3
Pricing ENSO Derivatives

In chapter 2 I showed that ENSO is associated with the volume and patterns of economic damage that may justify active trading. This chapter estimates how much ENSO risk protection costs. As part of that estimate, this chapter includes:

- an introduction to data quality issues surrounding ENSO;
- month-by-month analysis of the long-term average index values that could be used to price insurance against catastrophic El Niño/La Niña;
- modeling and analysis that links emerging ENSO forecasts to index values;
- futures and options prices conditional on ENSO forecasts (see the Pricing Appendix in part IV for full conditional prices by month); and,
- suggestions about how those theoretical prices will be modified for actual trading.

All the pricing routines presented in this chapter are based on the data available[1] at:

[1] . . . as of May 2013

- http://www.cpc.ncep.noaa.gov/data/indices/ersst3b.Ni~no.mth.81-10.ascii
- http://www.cpc.ncep.noaa.gov/data/indices/sstoi.indices

The forecasts used for conditional pricing in this chapter come from Colombia University's International Research Institute for Climate and Society (IRI). The archive of those forecasts is available[2] at:

[2] . . . also as of May 2013

- http://iri.columbia.edu/climate/ENSO/currentinfo/archive/

Understanding NOAA's SST indexes

NOAA publishes two primary sea surface temperate indexes. By and large, those indexes tell the same story about El Niño/La Niña. However, they are compiled with different methodologies, over different horizons. Understanding those distinctions is an important first step in pricing El Niño/La Niña risk protection, especially insofar as it suggests challenges to reliable contract settlement.

NOAA's Extended Reconstructed Sea Surface Temperature Index (ERSST) dataset provides a longer record, while NOAA's Optimum Interpolation Sea Surface Temperature Index (OISST) offers finer resolution. As I discuss below, my analysis in this chapter is focused on ERSST data, because OISST's limited horizon (the index begins in the early 1980s) may unfairly bias derivative and insurance prices upward.

NOAA's Extended Reconstructed Sea Surface Temperature Index (ERSST)

Most of the pricing in this chapter uses the latest iteration of NOAA's older ENSO-related SST index, ERSST version 3b. The key factor distinguishing ERSST from OISST is the use of in-situ and satellite data. With the exception of version 3[3], all the ERSST iterations (1,2, and 3b, the iteration used here) use in-situ measurement exclusively[4,5,6]. The full methodology for ERSST version 3b is described in detail in Smith et al. [2008][7].

Monthly anomalies in the ERSST version 3b index are measured relative to a 1971-2000 base period[8]. NOAA releases monthly ERSST estimates with a resolution of two degrees across the four ENSO regions. While the primary index record that NOAA posts to its websites goes back to 1950, monthly ERSST data are available from 1854 on.

Historically, all in-situ measurements came from passing ships. Smith and Reynolds [2004] suggests that ship-based measurements pose challenges to researchers:

> ...the historic distribution of in situ SST data from ships has varied with time due to a variety of economic and political changes (the opening of new canals, world wars, improved communication, etc.). In addition, biases in the ship in situ data have occurred as observational techniques have changed, and those biases must be corrected [statistically] ...

The last decades' in-situ records have relied more heavily on dedicated buoys, as Reynolds et al. [2002] describes:

> SST observations from drifting and moored buoys were first used in the late 1970s. Buoy observations became more plentiful following the

[3] ERRST version 3 included infrared satellite data starting in 1985. NOAA determined that this addition introduced some biases into the index - it tended to suggest temperatures that were too cold by a factor of .01 deg C. NOAA consequently removed satellite data (although it retains in situ data collected via satellite) from the calculation of ERSST version 3b, the current standard.

[4] Thomas M Smith and Richard W Reynolds. Improved extended reconstruction of SST (1854-1997). *Journal of Climate*, 17(12):2466–2477, 2004

[5] Thomas M Smith and Richard W Reynolds. Extended reconstruction of global sea surface temperatures based on COADS data (1854-1997). *Journal of Climate*, 16(10):1495–1510, 2003

[6] Thomas M Smith, Richard W Reynolds, Thomas C Peterson, and Jay Lawrimore. Improvements to NOAA's historical merged land-ocean surface temperature analysis (1880-2006). *Journal of Climate*, 21(10): 2283–2296, 2008

[7] Note that this citation is actually for version 3

[8] Yan Xue, Thomas M Smith, and Richard W Reynolds. Interdecadal changes of 30-yr SST normals during 1871-2000. *Journal of Climate*, 16 (10):1601–1612, 2003

> start of the Tropical Ocean Global Atmosphere (TOGA) program in 1985 [9].[AUTHOR'S NOTE: This program began as a response to the 1982/1983 El Niño event.] These observations are typically made by thermistor or hull contact sensor and usually relayed in real time by satellites. Although the accuracy of the buoy SST observations varies, the random error is usually smaller than 0.5 °C and, thus, is better than ship error. In addition, typical depths of the measurements are roughly 0.5 m rather than the 1 m and deeper measurements from ships. ... [The] deployment of the buoys has been designed to fill in some regions with few ship observations. This process had the most impact in the tropical Pacific Ocean and the Southern Hemisphere.

[9] Michael J McPhaden, Antonio J Busalacchi, Robert Cheney, Jean-René Donguy, Kenneth S Gage, David Halpern, Ming Ji, Paul Julian, Gary Meyers, Gary T Mitchum, et al. The Tropical Ocean-Global Atmosphere observing system: A decade of progress. *Journal of Geophysical Research: Oceans (1978–2012)*, 103 (C7):14169–14240, 1998

Given the improvements to the SST record over time, someone pricing risk management contracts using ERSST might want to give different weights to different eras in the historical record. However, the period since 1970 has had a much greater incidence of extreme El Niño/La Niña events than any other in the historical record. So, prices that rely more heavily on recent data will almost certainly be more expensive.

NOAA's Optimum Interpolation Sea Surface Temperature Index (OISST)

By the early 1980s, in-situ measurements were complemented by direct satellite measurements. Again, Reynolds et al. [2002] explains:

> In late 1981, Advanced Very High Resolution Radiometer (AVHRR) satellite retrievals improved the data coverage over that of in situ observations alone. The satellite retrievals allowed better resolution of small-scale features such as Gulf Stream eddies. Because the AVHRR cannot see the surface in cloud-covered regions, the biggest challenge in retrieving SST is to eliminate cloud contamination

This satellite data became the basis of NOAA's alternative to ERSST, the Optimum Interpolation Sea Surface Temperature Index (OISST). NOAA releases OISST data weekly on a one-degree grid across the key ENSO regions. So it is available more often, and on a finer scale, than ERSST.

OISST, currently at version 2, combines in situ SST measurements, daytime and nighttime satellite data readings, and data from sea ice cover simulations. The satellite data is adjusted statistically for natural sources of bias, like cloud cover and atmospheric water vapor[10] [11] [12] [13].

[10] Richard W Reynolds, Nick A Rayner, Thomas M Smith, Diane C Stokes, and Wanqiu Wang. An improved in situ and satellite SST analysis for climate. *Journal of Climate*, 15(13):1609–1625, 2002

[11] Richard W Reynolds and Thomas M Smith. Improved global sea surface temperature analyses using optimum interpolation. *Journal of Climate*, 7(6):929–948, 1994

[12] Richard W Reynolds and Diane C Marsico. An improved real-time global sea surface temperature analysis. *Journal of Climate*, 6(1):114–119, 1993

[13] Richard W Reynolds. A real-time global sea surface temperature analysis. *Journal of Climate*, 1(1): 75–86, 1988

Additional dataset considerations

After picking a dataset for pricing, you must also decide on the region to price, whether to use absolute SST measurements (°C) or anomalies, and any other data cleaning routines, like adjusting the standard

deviation to match recent decades. In this section I walk through those considerations and settle on a baseline for pricing of:

- Niño 3.4;
- measured in absolute degrees Celsius;
- without any standard deviation adjustments.

Niño region

As discussed in chapter 2, the Niño 3.4 region (figure 2.4) is the de facto benchmark for identifying El Niño/La Niña events world wide. It is consequently the basis of the pricing below.

In chapter 2, I also discussed the relative merits of Niño 1.2, the region with the tightest connection to Peruvian flooding, which GlobalAgRisk used for its El Niño insurance. The main disadvantages of that index are:

- Niño 1.2 anomalies arise later in the year, so the index has a smaller window for advanced payments; and
- Niño 1.2 may do poor job of representing El Niño/La Niña events with the same geographic signature as the 1972/1973 El Niño - where warming was focused on the central Pacific.

If Niño 1.2 is the best proxy for Peru, does that mean that Niño 4, on the other side of the Pacific, is the best proxy for Australia. Andrew Watkins of the Australian Bureau of Meteorology (ABM) suggests no. Both indexes represent Australian risk quite well. So, unlike Niño 1.2, Niño 4 is unlikely to find a niche of devoted specialized hedgers.

Anomalies vs. absolute SST measurements

NOAA releases each of its datasets as departures from monthly averages (anomalies) and absolute degrees Celsius. Its not immediately clear which format is better for financial contracts.

Presenting contracts in terms of anomalies facilitates interpretation of actual El Niño/La Niña events, since most major meteorological organizations define those events in terms of persistent monthly anomalies. Indeed, many forecasts of SSTs (like those from the ABM and IRI) are only provided in terms of anomalies.

The primary disadvantage of anomalies is that they have been, and will continue to be, subject to revision as underlying SSTs drift over time. In chapter 2 I briefly discussed the possible link between climate change and higher Pacific SSTs. To the extent that such trends continue, the index may revise its baseline and the interpretation

of anomalies may become less clear. The ONI index, which NOAA uses to define El Niño/La Niña already uses a rolling window for its monthly base periods.

The weather traders I interviewed for chapter 7 suggested that the temperature derivatives are currently subject to annual revision. The practice has not been a problem for traders. Nevertheless, there may be advantages to using absolute SSTs. Absolute measurements will directly incorporate any underlying shifts in the index, allowing, for example, traders to simply express theories about the long-term trends in the index. Those theories and, by proxy, the market's judgment of long-term climate change might be obscured in an anomaly-based contract.

Miscellaneous data preparation

As mentioned above, many of the strongest El Niño/La Niña events on record have happened since 1970. If you believe that the more recent record is indicative of a regime change and is likely to be a better guide to the future than all the other decades on record, then there is no problem pricing risk based on that subset of the data. However, if you believe that the clustering of anomalies in last 30 years was primarily a function of random chance, then you can use the full ERSST dataset.

If you are uncertain, then some statistical adjustment of the data may provide a middle ground. For example, you can increase the volatility in the earlier part of the ERSST record to match that from the more recent record.

Comparing and choosing a dataset to price

Figure 3.1 provides the baseline monthly values that NOAA uses to calibrate anomalies in OISST and ERSST. Note OISSTs tendency toward colder SSTs. The cold bias in satellite data is a great concern in the climate literature and is noted in all the index construction papers on ERSST and OISST cited above.

Figure 3.1 also shows that the winter months (in the Northern Hemisphere) are the coldest in both indexes. This is interesting, given that the ENSO phenomenon takes place in the tropics and its most dramatic human impacts are in the southern hemisphere.

Finally, February/March and June/July are inflection periods, moving both indexes from cold to warm phases (the former months) and back (the latter months). The baseline SST fluctuations over these two windows is dramatic. I suspect that those months will consequently host very active trading, if traded ENSO markets launch. Those are

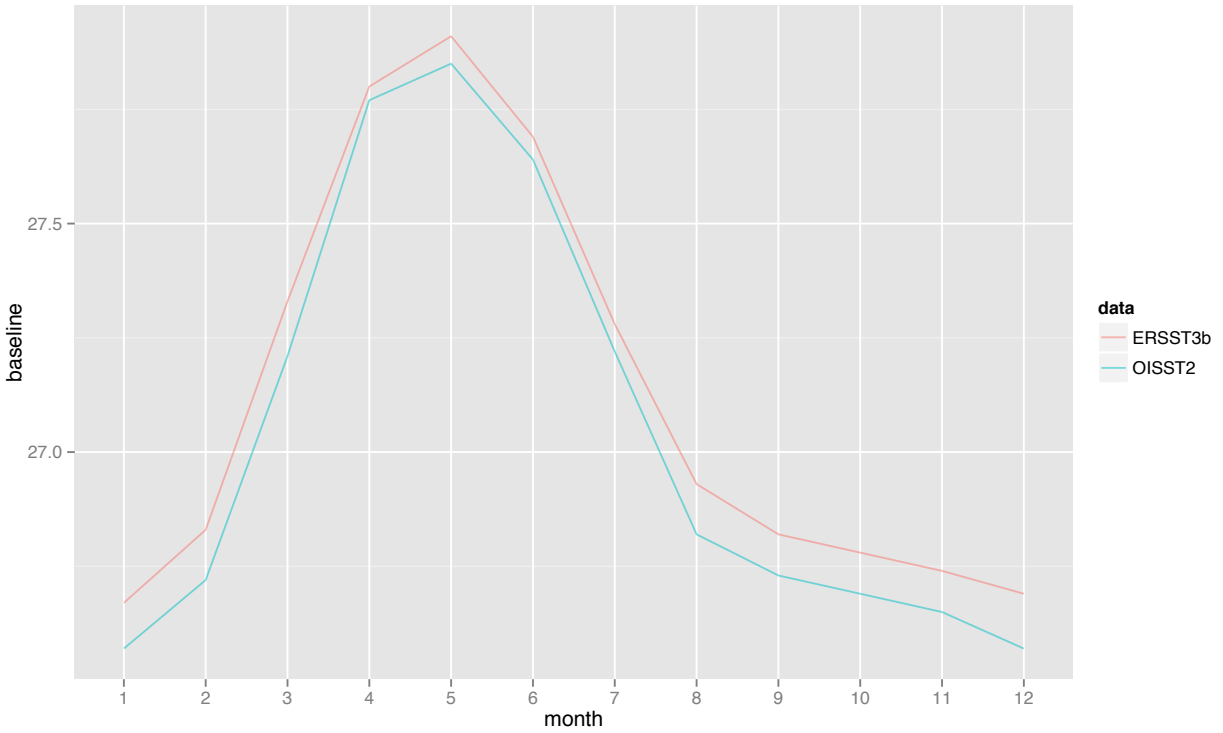

Figure 3.1: Comparing OISST and ERSST monthly baselines

also likely to be the months where climate expertise and proprietary data will provide the largest edge to traders. The possibility of information asymmetries in those months may undermine the volume boost that traded markets might otherwise get from increased volatility.

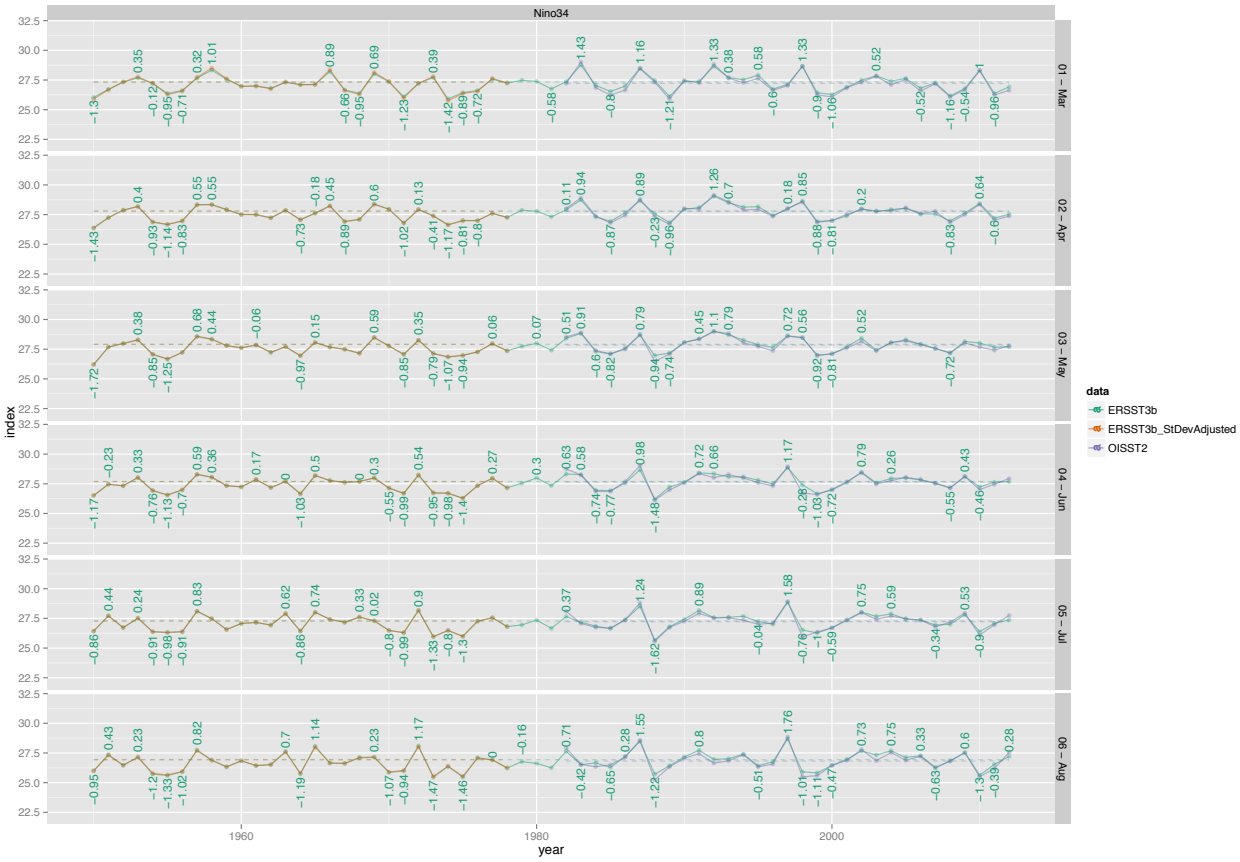

Figure 3.2: Comparing SST dataset options: March through August

Figures 3.2 and 3.3 show monthly time series for ERSST, ERSST with a standard deviation adjusted upward for pre-1979 data, and OISST. Looking at these graphs we can see a few trends that may be important in pricing risk coverage.

First and foremost, we need to look for systemic divergences between any of the indexes in the tails of their distributions, since that hedgers and speculators will be most interested in those extreme SST measurements that trigger payments. At first glance, there does appear to be a trend in tail behavior. On the La Niña side (when the indexes are in La Niña anomaly territory according to NOAA's ONI index discussed in chapter 2), OISST appears to show lower numbers than ERSST in any given month. On the El Niño anomaly side, the link is less clear. This raises the possibility that the OISST measure-

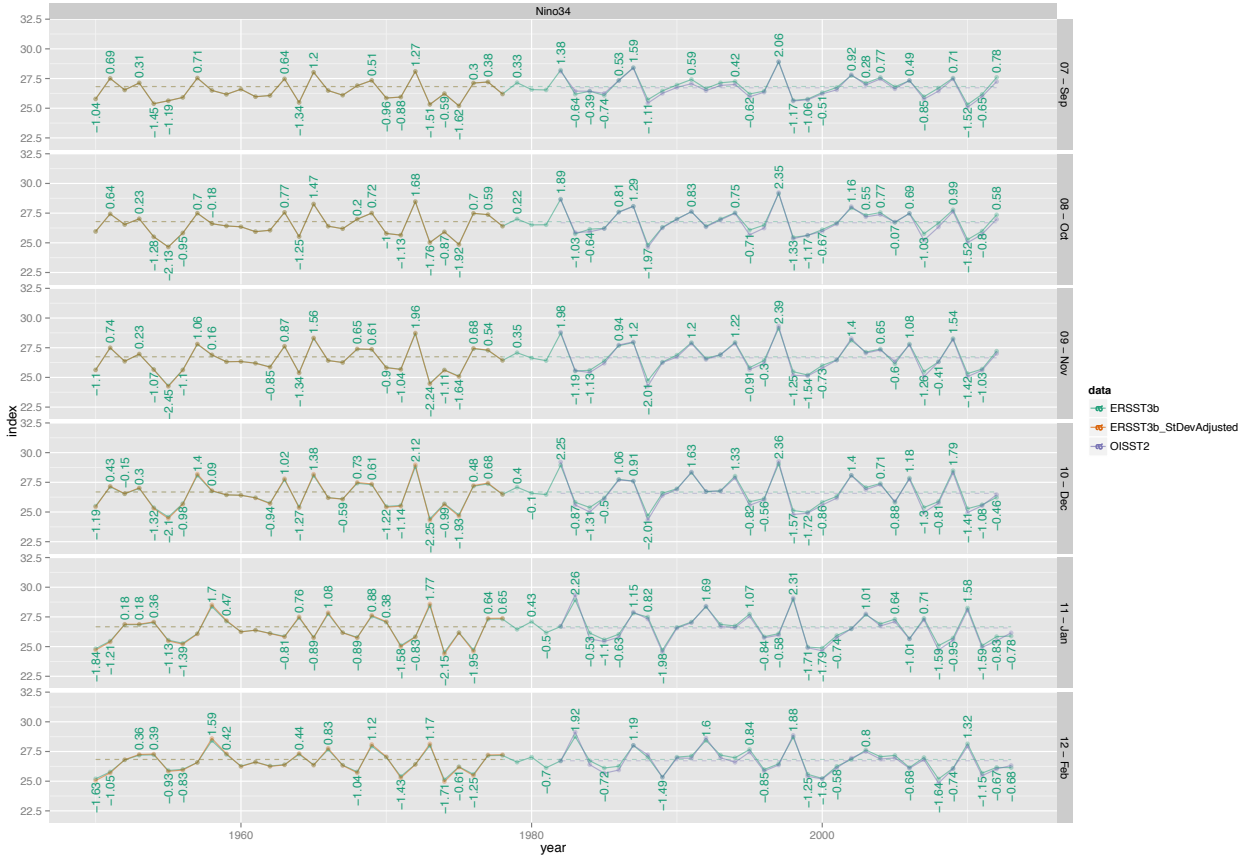

Figure 3.3: Comparing SST dataset options: September through February

ment has higher volatility on the cold side than on the warm side.

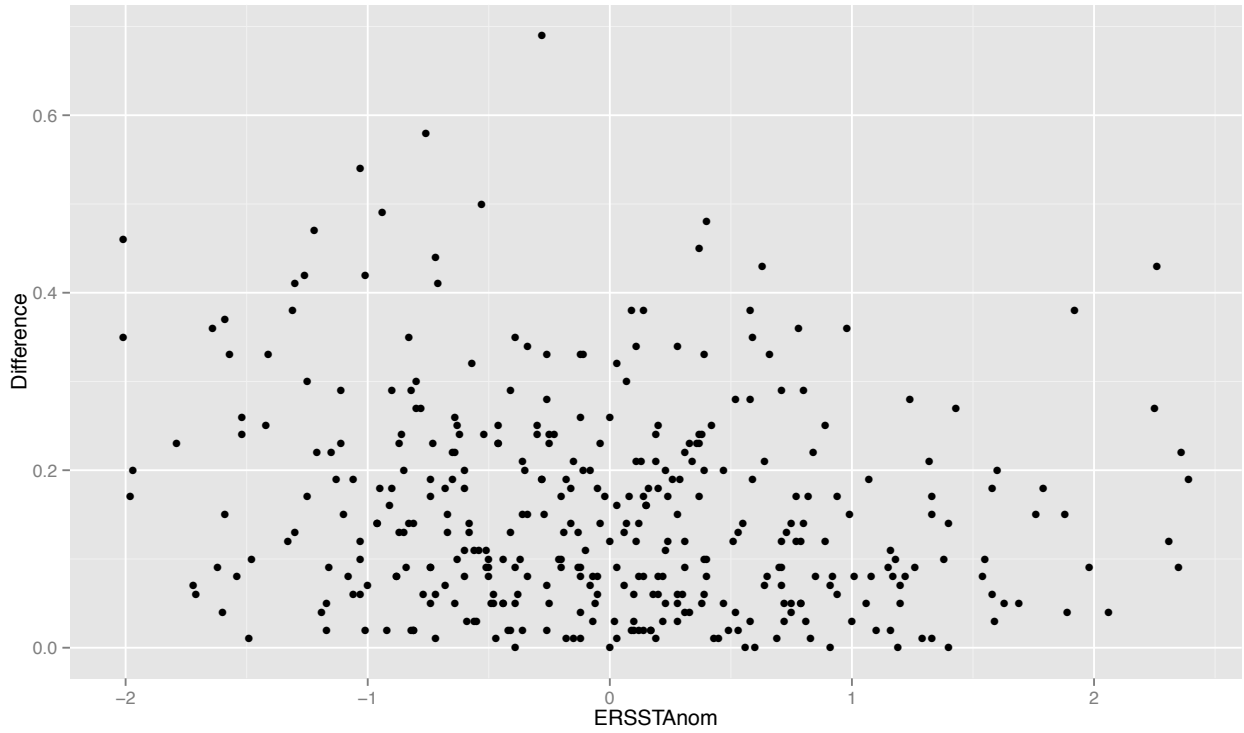

Figure 3.4: Comparing absolute value of the difference between OISST and ERSST (y-axis) and ERSST (x-axis)

Figure 3.4 is a scatter plot showing the difference between the OISST time series and the ERSST time series on the y-axis and the corresponding anomaly in the ERSST data on the x-axis. Are there patterns that suggest index choice matters a great deal in the tails of the anomaly distribution? Do the indexes agree least when we need reliable index measurements most - in a disaster?

Negative anomalies in the ERSST dataset are loosely associated with larger discrepancies between OISST and ERSST - indicated by a clump of points in the upper left-hand quadrant of figure 3.4. This is exactly what we saw in the raw time series comparison above. The implication of this trend for pricing is that simulations used for pricing OISST-based risk cover should explicitly model this downside volatility clustering.

Turning to the adjusted ERSST dataset, we see that its deviations from the basic ERSST dataset are smaller than those of OISST. Consequently, the choice between OISST vs ERSST seems more important for the ultimate pricing decision than the choice of whether or not to adjust the standard deviation of the earlier piece of the ERSST

dataset.

Figure 3.5 shows the relationship between those differences and extreme values in the underlying ERSST dataset. The figure shows the distinct pattern, that the adjustment itself was intended to produce - the adjusted dataset produces its largest discrepancies when the underlying index is high. The anomalies for any given month clump in straight line patterns.

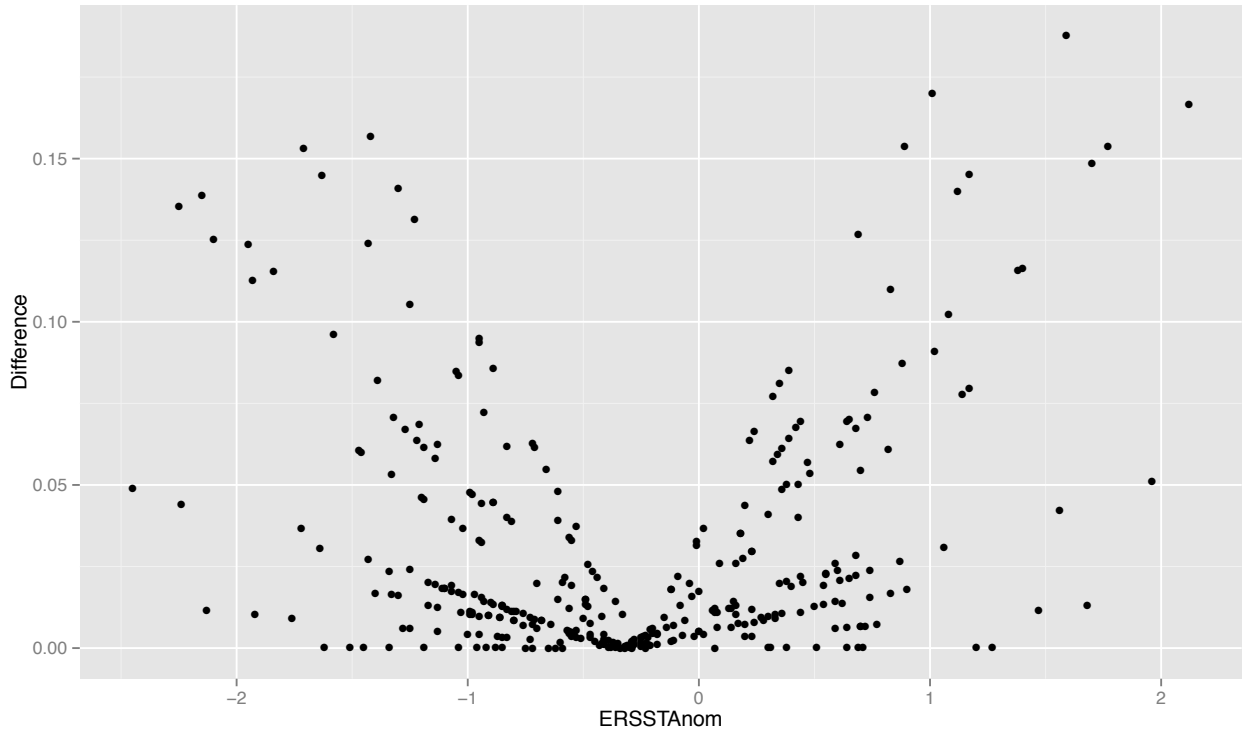

Figure 3.5: Comparing absolute value of the difference between ERSST with standard deviation adjustments and ERSST (y-axis) and ERSST (x-axis)

Arbitrage vs. expectations pricing

In a market such as corn or equities, derivatives prices come from arbitrage, buying and selling in two or more closely related markets to take advantage of a price differences across those markets. If, for example, call options on corn appear too expensive relative to today's price on spot markets, a trader can:

- sell calls (collecting a premium);
- borrow cash;
- use the cash to buy physical corn on the spot market in proportion to those calls; and

- store that physical corn.

In this case, the calls sold represent an obligation to deliver corn. But as long as the trader has corn in storage (potentially over the full term of the call option) in equal proportion to that obligation, then she can *deliver* on those calls (i.e. provide all the corn implied) at a moments' notice. This means that subsequent moves in the spot price of corn will not cause the trader to incur any additional profits or losses. All that matters to the trader is whether the premium she collected by selling her calls is greater than the cost of borrowing and storing the corn.

In an efficient market, traders specializing in arbitrage, make similar trades across all clearly linked markets. By buying corn, they nudge up the price in the physical market and by selling calls they nudge down the implied price of corn in options markets. Over time, price discrepancies between markets fall until the margins on these *risk-less*[14] trades disappear.

Black, Scholes, and Merton's Nobel-Prize-winning work on options pricing demonstrated that, given ideal conditions, the simple existence of these arbitrage strategies suggests rational option prices that do not include traders' expectations for the price of the underlying good[15] [16]. In other words, you can rationally price an option that would pay if corn goes to USD 8 per bushel without guessing about the probability of corn actually going to USD 8 per bushel.

By contrast, SSTs cannot be arbitraged directly. There is no opportunity to take a chunk of the Pacific Ocean today and deliver it in the future at a pre-arranged temperature. I'll discuss later how it may be possible to piece together rough arbitrage strategies by looking at the prices of related markets, such as a basket of more localized weather derivatives. But even with those opportunities, arbitrage is not going to provide theoretically definitive ENSO derivatives prices. Instead, those prices will have to come from reasonable guesses about traders' expectations for future SST.

[14] *Risk-less* in this context refers to the fact that these arbitrage trades attempt to profit from markets without taking on the price risk associated with the underlying good in the market. They are not free from risk.

[15] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *The Journal of Political Economy*, pages 637–654, 1973

[16] Robert C Merton. Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, pages 141–183, 1973

Pricing outside the predictive window

In this section, I walk through the process of pricing El Niño/La Niña derivatives/insurance outside the predictive window. That window, before which predictions of El Niño/La Niña for the upcoming year are little better than long term averages, is marked by what climate scientists call the "spring predictive barrier". As recently as 2010, many climate scientists placed that barrier in March. However, it may already have moved into February or January as El Niño/La Niña prediction has steadily improved[17].

[17] In fact, even if February predictions are not valuable, they may still produce herding behavior - with hedgers buying based on the assumption that they are valuable. Indeed, in 2010 a large potential buyer of GlobalAgRisk Extreme El Niño insurance decided against signing a reinsurance agreement after they saw early forecasts of La Niña/neutral conditions in 2010. While it is difficult to know how

Outside the yearly predictive window, SST expectations are relatively straightforward. Traders should generally expect that SSTs in any given month in the future will look like SSTs in the past[18]. When historical averages are the only basis for pricing, options are functionally equivalent to index-based insurance. One of the key theoretical factors that distinguishes insurance from derivatives is the fact that *insurance risk* is limited to situations where neither party has private information relative to the settlement of the contract. By extension, that means that historical averages are the only basis for pricing in insurance. As SST predictions improve over time, insurers will need to push back the sales closing date on their coverage to make certain that they and their clients are making purchasing decisions exclusively based on historical information.

[18] I discuss non-stationarity in the following section.

Modeling the index

I begin pricing with basic exploratory data analysis, graphing the historical record of monthly SST in various ways that might suggest statistical properties of the underlying phenomenon.

Figures 3.6, 3.7, and 3.8 are a histogram, an empirical cumulative distribution function (ECDF), and a kernel density estimate respectively, for absolute SSTs in October[19] from ERSST.3b's Niño 3.4, running from 1950 to the present. In particular, I'm interested to see if these graphs suggest skewness, a bi-modal distribution, or other features that will be important for simulation and modeling in later sections.

[19] According to Dr. Andrew Watkins of the Australian Bureau of Meteorology (ABM), October is the single most decisive month for El Niño/La Niña worldwide. It is consequently the month I use for most of the examples in this chapter.

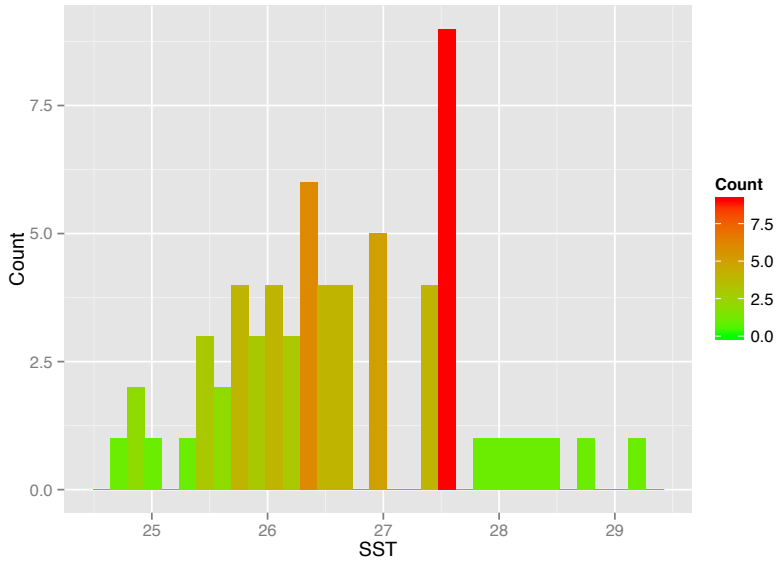

Figure 3.6: Histogram of October SST for Niño 3.4 ERSST.3b

The histogram of October SSTs suggests two important characteristics:

- A barrier at roughly 27.5°C - values running up to that level were progressively more frequent while values above that levels were relatively infrequent.
- Two distinct peaks of frequency - the biggest is at roughly 27.5°C and a second smaller peak at roughly 26.3°C.

The ECDF tells a similar story as the histogram. Its S-shape breaks between 27.5°C and 26.3°C. Alone, the ECDF and the histogram might indicate that October SSTs should be fit using a mixture model, combining draws from two distributions.

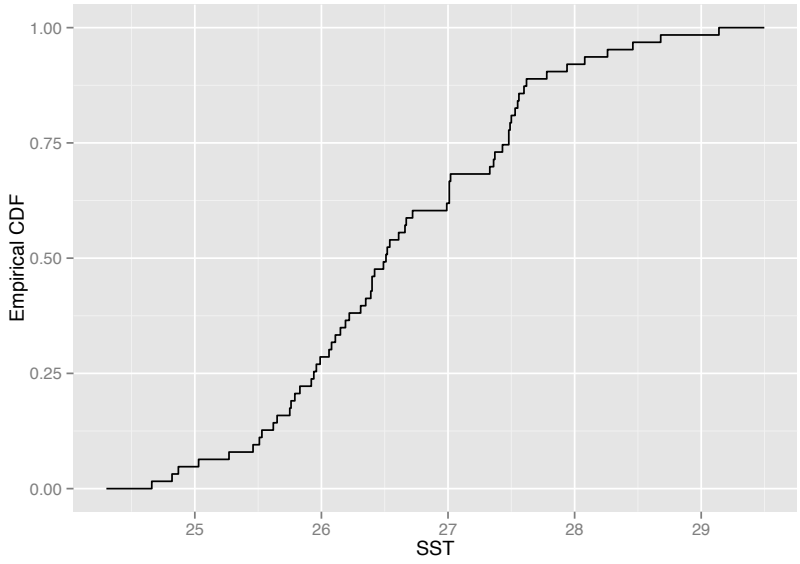

Figure 3.7: ECDF of October SST for Niño 3.4 ERSST.3b

However, these features are less prominent in the kernel density smoothed estimate in figure 3.8. The distribution is not clearly asymmetric. While there are two peaks, those peaks are relatively close together and of similar magnitude.

Satisfied by figure 116 that SSTs do not require a mixture model, I chose four parametric distributions to fit to the data: normal, log-normal, gamma, and Weibull. I fit both the normal and log-normal distributions using closed-form minimum least squares routines. I fit the gamma and Weibull distributions using direct optimization of the log-likelihoods [20]. I then generated one million random samples from each of the resulting distributions. QQ plots are shown in figure 3.9. They compare the quantiles of the ECDFs of the randomly generated samples to those from the actual historical SSTs for October.

[20] Brian D Ripley. *Modern applied statistics with S*. Springer, 2002

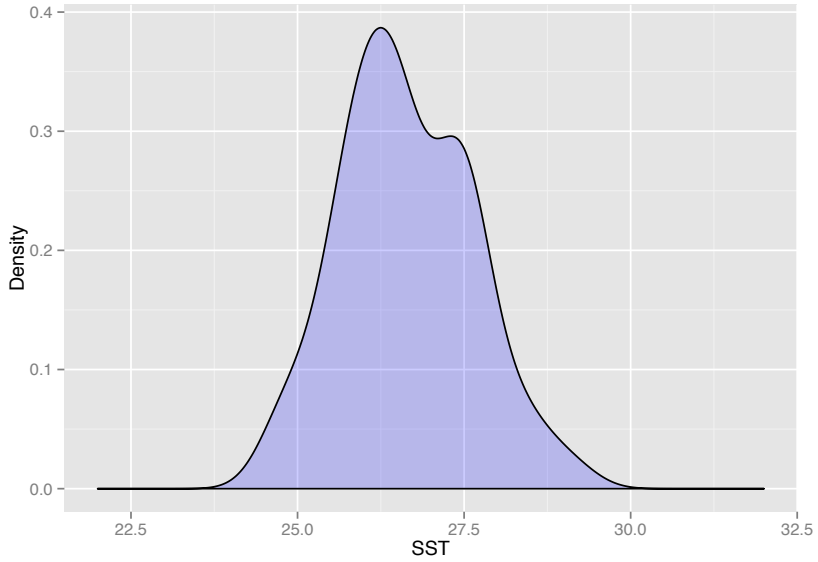

Figure 3.8: Kernel density estimate of October SST for Niño 3.4 ERSST.3b

The Weibull, a member of the family of extreme value distributions, is the only distribution that distinguishes itself in figure 3.9. It shows a poor fit to the data, generating too many extreme draws at both ends of the distribution.

Figure 3.9, suggests that the normal, log-normal, and gamma provide comparable fits to the data. In addition to looking at QQ-plots, I also performed a two-sample Kolmogorov-Smirnov (KS) test on each of the randomly generated samples to indicate whether each generated sample came from the same distribution as the historical record. The KS test uses the null hypothesis that the data from both samples follow the same distribution and we fail to accept the null (suggesting that the samples are indeed from distinct distributions) if the p-value of the KS test falls below our chosen threshold (generally 0.05.) The KS tests were inconclusive, unable to distinguish between the historical record and any of the generated samples at any reasonable level of statistical significance.

If ENSO indexes are non-stationary, as implied by figure 2.5, then the parametric modeling in this section must be adjusted to reflect the possibility that Pacific SST anomalies are slowly increasing in strength and/or frequency.

To some extent, NOAA's indexes of SSTs already adjust for that trend, by extrapolating historic measurements using a running window of baseline monthly temperatures. If modelers, believe that the index this adjustment is insufficient, they may attempt to model the index's gradual change directly. Alternatively, insurers who believe that non-

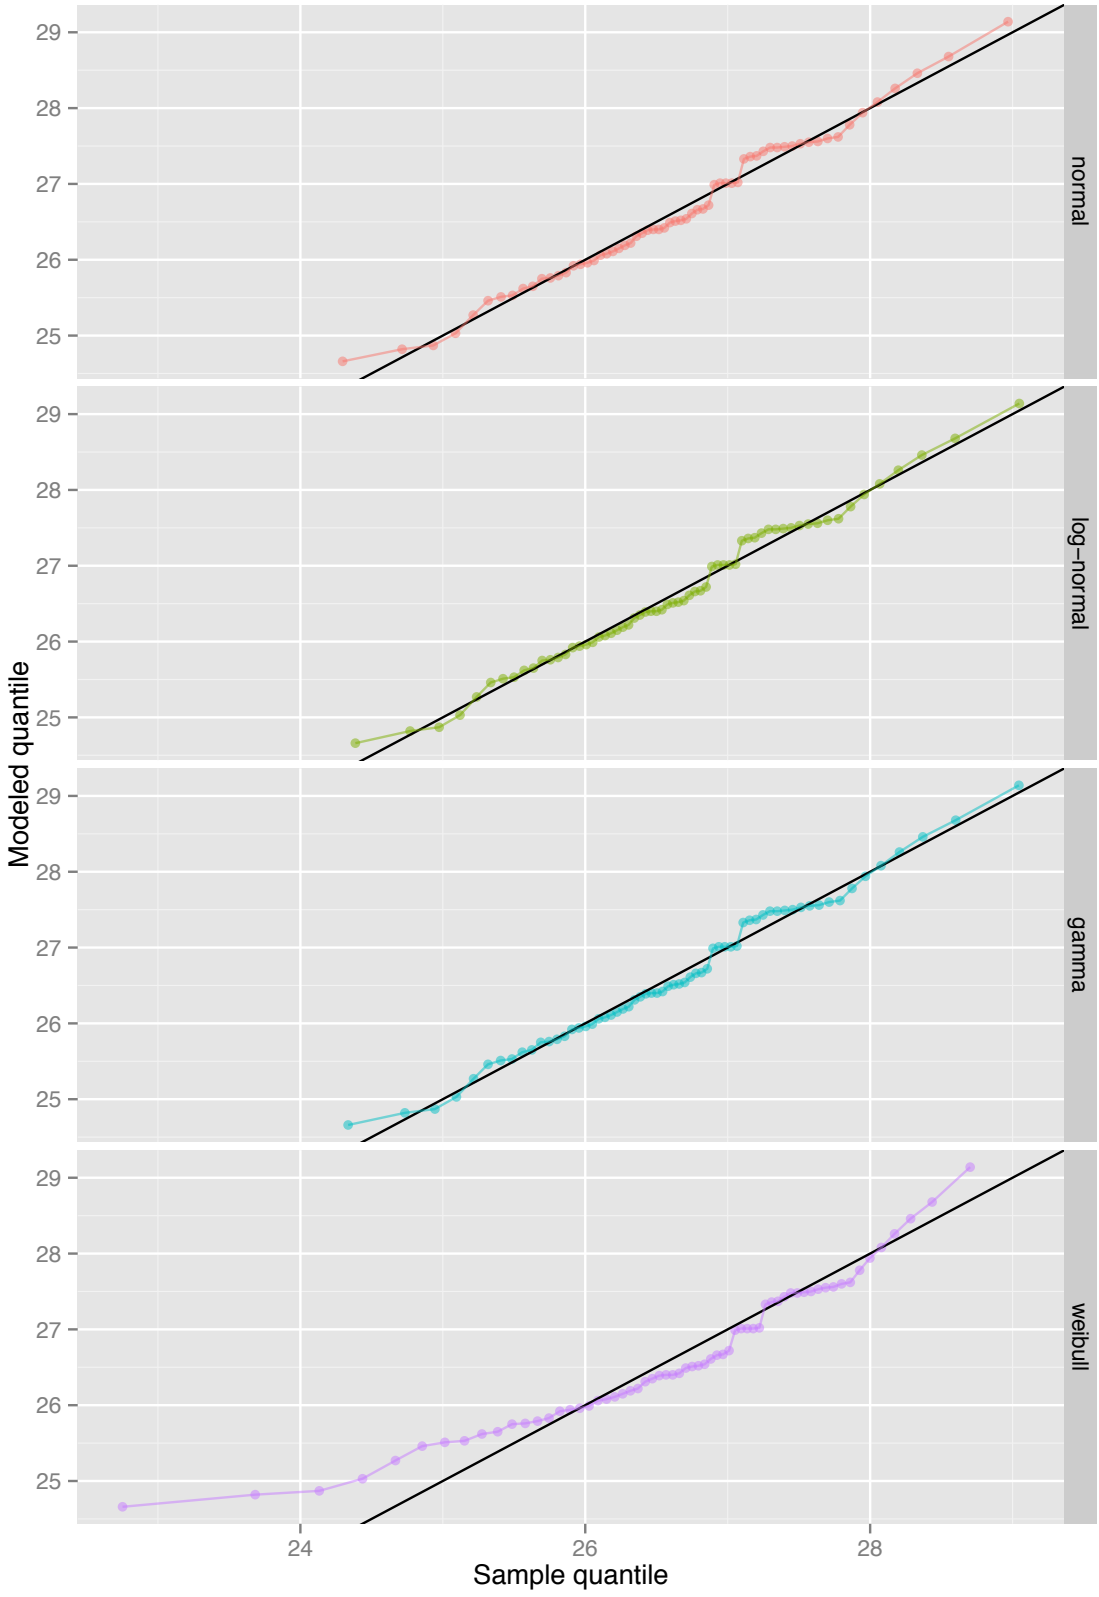

Figure 3.9: QQ plots of October SST for Niño 3.4 ERSST.3b compared to samples from various distributions (n=2 million)

month	normal	lognormal	gamma	weibull
1	0.57	0.66	0.65	0.33
2	0.77	0.82	0.78	0.18
3	0.96	0.97	0.96	0.36
4	0.81	0.80	0.83	0.76
5	0.90	0.90	0.92	0.70
6	0.91	0.90	0.92	0.82
7	0.96	0.94	0.95	0.62
8	0.97	0.97	0.98	0.58
9	0.77	0.81	0.81	0.47
10	0.80	0.75	0.79	0.37
11	0.53	0.63	0.62	0.18
12	0.75	0.80	0.81	0.36

Table 3.1: Kolmogorov-Smirnov test statistics comparing fitted distributional samples to historical SSTs

stationarity is a real, but difficult to quantify, phenomenon may choose not to model it at all, but to demand a slightly larger risk premium on their insurance.

Defining a payout function

Prior to pricing, simulated ENSO-SSTs need to be translated into payments for hedgers using a payout function.

I am most concerned with extreme El Niño/La Niña, so I've chosen to structure the payout functions for my example options around events between one and three standard deviations away from the monthly mean. More specifically, payments on the options begin at one standard deviation[21] above or below the monthly average (for El Niño coverage/calls and La Niña coverage/puts respectively) and payments reach one hundred percent of the notional value (or sum insured) at three standard deviations above or below the monthly average. Figure 3.10 shows the average monthly value for Niño 3.4 in black. The red and blue bands show the index values for each month that would trigger a payment on calls and puts respectively.

[21] This is also called the trigger or attachment point.

Within those ranges, I use linear pricing such that an index value halfway across the red band in figure 3.10 (i.e. halfway between the the trigger and max payout point) would obligate a payout that is half of the sum insured on a call/El Niño contract. The full linear function for October El Niño is shown in figure 3.11.

As an example, suppose that I bought USD 100 of coverage for USD 10 against October El Niño. If actual October SST was halfway across the red band, or 28.74°C, I would receive USD 50.

In practice, GlobalAgRisk found that hedgers (and speculators) prefer a payout function that offers a minimum payout in the event that the index reaches just above the trigger. For example, an index value

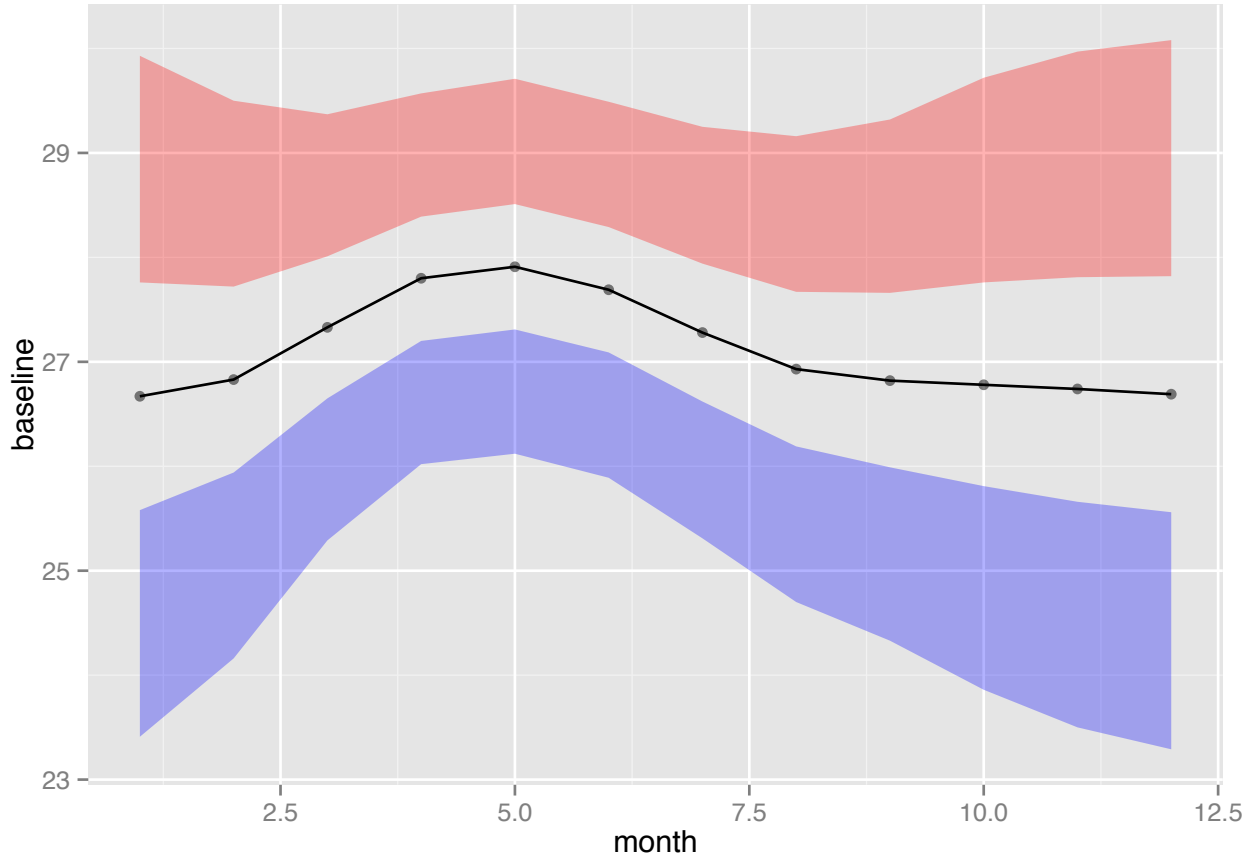

Figure 3.10: Index values for El Niño (red) and La Niña (blue) events between one and three standard deviations away from monthly average

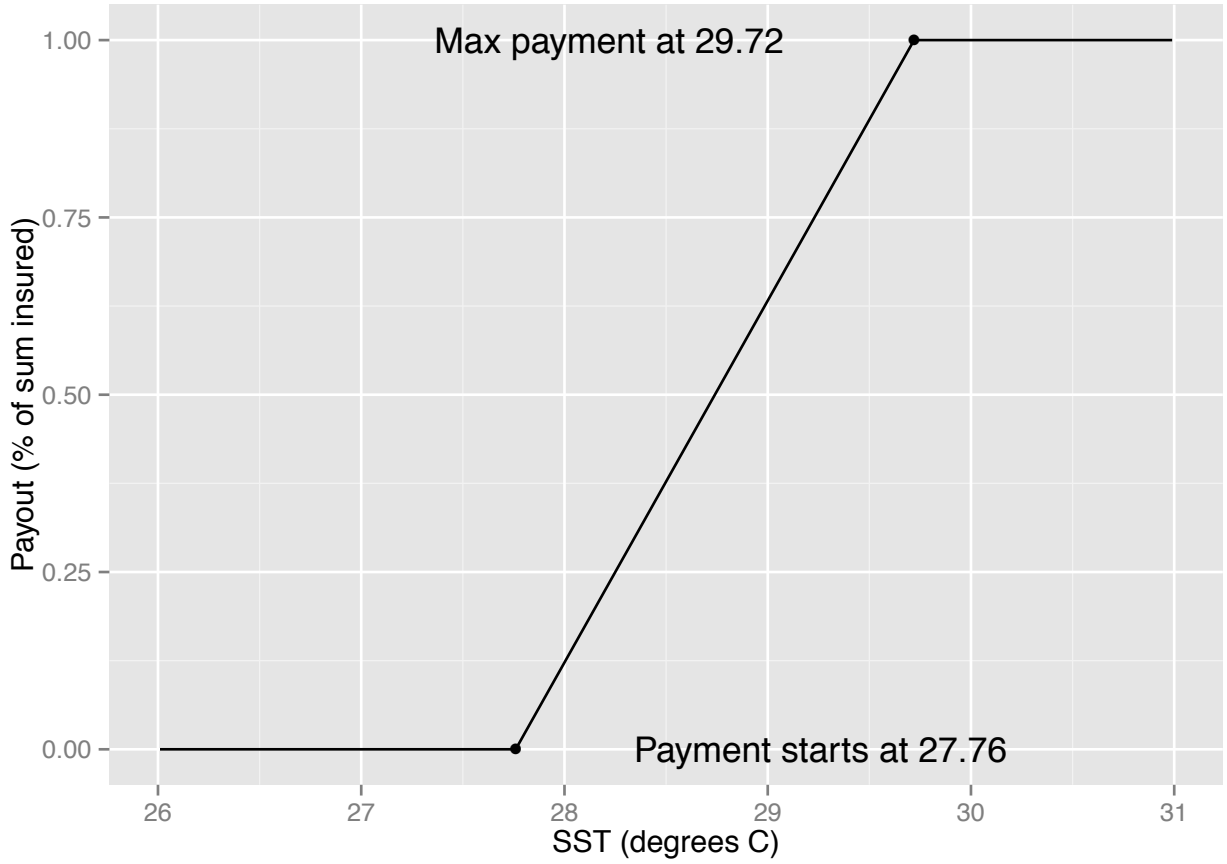

Figure 3.11: Payout function for call option on October SST for Niño 3.4 ERSST.3b covering index values between one and three standard deviations above the baseline

that just barely crosses into the red in 3.10 might trigger a payout of 5 percent on an El Niño/call contract, rather than the tiny payout suggested the kind of linear function in figure 3.11.

Some potential clients also expressed interest in a more customized payout function consisting of steps usually shaped around historical events e.g. a 25 percent payout for the 1972/1973 magnitude event and a 75 percent for a 1997/1998 magnitude event.

Static pricing

Given the option parameters, pricing function, and random samples from fit distributions discussed above, I can now price derivatives outside the predictive window.

Initially, I display the payouts generated just by historical October SSTs. The average of these is a starting point for the derivative price. This type of historical pricing is called burn analysis in (re)insurance. Figures 3.12 and 3.13 show burns and average payouts on calls and puts respectively.

As we discussed above, the last 30 years has been very active for El Niño/La Niña events. Earlier decades (going back to the mid-nineteenth century) hosted many fewer extreme El Niños/La Niñas. Consequently, the burn price displayed here may in fact be high relative to our future expectations for October SSTs.

In figure 3.14, I show the prices generated (in USD of premium per USD 100 of nominal coverage) from the random samples from fit distributions. The figure includes burn prices and prices from samples taken from kernel density smoothers fit over each month.

The prices from the various distributions are, with one prominent exception, close together. On the El Niño side, the highest and lowest prices are mostly within 125 basis points of one another in any given month. On the La Niña side, that spread is slightly larger at roughly 150 basis point, but only between April and June.

The Weibull, is the one model challenging this consensus. The prices from the Weibull samples are clearly distinct from the rest of the group - almost doubling the price of La Niña coverage relative to the rest of the group. The Weibull sample suggested the lowest prices for El Niño coverage, albeit by a much smaller margin than for La Niña. That is understandable given the distribution's heavy left tail.

Apart from the Weibull, the samples drawn from the kernel density smoother suggests the second highest prices for both El Niño and La Niña coverage. The burn prices are in the middle of the pack.

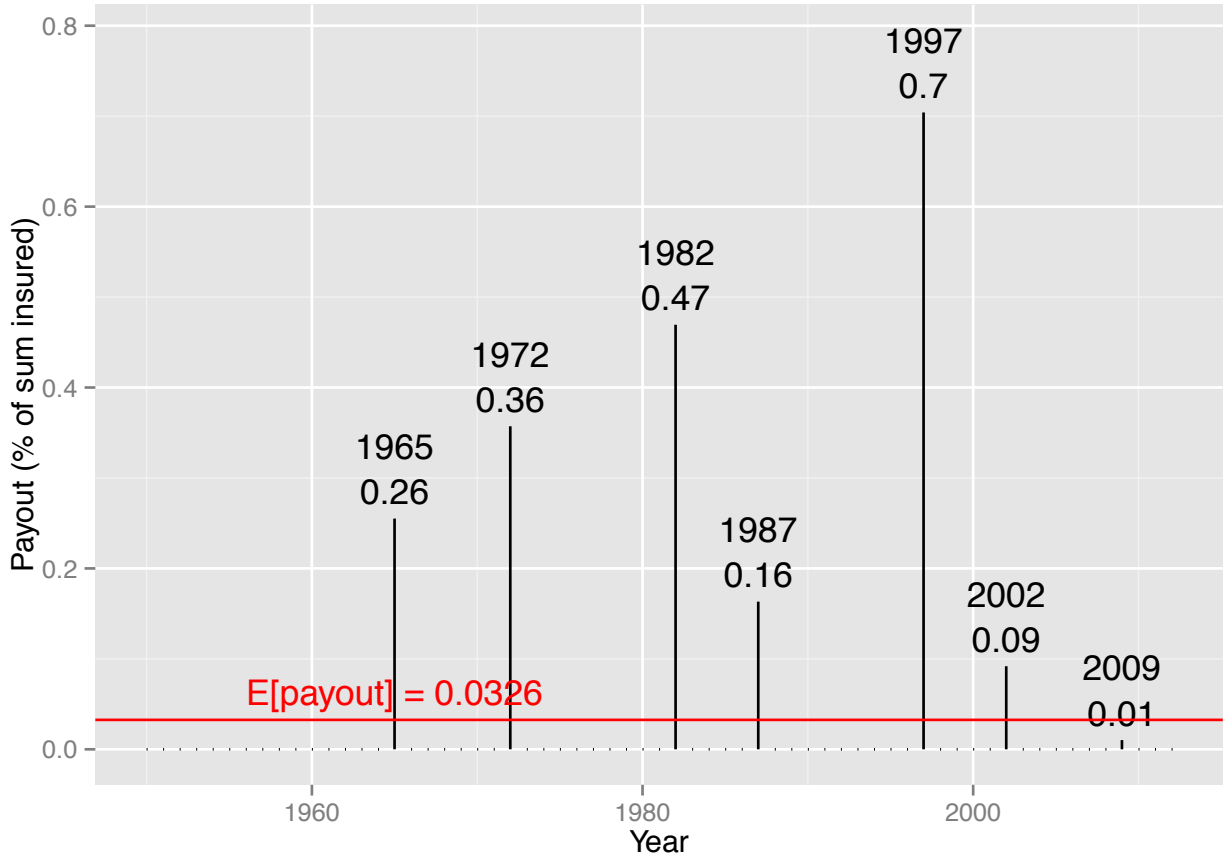

Figure 3.12: Historical burn on call option for October SST for Niño 3.4 ERSST.3b covering index values between one and three standard deviations below the baseline

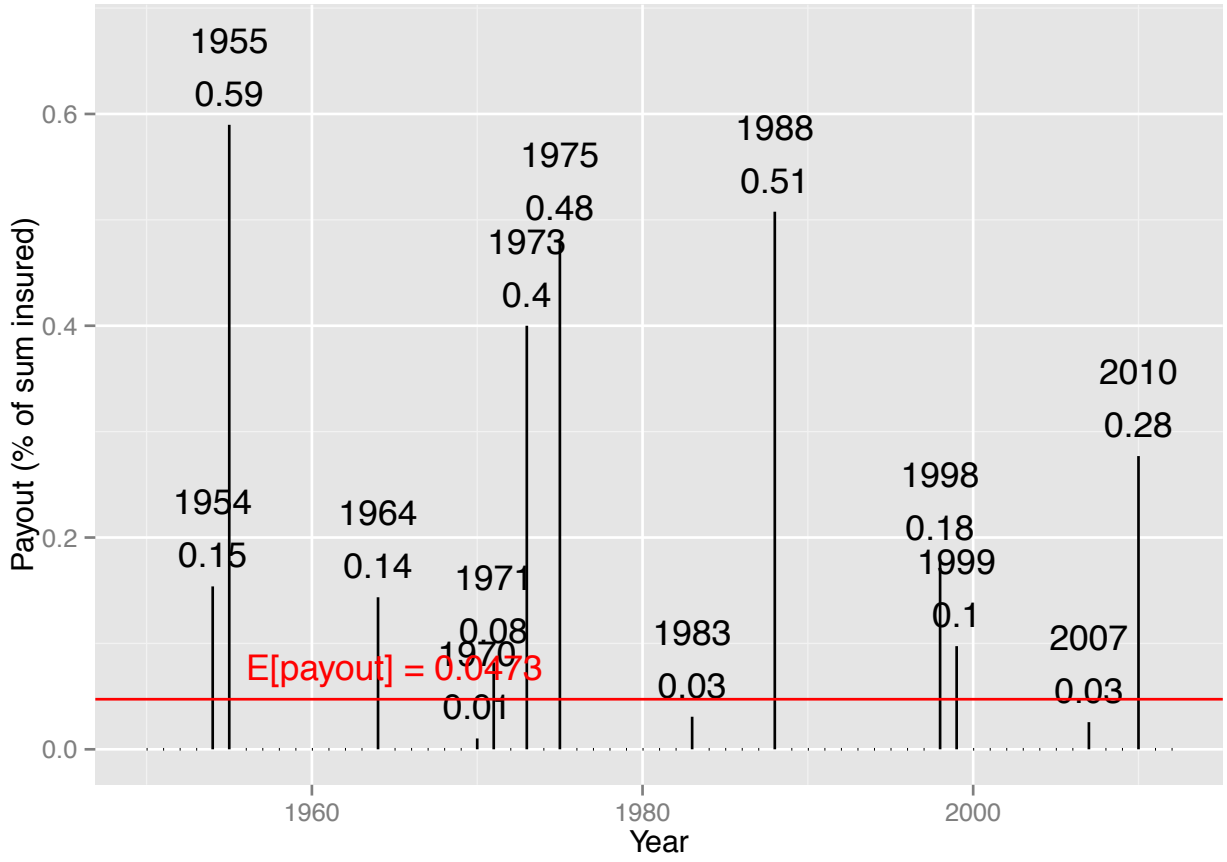

Figure 3.13: Historical burn on put option on October SST for Niño 3.4 ERSST.3b covering index values between one and three standard deviations below the baseline

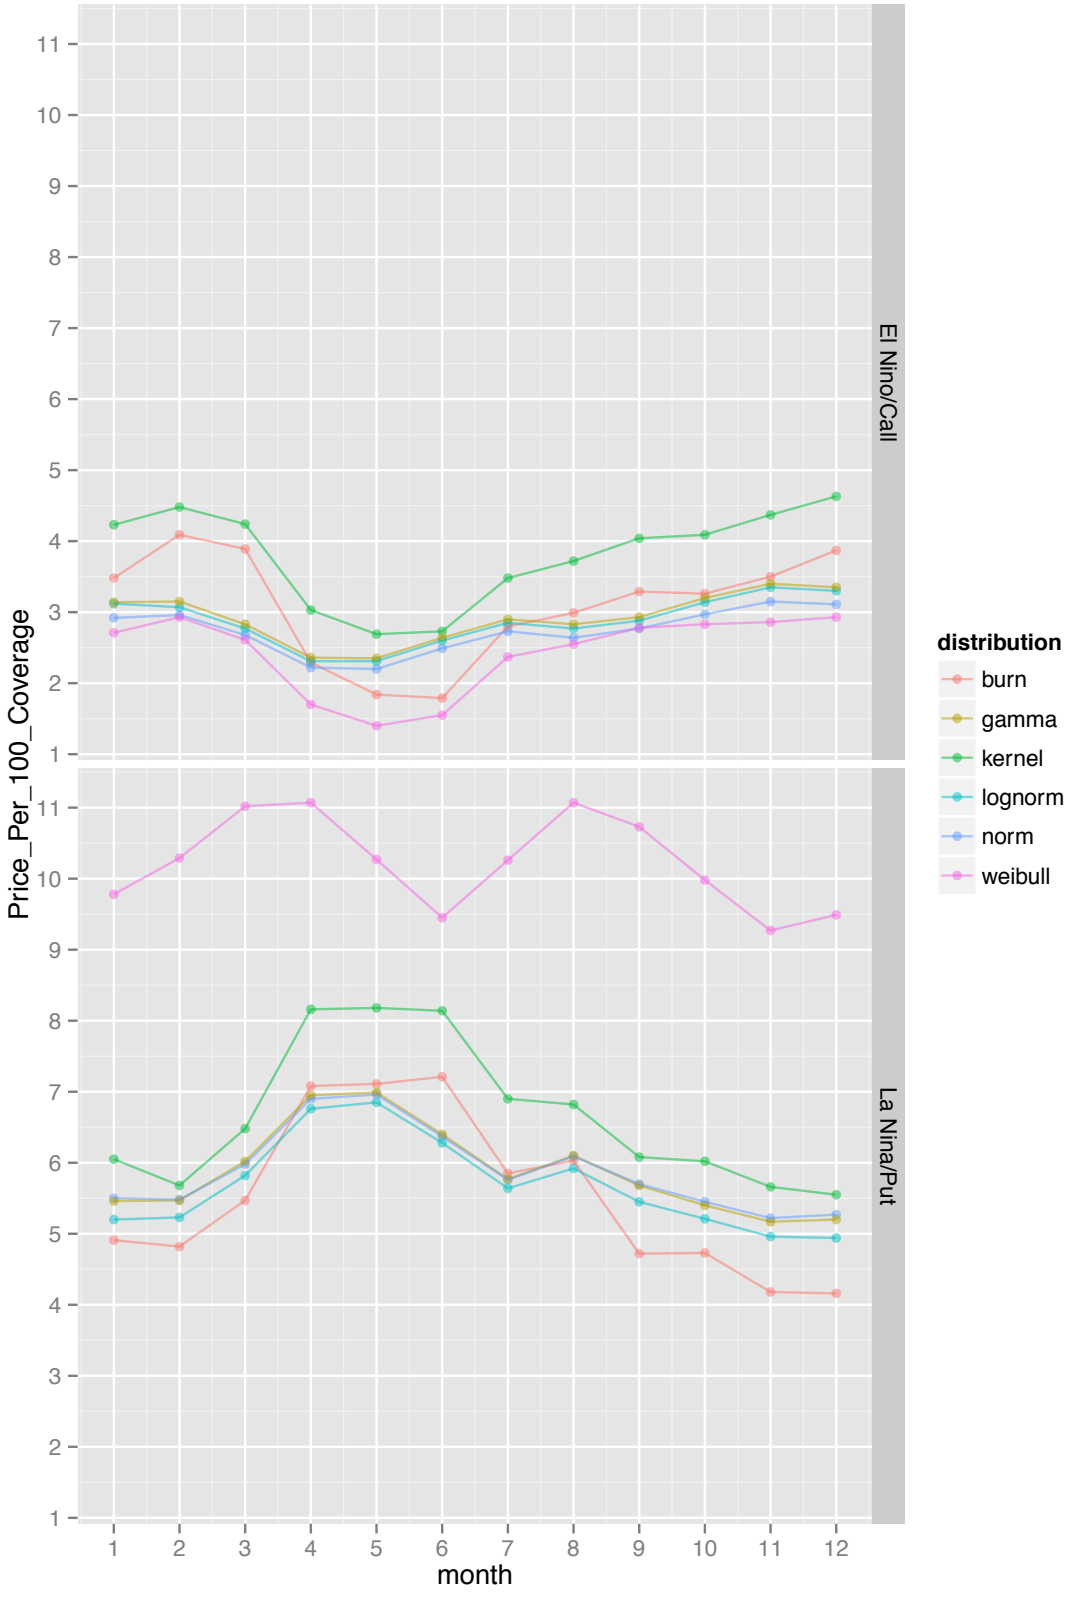

Figure 3.14: Expected price for options on Niño 3.4 by month, based on simulations from various distributions

Pricing inside the predictive window

Extreme El Niño/La Niña events emerge over time, with forecasts giving us even more useful hints in the months leading up to a given event. As those hints emerge, we change our beliefs around the likelihood of an event. The price of El Niño/La Niña risk protection should change to reflect those beliefs.

In this section, I present pricing analysis conditioned on SST forecasts released by Colombia University's International Research Institute for Climate and Society (IRI). Every month since mid-2002, IRI has collected forecasts issued by major centers of climatological research. Figure 3.15 shows IRI the forecasts as of March 2013.

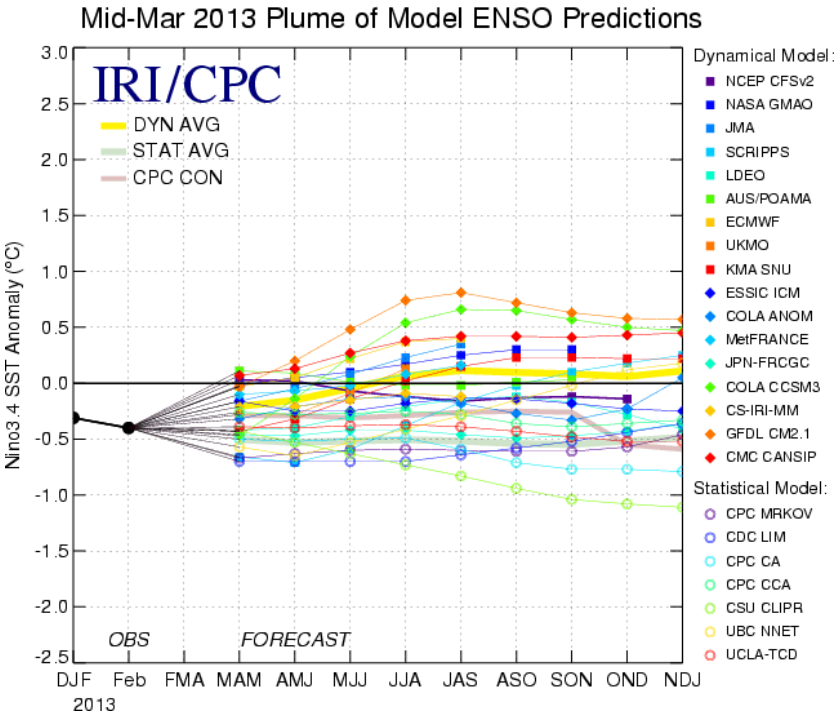

Figure 3.15: Example of IRI's collected forecasts - March 2013

I link forecasts and observed SSTs through a Bayesian regression that uses the long terms climate record as a prior. If the regression indicates that the forecasts have no predictive power, then all the simulated SSTs from the regression will simply reflect monthly historical averages.

Modeling the link between forecasts and SSTs

As an example, imagine that it is March and I am interested in predicting October Niño 3.4 SST. IRI's forecasts (given in terms of anomalies) are smoothed using three-month blocks, as in figure 3.15.

In that figure, there are three forecasts that contain information relevant to October SSTs - *ASO*, *SON*, and *OND*.

There are myriad ways of combining both individual and average forecasts for those three windows in a regression, but in this section I use as my predictive variable the IRI model average. So, in the above example, I would look at all the model averages made in March for *ASO*, *SON*, and *OND*, taking the average of those three numbers in any given year. I did the same for every month across that months valuable forecasts. That forecast average then conditions the long-term average anomaly for October[22]. IRI issues forecasts between 2 and 10 months prior to any given target month. For example, October SST forecasts begin in December and end in September. Since I want pricing for every month, from the vantage-point of every preceding month with IRI forecasts, I need to run a total of 108 separate regressions.

[22] I used anomalies rather than absolute SSTs to match IRI's convention.

$$\begin{aligned} \text{Monthly Niño 3.4 ERSST.3b anomalies}_{month,year} &\sim \mathcal{N}(\hat{y}_{month,forecastmonth,year}, \sigma^2_{y_{month,forecastmonth}}) \\ \hat{y}_{month,forecastmonth,year} &= a_{month,forecastmonth} \\ &\quad + b_{month,forecastmonth} * \\ &\quad \text{average of IRI average forecasts}_{month,forecastmonth} \end{aligned} \tag{3.1}$$

Those regressions, specified in equation 3.1, are a simplified version of a procedure that climate scientists and statisticians have recently used to merge ENSO forecasts[23,24]. Note first that I do not know the predictive power of IRI average forecasts. The parameter $\sigma^2_{y_{month,forecastmonth}}$ accounts for that forecasting uncertainty. It will be large where IRI average forecasts have shown low historical predictive power. Note also that this Bayesian regression will not be biased by non-stationarity. The underlying parameters are not assumed to be stationary, since they are realizations of an unknown distribution.

[23] Lifeng Luo, Eric F Wood, and Ming Pan. Bayesian merging of multiple climate model forecasts for seasonal hydrological predictions. *Journal of Geophysical Research: Atmospheres*, 112(D10), 2007

[24] CAS Coelho, S Pezzulli, M Balmaseda, FJ Doblas-Reyes, and DB StephENSOn. Forecast calibration and combination: A simple Bayesian approach for ENSO. *Journal of Climate*, 17(7):1504–1516, 2004

The prior probabilities I placed on model parameters are shown in equation set 3.2. There are weakly informative priors on b and σ_y, allowing them to move easily across a wide range of possible values in response to the data. a by contrast has a strongly informative prior based on historical data. This means that if b, the parameter indicating the predictive power of IRI's average forecasts, is at or near zero, then the resulting simulations from the posterior distribution will simply reflect long term trends in monthly SSTs.

$$\begin{aligned} a_{month,forecastmonth} &\sim \mathcal{N}(\text{mean anomalies}_{month}, \text{st dev anomalies}_{month}) \\ b_{month,forecastmonth} &\sim \mathcal{N}(0, 100) \\ \sigma^2_{y_{month,forecastmonth}} &\sim \text{Inv gamma}(0.001, 0.001) \end{aligned} \tag{3.2}$$

Dynamic pricing based on model results

The table below contains regression results for October SSTs, predicted between the preceding December and August. The regressions were all estimated using parallel Markov Chain Monte Carlo (MCMC) chains, each with 100,000 iterations, 50,000 of which were discarded as a warm-up[25]. The $\hat{R}$ on all parameters below and in part IV's Pricing Appendix were 1, indicating convergence on the simulation.

[25] Stan Development Team. Stan: A C++ library for probability and sampling, version 1.3, 2013. URL http://mc-stan.org/

	mean	sd	2.5th q	25th q	50th q	75th q	97.5th q	n_eff	Rhat
August forecast average covering October Niño 3.4 SST anomalies									
α	-0.10	0.10	-0.40	-0.20	-0.10	-0.10	0.10	91045	1
β	1.10	0.20	0.80	1.00	1.10	1.20	1.50	88920	1
σ_y^2	0.10	0.10	0.10	0.10	0.10	0.20	0.40	56829	1
July forecast average covering October Niño 3.4 SST anomalies									
α	-0.10	0.20	-0.50	-0.20	-0.10	0.00	0.20	92218	1
β	1.20	0.30	0.60	1.00	1.20	1.30	1.70	93712	1
σ_y^2	0.30	0.20	0.10	0.20	0.30	0.40	0.90	54297	1
June forecast average covering October Niño 3.4 SST anomalies									
α	-0.10	0.20	-0.40	-0.20	-0.10	0.00	0.30	95908	1
β	1.40	0.30	0.70	1.20	1.40	1.60	2.10	91107	1
σ_y^2	0.30	0.20	0.10	0.20	0.30	0.40	0.90	55596	1
May forecast average covering October Niño 3.4 SST anomalies									
α	-0.10	0.20	-0.50	-0.20	-0.10	0.10	0.40	92919	1
β	1.50	0.60	0.40	1.20	1.50	1.90	2.60	90255	1
σ_y^2	0.50	0.30	0.20	0.30	0.50	0.60	1.40	59205	1
April forecast average covering October Niño 3.4 SST anomalies									
α	-0.10	0.20	-0.50	-0.30	-0.10	0.00	0.30	88326	1
β	1.90	0.60	0.70	1.50	1.90	2.30	3.00	83902	1
σ_y^2	0.40	0.30	0.20	0.30	0.40	0.50	1.10	57674	1
March forecast average covering October Niño 3.4 SST anomalies									
α	0.00	0.20	-0.50	-0.10	0.00	0.20	0.50	101040	1
β	1.80	0.90	0.00	1.20	1.80	2.30	3.50	96782	1
σ_y^2	0.70	0.50	0.30	0.50	0.60	0.90	1.90	59539	1
February forecast average covering October Niño 3.4 SST anomalies									
α	-0.10	0.30	-0.70	-0.30	-0.10	0.10	0.60	98192	1
β	0.80	1.30	-1.80	0.00	0.80	1.60	3.40	88684	1
σ_y^2	1.10	0.80	0.40	0.60	0.90	1.30	3.20	54912	1
January forecast average covering October Niño 3.4 SST anomalies									
α	0.00	0.30	-0.60	-0.20	0.00	0.20	0.60	99518	1
β	1.00	1.60	-2.30	0.00	1.00	2.00	4.20	92225	1
σ_y^2	1.00	0.70	0.40	0.60	0.80	1.20	2.80	55715	1
December forecast average covering October Niño 3.4 SST anomalies									
α	0.00	0.30	-0.60	-0.20	0.00	0.30	0.70	80946	1
β	-0.30	1.90	-4.00	-1.40	-0.30	0.90	3.50	76663	1
σ_y^2	1.10	0.70	0.40	0.60	0.90	1.30	2.90	56323	1

Table 3.2: Bayesian regression linking October Niño 3.4 SST anomalies to average of relevant IRI ensemble forecasts

Looking at the 2.5th and 97.5th percentile of the distributions for b, its clear that the forecasts become more valuable predictors as the year goes on. Going from December to August, the 95 percent probability interval for the forecast parameter, b steadily tightens to a range including 1. This suggest that the correlation between forecasts and eventual SSTs increases throughout the predictive window. As the

explanatory value of b increases, a decreases. Just as climate scientists suggested, a's 95 percent probability tightening around 0 after March.

Using the posterior draws of parameter values from these 108 regressions, I simulated SSTs predicted by each possible forecast value between -2 and 2 (forecasts are rounded to one decimal). For example, I took 50,000 posterior draws of a, b, and σ_y^2 from the regression corresponding to October SSTs predicted by April forecasts. I used each of those 50,000 vectors of three parameters to randomly generate one October SSTs, based on an average April forecast of mild El Niño conditions in the coming October (a forecast value of 0.5.) That left me with 50,000 October SST conditioned on a forecast of 0.5 made in April. I repeated that procedure to produce conditional distributions for SSTs for each month of the year, predicted by a wide range of forecast values, from all possible forecast months. The resulting stochastic catalog allowed me to price El Niño/La Niña risk for any month given any IRI average forecast.

The empirical distribution functions of those posterior simulations, converted back into absolute SSTs, are shown in figures 3.16 and 3.16. In those figures, deeper blue lines indicate colder forecast averages from IRI and deeper red lines indicate warmer forecasts.

Notice how the blue and red lines are tightly bound ten months prior to any given target month (down the rightmost column) in figures 3.16 and 3.17. This indicates that forecasts had little or no predictive power, as warm forecasts were as closely associated with eventual warm conditions as cold forecasts, and visa versa. In some cases, where the blue lines peek above the red, the colder forecasts are actually associated with higher eventual SSTs. The fact that the red and blue lines bunch together as you move left to right across rows in figures 3.16 and 3.17 suggests that the signal from IRI's average forecasts deteriorates as we go further back in the predictive window.

By contrast, two months away from a target month (down the leftmost column of figures 3.16 and 3.17), forecasts are meaningful. Blue lines sit below red lines. So a warm forecast shifts the distribution of eventual SSTs warmer and visa versa.

The spring predictive barrier is also clear in the figures. The difference between April outcomes, conditioned on particularly cold and warm forecasts made just two months prior, is smaller than the same difference for February SSTs made ten months out. In visual terms, the ECDFs for row April, column t-2 months are more compact than the ECDFs for row February, column t-10 months. In other words, April SSTs show a weaker link to February predictions than February SSTs show to predictions from the preceding April.

In table 3.3, I translated these simulation results into pricing for October La Niña protection (put options on October SST). As before

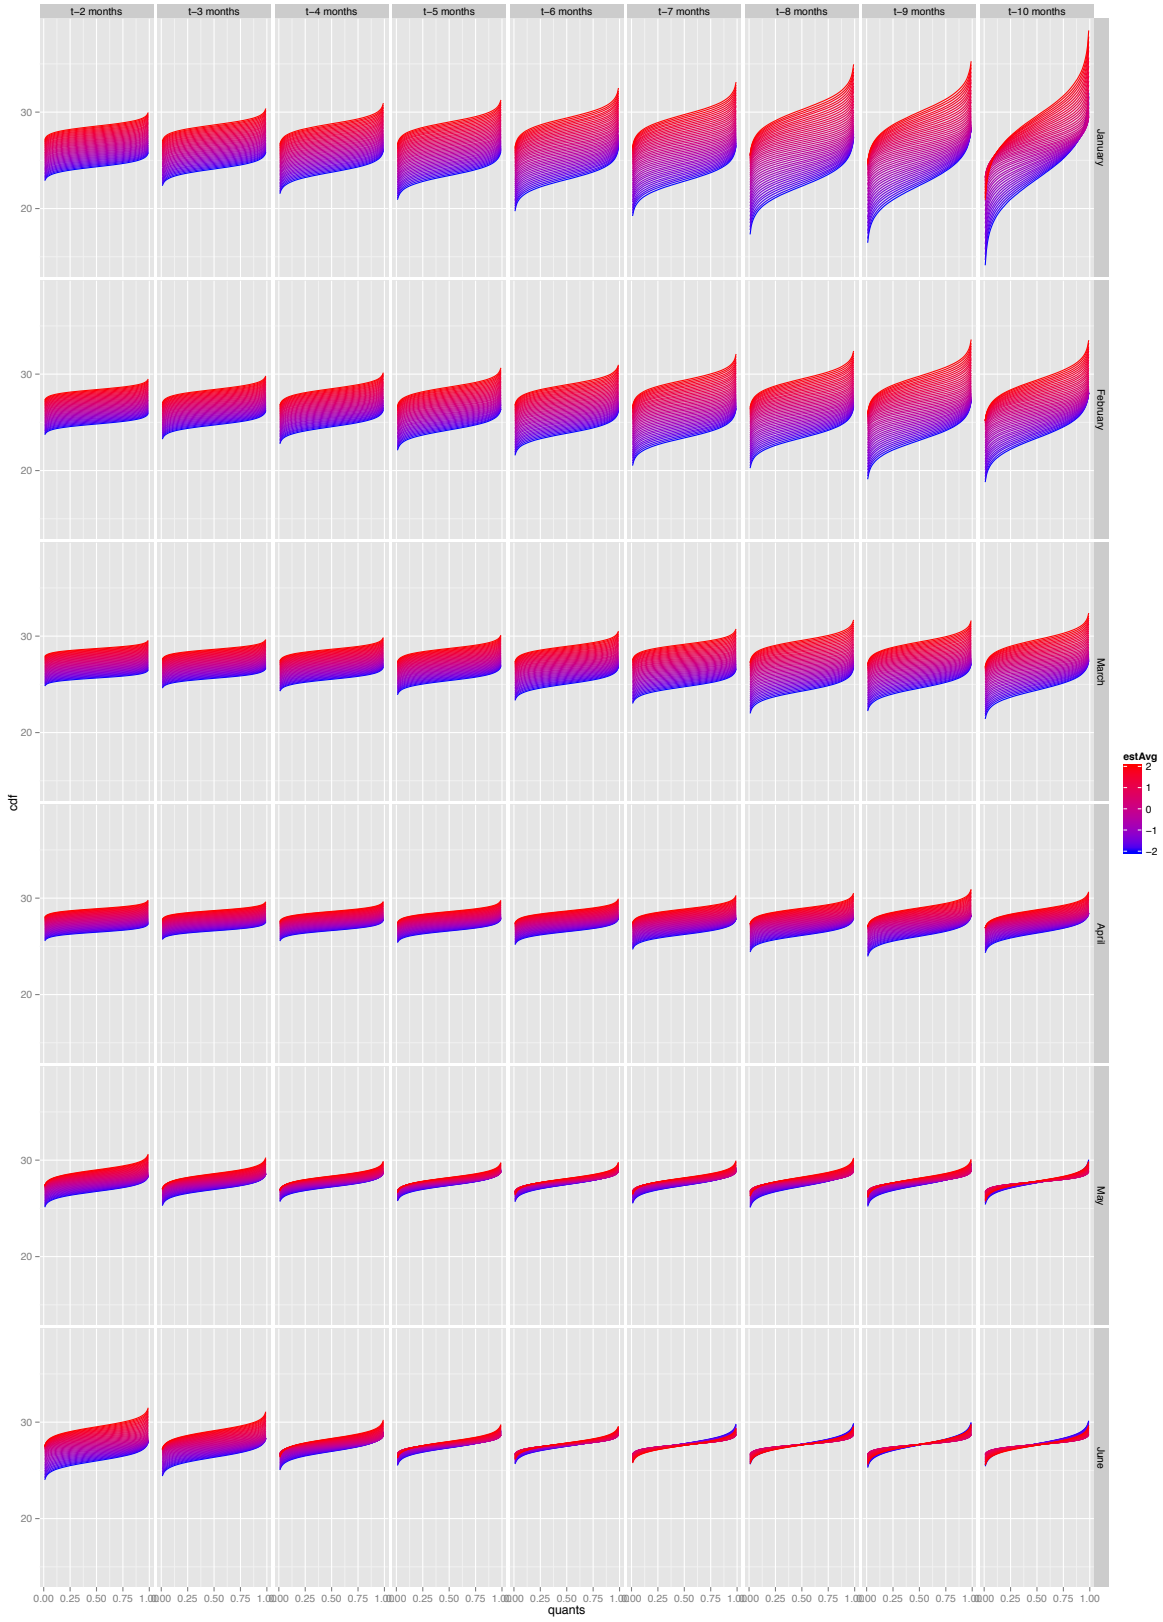

Figure 3.16: Cumulative distribution functions for realized January through June Niño 3.4 SST conditioned on average IRI ensemble forecasts for various months

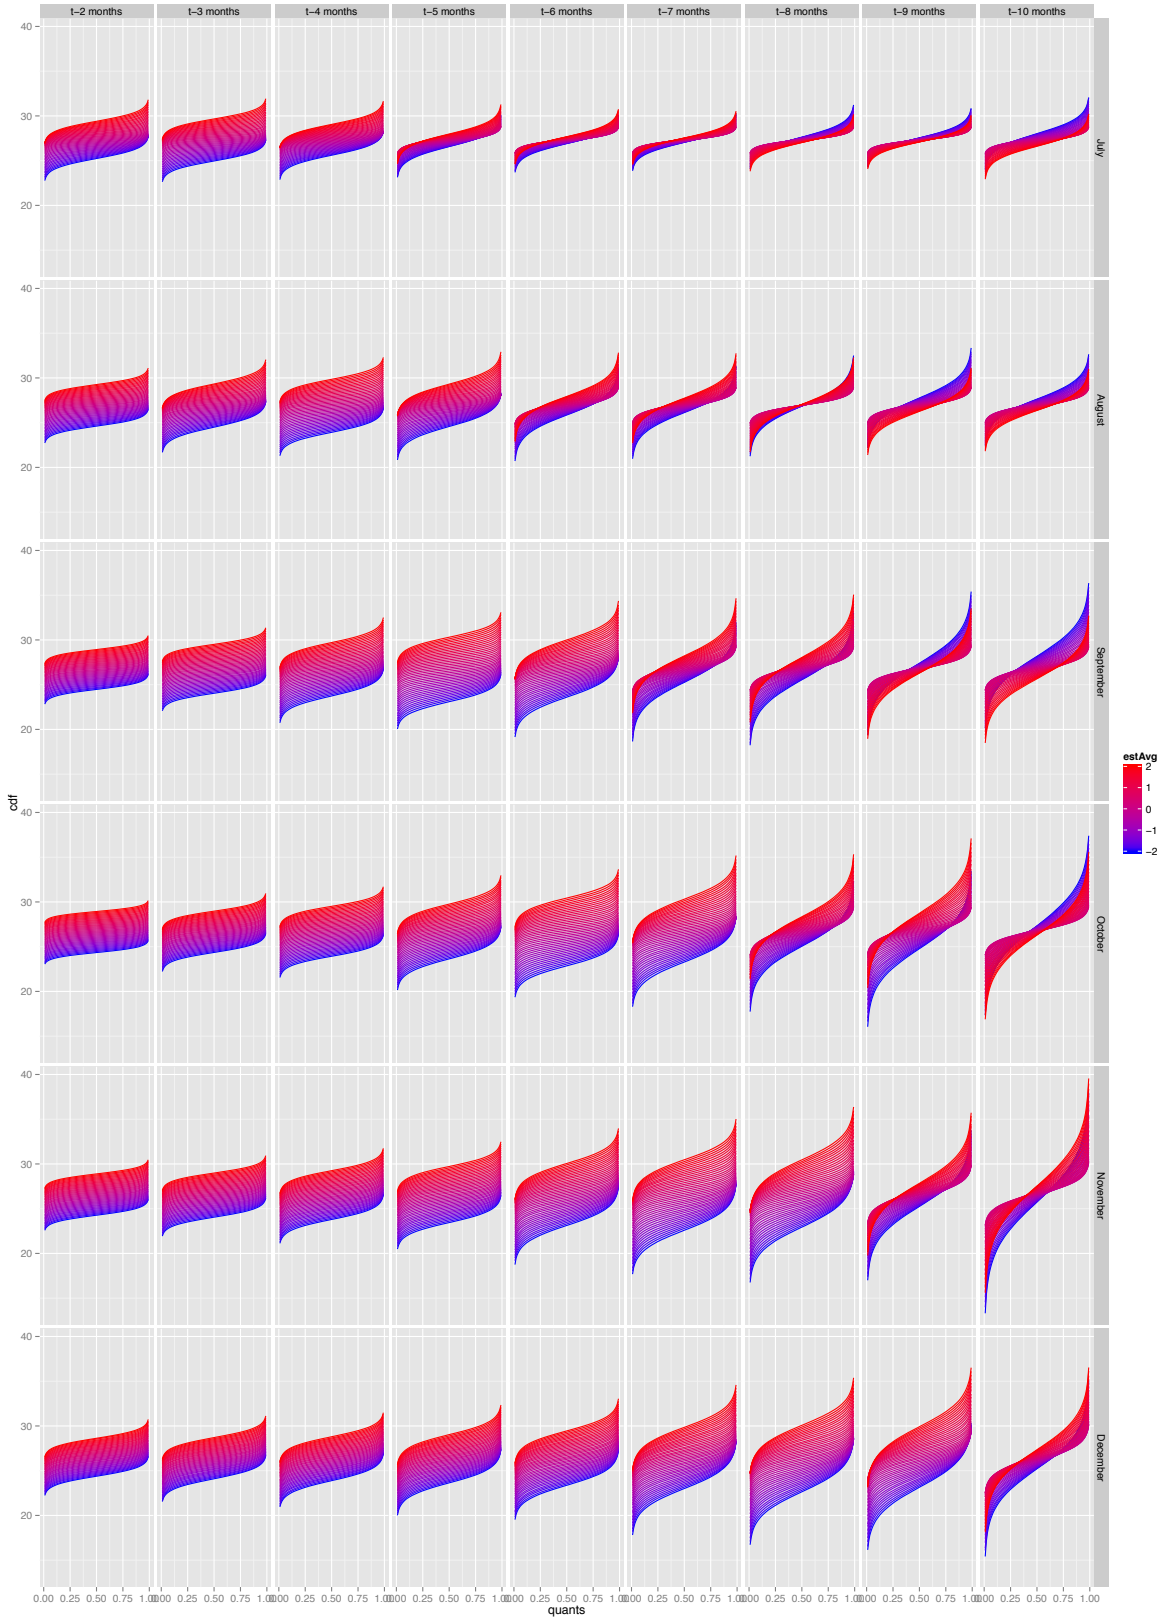

Figure 3.17: Cumulative distribution functions for realized July through December Niño 3.4 SST conditioned on average IRI ensemble forecasts for various months

IRI anom	price per USD	E[SST]	2.5th q	25th q	50th q	75th q	97.5th q
-2.00	0.80	23.93	0.00	0.66	0.96	1.00	1.00
-1.90	0.77	24.07	0.00	0.59	0.89	1.00	1.00
-1.80	0.73	24.21	0.00	0.54	0.82	1.00	1.00
-1.70	0.68	24.35	0.00	0.47	0.75	1.00	1.00
-1.60	0.64	24.49	0.00	0.41	0.68	0.95	1.00
-1.50	0.58	24.63	0.00	0.34	0.60	0.87	1.00
-1.40	0.53	24.77	0.00	0.28	0.54	0.79	1.00
-1.30	0.47	24.91	0.00	0.21	0.47	0.71	1.00
-1.20	0.41	25.05	0.00	0.15	0.39	0.63	1.00
-1.10	0.35	25.19	0.00	0.08	0.32	0.55	1.00
-1.00	0.30	25.33	0.00	0.02	0.25	0.48	0.99
-0.90	0.24	25.47	0.00	0.00	0.18	0.40	0.90
-0.80	0.19	25.60	0.00	0.00	0.11	0.33	0.81
-0.70	0.15	25.74	0.00	0.00	0.03	0.25	0.72
-0.60	0.11	25.88	0.00	0.00	0.00	0.17	0.63
-0.50	0.08	26.02	0.00	0.00	0.00	0.10	0.55
-0.40	0.06	26.16	0.00	0.00	0.00	0.02	0.46
-0.30	0.04	26.30	0.00	0.00	0.00	0.00	0.38
-0.20	0.02	26.44	0.00	0.00	0.00	0.00	0.31
-0.10	0.02	26.58	0.00	0.00	0.00	0.00	0.23
0.00	0.01	26.72	0.00	0.00	0.00	0.00	0.16
0.10	0.01	26.86	0.00	0.00	0.00	0.00	0.08
0.20	0.00	26.99	0.00	0.00	0.00	0.00	0.01
0.30	0.00	27.14	0.00	0.00	0.00	0.00	0.00
0.40	0.00	27.27	0.00	0.00	0.00	0.00	0.00
0.50	0.00	27.41	0.00	0.00	0.00	0.00	0.00
0.60	0.00	27.55	0.00	0.00	0.00	0.00	0.00
0.70	0.00	27.69	0.00	0.00	0.00	0.00	0.00
0.80	0.00	27.83	0.00	0.00	0.00	0.00	0.00
0.90	0.00	27.97	0.00	0.00	0.00	0.00	0.00
1.00	0.00	28.11	0.00	0.00	0.00	0.00	0.00
1.10	0.00	28.24	0.00	0.00	0.00	0.00	0.00
1.20	0.00	28.38	0.00	0.00	0.00	0.00	0.00
1.30	0.00	28.53	0.00	0.00	0.00	0.00	0.00
1.40	0.00	28.67	0.00	0.00	0.00	0.00	0.00
1.50	0.00	28.80	0.00	0.00	0.00	0.00	0.00
1.60	0.00	28.95	0.00	0.00	0.00	0.00	0.00
1.70	0.00	29.08	0.00	0.00	0.00	0.00	0.00
1.80	0.00	29.23	0.00	0.00	0.00	0.00	0.00
1.90	0.00	29.36	0.00	0.00	0.00	0.00	0.00
2.00	0.00	29.51	0.00	0.00	0.00	0.00	0.00

Table 3.3: Put option prices for October Niño 3.4 SST conditioned on IRI ensemble forecasts released in June

in this chapter, I used a payout function that began one standard deviation below normal and reached 100 percent of the nominal value of the agreement (sum insured) at three standard deviations below normal. The full conditional pricing tables for all months, covering both El Niño and La Niña, are available in the ENSO Pricing Appendix (part IV).

Adjusting risk prices for real transactions

The prices in table 3.3 and part IV only reflect the underlying risk of the index. In actual transactions, these pure risk prices will generally be:

- adjusted (downward) to reflect the time value of the premium paid by hedgers;
- subjected to some margining[26] rules, when applicable; and
- adjusted (upward) to allow for some reasonable expected profit for speculators.

[26] Margining refers to the process of setting aside collateral on financial trades. On exchange-traded derivatives there are clear, predictable rules for how much money must be set aside as collateral in a *margin account* as the trade's settlement index changes over time.

I won't address the first two procedures here. Time discounting is straight-forward and I have no expertise in margining. However, I will close this pricing chapter with a discussion on anchoring expectations about speculators' profits.

That anchoring is difficult because there are many reasonable benchmarks for profit expectations. Should speculators expect reinsurance-like returns? Futures-like returns? What have those returns been historically?

Figure 3.18 (discussed in greater detail in chapter 6) shows one reinsurance broker's estimates of returns to investment in CAT bonds and other insurance-linked securities (ILS). ILS markets show returns between 8 and 12 percent above LIBOR over the last decade for wind-exposed risk and between 3 and 9 percent for non-wind exposed risk. Most of the high prices on non-wind risk are clustered at the beginning of the observed period, so margins are at the low-end of that range.

These ILS returns are an important benchmark in their own right, since ENSO risk may trade in the form of CAT bonds. However, these returns are also as a standard proxy for returns in actual reinsurance markets.

Futures markets for major commodities provide another possible benchmark for ENSO risk. Gorton and Rouwenhorst [2004] compiled risk premium estimates for major commodities since 1959 by looking at futures prices relative to settlement at regular intervals. Those statistics are reproduced in the Miscellaneous Appendix's table 5. They include the average annualized arithmetic and geometric average returns. Note that these returns are not adjusted above a benchmark like LIBOR, so they are not directly comparable to the estimates in figure 3.18.

One interpretation of table 5 is that some futures markets offer speculators similar risk premiums as reinsurance markets, often above 10 percent per annum. That is surprising, given the general perception that futures are highly efficient.

However, the table also shows that some of the most liquid markets like corn, offer risk premiums of only a few percent. So clearly, the world of futures includes some highly efficient markets that sit along side many others that have provided speculators with high returns over the past five decades.

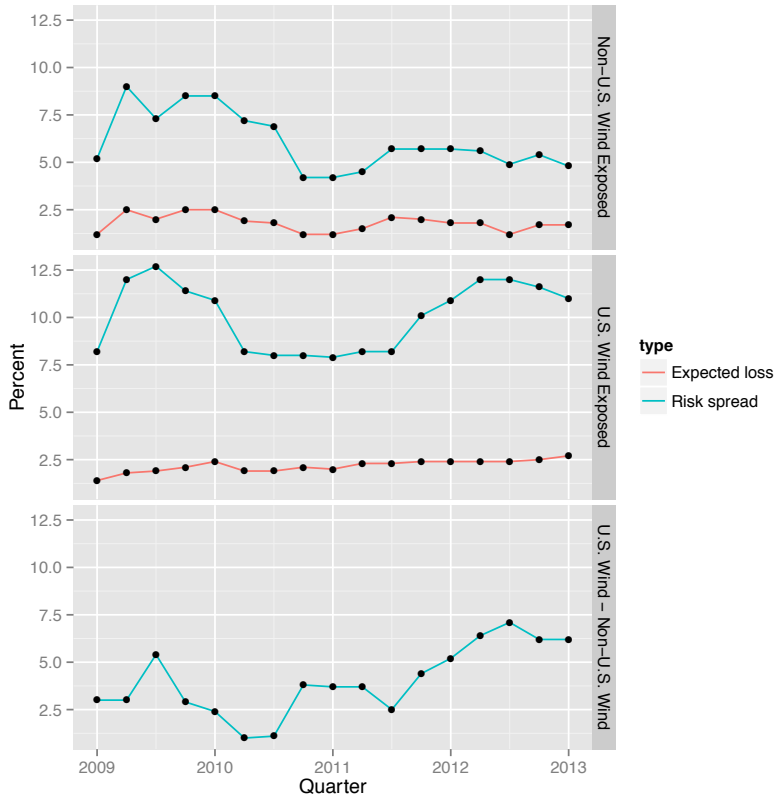

Figure 3.18: Weighted average risk premium and expected loss over last 12 Months on catastrophe bonds from "ILS Market Update" by Willis Capital Markets & Advisory (a large brokerage)

Regardless of their expected value, the risk premiums in futures markets are evidently more volatile than those in reinsurance markets. None of the markets above show returns that are more than two standard deviations away from zero. That may lead some to dismiss the observation of persistent positive risk premiums as a statistical anomaly.

Are there benchmarks within the world of exchange-traded derivatives that are more directly relevant to ENSO risk than the commodities markets profiled in Gorton and Rouwenhorst [2004]? Chincarini [2011] estimated risk premiums for on heating degree day (HDD) and cooling degree day (CDD) futures contracts on the Chicago Mercantile Exchange (CME). The results of that study are reprinted in Miscellaneous Appendix's tables 6, 7, 8, and 9. Only a small percentage of overall weather derivatives trading is in the form of on-exchange futures, so those estimates may not represent weather derivatives as a whole. Nevertheless, they suggest that recent speculative premiums for temperature risk have been on the low end of the range suggested by Gorton and Rouwenhorst [2004]. The high volatility of those historical premiums make it difficult to extrapolate about the long-term efficiency of those markets.

Finally, the CME's hurricane index derivatives the bring the reinsurance/ILS premiums noted in figure 3.18 into the context of exchange-traded derivatives. Surprisingly, the efficiency of hurricane derivatives markets relative to ILS is not obvious. The CME's marketing materials suggest a 9.64 percent risk spread over expected loss on an example hurricane contract. That is roughly 40 basis points below the mean risk spread since 2009. But. the risk spread was below that level in 4 out of 10 quarters since 2009 and averaged 150 basis points below the CME benchmark during those quarters.