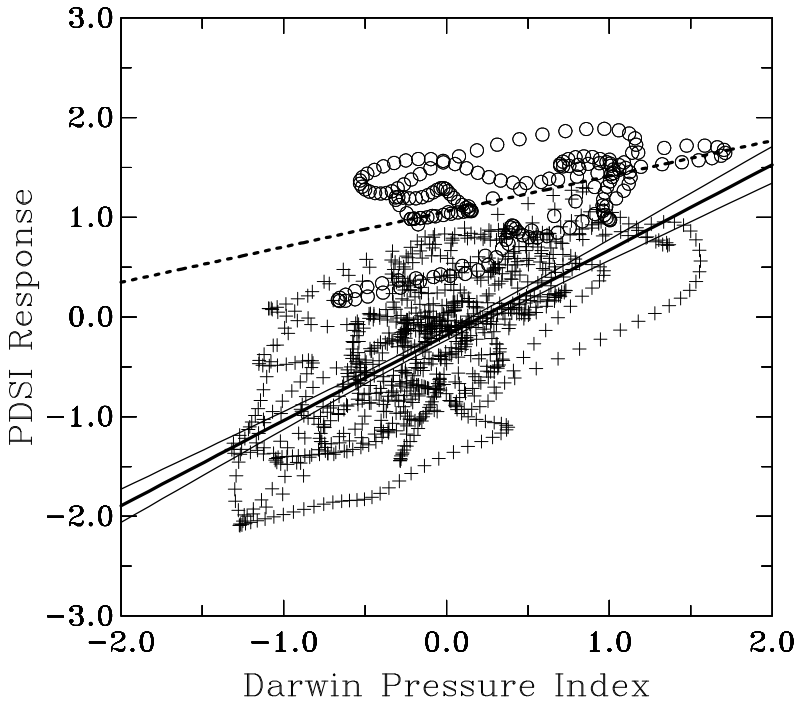

Figure 2.16: Reprinted from Dai et al. [1998] - a scatter plot of the first eigenvalue of the Empirical Orthogonal Function for a common drought index, the Palmer Drought Severity Index (PDSI) (response) plotted across the world versus the Darwin pressure index (a measure of ENSO strength) from six months previous. The crosses are monthly data points for 1900-1978 and the circles are for 1979-1995. The thick solid line is the linear regression for 1900-1978 and the thin lines are the 99% confidence intervals. The dashed line is the regression for 1979-1995.

The output from those regressions in table 2.4, place a slope coefficient of 0 well within the 95 percent probability interval. (See figure 2.17 for more detail.) This indicates that despite informative priors, the data in the EM-DAT database are too noisy to discern any relationship between ENSO and drought in these regions. The median economic burden of drought on the region is modest relative to those of the other peril/region groups studied here, so not only is the relationship statistically weak, but it is also of less economic consequence than the other relationships analyzed in this chapter.

Based on this regression I have decided against including economic damage from likely El Niño drought regions in my estimate of hedging

interest for an exchange traded ENSO market.

All ENSO phases	Observed seasons	51							
	mean	sd	2.50%	25.00%	50.00%	75.00%	97.50%	$\hat{R}$	n.eff
a	1.660	0.399	0.869	1.399	1.658	1.926	2.453	1.0009	11000
b	0.295	0.382	-0.455	0.039	0.290	0.548	1.056	1.0010	11000
sigma.y	2.916	0.303	2.390	2.704	2.893	3.097	3.599	1.0013	4300

Table 2.4: Diagnostics for Bayesian regression of economic damages in North America and the Caribbean from June to November on ONI August to October average

How do these markets compare to other widely traded commodities?

This statistical analysis approximates the hedging interest that could be generated on ENSO markets. The analysis deals exclusively with disaster damages - so my estimate of hedging interest is confined to the measurable losses that might otherwise be managed with insurance (e.g. a firm purchases a futures contract so that they have funds to rebuild critical infrastructure after a major anomaly). It does not cover the hedging interest that will come from firms or institutions using, for example, ENSO derivatives as a diversified asset that can improve their underlying portfolio of business. In other words, whereas my initial hedging interest estimate is based off of expected losses, much, perhaps most, of the hedging interest on a successful teleconnection index exchange will stem from firms and individuals anticipating lost opportunities.

Index/anomaly	Pr(event)	Peril	Region	Window	Median seasonal damage	Damage from large event	E[Damage (lg event)]
ENSO							
El Niño							
ONI index ≥ 2	0.075					for ONI = 2	
		Flood/ landslide/ epidemic	South America	Jan-Jun	0.259	3.4	0.257
		Storms	North America/ Caribbean	Jun-Nov	10.5	5.4	0.410
					from median	-5.0	-0.377
La Niña							
ONI index $\leq$ -1.85	0.05					for ONI = -1.85	
		Flood	Asia/Oceania	Sep-Aug	3.07	8.0	0.401

Table 2.5: Regression estimates of extreme ENSO events indicative of hedging interest (damages in USD b)

Table 2.5 presents aggregate findings from the damage regressions for historically large ENSO anomalies. These estimates are meant to provide general guidance about the hedging interest that might be

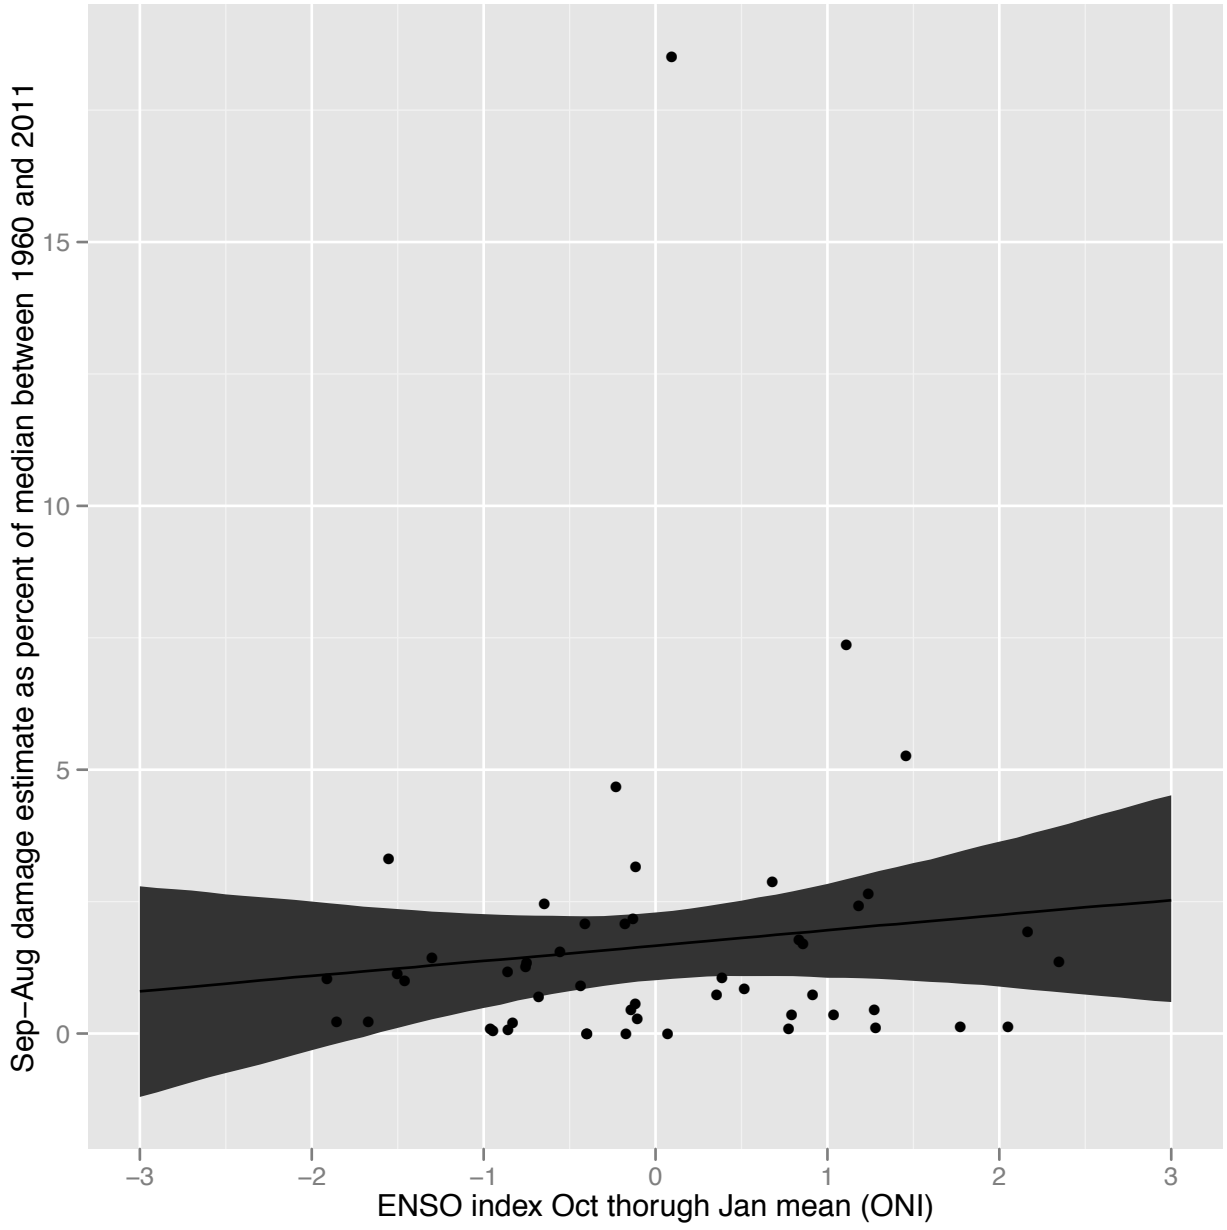

Figure 2.17: Bayesian regression of drought and wildfire damages estimates from likely El Niño drought regions, 1960-2011 predicted by ONI index

generated by individual large anomalies, conditional on their occurrence and weighted by their approximate historical probability.

Based on these results, ENSO anomalies, both high and low, could generate hedging interest in the range of a few billion dollars. The estimates are particularly promising for futures and options markets because they show a rough balance between interest in El Niño and La Niña coverage.

The damage associated with El Niño flooding in South America is entirely offset by a combination of savings to the insurance industry from El Niño's inverse correlation to Atlantic tropical storm damage and interest in hedging La Niña risk from Pacific Asia and Oceania. In fact, hedging interest may concentrate on the La Niña side of the market, although this is difficult to assess without additional analysis of damages in the regions likely to suffer from droughts during El Niño.

Of course, the figures for ENSO in table 2.5 are more valuable relative to a benchmark showing how similar analyses would apply to the indexes underlying successful futures and options contracts. In table 2.7, I present estimates of the impact of one and two standard deviation falls in the annual average crop price index from the US Department of Agriculture in terms of the percentage change in the total value of the US's annual crop for corn, wheat, and soybeans[21].[22] These benchmarks place ENSO risk in the context of indexes that are already the basis of successful exchange-traded derivatives markets.

For the sake of comparison, I've included a similar table coving anomalies in the Arctic Oscillation (AO) in table 2.6. The analysis underlying those benchmarks is available in Arctic Oscillation (AO).

The regressions indicate that large anomalies in ENSO and AO indexes could generate hedging interest of a comparable magnitude to large changes in price indexes for major US crops. It is difficult to compare risks as distinct as price changes in corn and extreme El Niño. But tables 2.5, 2.6, and 2.7 show that events that would be considered "extreme" (approximately two standard deviation events) in both AO and ENSO indexes are comparable to two standard deviation events in major crop price indexes, both in terms of absolute and expected losses.

[21] National Agricultural Statistics Service. National Agricultural Statistics Service Archive: Crop Values Annual Summary. http://usda.mannlib.cornell.edu/MannUsda/viewDocumentInfo.do?documentID=1050, 2012

[22] The figures in table 2.7 come from regressions of annual percentage changes in each index (price and production value) which are not discussed here.

Index/anomaly	Pr(event)	Peril	Region	Window	Median seasonal damage	Damage from large event	E[Damage (lg event)]
AO							
Low							
$-3.5 <$ AO index ≤ -3	0.025					for AO $= -3$	
		Storms/ extreme temps	Countries/ with territory above 45°N	Dec	0.181	1.8	0.045
				Jan	0.674	6.7	0.1675
				Feb	0.165	1.6	0.04
				Mar	0.202	2	0.05
				Dec-Mar	1.222		0.3025
AO index ≤ -3.5	0.01					for AO $= -3.5$	
		Storms/ extreme temps	Countries/ with territory above 45°N	Dec	0.181	2.3	0.0575
				Jan	0.674	8.5	0.2125
				Feb	0.165	2.1	0.0525
				Mar	0.202	2.6	0.065
				Dec-Mar	1.222		0.3875

Table 2.6: Regression estimates of extreme AO events indicative of hedging interest (damages in USD b)

Index/anomaly	Pr(event)	Fall in annual production value	E[Fall]
Corn, % change in price received			
23% $\leq$ percentage fall <46%	0.12	23% fall from the 2011 price 3.752	0.45024
percentage fall $\geq$ 46%	0.1	46% fall from the 2011 price 7.504	0.7504
Wheat, % change in price received			
18% $\leq$ percentage fall <36%	0.22	18% fall from the 2011 price 2.018	0.44396
percentage fall $\geq$ 36%	0.05	36% fall from the 2011 price 4.037	0.20185
Soybean, % change in price received			
17% $\leq$ percentage fall <34%	0.24	17% fall from the 2011 price 2.985	0.7164
percentage fall $\geq$ 34%	0.06	34% fall from the 2011 price 5.971	0.35826

Table 2.7: Regression estimates of change in total annual US crop value (1908-2011) based on percentage price change (USD b), for approximately one and two standard deviation moves.