INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

Bell & Howell Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA

UMI®
800-521-0600

University of California, San Diego

Climate Change and Variability,

and the Role of Information in Catastrophe Insurance Markets

A dissertation submitted in partial satisfaction of the

requirements for the degree Doctor of Philosophy

in Economics and International Affairs

by

Anthony Westerling

Committee in charge:

>Professor Richard T. Carson, Chair
>Professor Clive W. J. Granger
>Professor Barry J. Naughton
>Professor Jason M. Shachat
>Professor Richard C. J. Somerville

2000

Copyright

Anthony Westerling, 2000

All rights reserved.

The dissertation of Anthony Westerling is approved, and it is acceptable in quality and form for publication on microfilm:

Chair

University of California, San Diego

2000

iii

To Naoko Kada, who makes it all possible and worthwhile.

Limits of survival are set by climate, those long drifts of change which a generation may fail to notice. And it is the extremes of climate which set the pattern. Lonely, finite humans may observe climatic provinces, fluctuations of annual weather and, occasionally may observe such things as "This is a colder year than I've ever known." Such things are sensible. But humans are seldom alerted to the shifting average through a great span of years. And it is precisely in this alerting that humans learn how to survive on any planet. They must learn climate.

<div align="right">

– Frank Herbert
Children of Dune

</div>

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Acknowledgments

I would like to thank the following for their time and advice during the preparation of Chapter I, while retaining all blame for any errors and omissions: Richard Carson, Dan Cayan, Michael Dettinger, Reinhardt Flick, Clive Granger, Gordon J. MacDonald, and Richard Somerville.

I gratefully acknowledge the hard work and guidance of my coauthor for the second chapter, Jason Shachat. Jason Shachat and I acknowledge financial support for Chapter II from the University of California, San Diego's Committee on Research, Hewlett Packard Laboratories, and the National Science Foundation. We thank Naoko Kada, Sivan Ritz, and Jeff Bethard for their capable research assistance, the participants of the microeconomic theory workshop at UCSD, and Mark Machina for his insightful comments.

Thanks to Christopher DeMarco and his assistant Juan for keeping me on my feet while I finished my dissertation. Thanks also to David Noelle and Jeanne Milostan for inspiring me by finishing their dissertations. And thanks most of all to Naoko Kada for her unstinting love and support which made it all possible and very worthwhile.

x

VITA

1987   Bachelor of Arts
Economics/International Area Studies
University of California, Los Angeles

1987   Summer Intern
Bureau of Economics
U.S. Department of State

1988-89   English Instructor
Hangzhou Commercial Institute
Hangzhou, People's Republic of China

1989-90   Research Associate
Economics Analysis Corporation
Century City, California

1991-93   Assistant General Manager and Associate Editor
Journal of Environment and Development
University of California, San Diego

1991-96   Research Assistant
Graduate School of International Relations and Pacific Studies
University of California, San Diego

1992   Teaching Assistant
Thurgood Marshall College
University of California, San Diego

1993   Visiting Fellow
National Center for Environmental Health
Centers for Disease Control and Prevention

1993-97   Teaching Assistant
Graduate School of International Relations and Pacific Studies
University of California, San Diego

1994-95   Staff Research Associate II
California Environmental Technology Center
Scripps Institution of Oceanography

1996	Young Scientist
	International Institute of Applied Systems Analysis
	Laxenburg, Austria

1996   Young Scientist
International Institute of Applied Systems Analysis
Laxenburg, Austria

1997-98   Instructor, Summer Session
International Career Associates Program
Graduate School of International Relations and Pacific Studies
University of California, San Diego

2000   Doctor of Philosophy
Economics and International Affairs
University of California, San Diego

PUBLICATIONS

Westerling, A., (forthcoming abstract) "El Niño and One Hundred Years of Storm Surge in the Eastern North Pacific," In Proceedings of the 16th Annual Pacific Climate Workshop, Santa Catalina Island, California , 1999.

Westerling, A., (1992) "Climate Change and Chinese Energy Policy," Woodrow Wilson School *Journal of Public and International Affairs.*

FIELDS OF STUDY

Major Field: Economics

Studies in Econometrics, Time Series Analysis, and Numerical Methods
Professors Richard T. Carson, Wouter Den Haan, Robert F. Engle
Clive W. J. Granger, James Hamilton, Ramu Ramanathan and Halbert White

Studies in Microeconomic Theory, Game Theory and Experimental Economics
Professors Vincent P. Crawford, John McMillan, Garey Ramey and
Jason M. Shachat

Studies in Comparative Policymaking
Professors Stephan Haggard, Matthew D. McCubbins, Barry J. Naughton
Francis Rosenbluth and Matthew S. Shugart

Studies in Climate and Environmental Science and Policy
Professor Gordon J. MacDonald

Abstract of the Dissertation

Climate Change and Variability,

and the Role of Information in Catastrophe Insurance Markets

by

Anthony Westerling

Doctor of Philosophy in Economics and International Affairs

University of California, San Diego, 2000

Professor Richard T. Carson, Chair

This work is comprised of three loosely related papers, each of which

considers some aspect of catastrophe risk deriving from climate

variability. In *El Niño and One Hundred Years of Storm Surge in the*

Eastern North Pacific, I analyze roughly one hundred years of sea-level

height data in Honolulu, San Diego, San Francisco, and Seattle. The

results indicate that the frequency of large storm surges has increased

greatly overall, with strong increases in Honolulu and San Diego, a less

significant increase in San Francisco, and no significant change in

Seattle. In *Information Aggregation in Catastrophe Reinsurance*

Markets, Jason Shachat and I test whether experimental catastrophe

futures markets can aggregate diverse risk information. We conclude

that our markets equilibria reflect participants' primarily acting on prior

xiii

information only, with buyers of reinsurance underestimating the probability of catastrophic losses. Finally, in *The Value of Extended Climate Forecasts in Insurance Markets: Heterogeneous Risk Beliefs, Market Power and Regulation*, I ask whether twelve-to-eighteen month climate forecasts generate positive utility for consumers of catastrophe insurance in simple insurance models. I find that benefits to consumers depend upon additional model specifications such as solvency regulations, market power, or heterogenous risk beliefs. I also find that monopoly profits decrease and competitive profits remain the same when forecasts are introduced unless heterogenous risk beliefs are specified.

Chapter I

El Niño and One Hundred Years of Storm Surge in the Eastern North Pacific

A. Abstract

In this paper I analyze daily mean sea-level height measurements for one hundred years from four National Oceanographic and Atmospheric Administration (NOAA) tide stations in the Eastern North Pacific. The largest winter storm surges have increased in frequency since early in this century. Linear and negative binomial regression models for Honolulu and San Diego imply the frequencies of the largest storm surges have increased roughly 280% and 150%, respectively, while San Francisco shows a somewhat less significant increase of 40%. There is no statistically significant change in Seattle. Increased storm surge frequency in San Diego and San Francisco appears to be positively associated with a greater frequency of intense El Niño events in recent decades. With the possible exception of the very strong El Niño winters of 1983 and 1998, the incidence of El Niños appears to be negatively correlated with that of large storm surges in Seattle and Honolulu. To illustrate the potential impact of a continued increase in mean sea level, de-trended daily mean sea level series for four quarter-century time periods in San Francisco are projected onto an increasing sea level trend for the period 2001-25. The results indicate that the return times for

1

earlier benchmarks increase by a factor of 62 to 124 times.

B.   Introduction

The Intergovernmental Panel on Climate Change (IPCC) estimates that global mean sea levels have risen 10-25 cm in the past 100 years, and projects a further rise of 50 cm by 2100.[1] The sea-level records from three of the NOAA tide stations used in this analysis–San Diego, San Francisco, and Seattle–indicate a sea-level rise of roughly 20 cm this century. This is somewhat smaller than either the maximum storm surge or the variation due to tides experienced on the U.S. West Coast.[2] For example, the maximum storm surges recorded at San Francisco can exceed 60 cm, while the maximum sea-level variation off the California coast due to tides is nearly five times larger.[3] El Niño events can also elevate monthly mean sea levels by up to 30 cm or more off the California coast.[4]

When extreme high tides, storm surges, ENSO-related sea-level increases and long-term sea-level increases come together, the results can be costly. This is what happened during the 1982-83 El Niño, when elevated sea-levels from winter storms resulted in more than $100 million in damages to coastal property in California, in

---

[1]IPCC 1995: 50cm is the best estimate, with a low of 15 cm and a high of 95 cm.

[2]Storm surge here refers to the "local, instantaneous sea-level elevation that exceeds the predicted tide and which is attributable to the effects of low barometric pressure and high wind associated with storms...excluding the effect of waves." (Flick 1991)

[3]Flick 1998.

[4]Ibid.

addition to losses from high winds and intense precipitation.[5] Flick and Cayan (1984) note that this was a relatively rare event in that large storm surges coincided on three separate occasions with high tides near a four-year maxima in the monthly mean high tide, on top of a sustained sea-level elevation associated with strong El Niño conditions. By comparison, Flick (1998) finds peak storm surges and wave sizes in California set new records in 1998, but did not coincide with extreme high tides, and damages were reduced.[6]

Clearly, any changes in the frequency and intensity of the largest winter storms, or in their tendency to coincide with ENSO-related fluctuations in mean sea level, could be as or more important than the IPCC's anticipated rise in mean sea levels in determining the risk of damage to coastal property. While anticipated effects of climate change include an increased incidence of diverse climatic extremes, it is less clear what effect, if any, there will be on mid-latitude cyclones or El Niño events. Model simulations of mid-latitude storms under enhanced-CO2 climate regimes vary considerably with the model and region considered. As the IPCC (1992) notes, their conclusions are not easily comparable because of differences in the variables used to indicate storm activity. Keeping these caveats in mind, we note that while several studies of simulated winter cyclone activity in the northern hemisphere with enhanced

[5]See Flick and Cayan 1984. Note that this is damage due to coastal flooding, erosion, and large waves.

[6]It should be noted that part of this reduction in damages may also be due to the destruction and/or hardening of vulnerable structures after the strong El Niño winter of 1982-83 (See Griggs and Brown 1998).

CO_2 show no change or a decrease in cyclone frequency or activity, several also suggest an increase in cyclone intensity. Zhang and Wang (1997) and Bates and Meehl (1986) find a decrease in measures of winter cyclone activity in the northern hemisphere. Konig et al (1993) find no change in frequency of winter cyclones. Carnell et al (1996) and Hall et al (1994) find an increase in storm intensity, and Lambert (1995) finds both a decrease in northern hemisphere winter cyclone frequency and an increase in cyclone intensity. More recently, Timmermann et al (1999) find indications of an increase in the frequency of El Niño conditions and an intensification of La Niña events.

With respect to the effect of winter storms on coastal property, the frequency of the most intense cyclones is probably much more pertinent than that of winter storms overall. Since the IPCC Second Assessment has found "the balance of evidence suggests that there is a discernable human influence on global climate," including increases in global mean surface temperature, precipitation, and rainfall extremes, it seems reasonable to look for evidence of changes in winter storm frequency and intensity on a regional scale similar to that of model simulations.

In this analysis, we examine nearly 100 years of daily mean sea-levels constructed from hourly sea-level height data collected at NOAA tide stations in Honolulu, San Diego, San Francisco and Seattle. Effects of the tides and low-frequency climate fluctuations are removed. Large winter storms, as proxied by peak storm surges in the daily mean sea-level residuals, are tested for significant changes in frequency and maximum intensity. We also consider their propensity to coincide with sea-levels

additionally elevated by ENSO events. We are particularly interested in whether there is a detectable increase in the frequency of the most intense storms.

Effects of tide cycles of one day or less in period on the daily mean are, as we will discuss below, too small to affect this analysis. A long term increasing trend in sea level, changes in sea level correlated with climate fluctuations like ENSO events, and tide cycles of one month or greater in period are filtered from the data using estimated median sea levels. Median sea level is calculated using a median-smoothing procedure with a 31-day window on the sea level data, controlling for the estimated effects of deviations in atmospheric pressure at sea level. As we will show, this is effective in removing most of the effect of the tides since the most important tide cycles for daily means, the Solar Annual and Solar Semi-annual, are greater than one month in period.[7] ENSO-related sea-level changes typically persist for about a month or more as well. The result of this filtering is a time series of residuals for each tide station, containing fluctuations of relatively high frequency sea level activity. Positive excursions from the mean in these residuals are defined as storm surges.[8] Extreme peaks in the residuals for October through March are used as a proxy for the frequency of large winter storms.[9] Linear and negative binomial regression models are used to estimate the

---

[7]Schureman 1958.

[8]As in Flick 1991.

[9]Flick and Cayan (1984) find that storm episodes they identify by increased average daily wind speed and decreased barometric pressure are highly correlated with sea level anomalies in that part of the San Diego sea level height series they examine. We
(continued...)

relationship between annual large winter storm frequencies, a time trend, and the incidence of ENSO events. A joint Mantel-Haenszel chi-square statistic is used to test for a progressive relationship between winter storm frequency and time over multiple locations. The hypergeometric distribution is used to test for changes in the maximum intensity of storms over the observed period.

In the next section the data and summary statistics are presented, followed by a discussion of the trends and cycles apparent in the data and special problems presented by some of the data. In section 3 we describe the method used to filter annual and shorter period cycles from the data, as well as longer term fluctuations and the overall trend in sea level. In section 4, we motivate a definition of a storm in terms of the sea-level height record and compare the results to the identification of storms using wind speed and barometric pressure data. In section 5 the results of statistical tests are presented. In section 6 we conclude with a discussion of return times for extreme sea levels in San Francisco.

C.   The Data

The original data, supplied by the Oceanographic and Lake Levels Division (OLLD)of the National Oceanographic Service at NOAA, consist of hourly mean sea levels measured relative to a fixed point. Data are from four tide stations: Honolulu, Hawaii; San Diego and San Francisco, California; and Seattle, Washington. Their particulars are presented in Table 1.1. Daily mean series were constructed from the

---

[9](...continued)
assume that this relationship holds generally.

hourly data for each location.[10]

Table 1.1: Summary of daily mean tide series

	San Diego	San Francisco	Seattle	Honolulu
NOAA Tide Stn.	9410170	9414290	9447130	1612340
Start Date:	1906 : 01 : 21	1901 : 01 : 01	1899 : 01 : 01	1905 : 01 : 01
End Date:	1998 : 10 : 31	1998 : 10 : 31	1998 : 10 : 31	1998 : 10 : 31
Observations:	32884	35549	36264	33539
Missing (%):	1003 (3.0)	184 (0.5)	199 (0.5)	733 (2.1)
Winter Obs:	16287	17548	17978	16663
Winter Missing(%):	479 (2.9)	129 (0.7)	68 (0.4)	285 (1.7)
Trend mm / yr:	2.3	2.1	2.1	1.5

The unfiltered daily mean data for all four tide stations are in Figure 1.1. Note that the number of missing observations varies widely over the four data sets, but in any case does not exceed three percent of the days in any series. Fortunately, many of the gaps occur in summer and do not greatly affect this analysis. There is no apparent pattern to the remaining gaps in the series for the months of interest that would lead us to believe they will bias the results of statistical tests for changes in winter storm surge frequency in any particular direction. We will consider the effects of the missing data on filtering techniques further below.

---

[10]Seattle data through 1988 were provided as daily means by Dr. Gordon J. MacDonald. Other series were obtained directly from NOAA's National Oceanographic Service as hourly data.

San Diego

San Francisco

Seattle

Honolulu
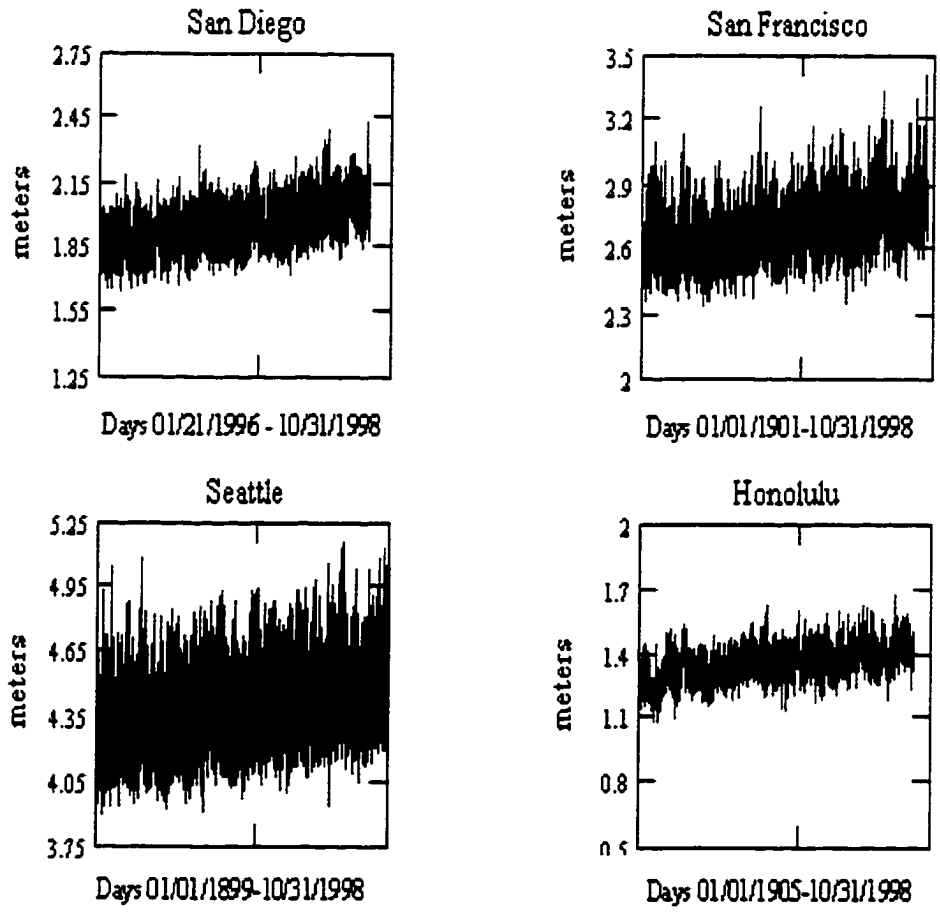
Figure 1.1: Raw daily mean sea level, station datum.
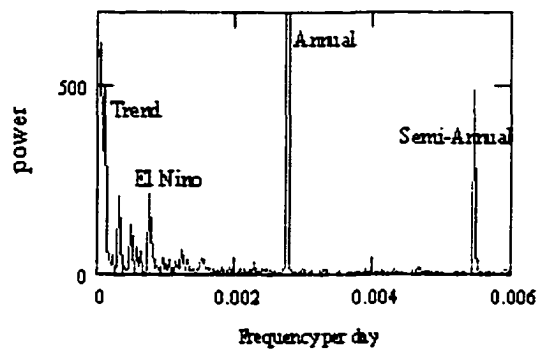
Figure 1.2: San Diego Lomb Periodogram
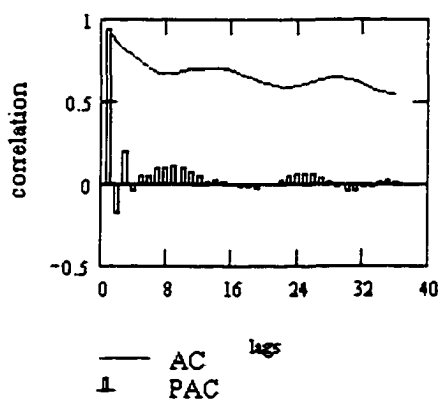
Figure 1.3: Autocorrelation (AC) and Partial Autocorrelation (PAC) for unfiltered San Diego daily mean data.

Power spectra for these series were obtained using the Lomb method for unevenly spaced data.[11] The Lomb normalized periodogram for San Diego, Figure 1.2, is typical for these data. The primary cyclical components are at the annual and semi-annual frequencies. The secular trend and multi-year shifts in sea level, probably associated with ENSO events,[12] are apparent in the low-frequency spectra in the periodogram in Figure 1.2. Some red noise is also evident in the periodogram, consistent with autocorrelation in sea level. Examination of the correlogram for the San Diego series, Figure 1.3, confirms the presence of autocorrelation in daily mean tide heights. An auto-regressive model with one lag was fit with an inverted root of 0.93 (see Table 1.2). This is indicative of the persistence of sea level effects introduced by tidal and climatological phenomena.

Table 1.2: Auto-regressive model estimation for San Diego daily means, 1 lag.

Variable	Coefficient	Std. Error	T-Stat.	Prob.
C	1.855	4.53E-03	409	0.00
Date	6.326E-06	2.31E-07	27	0.00
AR(1)	0.93	2.03E-03	458	0.00

[11]See Press et. al., 1992. This method allows us to estimate a periodogram for each data series despite the presence of gaps from missing observations. A common alternative is to interpolate missing values, but this can introduce distortions into the estimated power spectra. The Lomb method avoids these problems.

[12]For discussions of the effect of ENSO events on sea level, see Allan et al 1996, Clarke and Van Gorder 1993, Flick 1998, Flick and Cayan 1984, Quinn et al 1987, Quinn and Neal 1989, to name but a few.

These data generally accord with our expectations. In tide models, there are five primary long-period harmonic constituents which affect the daily mean tide.[13] Of these, the two with the largest amplitudes are the solar annual and solar semiannual cycles. According to Schureman (1994), the others–the lunar monthly, luni-solar synodic fortnightly and the lunar fortnightly–are usually too small to be of practical importance. The two fortnightly cycles appear in periodograms for some of the data series used here, but their amplitude appears small compared to the storm surges of interest. In the San Diego series, for example, the amplitude estimated for the luni-solar fortnightly cycle is an order of magnitude less than the maximum storm surge in the filtered daily mean sea level.[14] The lunar monthly cycle also appears to have little effect on the daily mean tide in these series. Considering the short periods of these cycles compared to the time scale of this analysis, their effect on the results should be negligible.

The secular trend for San Diego reported in Table 1.1, 2.3 mm/year, is the slope of the ordinary least squares fit to the daily mean sea level. This is slightly above the 0.7 ft per century (2.1 mm/year) rise for San Diego for 1906 to 1983 found by Flick and Cayan (1984). Flick and Cayan (1984) state that

> this is comparable to the global rate of rise and is typical of all long-term California stations, except for a few where local uplift or

---

[13]Schureman 1994 gives an excellent summary of harmonic tide models and the principal constituents of the daily mean tide.

[14]Some tide constituents less than one day in period can also disturb the daily mean tide heights (Schureman 1994). The strongest of these disturbances has the same period as the luni-solar fortnightly constituent, so any effect it has upon the daily mean values is subsumed in the amplitude estimated for the luni-solar fortnightly constituent.

subsidence dominate the relative sea level signal.

The long term trends in sea-level height for all stations fall between 1.5 mm/year and

2.4 mm/year and are consistent with the IPCC's (1995) estimated increase in global

mean sea level of 1 to 2.5 mm/year.

In addition to the tides and long term trend, the sea-level height data exhibit

substantial multi-year departures from the trend line. Much of this inter-annual

variability, as cited above, is related to ENSO events. Figure 1.4 demonstrates a clear

association between large positive departures from the trend in annual mean sea level

and moderate to strong El Niño events.
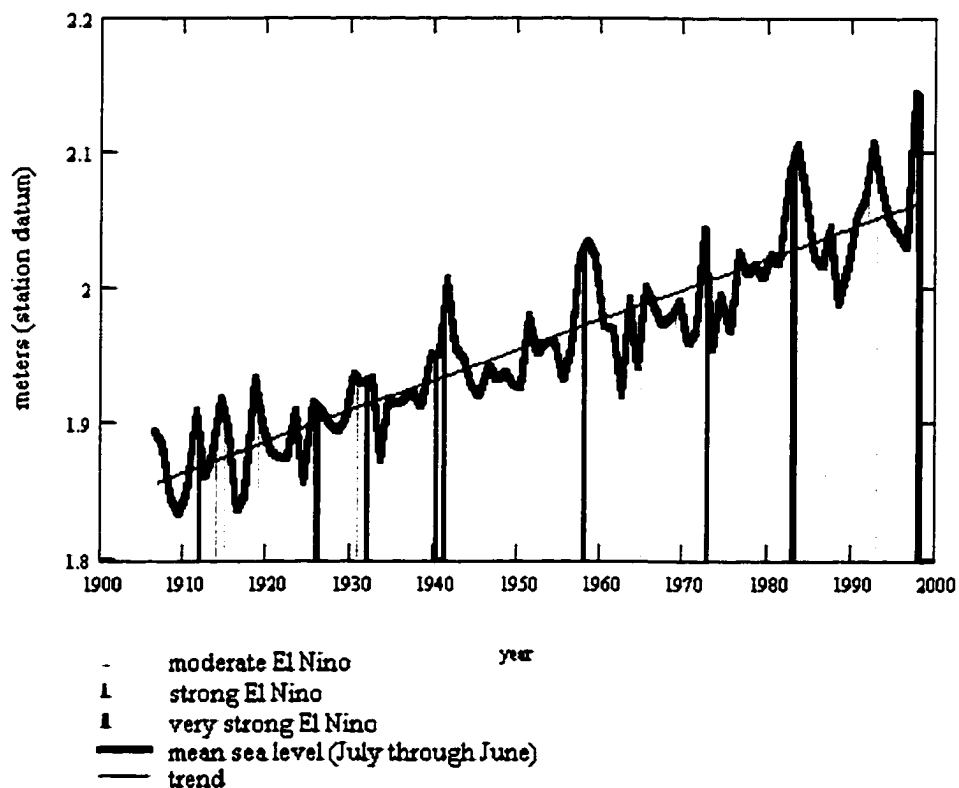
Figure 1.4: San Diego annual mean sea level and ENSO

Two of the data series presented special problems for this analysis. In the case of San Diego, the tide station was moved in 1926. In the Seattle data, a period of several hundred days beginning in 1964 is downshifted, possibly due to equipment mis-calibration. The San Diego data recorded before the tide station was moved and the down-shifted Seattle data are adjusted to the long term trend line for the respective series. Summary statistics are based on the adjusted series, except where otherwise noted. These adjustments do not have any effect on the residual series used for the analysis of storm frequency, since these data are filtered with median values. For San Diego, the unadjusted daily means were regressed on a constant, trend and dummy variable for the station move in 1926. While the move resulted in a change in mean, variability does not change over the two periods. Thus, it seems unlikely that these measurement problems have introduced a significant, systematic change in storm surge characteristics at either of the two locations.

There are frequent small gaps in the hourly data for all four stations. Wherever one or more hours is missing, the day is counted as a missing observation in the daily mean sea level series. In addition, there are occasional timing errors in the data, where measured values in the series are shifted forward or backward by one or more hours, and various measurement errors. The worst of the measurement errors are readily obvious (the city adjacent to the tide station would have been swept away...) and are treated as missing values. The filtering techniques used here were selected to avoid accentuating the effects of the remaining small errors and gaps in the data. In addition, there are four winters with a month or more of missing data in San Diego and

Honolulu, and one in San Francisco (see Table 1.3). The majority of the missing observations occur in Honolulu and San Diego. As we will see below, the results of our analysis are very strong for these two stations. The missing values would have to have been very different indeed from the rest of the winters in the two series in order to qualitatively change the results of this analysis.

Table 1.3: Winters with a month or more of missing observations.

	Honolulu year (days)	San Diego year (days)	San Francisco year (days)	Seattle year (days)
Winters with a month or more of missing observations	1973-74 (44) 1976-77 (56) 1993-94 (91) 1994-95 (31)	1925-26 (87) 1962-63 (116) 1970-71 (34) 1986-87 (52)	1977-78 (59)	none

D.   Filtering Techniques

We are interested in identifying and ranking by height large storm surges, which typically elevate the local sea level significantly for one to five days. Consequently, it is necessary to filter out changes in sea level from other causes to the greatest extent possible.   We do this primarily by removing sea-level effects with frequencies equal to or less than one day or greater than or equal to one month.   Effects of higher-frequency changes in sea level, such as those from tide cycles one day or less in duration, are effectively removed by the use of daily means. As we note above in the

Data section, there may be some residual effects of these high-frequency tide cycles on the daily mean sea level, but their magnitude is very small compared to that of a large storm surge. Likewise, there may be some effect on the daily mean sea-level from fortnightly tide cycles, but these are also small in magnitude in our data. Lower frequency effects of the tides and climatological phenomena are largely removed by subtracting the smoothed, 31-day median sea level, controlling for the effects of changes in barometric pressure.

We must control for the effects of changes in barometric pressure when determining the median sea level. In winters with a few small storms and many days of high barometric pressure, a relatively minor storm surge may appear extremely large compared to the unadjusted median sea level. Similarly, in a winter with many large storms and persistent low sea level pressure, the result is to reduce the apparent height of some large storm surges relative to the median sea level. To avoid introducing these biases, we adjust the median sea level for changes in barometric pressure. Daily average barometric pressure values are generated using the National Center for Atmospheric Research's (NCAR's) daily 5-degree lat-lon sea level pressure (SLP) grids and averaging nearby grid points for each station.[15] A regression of daily mean sea level heights on the daily mean SLP differential (series mean minus daily mean pressure) yields a relation of 1 cm per millibar for San Francisco and 1.2 cm per millibar for the other three stations (for example see Figure 1.5). The residual sea levels are still

---

[15]These data are available courtesy of NCAR and D. Cayan at the Climate Research Division, Scripps Institution of Oceanography.

affected by high winds and other sources of natural variability. These residuals are then smoothed with multiple passes of a median-smoothing procedure with a 31-day window to generate a smoothed, pressure-adjusted 31-day median sea level series (see Figure 1.6). Daily mean sea levels for which pressure data are unavailable are filtered with median values, but are otherwise unadjusted. The residual sea level used in the subsequent analysis is the difference between the daily mean sea level and the smoothed, pressure-adjusted 31-day median sea level. Lomb periodograms were calculated for the daily mean sea level residuals, and exhibit no significant red noise, trend, ENSO effects or tide cycles other than the small-amplitude lunar and lunisolar synodic fortnightly cycles discussed earlier.

Small runs of missing data are ignored in the median-smoothing procedure, and have little or no effect on the estimated median sea level. Where the end-point of a large gap in the data falls during the six months covered by this analysis, daily mean sea level values for the adjacent 15 days are not smoothed. Each of these cases in the four sea level records was individually examined and none appear to influence the analysis of the most extreme storm surges. There will likely be some effect on the relative heights estimated for some less extreme storm surges. Given the relatively large numbers of storm surges in the four records compared to the instances of missing data, we should expect any effect on the analysis to be trivial.

+++ daily mean sea level - series mean
- - - normalized sea level pressure   Days
—— change in sea level from pressure
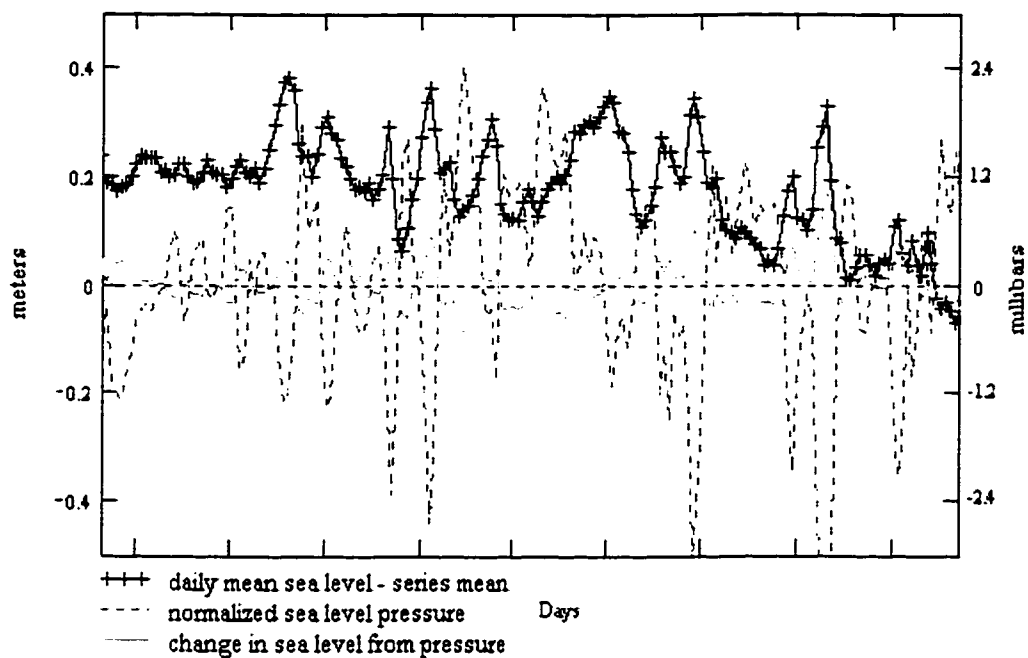
Figure 1.5: San Diego winter 1982-83, sea level pressure

+++ raw daily mean sea level   Days
—— daily mean sea level - 1.2 cm/mb SLP
- — smoothed 31-day median sea level
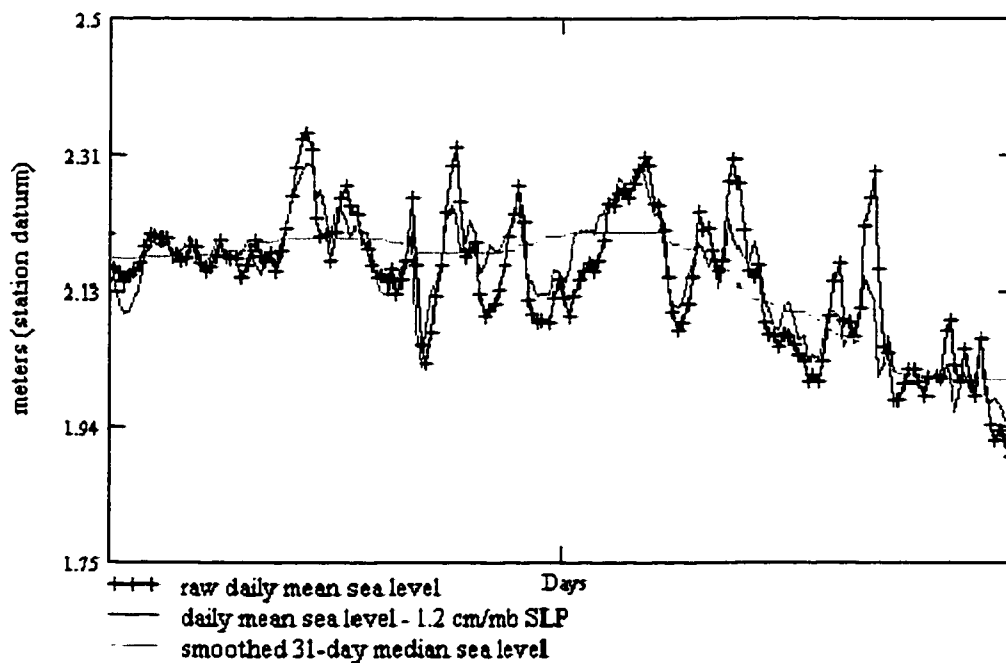
Figure 1.6: San Diego Winter 1982-83, smoothed 31-day median sea level

E.   Definition of Storms

In this analysis, storm surge in the sea level record is used as a proxy for storm frequency. According to Flick (1991),

> Storm surge is usually defined as that portion of the local, instantaneous sea level elevation that exceeds the predicted tide, and which is attributable to the effects of low barometric pressure and high wind associated with storms.

A correlation between storm surge height and storm strength is a crucial assumption for this analysis. It is motivated by the following relation from Pugh (1987):

$$E_w = \rho g \int_0^\lambda \zeta^2 dx$$

where E_w, the total energy per wavelength λ for an ocean wave, is proportional to the sum of squares of ζ, the height the wave displaces the sea level from its undisturbed mean over the wavelength.[16] The squared daily mean sea-level anomaly over a point is thus proportional to the total day's energy density from the storm surge at that point.

For a storm which passes near a tide station, daily mean surge height should be a good indicator of that storm's strength relative to others measured in the same manner. A small storm surge detected at the San Diego station could be from a weak storm passing through the immediate area, or could be due to a much more intense storm at some remove along the coast. So, surges detected at one location cannot

---

[16]Where ρ is the density of the sea water and g is gravitational acceleration. Pugh (1987) notes in his appendix that these "need not be simple harmonic" waves. Reinhard Flick (1998, personal correspondence) points out that the energy equations for a harmonic wave can also be applied to storm surge, since "non-sinusoids can be Fourier-decomposed into a unique bunch of sinusoids, and the energy for each calculated separately."

accurately characterize all storm activity for that location. However, for the largest storm surges, it is less likely that they indicate the passage of an even more intense storm further up the coast. Thus we would expect this measure to be most indicative of the relative strength of the largest storms.

Additionally, these data give us four widely spaced records of ocean storm activity that provide a measure of the sea level activity over a significant portion of the North Eastern Pacific. Considered together, they may be able to give us an indication of changes in the temporal and spatial pattern of large storms' distribution over the region as a whole. If sea level anomalies' correlation to total storm surge energy varies significantly across tide stations (due to vagaries of placement, depth, etc.), then comparisons between tide stations may be misleading. Our definition of a storm (detailed below) is a proportional measure which varies by site, so this may not be a concern. Assuming that these tide stations are sensitive to regional, and not just local, phenomena, the daily mean sea level anomaly at the four sites should be a good indicator of regional storm characteristics, since its square is proportional to energy derived from storms' low barometric pressure and high winds.

An examination of the tide record corresponding to known storm episodes indicates that large storms typically elevate the sea level for a day or more. Thus, we do not expect daily mean values to unduly dilute the strength of the signal we are trying to detect. In fact, the daily mean may help to ensure that our analysis is not contaminated by minor, local events. In Figure 1.7, storm episodes from the winter of 1982-83 in San Diego closely track peak elevations in the filtered daily mean sea level.

The filtered daily mean sea level allows for direct comparisons between storm episodes to be more easily taken because changes in sea level of a month or greater duration have been removed. Because sea level height data persist, extreme values tend to be clustered. We seek a definition of storms in terms of sea level heights that will identify peaks in the record and distinguish peaks which are unambiguously associated with storm episodes. For the purposes of this analysis, peaks are defined as five day maxima in the residual daily mean sea level series.

Peaks are identified as exceeding a threshold corresponding to a percentage of observed values in a filtered daily mean sea level series. For example, "1-percent" or "99^{th}-percentile" peaks are peaks above a threshold equal to the height of the n^{th} ranked peak, where n is 1% of the total observations + 1. If there were no missing observations, this would average 3.65 peaks per year for 1-percent peaks. For this analysis, we will restrict our attention to peaks which occur between October 1 and March 31. Figure **1.7** shows five 1-percent peaks in the filtered daily mean sea-level record for San Diego. Storm episodes from November 1, 1982, through March 31, 1983, denoted by the shaded columns, are identified by Flick and Cayan (1984) from increases in average daily wind speed and anomalously low barometric pressure.

The "El Niño" winter of 1982-83 was particularly noted for its intense winter storms in California, and includes the highest peak in the filtered sea level series for San Diego. Five of the filtered daily mean sea-level values during what is arguably the most extreme winter of the series are 1-percent peaks. These peaks correspond to 5 out

days, Nov. 1, 1982 - Mar. 31, 1983

——— Residual Daily Mean Sea Level
‐‐   Storm Episodes
●○   Peaks, 95-percent
●   Peaks, 99-percent
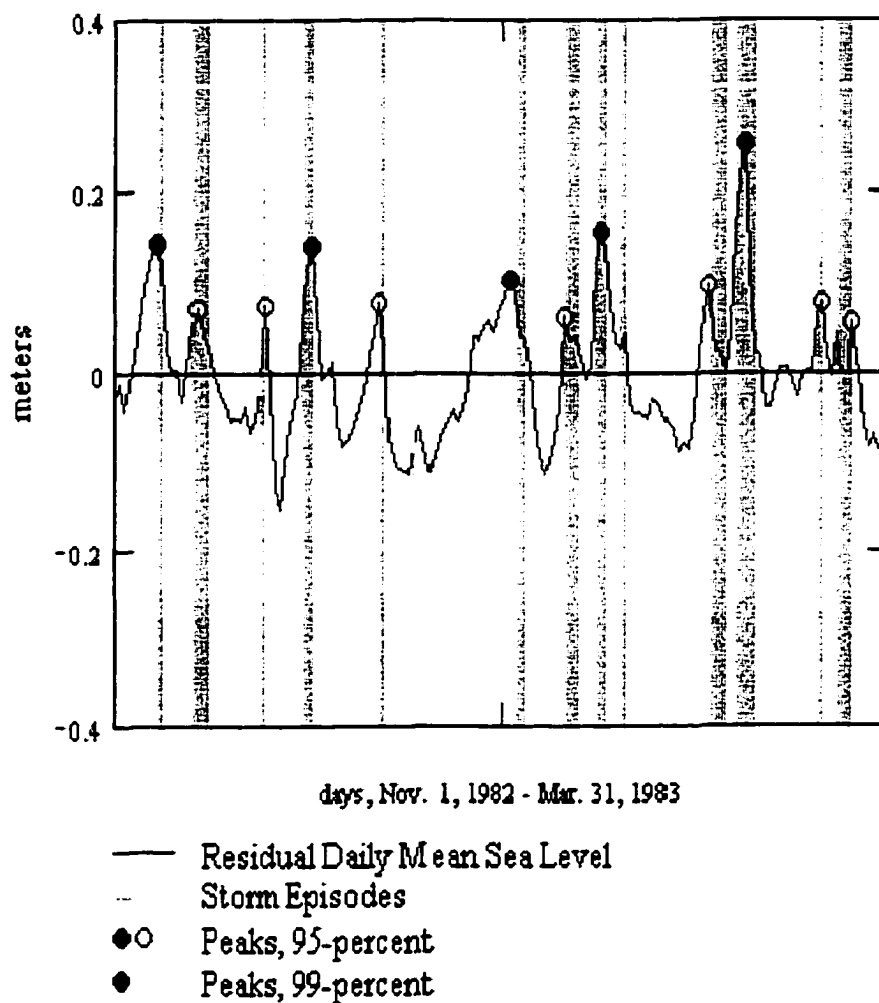
Figure 1.7: San Diego Winter 1982-83, storm episodes

of the 14 storms identified for the period. Figure 1.7 shows peaks in the 95[th] percentile corresponding to 12 out of the 14 storms. Once the threshold is moved lower than the 94[th] percentile of values we begin to capture peaks in these five months of the series which do not correspond to separate storm episodes as identified with local barometric and wind speed data. This analysis will be limited to winter peaks above thresholds corresponding to the 95[th] percentile of values or higher.[17] In general, the lower the threshold, the less confidence we can have that the results of this analysis describe characteristics of storm episodes only.

F.   Analysis

Linear regression models were fit to the annual winter storm frequencies for each station, regressing winter storm counts on a time trend. The results in Tables 1.4-1.7 clearly show a positive trend in annual winter storm frequencies in Honolulu, San Diego and, to a lesser extent, San Francisco. The increase in storm surge frequency is greatest, in percentage terms, for the largest storm surges. The modeled trend in storm frequency for Honolulu 1-percent peaks, 0.037 storms per year, represents a 280 percent increase in the frequency of these storm surges over the 93 years of the series, from 1.22 to 4.67 per year. Similarly for San Diego, the trend of 0.016 storms per year is equivalent to a 150 percent increase in the modeled frequency of 1-percent peak storm surges from a base of 0.96 to 2.38 per year. San Francisco's trend of 0.011 implies a

---

[17]Note that the average number of winter peaks over a threshold corresponding to any given percentile may vary widely from series to series, since in some locations large storm surges are more common in summer months than in other locations.

40 percent increase in 1-percent peak storm surges over the 97 years of the series, though at a 15% confidence level, from 2.45 to 3.46 per year. The longest series, Seattle, with 99 winters, does not demonstrate a trend at anything approaching a reasonable confidence level.

The histogram in Figure 1.8, however, is typical of the residuals from these regressions. These residuals are in general heteroskedastic[18] and non-normally distributed. In the presence of heteroskedasticity, our usual estimates of the standard error for the ordinary least squares regression—and thus our significance test statistics—may be biased. The test statistics reported in Tables 1.4-1.7 use White's (1980,1984) heteroskedasticity-consistent covariance matrix estimator for the least squares estimator, which allows for unbiased significance tests despite our inability to precisely specify the nature of the heteroskedasticity.

The storm surge frequency data approximate the poisson distribution, but are overdispersed. That is, the variance is typically higher than the mean, whereas the poisson model implies that they should be equal. The negative binomial (NB) model is a common correction to the poisson model allowing for over-dispersion. This specification assumes the conditional variance is $\text{var}(Y_t|X_t) = \lambda_t(1 + \alpha\lambda_t)$, where Y_t and X_t are the regressand and regressors. λ_t is the mean of Y_t conditional on the regressors X_t observed at time t, defined as $\lambda_t = \exp(X_t\beta)$. α is a nuisance parameter which is

---

[18]By heteroskedasticity we mean that the variance of the error term in the regression is not constant. In this case, the implication is that the variance of annual storm counts is changing over time, perhaps correlated with phase changes in climate like ENSO events.

Table 1.4:   Honolulu OLS regressions. (P-value is probability of type I error for two-sided test of H0: coefficient = 0. t-statistics are calculated using White's heteroskedasticity-consistent covariance estimator.)

	OLS reg. on year: Coefficient t-stat (p-value)	Percent increase 1906-98	OLS reg. on year & -|SOI|*SOI*100 t-stat (p-value)
1-percent peaks	*year:* 0.037 4.06 (0.00)	282%	*year:* 0.041 4.11 (0.00) *SOI:* -0.002 -0.92 (0.36)
2-percent "	*year:* 0.053 4.20 (0.00)	153%	*year:* 0.056 4.36 (0.00) *SOI:* -0.002 -0.77 (0.44)
3-percent "	*year:* 0.069 5.27 (0.00)	149%	*year:* 0.072 5.15 (0.00) *SOI:* -0.001 -0.75 (0.46)
4-percent "	*year:* 0.086 6.11 (0.00)	132%	*year:* 0.087 5.9 (0.00) *SOI:* -0.000 -0.13 (0.90)
5-percent "	*year:* 0.092 5.99 (0.00)	120%	*year:* 0.092 5.61 (0.00) *SOI:* -0.001 -0.68 (0.50)

Table 1.5:   San Diego OLS regressions. (P-value is probability of type I error for two-sided test of H0: coefficient = 0. t-statistics are calculated using White's heteroskedasticity-consistent covariance estimator.)

	OLS reg. on year: Coefficient	Percent increase 1907-98	OLS reg. on year & -|SOI|*SOI*100
	t-stat (p-value)		t-stat (p-value)
1-percent peaks	*year:* 0.016 2.38 (0.02)	148%	*year:* 0.010 1.91 (0.06) *SOI:* 0.004 4.18 (0.00)
2-percent "	*year:* 0.016 1.87 (0.07)	74%	*year:* 0.010 1.31 (0.19) *SOI:* 0.005 3.45 (0.00)
3-percent "	*year:* 0.023 1.98 (0.05)	71%	*year:* 0.014 1.51 (0.13) *SOI:* 0.008 4.09 (0.00)
4-percent "	*year:* 0.026 2.15 (0.03)	59%	*year:* 0.018 1.94 (0.06) *SOI:* 0.009 3.81 (0.00)
5-percent "	*year:* 0.034 2.62 (0.01)	64%	*year:* 0.028 2.61 (0.01) *SOI:* 0.008 3.17 (0.00)

Table 1.6:   San Francisco OLS regressions. (P-value is probability of type I error for two-sided test of H0: coefficient = 0. t-statistics are calculated using White's heteroskedasticity-consistent covariance estimator.)

	OLS reg. on year: Coefficient	Percent increase 1902-98	OLS reg. on year & -|SOI|*SOI*100
	t-stat (p-value)		t-stat (p-value)
1-percent peaks	*year:* 0.011 1.45 (0.15)	41%	*year:* 0.007 1.03 (0.30) *SOI:* 0.005 5.29 (0.00)
2-percent ··	*year:* 0.019 1.89 (0.06)	46%	*year:* 0.014 1.48 (0.14) *SOI:* 0.006 3.61 (0.00)
3-percent ··	*year:* 0.018 1.85 (0.07)	32%	*year:* 0.016 1.60 (0.11) *SOI:* 0.005 3.45 (0.00)
4-percent ··	*year:* 0.017 1.59 (0.11)	24%	*year:* 0.017 1.59 (0.12) *SOI:* 0.005 3.2 (0.00)
5-percent ··	*year:* 0.018 1.48 (0.14)	21%	*year:* 0.016 1.36 (0.18) *SOI:* 0.005 3.72 (0.00)

Table 1.7: Seattle OLS regressions. (P-value is probability of type I error for two-sided test of H0: coefficient = 0. t-statistics are calculated using White's heteroskedasticity-consistent covariance estimator.)

	OLS reg. on year: Coefficient	Percent increase 1900-98	OLS reg. on year & -|SOI|*SOI*100
	t-stat (p-value)		t-stat (p-value)
1-percent peaks	year: 0.004 0.55 (0.58)	12%	year: 0.001 0.10 (0.92) SOI: 0.002 0.90 (0.37)
2-percent "	year: 0.007 0.66 (0.51)	11%	year: -0.001 -0.09 (0.93) SOI: 0.005 2.18 (0.03)
3-percent "	year: -0.001 -0.07 (0.94)	-0.70%	year: -0.006 -0.61 (0.54) SOI: 0.004 1.76 (0.08)
4-percent "	year: -0.001 -0.06 (0.95)	-0.70%	year: -0.007 -0.54 (0.59) SOI: 0.004 1.98 (0.05)
5-percent "	year: -0.003 0.26 (0.79)	-2.30%	year: -0.004 -0.36 (0.72) SOI: 0.003 1.48 (0.14)
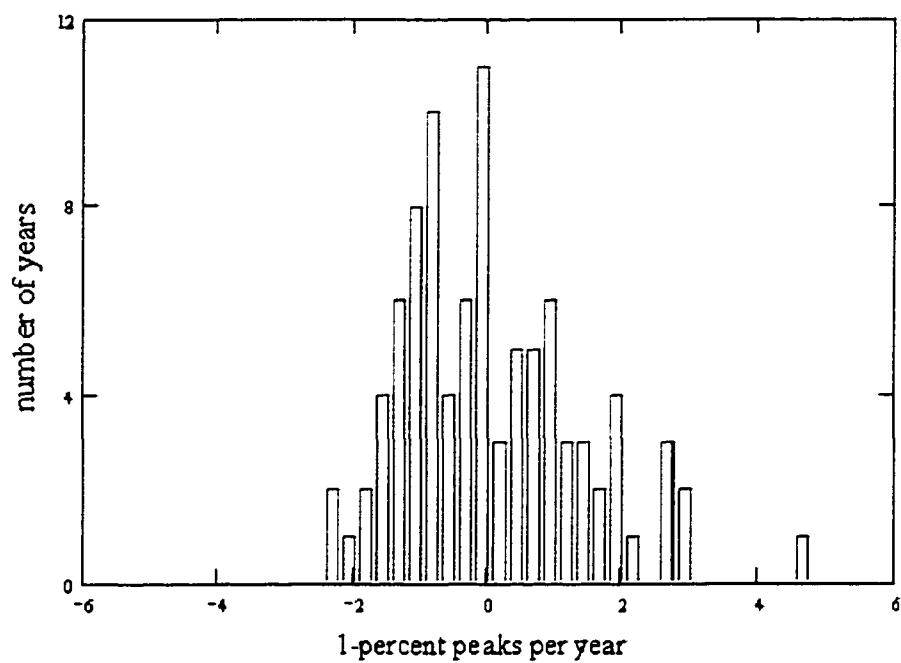
Figure 1.8: Histogram, San Diego OLS residuals

estimated along with the coefficient β.[19] Negative binomial models of a trend in winter storm frequencies were also fit for peaks in each data series, with results very close to those of the OLS regressions. The results of the NB regressions for 1-percent peaks are given in Table 1.8.

OLS and NB regressions were also run using a term for the NOAA Climate Analysis Center's seasonalized Southern Oscillation Index for December through February (SOI) of each winter as a regressor.[20] The SOI is the difference between the standardized Tahiti SLP and the standardized Darwin SLP measurements. Strong El Niño conditions are indicated by large negative values of the index. For example, the seasonalized December-February SOI for 1982-83 was -3.13, compared to a mean value of -0.9 (1909-1998). Likewise, positive values of the index are associated with the La Niña phase of the ENSO. The exact regressor used is $-|10 \times SOI| \times (10 \times SOI)$. The form of the regressor (the squared index with original sign reversed) was chosen merely to demonstrate the non-linear nature of the relationship between the SOI and storm surge frequency (the retained sign is reversed so that a positive coefficient implies an increase in storms during an El Niño winter). The results of these regressions are in column 4 of Tables 1.4-1.7. These regressions indicate a positive correlation between storm frequency and the occurrence of an El Niño in San Diego and San Francisco. In Seattle, the 1983 and 1998 El Niños were anomalous in that they coincided with above-average

---

[19]See Grogger and Carson 1991 and Hausman et al 1984.

[20]We use the seasonalized SOI as defined by Ropelewski and Jones, 1987.

Table 1.8:   Negative Binomial regressions of annual 1-percent peak frequencies on year and seasonalized SOI. Reported trend and ENSO effects are average change in conditional expectation of storm surge frequency given the explanatory variables as a function of the estimated NB coefficient, or $dE(y|X)/dX = \beta \times E(y|X)$. P-value is probability of type I error for two-sided test of H0: coefficient = 0.

	Negative Binomial regression on year: average trend t-statistic (& assoc. p-value)	Negative Binomial regression on year and - |SOI|*SOI*100: average trend and ENSO effect t-statistic (& assoc. p-value)	
Honolulu	0.041 4.05 (0.00)	*year:* 0.044 4.06 (0.00) *SOI:* -0.002 -1.17 (0.25)	
San Diego	0.016 2.82 (0.01)	*year:* 0.010 1.98 (0.05) *SOI:* 0.002 7.53 (0.00)	
San Francisco	0.010 1.47 (0.14)	*year:* 0.007 1.01 (0.32) *SOI:* 0.003 5.57 (0.00)	
Seattle	0.004 0.55 (0.58)	*year:* 0.001 0.08 (0.93) *SOI:* 0.002 1.75 (0.08)	

storm frequencies. When these winters are dropped from the sample in Seattle, the correlation between storm frequency and El Niños becomes negative and much less significant, while the secular trend is still not significant. In the case of Honolulu the 1983 El Niño is similarly anomalous. The already negative correlation becomes significant when 1983 is dropped from the sample (the t-statistic is -2.18), while the positive secular trend strengthens. Once again, the results of the NB regression were very similar to those of the OLS regressions.

We would like to test for an overall increase in storm surge frequency across all four tide stations as well. The Mantel-Haenszel summary chi-square statistic is a test for progressive association between multiple series of count data, drawn from separate populations, and scores based on a treatment variable. In this case, it is quite similar to regressing the sum of the four storm surge frequency series on time. Instead of testing for the regression coefficient's deviation from a hypothesized value, we test the deviation of

$$\sum_{y} A_{y,i} y$$

from its expectation conditional on the number of observations each period, where $A_{y,i}$ is the number of storms in year y for station i.. The test based on the Mantel-Haenszel statistic has the advantage of not imposing a linear model. It also makes an adjustment for differences in the number of missing observations across series, subtracting from each period's count of storm surges its expectation conditional on the number of observations in that period. The chi square statistic, with one degree of freedom, is

$$\chi^2 = \frac{\left[\sum_i \sum_y A_{y,i} y - \sum_i E\left(\sum_y A_{y,i} y\right)\right]^2}{\sum_i V\left(\sum_y A_{y,i} y\right)}.$$

The expectation of the statistic of interest is

$$E\left(\sum_y A_{y,i} y\right) = \frac{N_i}{T_i} \sum_y M_{y,i} y$$

where N_i is the total number of storms over series i, T_i is the total number of observations, and $M_{y,i}$ is the number of observations in year y for station i. Likewise, the conditional variance is

$$V\left(\sum_y A_{y,i} y\right) = \frac{N_i(T_i - N_i)}{T_i^2(T_i - 1)}\left[T_i \sum_y M_{y,i} y^2 - \left(\sum_y M_{y,i} y\right)^2\right].$$

The chi-square statistics from these tests in Table 1.9 below indicate a very significant increase in storm frequency over time. Note that the normal approximation for the statistic $A_{y,i} y$ is very good for large samples. In calculating the chi-square statistics, we pooled annual counts into four roughly equal periods to insure the large-sample properties would apply even to frequencies of the most extreme storm surges.

Using the properties of the hypergeometric distribution of extremes, we can also construct a nonparametric test of changes in the intensity of the most extreme winter storms. So long as we are interested in estimating the probability of exceeding a prior maximum without estimating by how much that threshold may be exceeded, we can use the hypergeometric distribution without having to model the distributions of the sea-

Table 1.9:   Mantel-Haenszel chi-square test of progressive association, Honolulu, San Diego, San Francisco, and Seattle combined longest common series (1907-98).

	summary Mantel Haenszel chi-square statistic	1 - Prob(chi-square)
1-percent peaks	34.3	0
2-percent "	43.4	0
3-percent "	43.4	0
4-percent "	47.4	0
5-percent "	47.0	0

Table 1.10:   Augmented Dickey-Fuller test for stationarity. 1% critical value: -2.6, 5% critical value: -1.9, 10% critical value: -1.7

	San Diego	San Francisco	Seattle	Honolulu
ADF test statistic	-48.749	-48.651	-49.8	-49.461

Table 1.11:   Hypergeometric test for increase in amplitude of largest storm surges

Station	Post-1950 exceedences of prior maximum	Probability of observed or greater exceedences
San Diego	1	0.51
San Francisco	2	0.24
Seattle	0	0.52
Honolulu	1	0.50

level series themselves.[21] We need only establish that the series are stationary and independent. As we saw with San Diego in section 1.3 above, the sea level series are auto-regressive. This was true for the data from all four tide stations. However, by restricting our analysis to peaks defined as maxima of 5-day series, the dependence is effectively removed. The series are also stationary. Augmented Dickey-Fuller tests for stationarity of the filtered series using 14 lags are in Table 1.10.[22]

Our analysis of extreme sea level peaks shows no significant trend in maximum storm surge intensity across the four series. Table 1.11 shows the number of post-1950 filtered mean sea level peaks that exceed the pre-1951 maximum of each series. In no case does the probability of the observed or greater number of exceedences fall below 24 percent. While the results of the trend analysis indicate that the frequency of the largest storms may have increased over the last century, the intensity of the most extreme storms does not appear to have been significantly affected.

G. Discussion

The foregoing analysis indicates that the frequency of the largest winter storms has increased on the order of 280 percent since 1906 for Honolulu and 150 percent since 1907 for San Diego. Further north, San Francisco saw an increase of 40 percent, though at a low level of significance, while Seattle remained essentially unchanged.

---

[21]Abarbanel et al (1992).

[22]The Augmented Dickey-Fuller test for stationarity of a series essentially consists of regressing the first difference of the series on its lagged differences. If the coefficients are significantly different from zero, the series is stationary.

Storm surges in general show a similar pattern, though the increase at the two southern stations is less dramatic. Recall that the storm surges analyzed here ride on top of the secular trend and inter-annual fluctuations in sea level, as well as on top of the tide. Thus, increases in large storm surge frequency combined with a rising mean sea level may result in a much greater frequency of extreme sea levels.

Regression models including a squared term for the SOI indicate that some, perhaps all, of the increase in annual winter storm frequency in San Diego and San Francisco may be due to an apparent intensification of the ENSO in the last quarter century. The models fit here are too crude to determine how much of the total increase in storm frequency can be attributed to a change in the frequency of moderate to strong El Niños. However, they are sufficient to indicate that much of the increase in large storm surge frequency coincides with El Niño events, which also further elevate mean sea levels along the California coast. The combination can only accentuate the increased likelihood of extremely high sea level events on the California coast.

The pattern of some of the changes in storm surge frequency observed—increases in San Diego and San Francisco and small, possibly negative changes in Seattle—may also be consistent with an increase in frequency of El Niño events. San Diego and San Francisco are crossed by an amplified storm track in many El Niño winters, while Seattle and Honolulu are not. The fact that the Honolulu record also sees a strong increase in storm frequency, while the ENSO cycle's effects on Honolulu appear to be small and possibly of the opposite sign as in the Californian stations, might indicate that the frequency of large storms in the Eastern North Pacific is increasing independent of

ENSO effects.

Local increases in winter storm frequency combined with rising sea levels and the short term sea level increases associated with El Niños can have serious consequences for coastal infrastructure designed for a less active winter storm season and a lower average sea level. Consider the historical sea level height frequencies presented in Figure 1.9 for San Francisco. These frequencies are computed for hourly sea level heights for San Francisco for four roughly equal periods spanning 1901-98. The right tails of these distributions correspond to winter storm events.

Despite the small trend in annual winter storm frequency for San Francisco, the pre-1926 maximum of 143 cm above mean sea level (MSL) is exceeded for a total 36 hours, corresponding to 15 separate winter storm events, after 1975. All but one of these storms occurred when sea levels were elevated due to an El Niño event. De-trending these samples and projecting them onto an annual sea level rise of 3.5 mm/year for 2001-25 indicates an annual frequency of 2.5 to 5 hours a year in which the pre-1926 maximum sea-level is attained or exceeded.[23] This is an increase of between 62 and 124 times for the return time for the pre-1926 maximum. If future climate change results in changes in storm frequency and the inter-annual variability in sea level and storm frequency, the impact on coastal infrastructure may be far different from what we

---

[23]These frequencies represent what might be expected if we were to experience the same sequence of weather and climate events that occurred in each of the four samples considered here—1901-25, 1926-50, 1951-75, and 1976-98—at mean sea levels similar to what the IPCC anticipates for the 25 years 2001-25. The rate of sea level rise projected here is a continuation of the rate of rise experienced in San Francisco since 1976.

would anticipate based on long term sea level rise alone.
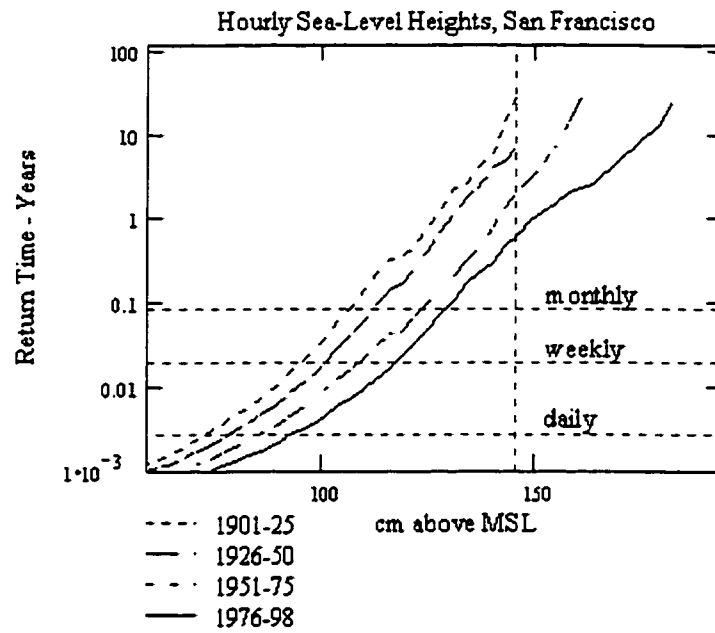
Figure 1.9: Hourly sea-level heights, San Francisco

H. Bibliography

Abarbanel, H., S. Koonin, H. Levine, G. MacDonald, and O. Rothaus (1992) *Statistics of Extreme Events with Application to Climate*, JASON, Mitre Corporation, McLean, VA.

Bates G. T. and G. A. Meehl (1986) "The effect of CO2 concentration on the frequency of blocking in a general circulation model coupled to a simple mixed layer ocean model," *Monthly Weather Review*,114:687-701.

Climate Analysis Center (CAC) (1995) *CPC: Special Climate Summary - 95/1*, National Oceanographic and Atmospheric Administration, Climate Analysis Center: February 1995.

Carnell, R.E., C.A. Seiñor, and J.F.B. Mitchell (1996) "Simulated changes in Northern Hemisphere winter storminess due to increasing CO2," *Climate Dynamics*,12(7), 467-476.

Clarke, A. J. and S. Van Gorder (1993) "On ENSO coastal currents and sea levels," *Journal of Physical Oceanography*, vol. 24, pp. 661-680.

Flick, R. E. and D. R. Cayan (1984) "Extreme sea levels on the coast of California," *19th Coastal Engineering Conference Proceedings*, ASCE/September 3-7, 1984, Houston, Texas.

Flick, R. E. (1991) "Joint Occurrence of High Tide and Storm Surge in California," *World Marina '91*, 1st International Conference, ASCE/Southern California Advisory Group, Long Beach, CA/September 4-8, 1991.

Flick, R. E. (1998) "Comparison of California tides, storm surges, and mean sea level during the El Niño winters of 1982-83 and 1997-98," *Shore and Beach*, 66(3) pp. 7-11.

Greene, W. H. (1993) *Econometric Analysis*, Macmillan Publishing Company, New York, NY.

Griggs, Gary B. and Kristin M. Brown (1998) "Erosion and Shoreline Damage along

the Central California Coast: A Comparison Between the 1997-98 and 1982-83 ENSO Winters," *Shore and Beach*, (July 1998) 66.3:18-23.

Grogger, J. T. and R. T. Carson (1991) "Models for truncated counts," Journal of *Applied Econometrics*, vol. 6, pp. 225-238.

Hall, N.M.J., B.J. Hooskins, P.J. Valdes, and C.A. Seiñor (1994) "Storm tracks in a high resolution GCM with doubled CO2," *Quarterly Journal of the Royal Met. Society*, vol. 120, pp.1209-1230.

Hausman, J., B. H. Hull and Z. Griliches (1984) "Econometric models for count data with application to the patents-R & D relationship," *Econometrica*, 52(4), pp. 909-938.

IPCC (1990) *Climate Change: The IPCC Scientific Assessment*, University Press, Cambridge, Great Britain.

IPCC (1992) *Climate Change 1992: The Supplemental Report to the IPCC Scientific Assessment*, University Press, Cambridge, Great Britain.

IPCC (1995) *Climate Change 1995: IPCC Second Assessment Report*, University Press, Cambridge, Great Britain.

Konig, W., R. Sausen and F. Sielmann (1993) "Objective identification of cyclones in GCM simulations," *Journal of Climate*, vol. 6, pp. 2217-2231.

Lambert, S.J. (1995) "The effect of enhanced greenhouse warming on winter cyclone frequencies and strengths," *Journal of Climate*, vol. 8, pp. 1447-1452.

Press, William H., Saul A. Teukolsky, William T. Vetterling and Brian Flannery (1992) *Numerical Recipes in Fortran 77: The Art of Scientific Computing*, Reprinted 1996. Cambridge University Press, New York, N.Y.

Pugh, David T. (1987) *Tides, Surges and Mean Sea Level*. John Whiley & Sons, New York, NY.

Quinn, W. H. and V. T. Neal (1989) *The Historical Record of El Niño Events*, manuscript.

Quinn, W. H. and V. T. Neal (1987) "El Niño occurrences over the past four and a half centuries," *Journal of Geophysics Research*, vol. 92, pp. 14449-14461.

Ropelewski, C. F. and P. D. Jones (1987) "An extension of the Tahiti-Darwin Southern

Oscillation Index," *Monthly Weather Review*, vol. 115, p. 2161-65.

Schureman, Paul (1994) *Manual of Harmonic Analysis and Prediction of Tides*, reprint of the 1940 edition, United States Department of Commerce, Washington, DC.

Stephenson, D. B., and I. M. Held (1993) "GCM response of northern winter stationary waves and storm tracks to increasing amounts of carbon dioxide," *Journal of Climate*, vol. 6, 1859-1870.

Timmermann, A., J. Oberhuber, A. Bacher, M. Esch, M. Latif, and E. Roeckner (1999) "Increased El Niño frequency in a climate model forced by future greenhouse warming," *Nature*, vol. 398, pp.694-697.

White, H. (1980) "A heteroskedasticity-consistent covariance estimator and a direct test for heteroskedasticity," *Econometrica*, vol. 48, pp. 817-838.

White, H. (1984) *Asymptotic Theory for Econometricians*, Academic Press, San Diego, CA.

Zetler, B. D. and R. E. Flick (1985) "Predicted extreme high tides for mixed-tide regimes," *Journal of Physical Oceanography*, vol. 15, pp. 357-359.

Zhang, Y. and W. C. Wang (1997) "Model-simulated northern winter cyclone and anticyclone activity under a greenhouse warming scenario," *Journal of Climate*, vol. 10, pp. 1616-1634.

Chapter II

Information Aggregation in Catastrophe Reinsurance Markets

A. Abstract

We experimentally examine a reinsurance market in which participants have differing information regarding the probability distribution over losses. The key question is whether the market equilibrium reflects traders maximizing value with respect to their different priors, or whether the equilibrium is one based on a common belief incorporating all participants' information. When assuming subjects are expected value maximizers, we reject both full information aggregation and no information aggregation equilibria. We discover, as in past individual choice insurance experiments, that buyers under-assess the probabilities of large loss states, or alternatively, subjects assign larger utility values to losses than to comparable gains. After accounting for these decision theoretic concerns, the non-aggregation of information hypothesis explains the data better than full information aggregation.

B. Introduction

It is commonly thought that insurance markets facilitate the efficient sharing of risk, but whether they facilitate the efficient sharing of information is an open question.

41

A defining feature of an insurance market is its underlying uncertainty. It is reasonable to assume that market participants possess differing information regarding the objective probabilities governing states of nature. When these agents participate in a market there are two natural conjectures regarding the nature of the arising competitive equilibrium. First, agents maximize their objectives (holding their priors constant) and the resulting market prices and allocations reflect efficiency with respect to these initial beliefs. Second, market prices and allocations arise that reflect a competitive outcome of agents maximizing their objectives conditional upon a common belief formed by the pooling of the agents' differing information. In the first conjecture, the invisible hand only optimally coordinates activity treating the initial beliefs as exogenous parameters, while in the second conjecture the invisible hand does substantially more. The process of market feedback aggregates disparate information and generates individually optimal outcomes with respect to the most informed sets of beliefs possible. Such a feature is highly desirable within an insurance market.

The study of whether markets efficiently aggregate information is well suited for an experimental approach. A laboratory experiment allows for the control of preferences, endowments, and information structures that are essential in identifying when a market achieves a non-information aggregation (NA) equilibrium or a full information aggregation (FA) equilibrium. Several past experimental studies have addressed this question in the context of basic asset markets with mixed results. Plott

and Sunder (1988) find aggregation can occur when market participants have a complete set of Arrow-Debreu securities to trade, or when there are homogeneous preferences. In Forsythe and Lundholm (1990) information aggregation occurs only when traders have experience with market institutions and common knowledge of each others' dividends. Plott and Wit and Yang (1997) find some success for information aggregation in parimutuel markets for situations where Bayes' Law is not needed.
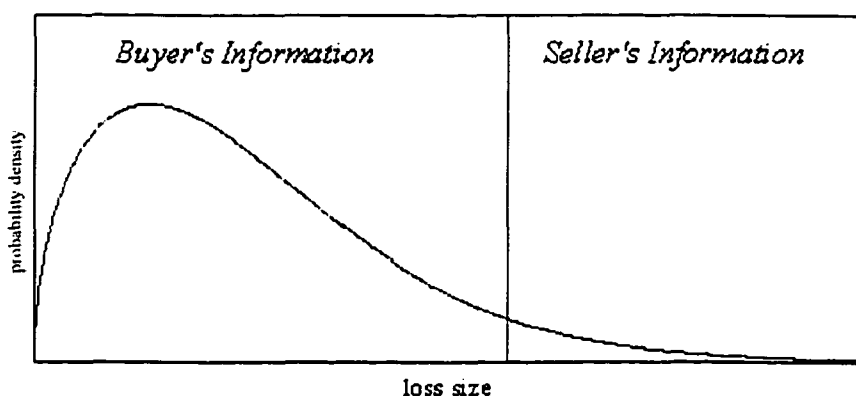
Figure 2.1: Reinsurance market risk and information structure.

Unfortunately, these experiments' designs and results do not lend sufficient insight into how effectively information aggregates in an insurance market because of the strikingly different information structure. In this study we consider a property reinsurance market. It is natural to suppose a risk and information structure like that in Figure 2.1. Purchasers of reinsurance have considerable experience with the high-frequency, low-value claims processes represented by the left side of the figure. Sellers of reinsurance, on the other hand, with a long history of business in multiple regions

and lines of reinsurance, have better information about the large less likely catastrophe risks represented by the right tail of the probability density in Figure 2.1.[1]

The presence of low-probability, large-loss states also is not captured in previous experimental market studies, but is an integral part of an insurance market. However, there is an extensive body of survey and experimental work addressing how *individuals* make insurance decisions when faced with low-probability, high-value risks. Slovic et al (1977) and Kunreuther et al (1978) find evidence of either persistent probability biases or convex utility over losses in insurance experiments. McClelland, Schulze and Coursey (1993) find, when agents purchase insurance from the experimenter in a Vickrey auction, evidence of a bimodal response to very low probability risks, with some participants disregarding very small risks and others highly sensitive to small risks. None of these experiments are conducted in a bilateral-market context (i.e., subjects only perform the task of buying insurance). Also these experiments do not consider the situation of differential information.

An empirical example motivates us to draw distinct elements from the two literatures: a recent innovation in the U.S market for catastrophe reinsurance. After three recent low probability large loss events, Hurricane Hugo ($4.2 billion in insured claims), Hurricane Andrew (claims over $16 billion), and the Northridge Earthquake

---

[1]The property insurance market here is assumed to have little in the way of moral hazard. We believe this would muddle the central issue of information aggregation. Moreover, we feel secure in assuming that the market participant's actions do not exert significant influence over the probabilities of catastrophic events such as hurricanes, earthquakes, and floods.

(claims over $12.5 billion), many insurers tried to withdraw from the catastrophe insurance market for earthquake risk in California and wind risk in Florida.[2] However, regulatory measures kept firms from fleeing these markets. At the same time, available reinsurance coverage grew increasingly scarce,[3] as the reinsurance market did not face the same regulations. These changes created an opportunity for new and innovative entrants to the reinsurance industry.[4] The Chicago Board of Trade (CBOT) was one of the first non-traditional entrants, inaugurating trading in Catastrophe Futures and Options in December 1992. CBOT officials were particularly enthusiastic about the potential success of catastrophe insurance futures. Numerous members of the academic community shared this enthusiasm. There were many anticipated benefits of catastrophe insurance futures and one of the strongest was the reduction of information asymmetries.[5] Despite the initial optimism, trading in the CBOT's catastrophe futures never amounted to much,[6] and they are no longer traded today. We hope our experiment sheds some light into this lack of success, and give insights into whether

---

[2]See Nutter 1994, and Marlett and Eastman 1997, Lecomte 1998, Roth 1998.

[3]O'Hare 1994 and Kunruether 1996.

[4]See Doherty 1997 for a good review of conditions in the insurance industry at the time.

[5]D'Arcy and France (1992), Niehaus and Mann (1992), Harrington, Mann and Niehaus (1995), Doherty (1996 and 1997) discuss benefits of trading in catastrophe futures and insurance derivatives in general. Cox and Schwebach (1992), Cummins and Geman (1995) and Doherty 1997 also address the role of catastrophe futures markets in resolving information asymmetries.

[6]Harrington and Niehaus (1997).

any market of this structure leads to information aggregation.

The results of our experiments do not offer much hope in this regard. First, when we assume individuals are expected value maximizers, the market price and quantity data do not support either an NA equilibrium or an FA equilibrium. However, there is strong evidence that prices and quantities rely more heavily upon the realization of the buyer's prior information regarding high-probability, low-loss events than the seller's prior information regarding low-probability, high-loss events. This leads us to investigate the impact that subjective probability biases and risk aversions, found in individual choice insurance experiments, could be having in our markets. We find that buyers tend to underestimate the probability of disasters while sellers on average assess these probabilities correctly. This finding is also consistent with an agent model where the correct probabilities are used by both buyers and sellers but subjects' preferences are those given in Prospect theory (Kahneman and Tversky 1979) in which losses loom larger than gains. Once controlling for these preferences, we find that an NA equilibrium typically explains the data more robustly than does an FA equilibrium.

In the next section we review the experimental paradigm for information aggregation in markets, and introduce our modifications in more detail. Subsequently we lay out the environment for our experimental design. In section 4 we provide a theoretical basis for our induced supply and demand schedules. In section 5 we present hypotheses for information aggregation and non-aggregation testable in equilibrium price and quantity. The results of our experimental markets are analyzed in section 6.

C.   Design of Information Aggregation Experiments

The standard paradigm used in experimental asset markets with diverse information is introduced by Plott and Sunder (1982, 1988). In a typical treatment of these innovative experimental designs, there are three equally probable states of the world (x, y, and z) and a security which pays a single dividend in each state. Traders are each endowed with two units of the security and working capital. These securities are traded in an oral double auction. There are two types of agents with heterogeneous dividend values. The dividend values in Table 2.1 are an example of the heterogenous payoff structures used in these experiments.

Table 2.1: Plott and Sunder Dividend Values

| | Dividend Value | | |
	State X	State Y	State Z
Type 1	70	160	300
Type 2	230	90	60

Agents in these experiments have diverse prior information. After the state of the world is determined, but prior to securities trading, equal numbers of randomly selected agents are informed that certain states of the world are not possible outcomes. For example, if the state of the world is y, then half of the traders are informed that the state is "not x" and half that it is "not z."
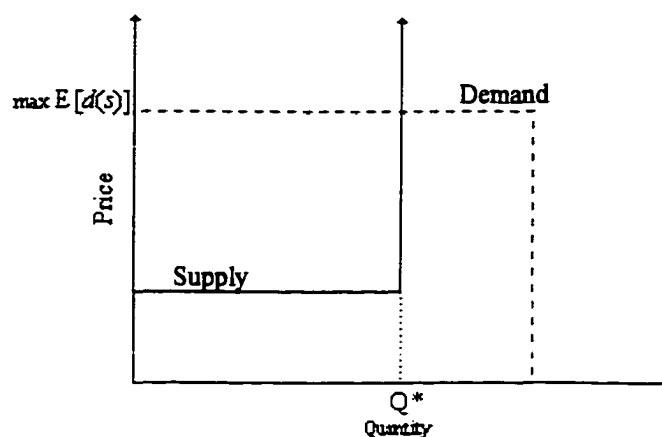
Figure 2.2: Supply and demand for a fixed set of beliefs

The implication of a competitive outcome is that the individuals with the highest expected dividend value will purchase the other agent type's holdings of the security at a price equal to their expected dividend. It is helpful to view the competitive outcome in the context of "box-design" supply and demand functions like those in Figure 2.2. The perfectly elastic demand curve's location is determined by the value of the highest expected dividend. Likewise, the perfectly elastic supply curve's location is determined by the value of the lower expected dividend, and the elbow occurs where total quantity equals two times the total number of agents in the market of the type with the lower expected dividend value. Information aggregation in these experiments is defined as all agents knowing the true state. In the absence of information aggregation, each agent's expected dividend value is the average of the two states that they know have not been eliminated. Thus prices at which transactions take place are affected by information aggregation. In an equilibrium, the buyers are indifferent between holding

a unit of the security and holding currency equal to the price of the security. The testable implications for quantity are weak as a result.

In a catastrophe futures market, the structure of the uncertainty, dividend values, and information have a strong ordinal property which cannot be forced to fit into the Plott and Sunder formulation. We can give a strong ordinal ranking to states according to the monetary amount of the state-dependent dividend (ie., as determined by values of the catastrophe futures index in the real market). Furthermore, in a given state, market participants' values of the dividend will be highly correlated. A natural assumption is that market participants' valuations of the state-dependent dividends will follow the same ordinal structure as the state space. To summarize, participants' will all place a greater utility on dividends in periods with a high futures index value than in periods with a low value. Information diversity is also related to the ordinality of the state space. Specifically, we assume that those who have the best information regarding the right tail of the distribution of losses are more likely to be sellers in this market (eg. reinsurers). Likewise, those with higher quality information regarding high-probability, low-value losses are more frequently buyers in this market. We should note that Plott and Sunder (1988) conducted three experiments with homogenous (and thus ordinal) dividends. However, the information and inherent risk properties of these treatments do not correspond to those of a reinsurance market.

In developing a design suitable for the economic environment of interest, we use a random supply and demand framework to induce preferences which correspond to a

quasi-linear utility representation (see Smith 1982 and Gjerstad and Shachat 1996). This allows us to generate testable predictions in price and quantity which distinguish between models of traders' behavior with and without information aggregation. This is unlike Plott and Sunder's formulation in that instead of a corner solution equilibrium, we generate interior solutions in price and quantity. We hope that we will observe quicker convergence to an interior equilibrium than to a corner equilibrium.[7]

In inducing our preferences, we use an induced supply and demand framing, which is a time-tested tool of experimental markets and auctions. In the next section we give a description of the environment. In this description our market will not be framed as a reinsurance market. Subsequently we will demonstrate that our environment corresponds to that of a catastrophe futures market.

D.   The Experimental Design: Institutions, Assets, and Information

In these experiments we have 12 participants randomly partitioned into a group of six Buyers and a group of six Sellers. Participants are Buyers and Sellers of a single asset in an oral double auction. Trading is conducted in successive seven minute periods. In each period of the auction, Buyers may purchase from zero to four units of the security (henceforth "units"). For each unit a Buyer purchases, there is unique random value. Likewise, each Seller may choose to sell from zero to four units of the security. For each unit sold, the seller incurs a random cost. At the conclusion of each

---

[7]Our optimism comes from the results of Smith 1964, Smith 1981, Holt, Langan and Villamil 1986, and Smith and Williams 1989, all of which demonstrate slow convergence in markets without uncertainty and with box designs.

period, buyers are also subject to a random transfer.

Realizations of random values, transfers, and costs are determined by one of four possible states of nature. There are two distinct types of states, Normal (N) and Disaster (D). Within each of these types one state corresponds to "low" values for buyers and "low" costs for sellers, and the other corresponds to "high" values and costs. In Normal states the random transfer is positive and buyers receive money. In Disaster states it is negative, and buyers lose this amount. For simplicity we describe the state space as (N_L, N_H, D_L, D_H). The initial prior over the state space is given by the vector (.45, .45, .05, 05).
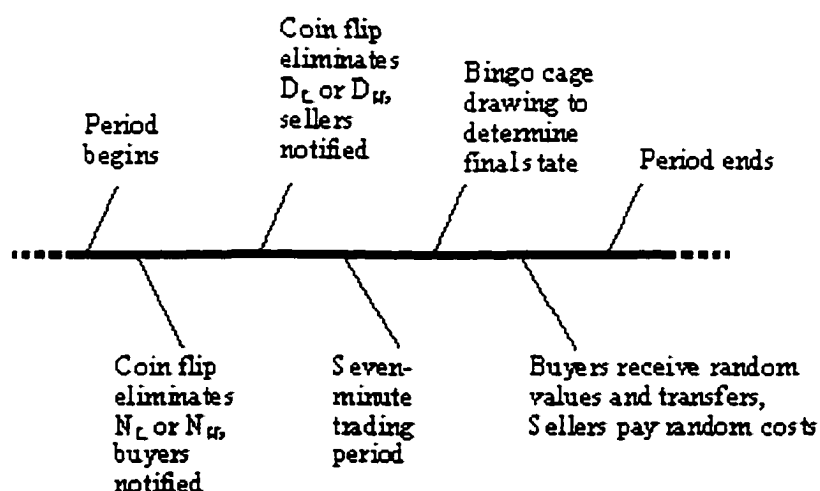
Figure 2.3 Time line for a trading period.

Consider the time line in Figure 2.3. Before the start of each trading period, the experimenter flips a coin. If the result is heads, the N_L state is eliminated. If the result is tails, the N_H state is eliminated. Buyers are privately informed of the value of the

remaining Normal state (N_L or N_H) with the use of a code sheet. Likewise, another coin flip is used to eliminate one of the Disaster states. Sellers are privately informed of the cost of the remaining Disaster state (D_L or D_H). Next a seven-minute trading period begins. Buyers may offer bids or accept asks, and sellers may make asks or accept bids in an oral double auction format. A valid bid or ask must improve upon any standing bid or ask. Once a bid or ask is accepted, bidding starts over; buyers are then free to open bidding at any non-negative price, and sellers are free to make an initial ask at any price between 0$ and $20. Bids, asks, and trades are displayed on an overhead projector as they are made. After the seven-minute trading period has expired, the final state of nature is resolved by drawing one ball from the bingo cage in view of the participants. If the ball is numbered "1" through "9", the result is the remaining Normal state. If a "10" is drawn, the result is the remaining Disaster state. The ball is returned to the bingo cage prior to the next trading period. Buyers then receive the random values of the units they purchased and the random transfers, and sellers pay the random costs of the units they sold.

Figure 2.4 below is a typical Buyer's Decision Sheet. In row number 1 Buyer 1 carries over cumulative earnings from the previous period ($0.00 since this is the first period). On the left side of the Buyer's Decision sheet are four columns labeled X1, X2, Y1, and Y2, corresponding to the state-space (N_L, N_H, D_L, D_H). In this period the statement "not White" would inform Buyers that X1 had been eliminated, and "not Blue" that X2 had been eliminated. There are no codes listed for the "Y" states (D_L,

Buyer Decision Sheet for Buyer # | 1 | Period: | 1 |

Name: _____

Probability of an X-state: 90%

Probability of a Y-state: 10%

State					Unit #	Row #		
X1	X2	Y1	Y2					
White	Blue	—	—			1	Cumulative Earnings	0.00
2.60	0.60	-5.40	-15.40			2	Random Transfer	
1.54	2.54	5.54	10.54		1	3	Unit Value	
						4	Purchasing Price	
						5	Unit Earnings (3 - 4)	
1.30	2.30	5.30	10.30		2	6	Unit Value	
						7	Purchasing Price	
						8	Unit Earnings (6 - 7)	
0.66	1.66	4.66	9.66		3	9	Unit Value	
						10	Purchasing Price	
						11	Unit Earnings (9 - 10)	
0.42	1.42	4.42	9.42		4	12	Unit Value	
						13	Purchasing Price	
						14	Unit Earnings (12 - 13)	

15 Total Unit Earnings

(5+8+11+14)

16 Period Net Earnings

(2 + 15)

17 Cumulative Earnings

(1 + 16)

Figure 2.4: Buyer's Decision Sheet

Seller Decision Sheet for Seller # [1]   Period: [1]

Name: _____

Probability of an X-state: 90%

Probability of an Y-state: 10%

	State				Unit #	Row #		
X1	X2	Y1	Y2			1	Cumulative Earnings	[0.00]
—	—	Mango	Grape					
						2	Selling Price	
0.46	1.46	4.46	9.46		1	3	Unit Cost	
						4	Unit Earnings (2 - 3)	
						5	Selling Price	
0.70	1.70	4.70	9.70		2	6	Unit Cost	
						7	Unit Earnings (5 - 6)	
						8	Selling Price	
1.34	2.34	5.34	10.34		3	9	Unit Cost	
						10	Unit Earnings (8 - 9)	
						11	Selling Price	
1.58	2.58	5.58	10.58		4	12	Unit Cost	
						13	Unit Earnings (11 - 12)	

14   Total Unit Earnings [ ]
      (4+7+10+13)

15   Cumulative Earnings [ ]

Figure 2.5: Seller's Decision Sheet

D_H), since the buyers are not privy to this information. The values in row number 2 are the random transfers which would apply in each state. Similarly, in rows three, six, nine and twelve the values for each of the four units that Buyer number 1 may purchase are listed for each of the four states. For each unit he purchases, Buyer 1 enters the purchase price in the appropriate space on the far right column. After trading is finished the final state is drawn. The Buyer copies the amount of the random transfer and the values of any units purchased from the column corresponding to the final state to the blank spaces at the far right of each column. For units not purchased, the buyer enters zeros. Then the earnings for each unit and for the period and the cumulative total are calculated as indicated. For the typical Seller's Decision Sheet in Figure 2.5, note that the primary difference is that the seller has no random transfer. The amounts listed in columns X1, X2, Y1, and Y2 are unit costs instead of values. For both buyers and sellers, cumulative earnings are carried over from period to period, and participants are paid the dollar amount of any positive cumulative earnings at the end of the experiment in cash.

The Buyer's random transfer in row 1 of Figure 2.4 is the difference between a fixed premium income of \$4.60 and a random loss of (\$2, \$4, \$10, \$20) over the state-space. This is the same for every Buyer. Similarly, the random value of each unit is equal to the sum of a fixed amount unique to each buyer and unit and a random amount of (\$1, \$2, \$5, \$10) over the state-space. For example, values for unit 1 in row number 3 of Figure 2.4 are the sum of a fixed value of 54¢ and (\$1, \$2, \$5, \$10). Thus,

while this market is not framed to the participants as a reinsurance market, the purchase of two units of the security fully insures Buyers against losses from the primary insurance portfolio represented by the random transfer.

Market participants are inexperienced prior to their arrival for the experiment. They are trained in the procedure for resolving uncertainty and receiving private information by several repetitions of the procedure without trading, and by participating in one to three practice periods that include trading in the security.

Buyers and Sellers begin the experiment with zero cash endowments. They are permitted to run negative cash balances without being expelled from the experiment, but receive no compensation other than a non-salient show-up fee of five dollars if their cumulative earnings are negative at the end of the experiment. The number of periods over nine is randomly determined, and participants are not informed ahead of time which period will be the final period.

E.   Induced Supply and Demand

In our formulation of a catastrophe futures market there are insurers and reinsurers. We present the environment for a single trading period. Each trading period in the experiment is an independent replication of this environment. Conceptually, economic activity within a period occurs in two stages, before and after the resolution of insurable risks. There are two commodities in stage one (x_1 and y_1) and a single commodity (x_2) in stage two. In the context of our catastrophe futures market, insurers and reinsurers form a *net* stage-one cash holding (x_1) through the purchase or sale of

reinsurance instruments. These trades give each insurer and reinsurer a reinsurance position. Firms are motivated to enter into reinsurance positions by both their risk management strategies and regulatory requirements. Note that while risk management strategies may be determined by idiosyncratic risk preferences, restrictions imposed by regulators are independent of these preferences. These independent restrictions lead us to treat reinsurance positions as unique stage-one commodities (y_1), entering the utility function separately from the probabilistic dividends.

In stage two, the commodity (x_2) is the *net* change in cash holdings associated with the payment and receipt of catastrophe future dividends and the net of claims and premiums for primary insurance. We resolve the treatment of stage-one and stage-two cash as distinct commodities by specifying them as perfect substitutes in the utility functions, i.e. the joint value of these cash holdings is their sum.

Next we specify the states of nature and the dividend structure of the reinsurance instrument. In stage one there is a single possible state, and in stage two there are four possible states. These states are denoted s_i, where $i = 1,2,3,4$. Let w denote a probability distribution over the four states and w_i the probability of state s_i. A reinsurance contract is a security we denote z. The contract pays a dividend of a single unit of y_1 in stage one and a state dependent dividend of x_2 in stage two, $d(s_i)$.

In our specification, the market involves the exchange of good x_1 (the numeraire good) for units of the security z prior to stage-one consumption. In this market insurers are only permitted to purchase the security and we refer to them as buyers. Likewise

reinsurers can only sell security units and we refer to them as sellers. The trading institution we adopt in the market is a double oral auction. For institutional details, see Smith (1962) and the instructions to the participants in Appendices 2.1-2.3.

Since we are considering the net change in cash positions associated with activity in the market, we set buyers' and sellers' endowments of x_1 to zero. An agent who neither buys nor sells a future contract holds zero units of x_1 for consumption. The net cash holdings of x_2 for buyers are determined by their net state-dependent primary insurance receipts (premiums less claims) and reinsurance dividend values. We set each buyer's state-dependent endowment of x_2 equal to twice the amount of the state-dependent dividend of z, i.e. $2*d(s_i)$. In stage two, sellers do not hold any primary insurance, and consequently have a zero endowment of x_2. A seller who does not sell any futures contracts does not incur any stage two liabilities and consumes zero units of x_2.

We derive a buyer b's demand function by first stating their expected utility over the three commodities as

$$u_b(x_{1b}, y_{1b}, x_{2b}) = x_{1b} + x_{2b} + v_b(y_{1b}) - 2\sum_{i=1}^{4} \omega_i d(s_i).$$

In this expression $v_b(\cdot)$ is an increasing concave function. Since the holding of (y_{1b}, x_{2b}) is determined by the holding of the security z_b, we can rewrite the above expected utility function as

$$u_b(x_{1b}, z_b) = x_{1b} + z_b\sum_{i=1}^{4} w_i d(s_i) + v_b(z_b) - 2\sum_{i=1}^{4} w_i d(s_i).$$

Since the market does not permit a buyer to sell short, we can proceed by simply deriving his demand function for z. Also recall that if buyer b does not purchase any catastrophe futures then his expected utility is exactly negative two times the expected dividend. A buyer maximizes his expected utility function subject to the budget constraint

$$p_z z_b + x_{1b} \leq 0.$$

From this problem we see that buyer b's optimal choice of z_b^* satisfies

$$\sum_{i=1}^{4} w_i d(s_i) + v_b'(z_b^*) = p_z.$$

The implication of this equation is that buyer b will desire units of z up to the point that the expected dividend plus the individual specific marginal benefit of the stage one dividend is equal to the price of the security.

This structure is similar to that of induced demand as explicated in Smith (1982) and in Gjerstad and Shachat (1999). In these two formulations, the unit valuations in an induced demand schedule are exactly the private marginal benefits described by $v_i'(\cdot)$. It is worthwhile to point out the direct relationship between the consumer surplus earned from the purchase–the unit's valuation less its price–is exactly equal to the gain in utility. In typical induced demand environments, the utility of making no

purchases is normalized to zero, and then the total earnings of a buyer is exactly the total utility of the final period holdings (and because of the quasi-linearity of the utility function, it is also exactly the consumer surplus.)

In our experiment the buyer also earns a dollar amount equal to the total utility of the final holding. The only difference from the typical formulation is the utility of not trading is equal to the endowment of x_2, and thus a buyer's earnings is their consumer surplus minus the realization of their state dependent liabilities. To summarize, in the structure of this paper, the total unit valuation for a buyer of any unit is the sum of the *private* marginal benefit plus the expected value of the *common* dividend. The total period earnings is the sum of the unit valuations purchased less the sum of the payments made for the units and less the realized stage two liabilities.

Unlike previous market experiments with uncertainty, our structure permits us to use the time-tested experimental procedures of inducing preferences through the framing of induced supply and demand. In each period, every buyer is given positive private valuations for four units.[8] This implicitly assumes that for each buyer b, that for any n and $n' \geq 4$, $v_i(n) - v_i(n') = 0$. For any buyer we can describe his discrete demand function by the sum of the following four-element vector of private unit valuations and the expected dividend,

$$\left(v_b(1) - v_b(0), \; v_b(2) - v_b(1), \; v_b(3) - v_b(2), \; v_b(4) - v(3) \right).$$

---

[8]In addition, they were only permitted to purchase up to four units in any period.

A graphical interpretation of this setup is given below. The step function is an induced demand curve given an increasing vector of private valuations. The arrows above each of the steps indicate that this demand curve has a unknown vertical shift whose magnitude is the buyer's expected value of the stage two dividend z.
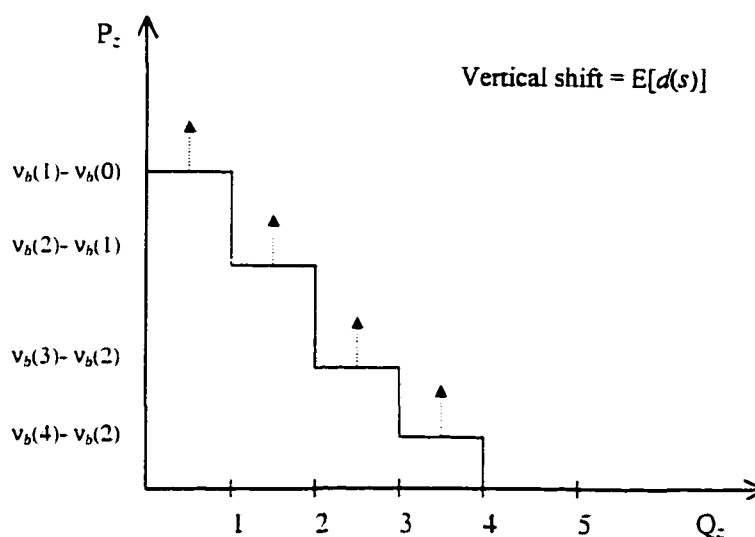
Figure 2.6: Induced demand with random shift

We now turn our attention to the sellers. We derive a seller s's demand function by first stating their expected utility over the three commodities as

$$u_s\left(x_{1s},\ y_{1s},\ x_{2s}\right) = x_{1s} + x_{2s} + c_s\left(y_{1s}\right).$$

In this expression $c_s(\cdot)$ is an increasing concave function. Recall that for sellers the value of y_1 is non-positive. This corresponds to an increasing administrative and regulatory cost of larger risk transfers. Given the relationship between a seller's sales

of security z_s and their allocation $\left(x_{1s}, y_{1s}, x_{2s}\right)$, we can rewrite their expected utility function solely as a function of z_s, recalling that z_s is non-positive for sellers:

$$u_s(z_s) = -p_z z_s + z_s \sum_{i=1}^{4} w_i d(s_i) + c_s(z_s).$$

Notice that if seller s does not sell any contracts then her expected utility is zero.

Since the market does not permit a seller to purchases units of z, we derive her supply function for z by maximizing her expected utility function subject to the budget constraint that z_s is non-positive. From first order condition one can see that seller s's optimal choice of z_s^* satisfies

$$\sum_{i=1}^{4} w_i d(s_i) + c_s'(z_s^*) = p_z.$$

The implication of this equation is that seller s will sell units of z up to the point that the expected common dividend plus the individual-specific marginal cost of taking on additional risk in their reinsurance position is greater than or equal to the price of the security.

We can use this formulation to construct an induced supply schedule for a seller. The seller's producer surplus will be exactly equal to the utility of the final commodity holdings. Each seller may sell up to a maximum capacity of four units of z. We impose this constraint by setting $c_s'(k) = -\infty, \ \forall k \leq -5$. In a discrete experiment we can calculate the private marginal cost array of selling units of z as

$$\left(c_s(0) - c_s(-1), c_s(-1) - c_s(-2), c_s(-2) - c_s(-3), c_s(-3) - c_s(-4)\right).$$

The marginal cost to the seller is the sum of this individual-specific marginal cost and the expected stage-two dividend. A graphical interpretation of seller s's supply curve is given below. Each step is the deterministic marginal cost given by the above array and the vertical arrow above each step represents the random vertical shift whose value is equal to the expected value of z's stage-two dividend. The market supply and demand curves are obtained by the horizontal sums of these curves. The missing element to fixing the values of the supply and demand curves is the specification of the stage two state probabilities and dividends.
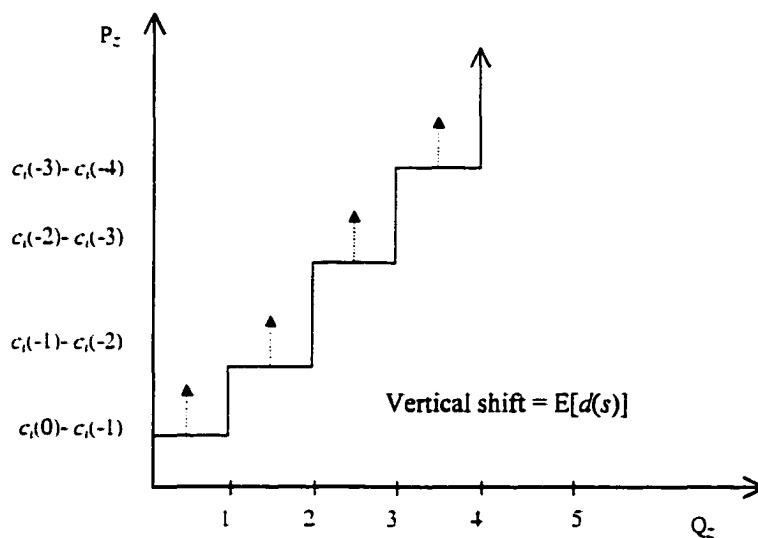

Figure 2.7: Induced supply with random shift

F.   Hypotheses and Testable Implications

The hypotheses of interest are full information aggregation (FA) versus non-information aggregation (NA). The basis of the FA hypothesis is the ability of a market to generate an information aggregation equilibrium, i.e. the market generates a

competitive outcome that reflects the pooling of all diverse information regarding the true state of nature. The NA hypothesis is generated by the conjecture that the market generates a competitive outcome reflecting the agents' prior beliefs regarding the true state of nature. In the previous section we showed that our experiment's induced supply and demand schedules coupled with a common random vertical shift constitute a reduced form of an insurance market with quasi-linear preferences. The impact diverse information has on competitive prices and quantities manifests itself through the buyers' and sellers' respective expectations of the size of the vertical shift in the demand and supply schedules. This vertical shift is the expected cash dividend. Testable implications of FA and the NA hypotheses are found by comparing the competitive equilibria for these two aggregation scenarios under the various information priors.

Let's recall the sequence by which information regarding state determination is produced and transmitted to participants in a market period. The initial prior over the state space (N_L, N_H, D_L, D_H) is given by the vector $(.45, .45, .05, 05)$. Then, before the market period, buyers learn which Normal state is eliminated and the sellers learn which Disaster state is eliminated. This results in four distinct *prior information regimes* which we denote LL, LH, HL, HH. The first letter in a pair refers to the remaining Normal state and the second letter refers to the remaining Disaster state. If information does not aggregate, then a pair of distributions characterizes a prior information regime: the buyers' and sellers' beliefs that incorporate their diverse

information. When information aggregates, a single distribution that incorporates all information characterizes a prior information regime. Table 2.2 gives these characterizations.

Table 2.2: Prior information regimes

Regime	Buyer	Seller	Aggregate
(LL)	(.9, 0, .05, .05)	(.45, .45, .1, 0)	(.9, 0, .1, 0)
(LH)	(.9, 0, .05, .05)	(.45, .45, 0, .1)	(.9, 0, 0, .1)
(HL)	(0, .9, .05, .05)	(.45, .45, .1, 0)	(0, .9, .1, 0)
(HH)	(0, .9, .05, .05)	(.45, .45, 0, .1)	(0, .9, 0, .1)

The impact of these two competing models is generated through differing expected dividend values. Under the FA conjecture, a competitive outcome reflects a common expected dividend value based on the pooling of buyers' and sellers' private information. The expected value is calculated as

$$E[d(s)] = 0.9(\text{remaining N-state's dividend})$$

$$+ 0.1(\text{remaining D-state's dividend}). \qquad (1)$$

On the other hand, if the NA conjecture holds true, the market outcome will reflect the following distinct expected dividends for the buyer and seller:

$$E[d(s)]_{buyer} = 0.9(\text{remaining N-state's dividend})$$

$$+ 0.1(\text{average D-state's dividend}) \qquad (2)$$

and

$$E[d(s)]_{seller} = 0.9(\text{average N-state's dividend})$$

$$+ 0.1(\text{remaining D-state's dividend}). \qquad (3)$$

Buyers' and Sellers' expectations of the dividend values determine the vertical location

of supply and demand curves. Hence, the implications of the comparative statics of FA

versus NA are obtained from the inspection of the competitive equilibrium for their

respective supply and demand curves. Table 2.3, and Figures 2.8-2.9 summarize the

equilibria for the two models in the four prior information regimes.

Table 2.3: Model predictions for equilibrium prices and quantities

			Disaster State	
			Low	**High**
			LL	**LH**
Normal State	**Low**	FA model:	12 units, $1.30-$1.50	12 units, $1.80-$2.00
		NA model:	12 units, $1.75	6 units, $1.81-$2.19
			HL	**HH**
	High	FA model:	12 units, $2.20-$2.40	12 units, $2.70-$2.90
		NA model:	18 units, $2.19-$2.21	12 units, $2.25-$2.65

Figure 2.8 shows the market supply and demand curves under the FA premise

for the four prior information regimes.[9] First notice that for all four prior information

regimes the equilibrium market quantity is twelve units. Private marginal value and cost

schedules are selected so that in these equilibria each buyer and seller engages in two

futures contracts. This allocation is equivalent to buyers fully reinsuring their endowed

---

[9]Market supply and demand curves were obtained by the "horizontal" summations of individual private marginal cost and private marginal benefits. To these aggregated schedules the appropriate expected dividend was added "vertically."
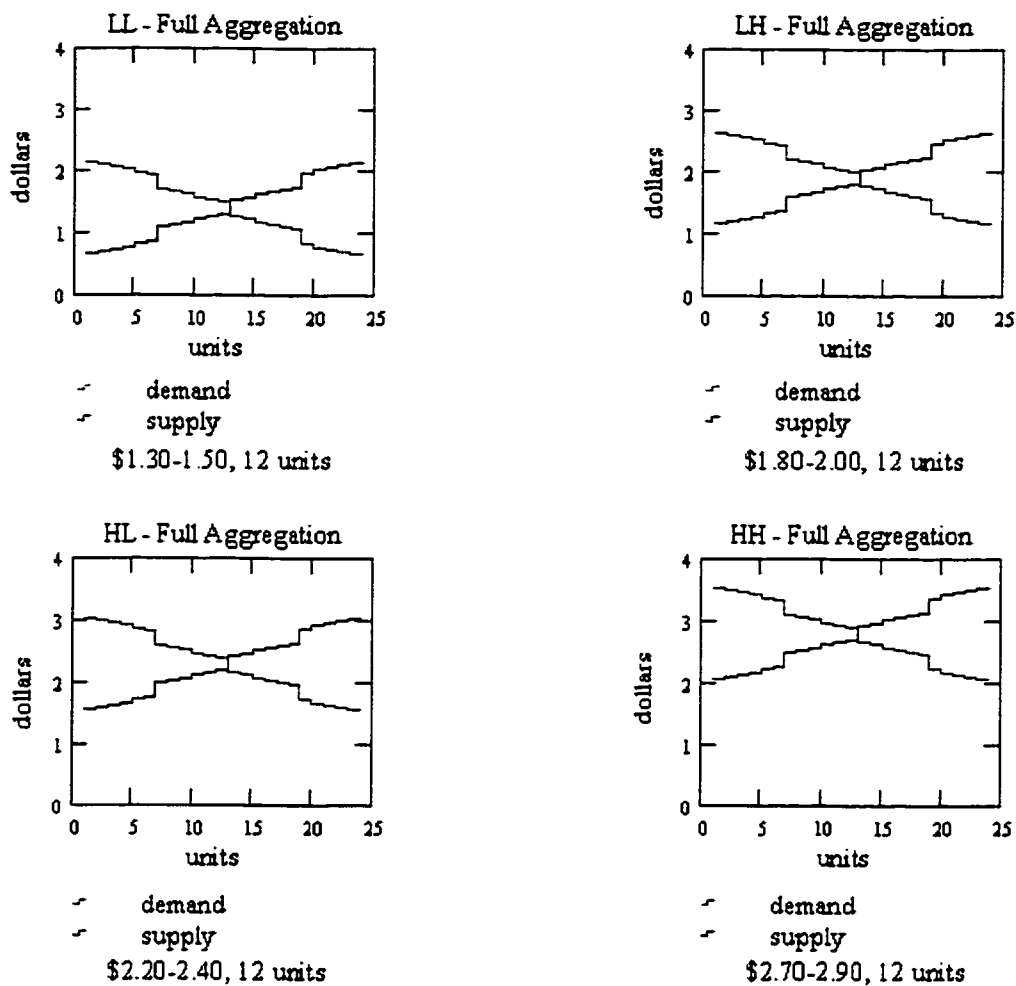
Figure 2.8: Full-aggregation equilibrium induced supply and demand.

portfolio risk. The quantity is independent of the prior information regime because the FA assumption dictates that the supply and demand curves have the same random vertical component. Turning our attention to price, the FA outcome generates distinct equilibrium price tunnels. The midpoints of these price tunnels represent actuarial fair premiums for reinsurance.

LL - No Aggregation

demand
supply
$1.75, 12 units

LH - No Aggregation

demand
supply
$1.81-2.19, 6 units

HL - No Aggregation

demand
supply
$2.19-2.21, 18 units

HH - No Aggregation

demand
supply
$2.25-2.65, 12 units
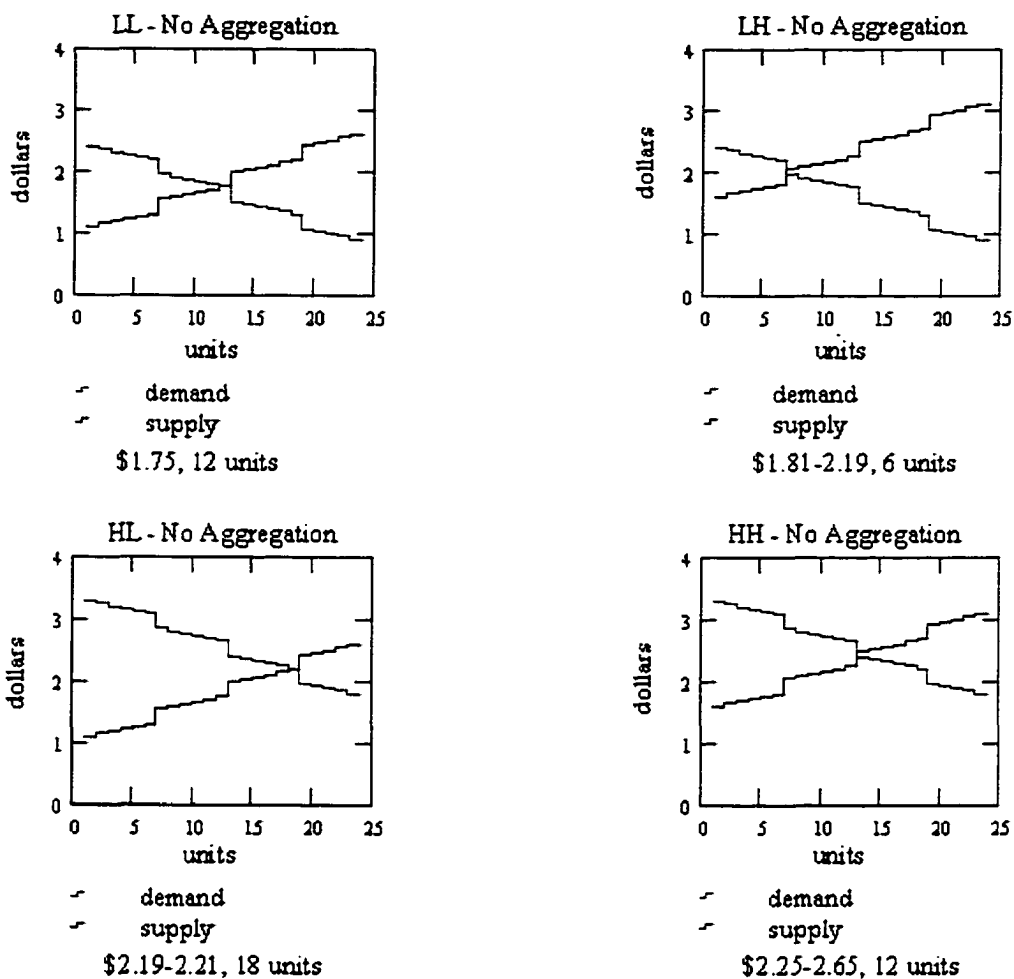
Figure 2.9: No-aggregation equilibrium induced supply and demand.

In any equilibrium the NA model will distinctly differ from the FA model in either the equilibrium price or quantity. In the LH and HL regimes, the NA and FA models only differ strongly in equilibrium quantities. The NA model predicts that in the LH regime only 6 units are traded, resulting in an under-provision of reinsurance; in the regime HL 18 units are traded, and there is an over-provision of reinsurance. One can also observe that under prior information regimes LL and HH the equilibrium

prices are distinct under the FA and NA hypotheses, but full reinsurance is achieved in both scenarios. However, in these two regimes the NA hypothesis does not generate actuarial fair reinsurance premiums: In HH, the midpoint of the price tunnel is below the actuarial fair rate and in LL the midpoint is above the actuarial fair rate.

G. Results and Analysis

We focus our analysis of the experimental data into two activities. First, we compare how well the data conforms to our interior predictions for price and quantity for the two competing models. Prices and quantities for units traded each period, with few exceptions, do not match the equilibrium predictions for either the full-aggregation or the no-aggregation model. Prices typically are lower than either models' predictions and market prices do not depend on the sellers' prior information. The volume of reinsurance contracts also does not reflect either model's predictions. We do observe that the impact of buyers' prior information is more influential on quantity than is the sellers' prior information.

Since prices are generally lower than either hypothesis predicts, and buyers' prior information has a greater than expected impact on both price and quantity, we consider alternative explanations. We turn to the experimental and survey research on disaster insurance for possible explanations. Given the subjective probability biases that underestimate the probability of disaster states found in these literatures, we explore the possibility that the buyers' and sellers' posses this bias in our experiment. From the experimental market data, we calculate implicit subjective probability beliefs of a

disaster for both buyers and sellers under the FA and NA hypotheses. The result of this exercise suggests there is a strong bias: the buyers' implicit beliefs are typically below the sellers' implicit beliefs (which are on average statistically indistinguishable from 10 percent.) Once we account for this bias, there is evidence that the NA assumption is more appropriate. We also point out that there is an alternative to our subjective probability bias conclusion: individuals use the objective probabilities but differ in the way they evaluate risky choice. In this scenario we conclude that the implications of prospect theory hold: sellers' losses loom larger than buyers' gains.

G.1 Data Preliminaries

We start by presenting the data from the five catastrophe futures markets in Figures 2.10-2.14 in Appendix 2.4. We show the transaction prices for each experiment in chronological order, separated by trading period. For each period, the shaded areas represent the quantities and the range of prices we would expect to observe if markets are in the FA model equilibrium. The NA model equilibrium prices and quantities are the clear areas; overlapping regions are cross-hatched. The x-axis gives the period, prior information regime, and the triple FA predicted quantity/ NA predicted quantity/ observed quantity.

For example, in the first period of Experiment 1 in Figure 2.10, the information set is LL. The no-aggregation model prediction of 12 units traded at $1.75 is represented as a horizontal line 12 units wide. The full-aggregation model prediction—12 units traded between $1.30 and $1.50—is represented by the shaded area.
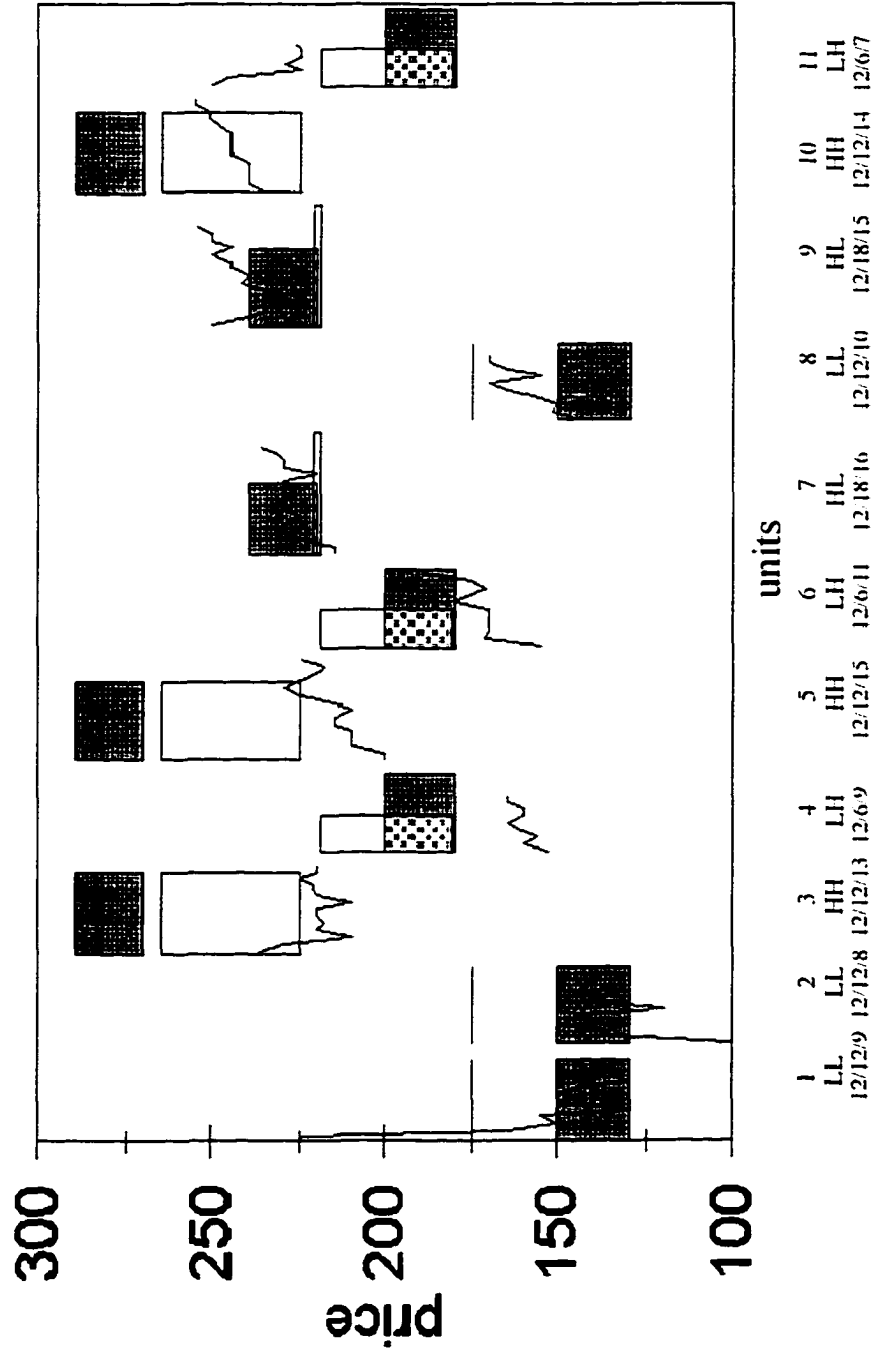
Figure 2.10: Market I
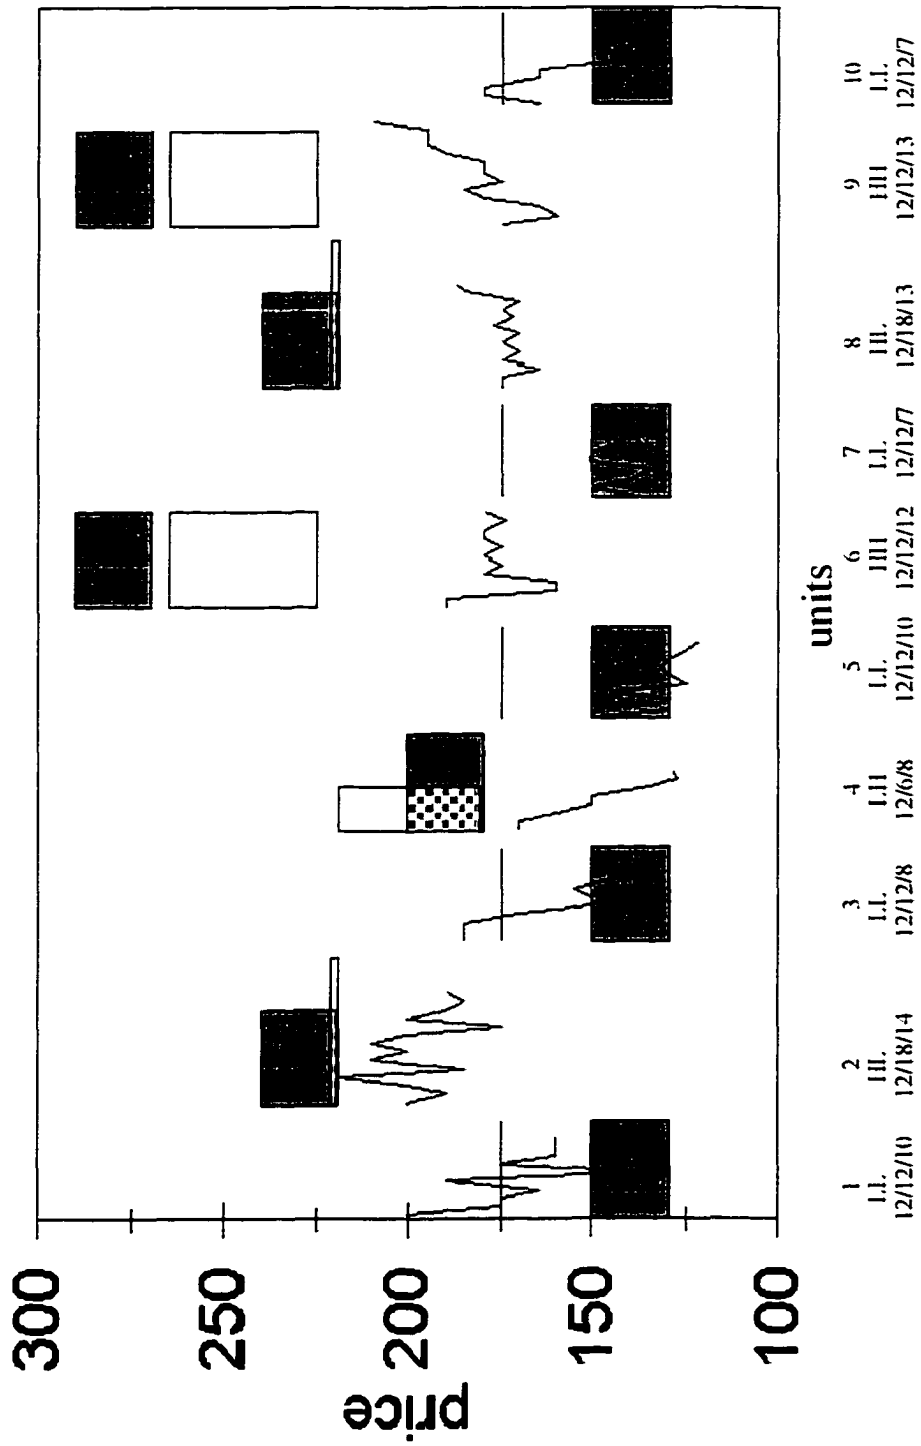
Figure 2.11: Market 2
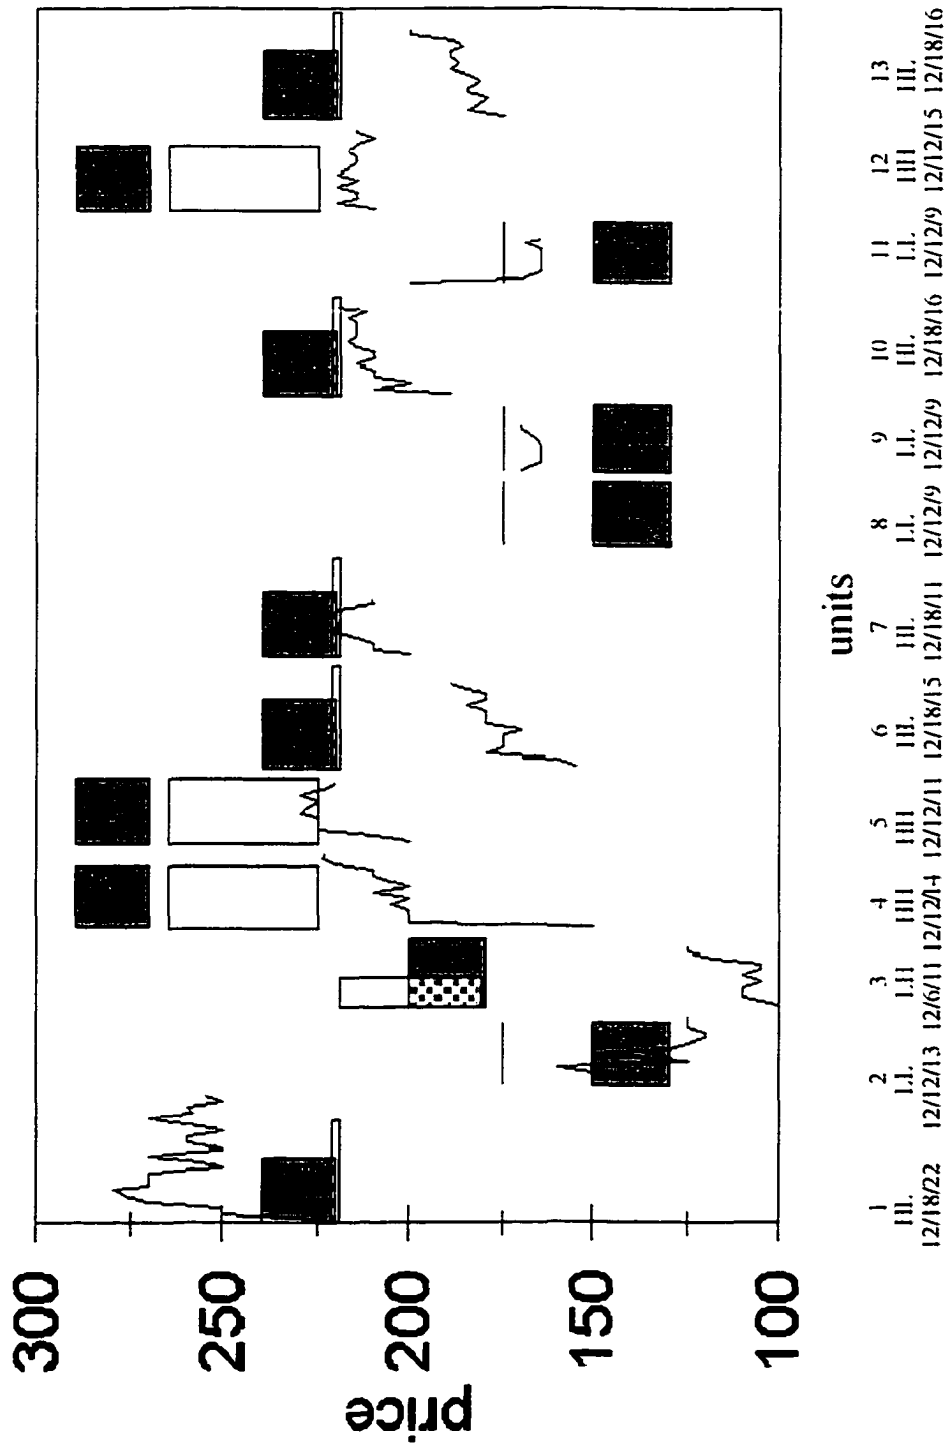
Figure 2.12: Market 3
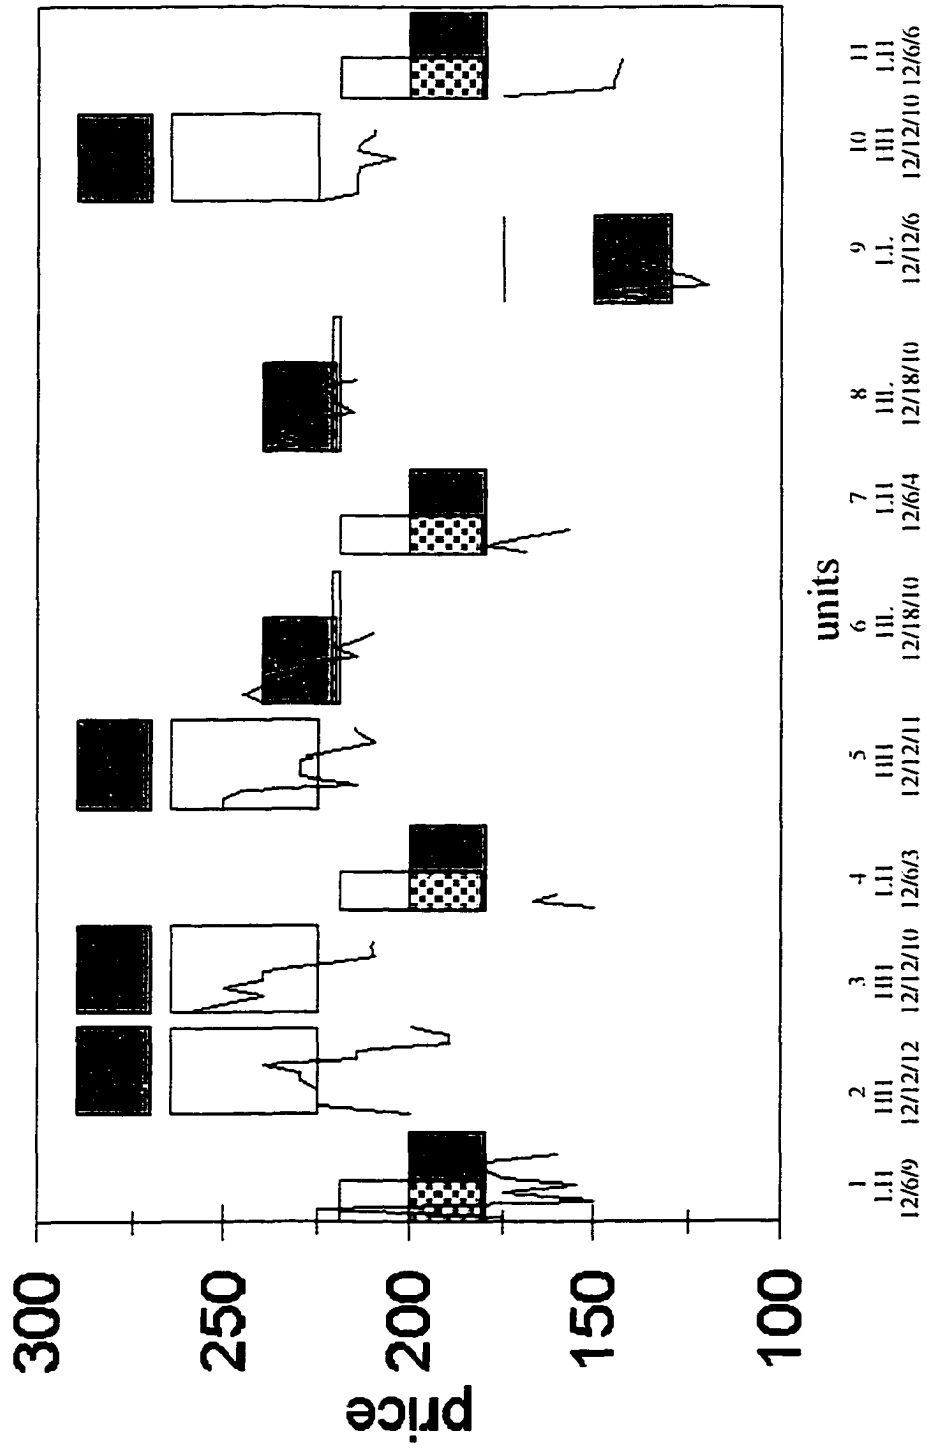
Figure 2.13: Market 4
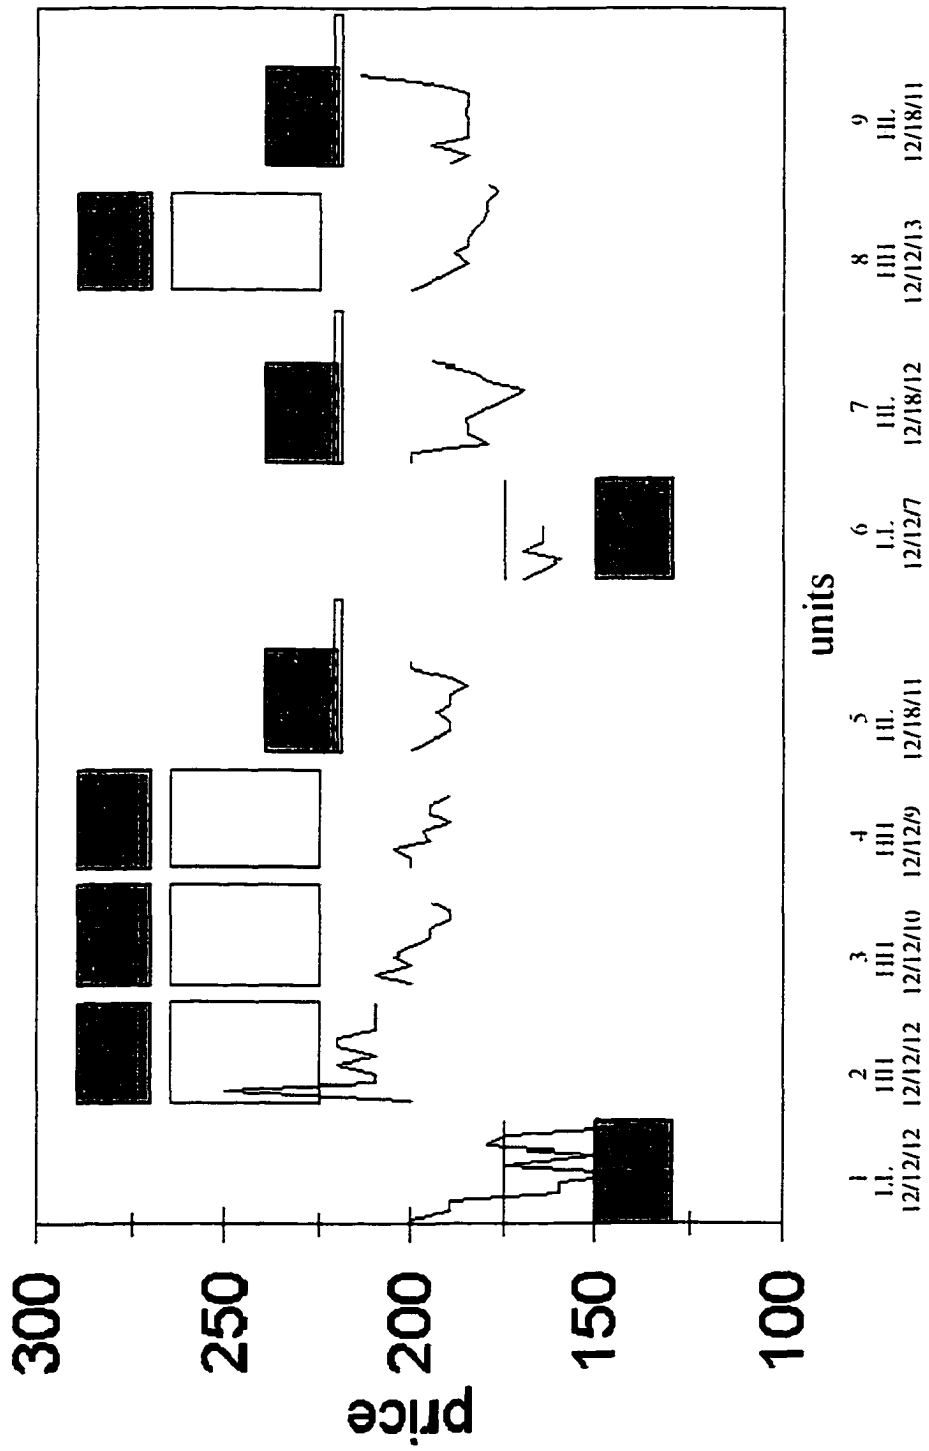
Figure 2.14: Market 5

The line representing actual trades shows the first unit traded at \$2.25. Subsequent prices fell rapidly to the full-aggregation price range, and the total quantity traded was 9 units.

G.2 Price and Quantity Data Analysis

A visual inspection of Figures 10-14 quickly reveals that the observed prices tend to lie outside the ranges predicted by either model. To assess the impact prior information has on prices we obtain the ordinary least squares estimate of the coefficients in the following dummy variable equation:

$$Price = a_1 LL + a_2 LH + \alpha_1 HL + \alpha_2 HH.$$

The results of this regression, along with NA and FA price predictions, are given in Table 2.4. First notice that mean price for each prior information regime falls below the predicted range except in the case of the FA prediction in the LL regime. The second striking result is that price seems to solely depend upon the buyer's prior information. Specifically, the mean prices in LL and LH are close and the mean prices in HL and HH are close. We conduct an F-test to confirm this observation. The F-statistic for the hypothesis that $a_1=a_2$ and $\alpha_3=\alpha_4$ is 2.63 with a p-value of .073.

These results regarding price are quite surprising given the results of similar treatments in Plott and Sunder (1988). In three of their experimental sessions, subjects are given homogeneous preferences over dividends, thus giving an ordinal ranking of states. Strong convergence to the FA predicted prices occurred by the end of each of

the three sessions.[10] The lack of price convergence in our experiment must result from one or some combination of the following: correlation of prior information with buyer and seller roles, pooled information does not reveal the true state of nature, the low probability of large loss states, and how individuals form assessments in the presence of this uncertainty.

Table 2.4: Dummy regression: Price = a_1LL + a_2LH + α_1HL + α_2HH

Variable	Coefficient	Standard Error	FA Prediction	NA Prediction
LL	153.28	2.189	130-150	175
LH	161.19	3.072	180-200	181-219
HL	210.35	1.828	220-240	219-221
HH	210.71	1.819	270-290	225-265

Before completely dismissing the applicability of either model, consider the effects a probability bias might have on the hypothesized prices. The price data imply that market participants may tend to under-weight the probability of a disaster state occurring. Note that under the FA model we would expect the difference in price between low and high normal-state information sets to average 90¢, and under the NA model, 45¢. As the probability of a disaster state goes to zero, these predictions approach $1.00 and 50¢, respectively. The difference we observe--about 53¢--is supportive of the NA hypothesis. We more vigorously pursue this idea below.

---

[10]For more details of the results from these three sessions see Plott and Sunder (1988) pages 1100-1102.

One of the attractive features of our induced supply approach is the ability to discriminate between models through the inspection of quantities. In the five futures markets, observed quantities tend to diverge from those predicted by either model. The lack of convergence in quantity is readily seen in the Figures 2.10-2.14. We now ask whether either model can explain the average market quantities. Recalling the quantity predictions of the two models summarized in Table 2.3, note that under the FA model we expect 12 units to be traded in each period. Also note that under the NA model the quantity prediction differs in two prior information regimes: in LH the quantity is six and in HL the quantity is eighteen. The FA and NA models both give testable implications in the following expression:

$$Q_t = \alpha + \nu \, Hx_t + \delta \, xH_t,$$

where Q_t is the market quantity in period t. Hx_t is dummy variable for the prior information regimes in which buyers are informed that the low Normal state is eliminated (i.e. regimes HL and HH), and xH_t is a dummy variable for the prior information regimes in which the seller has been informed that the low Disaster state is eliminated (i.e. LH and HH). Under the FA model, $\alpha = 12$ and $\nu = \delta = 0$ and under the NA model $\alpha = 12$ and $\nu = -\delta = 6$. The OLS estimates of these coefficients are presented in Table 2.5. The F-statistic for this regression (24.301) rejects the hypothesis that the mean quantity is independent of the prior information regime. This is a rejection of the FA coupled with symmetric subjective probability beliefs of a Disaster state. On the other hand, the estimated model coefficients do not follow the predictions

of the NA model either. The estimated value of α (9.0) is not the predicted 12 units, and a t-test indicates a 0.00 probability that $\alpha = 12$. While the estimated values of v and δ are significantly different from zero, and have the correct sign for the NA model, they are not equal to 6 and - 6, respectively. The probability that v, given an estimated value of 4.7, is equal to 6 is 0.059 and the probability that δ, given an estimated value of -1.5, is equal to - 6 is 0.000, again according to two-sided t-tests. The other notable result of this exercise is the magnitude of v is significantly greater than δ. This result is indicative of the more significant impact the buyers' information has than the sellers' information.

Table 2.5: Regression: $Q_t = \alpha + v\ Hx_t + \delta\ xH_t$

Variable	Coefficient	Standard Error	T-Statistic
Constant	8.99	0.563	0.000
Hx	4.69	0.681	0.000
xH	-1.50	0.679	0.032

In our analysis of prices we noted that observed biases were consistent with the buyers and sellers assigning a probability of a disaster state as less than ten percent. Is this consistent with the data on quantities? If buyers and sellers tend to under-weight the probability of a disaster, we would still expect under the FA model a quantity of 12 units traded in each period. Under the NA model, we would expect, as observed, a value for $|\delta|$ less than 6; as the probability of a disaster state goes to zero, δ goes to 0 as well. As the perceived probability of a disaster declines, however, the observed

value of v should increase under the NA model, converging to 7 as the probability of a disaster state goes to zero, contrary to our result. How then do we account for these results? Some possible explanations for our results are that the experimental subjects' perceived probability of a disaster state changes over time, that buyers' and sellers' beliefs may differ, or both.

G.3   Subjective Probability Biases

We assess whether subjective probability biases combined with either the FA or NA model can rationalize our market data. We start by assuming that the market prices and quantities we observe each period reflect a competitive equilibrium. This assumption relies upon the oral double auction's substantial history of robustly generating competitive outcomes in induced supply and demand experiments. Next we know that the schedules of private marginal valuations and costs give us the slopes of the demand and supply curves. What is not known is the vertical location of these curves as these are defined by the experimental subjects' subjective probability beliefs of a disaster state. We further assume that all buyers have the same belief and that all sellers have the same belief. The size of a vertical shift given a belief depends upon whether there is information aggregation or not. We proceed by calculating implicit beliefs under both the FA and NA hypotheses. To summarize, we have two parameters (the subjective size of the supply and demand curves' positive vertical shifts) whose values we can use to calibrate the observed market price and quantity.

The answer to the following question is not obvious; are there role-specific probability biases which can explain our results under these two models? To address this question, we perform a numerical exercise in which we deduce the implicit probability biases for buyers and for sellers using the FA and NA hypotheses. The are four main conclusions: the NA model most plausibly explains results in most periods, buyers' average implied beliefs of disaster under the NA hypothesis are below the actual ten percent probability, sellers' average probability beliefs of disaster under the NA hypothesis do not differ significantly from ten percent on average, and correspondingly sellers' implied probabilities are higher than buyers'.

Let p_b denote the buyers' perceived probability of a Disaster state and p_s denote the sellers' perceived probability of a Disaster state. Substituting into equations 1-3, we get

$$E(d)_{buyer} = (1 - p_b) \text{ (remaining N-state's dividend)}$$

$$+ p_b \text{ (remaining D-state's dividend)}$$

$$E(d)_{seller} = (1 - p_s) \text{ (remaining N-state's dividend)}$$

$$+ p_s \text{ (remaining D-state's dividend)}$$

for the expected values of the common dividend under the FA hypothesis, and

$$E(d)_{buyer} = (1 - p_b) \text{ (remaining N-state's dividend)}$$

$$+ p_b \text{(average of the D-states' dividends)}$$

$$E(d)_{seller} = (1 - p_s) \text{ (average of the N-states' dividends)}$$

$$+ p_s \text{ (remaining D-state's dividend)}$$

for the expected values of the common dividend under the NA hypothesis. Combining these equations with the private value and cost increments, we solve for market equilibrium prices and quantities for both models for all the combinations of probability beliefs (p_b, p_s) over $p_b = 0.01, 0.02, ..., 1$ and $p_s = 0.01, 0.02, ..., 1$. From these results we identify the range of probability beliefs of sellers and buyers in our experiments that could support the observed quantities and median prices for each period.

The median and range of probability beliefs for buyers and sellers supporting the observed quantities and median prices for each period's trades are shown in chronological order in Figures 2.15-2.16, separated by experiment. The dashed vertical lines mark occurrences of disaster states. Having added two degrees of freedom to our models, the choice between hypotheses becomes a matter of judgement and interpretation, rather than a test of predictions. Nevertheless, there are two features of these implied probability beliefs that tend to support the conclusion that the NA model has more explanatory power:

- The implied probability beliefs calculated for the full-aggregation model are much sparser than those calculated for the NA model. This is because no combination of buyers' and sellers' probability beliefs support the observed prices and quantities in 25 out of 54 periods for the FA model, while the same is true in just 14 out of 54 periods for the NA model.

- Buyers' and sellers' implied probabilities vary more, and more erratically, over time, and vary more from buyer to seller, under the full-aggregation model than is
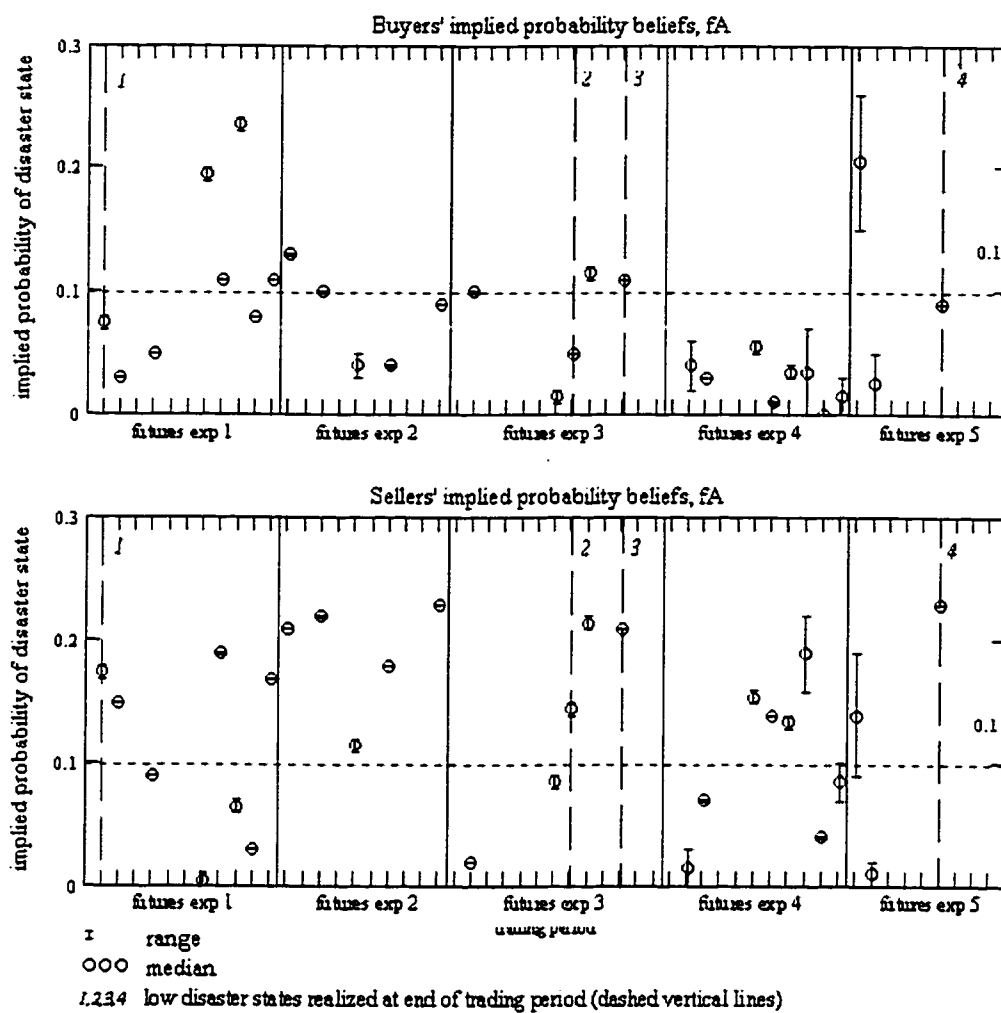
Figure 2.15: Buyers' and Sellers' implied probability beliefs under the full-aggregation model.

Buyers' implied probability beliefs, nA

Sellers' implied probability beliefs, nA

I range
OOO median
1,2,3,4 low disaster states realized at end of trading period (dashed vertical lines)
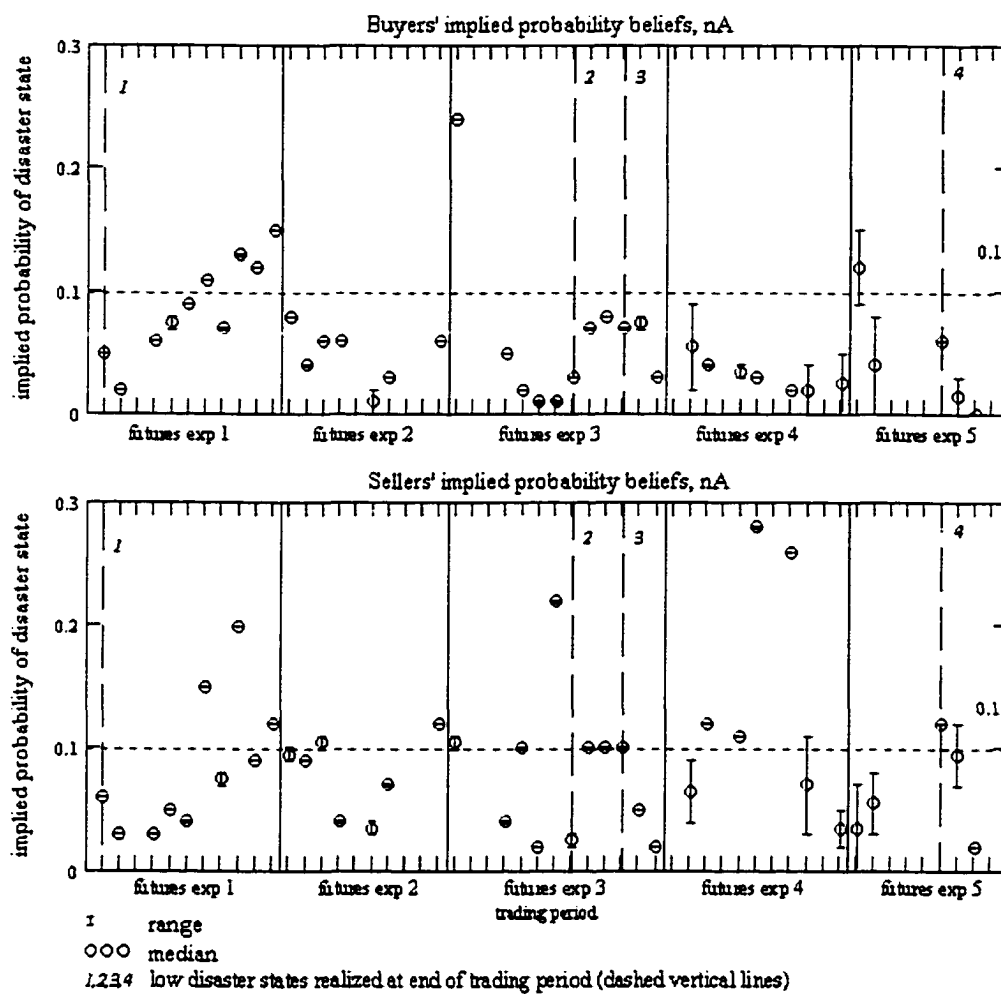
Figure 2.16: Buyers' and Sellers implied probability beliefs under the no-aggregation model.

the case under the no-aggregation model. This is likely an artifice of the data being forced to fit the model, rather than a true representation of the evolution of participants' probability beliefs. By contrast, the beliefs implied by the NA model tend to move together. Buyers' and sellers' implied beliefs tend to move in the same direction under the NA model, and period-on-period changes in beliefs tend to be much less extreme.

Clearly there is variation from period to period in both the buyers' and sellers' subjective beliefs. Table 2.6 gives some brief statistical analysis of the sets of beliefs under the NA hypothesis. For each statistic we conduct a hypothesis test that the mean is equal to ten percent versus the alternative that the mean is less than ten percent. For the sellers' beliefs we fail to reject the null at all typical levels of significance, however for the buyer we do reject the hypothesis. We also conduct a t-test for difference in means for the two sets of beliefs. Here we reject the null hypothesis that the means are equal in favor of the alternative that the sellers' mean is larger than the buyers' mean. (The t-statistic is 2.469, has 78 degrees of freedom and a p-value of 0.008.) The strong negative bias possessed by buyers corresponds to similar results found in individual choice experiments, for example Slovic et al (1977) and Kunreuther et al (1978), in which subjects purchase insurance from the experimenter against small-probability, large-loss events.

Tabe 2.6: Test of mean implied probability beliefs

Statistic	Mean	Standard Deviation	Mean Test Statistic	P-value
Seller Belief	0.089	.062	-1.152	0.125
Buyer Belief	0.059	.046	-5.590	0.000

Our experiment is the first in which some subjects *sell* insurance against small-probability, large losses. It also appears that changing the point of reference and the framing of the reinsurance task has eliminated this bias for sellers. However, there is another interesting perspective from which we can view these results. Instead of assuming that individuals are expected value maximizers who have probability biases, we could have assumed that the did not have subjective probability biases but that there preferences differ from risk neutrality. Under this interpretation we would conclude that the sellers give a greater assessment to the potential large losses of selling insurance contract than buyers give to the assessment of the large gains. This interpretation is consistent with the implications of the Kahneman and Tversky's (1979) prospect theory of decision making under uncertainty, where relative losses typically loom larger than relative gains.

H.   Conclusion

In this paper we examine an insurance market's ability to generate equilibria which reflect the union of market participants' diverse information regarding the probabilities that govern states of nature. The correlation of prior information with

market roles and the structure of uncertainty in these markets lead us to develop significant changes to the standard experimental design, introduced by Plott and Sunder (1988), used to test information aggregation. We found that the economic environment of a reinsurance market failed to generate the equilibrium predictions under either the FA model or the NA model. This is in contrast to Plott and Sunder's finding of information aggregation in simpler environments. In evaluating the hypotheses we found strong evidence that the value of the buyer's prior information had more impact on economic outcomes than did the seller's prior information. This suggested alternative explanations.

The uncertainty that characterizes insurance markets requires individuals to assess the value of small-probability, large-loss (gain) states. A plethora of past studies show that traditional expected utility theory's robustness falters in these situations, and that subjective probability biases or non-expected utility preferences can characterize behavior. In our setting one can not distinguish between a subjective probability bias and a utility phenomenon. After we calculate the implicit subjective probability beliefs in our experiment we conclude that buyers posses a strong subjective probability bias and sellers do not. The corresponding utility explanation is that sellers' potential losses from reinsurance contracts loom larger than buyers' gains from reinsurance. Finally, after we control for these decision theoretic aspects, we see that the NA hypothesis has more explanatory power than the FA hypothesis.

These results do not provide optimism that insurance markets, such as the catastrophe futures index introduced by the CBOT in 1992, can lead to outcomes in which information is aggregated and risk is efficiently shared. Given the strong desirability of the information aggregation property in insurance market, it is worthwhile to explore whether other financial instruments (e.g. PCS option spreads and Act of God Bonds) and other institutions (such as the long standing bilateral contractual relationships that governed the reinsurance market prior to 1990) fare better than the market we study here.

Our results also suggest future directions in the study of information aggregation in general. Specifically, can we explain why the challenging decision making under uncertainty environment of catastrophe insurance impedes the information aggregation process? If we can not answer this question, can we at least establish the boundary of this breakdown empirically? Furthermore, in previous experiments in which information aggregation occurs, the pooled information reveals the true state. In our experiments pooled information does not reveal the true state of nature, and it is of interest to assess the impact this has. Clearly, in most cases of interest, pooled information does not reveal the true state. Finally, we believe the introduction of the induced supply and demand approach to the study of markets with uncertainty is an innovation which may permit the performance of a wider class of experiments. The robustness of this approach needs to be more thoroughly tested.

Appendix 1

Oral Double Auction Instructions (Buyers)

Today we are going to set up a market in which some of you will be buyers and some of you will be sellers. The commodity to be traded is divided into distinct items or units. We will not specify a name for the commodity; we will simply refer to units.

Trading will occur in a sequence of trading periods. The prices that you negotiate in each period will determine your earnings. These earning will be paid to you in U.S. dollars.

We will proceed in the following way. First I will explain how buyers compute their earnings, and then I will explain how sales and purchases are arranged in the market. In today's market you are a buyer. Information specific to your role in today's market will be presented at the end of the instructions. After reading the instructions and reviewing your specific information, I will give you a chance to ask any questions you might have. Then we will begin the first trading period which will be for practice (i.e. no earnings) and then more periods for earnings.

Instructions for Buyers

Buyer decisions and earnings will be recorded on a Buyer Decision Sheet like the ones included with these instructions. Each trading period will be recorded on a

separate Buyer Decision Sheet. In each trading period, a Buyer may buy up to four units. For the first unit that may be bought during a period, the Buyer receives one of the four amounts listed on the left side of the Buyer Decision Sheet in row 3. The method for selecting which of the four values you will receive will be described later. If a second unit is purchased during the same period, the Buyer receives one of the four amounts listed on the left side of the Buyer Decision Sheet in row 6. If a third unit is purchased, the Buyer receives one of the four amounts listed in row 9. Likewise, if a fourth unit is bought during the same period, the Buyer receives one of the four amounts listed in row 12. The method for determining which of the four amounts is received each period will be explained on a separate sheet. This value will be entered by the buyer into the box to the right of the words "Unit Value" on rows 3, 6, 9 and 12 on the Buyer Decision Sheet. A Buyer may buy between 0 and 4 units and may buy these units from the same or different sellers.

Buyers earn money by purchasing units at prices that are **below** their Unit Values. Unit Earnings from the purchase of each unit are computed by taking the difference between the Unit Value and the Purchase Price. Total Unit Earnings for the period are computed by adding up the Unit Earnings from all units purchased.

In addition to their Total Unit Earnings from purchases of units, Buyers also receive one of the four Random Transfers listed on the left side of the Buyer Decision Sheet on row 2. Note that two of the possible Random Transfers are positive, and two are negative. If the Random Transfer is positive, then the Buyer receives Period Net

Earnings equal to the amount of the transfer plus their Total Unit Earnings. If the Random Transfer is negative, then the buyer receives Period Net Earnings equal to the amount of their Total Unit Earnings minus the amount of the Random Transfer. If Period Net Earnings are negative, then the Buyer has lost money for that period.

Subsequent periods are represented by separate Buyer Decision Sheets. The period number for each sheet is displayed in the upper right corner. All calculations for each period should be reflected on the Buyer Decision Sheet *for that period*.

Importantly, a buyer does not receive the value for a unit unless that unit is purchased. Thus earnings for each un-purchased unit in a period are zero (0). If you are a buyer, the first unit you buy during a trading period is *your* **1ˢᵗ unit**, regardless of whether or not other buyers have previously purchased units in the period. The purchase price of you first unit should be recorded in row 4 immediately after the purchase, and Unit Earnings should be recorded in row 5. Do the appropriate actions when you purchase your second, third or fourth unit in a period. You cannot purchase your second unit before your first, and therefore you will move *down a page* during a period. At the end of the period , record your Total Unit Earnings in row 15 of your decision sheet.

Note that a buyer receives (or pays) their Random Transfer every period, regardless of whether they purchase units or not. At the end of each period, write your Cumulative Earnings in row 17 of that period's decision sheet. You should carry over your Cumulative Earnings to row 1 of the next decision sheet.

Buyer's Example

Look at the sample Buyer's Decision Sheet that you received with these instructions. It should have a letter "A" in the box in the upper right hand corner. In this example, the buyer had $30 in cumulative earnings from previous periods, noted in row 1. The Unit Values and Random Transfer were determined to be the amounts in column Y2 of the box on the left side of the Buyer Decision Sheet. Find the Random Transfer amount in column Y2 and enter it in row 2. Two units were purchased—unit one and unit two. Suppose unit one was purchased for $7. Enter $7 in row 4. The Unit Value for unit one was determined to be the amount in column Y2, row 3: $10.66. Enter $10.66 in row 3. Thus the Unit Earnings for unit one were 10.66 - 7 = 3.66. Please enter $3.66 in row 5. Suppose unit two was purchased for $5. Enter the appropriate purchase price for unit two in row 7 and the appropriate Unit Value in row 6. Calculate the Unit Earnings for unit two and enter them in row 8. The Unit Earnings in row 8 should equal $5.18.

Calculate the Total Unit Earnings, Period Net Earnings and Cumulative Earnings for this sample trading period and enter them in rows 15, 16 and 17. Total Unit Earnings should equal 3.66 + 5.18 = 8.84. Period Net Earnings should equal the Random Transfer plus the Total Unit Earnings, or -15.40 + 8.84 = -6.56. Thus the Period Net Earnings are a loss of $6.56 for this sample period. The Cumulative Earnings in row 17 should be the sum of the prior Cumulative Earnings in row 1 and the Period Net Earnings in row 16, or 30 - 6.56 = 23.44.

Trading Rules

I will begin each 7-minute trading period with an announcement that the market is open. At any time during the period any buyer is free to raise his/her hand and when called on, to make a verbal bid to buy a unit at the price specified in the bid. Similarly, any seller is free to raise his/her hand and, when called on, to make a verbal offer to sell a unit at the price specified in the offer. All bids and offers pertain to one unit, it is not possible to sell two units as a package.

All buyers and sellers have identification numbers; your number is given in the upper left corner of the Decision Sheet attached to these instructions. These numbers must be used when making a bid or offer. Buyers should use the word bid, and sellers should use the word ask. For example, if Buyer 1 wants to make a bid of $120, then this person would raise their hand and, when recognized, say "Buyer 1 bids $120." I will repeat the buyer number and the bid to give the person at the computer time to record it. Similarly, if Seller 5 decides to offer a unit for sale at $250 , this seller should raise their ha d and, when recognized, say "Seller 5 asks $250." I will repeat this information while it is recorded. At this point the spreadsheet projected onto the overhead display will appear as follows

Bids		Asks	
B1	120	S5	250

We ask you to help us enforce a bid/ask improvement rule: All bids must be

higher than the lowest outstanding offer, should one exist. In the example above, the next bid must be above $120 and the next ask must be below $250.

For example, suppose that Buyer 1, the next person recognized, raises his/her own bid from $120 to $130, and then Seller 4 is called on and asks $165. I would repeat the bid and ask as they are recorded on the overhead display:

Bids		Asks	
B1	120	S5	250
B1	130	S4	165

To save space, the bids and asks will be written in small numbers, without the dollar signs and decimals. Please tell us if you cannot read the numbers recorded or if you think that a bid or asks was not recorded correctly.

Any seller is free to accept or not accept the bid of any buyer, and any buyer is free to accept or not accept the asking price of any seller. To accept a bid or ask, simply raise your hand. After you are recognized, announce your identity and indicate acceptance, e.g. "Buyer 2 accepts Seller 3's ask."

Suppose that buyer 3 bids $160 and that the next person recognized is Seller 5 who accepts this bid. I would repeat this acceptance, while the person at the computer enters the buyer number, seller number, transaction price and the letter "A" for accepts.

Bids		Asks	
B1	120	S5	250
B1	130	S4	165
B3	**160**	**S5**	**A**

If a bid or ask is accepted, a binding contract has been closed for a single unit, and the buyer and seller involved will immediately record the contract price and earnings for the unit. After each contract is closed, all previous bids and asks will be automatically withdrawn before any new ones can be made.

Following the acceptance of buyer 3's bid of $160, a horizontal line would have been drawn below the accepted contract. Subsequent bids need not be above $160 and in fact can be below any of the earlier bids. The horizontal line is to remind you that the contract invalidates previous bids and asks.

If seller 4 wished to ask $165 again, this seller would raise his/her hand and be recognized. Suppose that Buyer 1 bids $140 and Buyer 3 is then recognized and accepts Seller 4's asking price. The display will appear as below.

Bids		Asks	
B1	120	S5	250
B1	130	S4	165
B3	**160**	**S5**	**A**
B1	140	**S4**	**165**
B3	**A**		

Except for bids, asks and their acceptances, you are expected not to speak to any other person, even if there are many bids and offers that are not accepted. Also note that on your attached Decision Sheet their is your private costs or values. You should not share these with anyone. Once everyone has finished reading the instructions we will describe the method by which the values of the random Unit Earnings and Transfers is determined, then we will start with a 7 minute practice round, and then do multiple 7 minute rounds.

Appendix 2

Oral Double Auction Instructions (Sellers)

Today we are going to set up a market in which some of you will be buyers and some of you will be sellers. The commodity to be traded is divided into distinct items or units. We will not specify a name for the commodity; we will simply refer to units.

Trading will occur in a sequence of trading periods. The prices that you negotiate in each period will determine your earnings. These earning will be paid to you in U.S. dollars.

We will proceed in the following way. First I will explain how sellers compute their earnings, and then I will explain how sales and purchases are arranged in the market. In today's market you are a seller. Information specific to your role in today's market will be presented at the end of the instructions. After reading the instructions and reviewing your specific information, I will give you a chance to ask any questions you might have. Then we will begin the first trading period which will be for practice (i.e. no earnings) and then more periods for earnings.

Instructions for Sellers

Seller decisions and earnings will be recorded on a Seller Decision Sheet like the ones included with these instructions. Each trading period will be recorded on a

separate Seller Decision Sheet. In each trading period, a Seller may sell up to four units. For the first unit sold during a period, the Seller pays one of the four costs listed on the left side of the Seller Decision Sheet in row 3. The method for selecting which of the four costs you will pay will be described later. If a second unit is sold during the same period, the Seller pays one of the four costs listed on the left side of the Seller Decision Sheet in row 6. If a third unit is purchased, the Seller pays one of the four costs listed in row 9. Likewise, if a fourth unit is bought during the same period, the Seller pays one of the four costs listed in row 12. The method for determining which of the four costs is payed each period will be explained on a separate sheet. This cost will be entered by the seller into the box to the right of the words "Unit Cost" on rows 3, 6, 9 and 12 on the Seller Decision Sheet. A Seller may sell between 0 and 4 units and may sell these units to the same or different buyers.

Sellers earn money by selling units at prices that are **above** their Unit Costs. Unit Earnings from the sale of each unit are computed by taking the difference between the Selling Price and the Unit Cost. Total Unit Earnings for the period are computed by adding up the Unit Earnings from all units purchased.

Subsequent periods are represented by separate Seller Decision Sheets. The period number for each sheet is displayed in the upper right corner. All calculations for each period should be reflected on the Seller Decision Sheet *for that period*.

Importantly, a seller does not pay the cost for a unit unless that unit is sold. Thus Unit Costs (and Unit Earnings) for each unsold unit in a period are zero (0). If you

are a seller, the first unit you sell during a trading period is *your* **1ˢᵗ unit**, regardless of whether or not other sellers have previously sold units in the period. The selling price of you first unit should be recorded in row 2 immediately after the sale, and Unit Earnings should be recorded in row 4. Do the appropriate actions when you sell your second, third or fourth unit in a period. You cannot sell your second unit before your first, and therefore you will move *down a page* during a period. At the end of the period, record your Total Unit Earnings in row 14 of your decision sheet.

At the end of each period, write your Cumulative Earnings in row 15 of that period's decision sheet. You should carry over your Cumulative Earnings to row 1 of the next decision sheet.

Seller's Example

Look at the sample Seller's Decision Sheet that you received with these instructions. It should have a letter "A" in the box in the upper right hand corner. In this example, the seller had $30 in cumulative earnings from previous periods, noted in row 1. The Unit Costs were determined to be the amounts in column Y2 of the box on the left side of the Seller Decision Sheet. Two units were sold⬚unit one and unit two. Suppose unit one was sold for $7. Enter $7 in row 2. The Unit Cost for unit one was determined to be the amount in column Y2, row 3: $9.34. Enter $9.34 in row 3. Thus the Unit Earnings for unit one were 7 - 9.34 = -2.34. Please enter -$2.34 in row 4. Suppose unit two was sold for $5. Enter the appropriate selling price for unit two

in row 5 and the appropriate Unit Cost in row 6. Calculate the Unit Earnings for unit two and enter them in row 7. The Unit Earnings in row 7 should equal -$4.82.

Calculate the Total Unit Earnings and Cumulative Earnings for this sample trading period and enter them in rows 14 and 15. Total Unit Earnings should equal - 2.34 - 4.82 = -7.16. Thus the Total Unit Earnings for this example period are a loss of $7.16. The Cumulative Earnings in row 15 should be the sum of the prior Cumulative Earnings in row 1 and the Total Unit Earnings in row 14, or 30 - 7.16 = $22.84.

Trading Rules

I will begin each 7-minute trading period with an announcement that the market is open. At any time during the period any buyer is free to raise his/her hand and when called on, to make a verbal bid to buy a unit at the price specified in the bid. Similarly, any seller is free to raise his/her hand and, when called on, to make a verbal offer to sell a unit at the price specified in the offer. All bids and offers pertain to one unit, it is not possible to sell two units as a package.

All buyers and sellers have identification numbers; your number is given in the upper left corner of the Decision Sheet attached to these instructions. These numbers must be used when making a bid or offer. Buyers should use the word <u>bid</u>, and sellers should use the word <u>ask</u>. For example, if Buyer 1 wants to make a bid of $120, then this person would raise their hand and, when recognized, say "Buyer 1 bids $120." I will repeat the buyer number and the bid to give the person at the computer time to record it. Similarly, if Seller 5 decides to offer a unit for sale at $250 , this seller should raise

their hand and, when recognized, say "Seller 5 asks $250." I will repeat this information while it is recorded. At this point the spreadsheet projected onto the overhead display will appear as follows

Bids		Asks	
B1	120	S5	250

We ask you to help us enforce a bid/ask improvement rule: All bids must be higher than the lowest outstanding offer, should one exist. In the example above, the next bid must be above $120 and the next ask must be below $250.

For example, suppose that Buyer 1, the next person recognized, raises his/her own bid from $120 to $130, and then Seller 4 is called on and asks $165. I would repeat the bid and ask as they are recorded on the overhead display:

Bids		Asks	
B1	120	S5	250
B1	130	S4	165

To save space, the bids and asks will be written in small numbers, without the dollar signs and decimals. Please tell us if you cannot read the numbers recorded or if you think that a bid or asks was not recorded correctly.

Any seller is free to accept or not accept the bid of any buyer, and any buyer is free to accept or not accept the asking price of any seller. To accept a bid or ask, simply raise your hand. After you are recognized, announce your identity and indicate acceptance, e.g. "Buyer 2 accepts Seller 3's ask."

Suppose that buyer 3 bids $160 and that the next person recognized is Seller 5 who accepts this bid. I would repeat this acceptance, while the person at the computer enters the buyer number, seller number, transaction price and the letter "A" for accepts.

Bids		Asks	
B1	120	S5	250
B1	130	S4	165
B3	**160**	**S5**	**A**

If a bid or ask is accepted, a binding contract has been closed for a single unit, and the buyer and seller involved will immediately record the contract price and earnings for the unit. After each contract is closed, all previous bids and asks will be automatically withdrawn before any new ones can be made.

Following the acceptance of buyer 3's bid of $160, a horizontal line would have been drawn below the accepted contract. Subsequent bids need not be above $160 and in fact can be below any of the earlier bids. The horizontal line is to remind you that the contract invalidates previous bids and asks.

If seller 4 wished to ask $165 again, this seller would raise his/her hand and be recognized. Suppose that Buyer 1 bids $140 and Buyer 3 is then recognized and accepts Seller 4's asking price. The display will appear as below.

Bids		Asks	
B1	120	S5	250
B1	130	S4	165
B3	**160**	**S5**	**A**
B1	140	**S4**	**165**
B3	**A**		

Except for bids, asks and their acceptances, you are expected not to speak to any other person, even if there are many bids and offers that are not accepted. Also note that on your attached Decision Sheet their is your private costs or values. You should not share these with anyone. Once everyone has finished reading the instructions we will describe the method by which the values of the random Unit Earnings and Transfers is determined, then we will start with a 7 minute practice round, and then do multiple 7 minute rounds.

Appendix 3

Information about Unit Costs/Values

There are four possible outcomes–X1, X2, Y1 and Y2–for the random Unit Values, Transfers, and Costs. The Unit Values, Transfers, or Costs that are associated with each of the four outcomes in each period are your private information. Each year, the final outcome is determined in two stages. First, at the beginning of each market year, one of the X outcomes (X1 or X2) and one of the Y outcomes (Y1 or Y2) is eliminated as follows:

At the beginning of each year, before trading starts, the experimenter flips a coin to eliminate either X1 or X2. In this example, suppose that the "X1" outcome is eliminated. In each period, there is a different, randomly assigned code for "X1" and "X2" noted on the Buyer's Decision Sheet. For example, in the sample Buyer Decision Sheet below, the "X1" column on the left side of the Buyer Decision Sheet is labeled "Green" and the "X2" column is labeled "Purple." In this example, the experimenter would announce "not Green." The buyers would record this information by drawing a vertical line crossing out the column labled "Green" on the left side of their Buyers Decision Sheet. You must not discuss this information. This information is for buyers to know only.

Similarly, at the beginning of each year, before trading starts, the experimenter flips a coin to eliminate either Y1 or Y2. In this example, suppose that the "Y2" outcome is eliminated. In each period, there is a different, randomly assigned code for "Y1" and "Y2" noted on the Seller's Decision Sheet. For example, in the sample Seller's Decision Sheet below, the "Y1" column on the left side of the Seller's Decision Sheet is labeled "Guava" and the "Y2" column is labeled Mango." In this example, the experimenter would anounce "not Mango." The sellers would record this information by drawing a verticle line crossing out the column labeled "Mango". You must not discuss this information. This information is for sellers to know only.

In the second stage, after all transactions are completed, the experimenter at the end of the period draws a ball from a bingo cage containing ten balls numbered one through ten. If the ball drawn is numbered one through nine, the outcome is the remaining X outcome. If the ball drawn is numbered ten, the outcome is the remaining Y outcome. Thus, there is a 90% chance of an X outcome, and a 10% chance of a Y outcome. The experimenter then announces the final outcome to the whole room.

To summarize our example: Suppose in the first stage of year 1 the experimenter eliminates X1 and Y2. The experimenter announces "not Green" and "not Mango." Buyers look at their Buyer's Decision Sheets and see that "X1" in period 1 is labeled "Green." They should mark down this information on their Buyer Decision Sheet by drawing a vertical line through column X1. Sellers look at their Seller's Decision Sheets and see that "Y2" is labled "Mango." They should mark down

this information on their Sellers Decision Sheet   by drawing a vertical line through column Y2. Now the possible final outcomes are only X2 or Y1. At the end of the year, suppose the experimenter draws a ball from the bingo cage with the number ⊔6" on it. Then the outcome for the year is X2. For every unit he or she bought, a buyer's unit values are those shown in column X2 of their Buyers Decision Sheet. A buyer's Random Transfer is listed at the top of column X2. For every unit he or she sold, a seller's Unit Cost is listed in column X2 of their Sellers Decision Sheet.

J. Bibliography

Cox, Samuel H. and Robert G. Schwebach (1992) "Insurance Futures and Hedging Insurance Price Risk," *The Journal of Risk and Insurance*, 59.4:628-644.

Cummins, J. David and Hélyette Geman (1995) "Pricing Catastrophe Insurance Futures and Call Spreads: An Arbitrage Approach," *The Journal of Fixed Income*, 4:46-57.

D'Arcy, Stephen P. and Virgina G. France (1992) "Catastrophe Futures: A Better Hedge for Insurers," *The Journal of Risk and Insurance*, 59.4:575-601.

Doherty, Neil A. (1996) "Insurance Markets and Climate Change," Presented at IIASA 1996 conference *Climate Change: Cataclysmic Risk and Fairness*.

Doherty, Neil A. (1997) "Innovations in Managing Catastrophe Risk," *Journal of Risk and Insurance*, 64.4:713-718.

Forsythe, Robert and Russell Lundholm (1990) "Information Aggregation in an Experimental Market," *Econometrica*, 58.2:309-347.

Gjerstad, Steven and Jason Shachat (1996) "A General Equilibrium Structure for Induced Supply and Demand," UCSD Economics Discussion Paper 96-35.

Harrington, Scott E., Steven V. Mann and Greg Niehaus (1995) "Insurer Capital Structure Decisions and the Viability of Insurance Derivatives," *The Journal of Risk and Insurance*, 62.3:483-508.

Harrington, Scott and Greg Niehaus (1999) "Basis Risk with PCS Catastrophe Insurance Derivative Contracts," *Journal of Risk and Insurance*, 66:49-82.

Holt, Charles A., Loren Langan, and Anne P. Villamil (1986) "Market Power in Oral Double Auctions," *Economic Inquiry*, 24, 107-123.

Kahneman, Daniel and Amos Tversky (1979) "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica* 47:263-291.

Kunreuther, Howard (1996) "Mitigating Losses and Providing Protection against Catastrophic Risks: The Role of Insurance and Other Policy Instruments," The Geneva Lecture, presented at the *European Society for Risk Analysis*, June 4, 1996, Guilford, England.

Kunreuther, Howard (1998) "Preface," in *Paying the Price: The Status and Role of Insurance Against Natural Disasters in the United States*, Howard Kunreuther and Richard J. Roth, Eds. (Originally: 1996 Monograph for the Second Assessment of Research and Applications on Natural Hazards.) Washington D.C.: Joseph Henry Press.

Kunreuther, Howard, Ralph Ginsberg, Louis Miller, Philip Sagi, Paul Slovic, Bradley Borkan and Norman Katz (1978) *Disaster Insurance Protection: Public Policy Lessons*, John Whiley and Sons: New York.

Lecomte, Eugene (1996) "Hurricane Insurance Protection in Florida," chapter 5 in in *Paying the Price: The Status and Role of Insurance Against Natural Disasters in the United States*, Howard Kunreuther and Richard J. Roth, Eds. (Originally: 1996 Monograph for the Second Assessment of Research and Applications on Natural Hazards.) Washington D.C.: Joseph Henry Press.

Marlett, David C. and Alan Eastman (1997) "The Estimated Impact of Residual Market Assessments on Florida Property Insurers," *American Risk and Insurance Association 1997 Annual Meeting*.

McClelland, Gary H., William D. Schulze and Don Coursey (1993) "Insurance for Low-Probability Hazards: A Bimodal Response to Unlikely Events," *Journal of Risk and Uncertainty*, 7:95-116.

Niehaus, Greg and Steven V. Mann (1992) "The Trading of Underwriting Risk: An Analysis of Insurance Futures Contracts and Reinsurance," *The Journal of Risk and Insurance*, 59.4:601-627.

Nutter, Franklin W. (1994) "The Role of Government in the United States in Addressing Natural Catastrophes and Environmental Exposures," *The Geneva Papers on Risk and Insurance: Issues and Practice*, (no. 72) 19:244-256.

Ohare, Dean (1994) "The Need for Insurers to Change," *Geneva Papers on Risk and Insurance: Issues and Practice*, (no. 72) 19:357-364.

Plott, Charles R. and Shyam Sunder (1982) "Efficiency and Experimental Security Markets with Insider Information: An Application of Rational Expectations Models," Journal of Political Economy, 90:663-698.

Plott, Charles R. and Shyam Sunder (1988) "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica*, 56.5:1085-1118.

Roth, Richard J. Jr. (1996) "Earthquake Insurance Protection in California," chapter 4 in *Paying the Price: The Status and Role of Insurance Against Natural Disasters in the United States*, Howard Kunreuther and Richard J. Roth, Eds. (Originally: 1996 Monograph for the Second Assessment of Research and Applications on Natural Hazards.) Washington D.C.: Joseph Henry Press.

Slovic, Paul and Baruch Fischhoff, Sarah Lichtenstein, Bernard Corrigan and Barbara Combs (1977) "Preference for Insuring Against Probable Small Losses: Insurance Implications," *The Journal of Risk and Insurance*, 44.2:237-258.

Smith, Vernon L. (1962) "An Experimental Study of Competitive Market Behavior," *The Journal of Political Economy*, 60.2:111-137.

Smith, Vernon L. (1964) "The Effect of Market Organization on Competitive Equilibrium," *Quarterly Journal of Economics*, 78, 181-201.

Smith, Vernon L. (1981) "An Empirical Study of Decentralized Institutions of Monopoly Restraint," in J. Quirk and G. Horwich, eds., *Essays in Contemporary Fields of Economics in Honor of E.T. Weiler (1914-1979)*. West Lafayette: Purdue University Press, 83-106.

Smith, Vernon L., and Arlington W. Williams (1981) "The Boundaries of Competitive Price Theory: Convergence, Expectations, and Transactions Costs," in L. Green and J. Kagel, eds., *Advances in Behavioral Economics*, vol. 2. Norwood, N.J.:Ablex Publishing.

Chapter III

The Value of Extended Climate Forecasts in Insurance Markets: Heterogenous Risk Beliefs, Market Power and Regulation.

A. Introduction

In recent years, our ability to forecast regional climatic conditions twelve to eighteen months in advance has greatly improved. It seems natural to suppose that such information will be of considerable value to providers and consumers of property catastrophe insurance. In this study we identify conditions under which simple models of insurance transactions generate positive utility for the use of forecast information. We examine model insurance markets where consumers may only purchase full insurance contracts, as well as models like that of Mossin (1968) where consumers may choose less than full insurance (which we will also refer to as a variable insurance model). Under either formulation, positive expected utility of forecast information for consumers derives from insurers' market power, regulatory constraints, and differences in risk beliefs. In other words, consumers of insurance derive value from reductions in uncertainty only if they are not fully insured against loss, or if the price of insurance

includes a risk premium over and above the value of the expected loss.[1] When consumers may choose less than full insurance, these amount to much the same thing, since then consumers only fully insure when there is no risk premium.

In most of the cases examined here, the consumer's gains from the introduction of extended forecasts are greater if the consumer may only purchase full insurance or nothing. This is because the ability to purchase intermediate levels of insurance allows the consumer to better optimize their utility independent of the availability of forecast information. One area where the models' implications diverge is in the effects of different forms of regulation. Adding a regulatory constraint in the form of a fixed risk reserve to a competitive insurance market results in increased consumer utility from the use of forecasts if consumers can choose intermediate levels of insurance, but not under the full-insurance model. On the other hand, constraining the insurer's risk of ruin (bankruptcy) below a fixed probability results in gains for the consumer from the use of forecasts under the full insurance model.

Under both models, a monopolist insurer's expected profits are reduced by the use of forecasts. Differences in perceived risks between insurer and insured can, however, result in an increase in a monopolist insurer's profits from the use of forecast

---

[1]We do not consider the potential for gain or loss from consumers undertaking different activities separate from insurance policies due to forecasts. For example, consumers might take unusual preventative measures such as repairing their roof if a wet winter was forecast. Alterna-tively, they might schedule construction and maintenance activities during a dry winter when labor and materials costs may be lower.

information under certain circumstances. We consider two cases here. In the first, the parameters governing the risk processes, and any forecast information, are assumed to be public information, but some consumers of insurance have a low probability bias. That is, they underweight the probability of the less likely state of nature. Kunreuther et. al. (1978) and McClelland, Schulze and Coursey (1993) find evidence in laboratory experiments and field surveys that some buyers of insurance behave as though their subjective probability of infrequent, high value risks is zero. If a some consumers under-insure because they hold these biases, an insurer's monopoly profits are increased by the introduction of forecasts under certain circumstances.

In the second scenario, we suppose that insurers privately generate forecast information that subsequently becomes known to consumers only as a result of the operation of the market. The full-insurance model implies that a monopolistic insurer with private forecast information will increase profits by concealing his private information when the expected loss conditional on the forecast is lower than the unconditional expected loss, and revealing his information when the forecast loss is higher than the unconditional expected loss. Strictly maximizing profits under the variable-insurance model may, however, send clear signals to consumers regarding the insurer's private information. Similarly, an insurer in an otherwise competitive market can earn positive profits by concealing private forecast information when the expected loss conditional on the forecast is lower than the unconditional expected loss, and either revealing his information or withdrawing from the market when the forecast loss is

greater than the unconditional expected loss.

Finally, consumers in a competitive, unregulated market with common risk beliefs do not benefit from the introduction of forecasts. Consumers with a low probability bias, however, can benefit from forecasts in two ways. First, without changing their beliefs they can improve their expected utility by selectively insuring in periods when forecast losses are consistent with their beliefs. Second, their utility is further enhanced if the greater conditional probability of a rare, forecast state induces them to re-evaluate their risk beliefs.

Note that throughout this analysis we assume that insurance markets allow for price and quantity to adjust to new information each period. It is possible that traditional property insurance markets may not have demonstrated such flexibility in the past. However, extended climate forecasts that reliably predicted parameters of interest to property catastrophe markets were also not previously available, so we do not know whether these markets would have functioned in the way our models describe. In any event, the ongoing securitization of property catastrophe insurance risk in index futures and options markets and bond markets should in future provide opportunities to observe flexible insurance markets responding to accurate extended forecasts.

In the next sections I will review the related literature and some basic insurance market models. In subsequent sections I use these models to examine the effects of various market conditions on the utility of forecast information. We will start with

competitive and monopolistic examples in sections D and E, and then consider the effects of differences in risk beliefs in sections F through H. Finally, the effects of different regulatory constraints are modeled in section I.

B.   Review of the Literature

The vast majority of studies addressing the value of either current weather or extended climate forecasts are limited to the agricultural and energy sectors. There are very few studies of the value of climate forecasts to the insurance industry. This is understandable, considering that only recently have we developed the capacity to reliably predict climate on a timescale useful to the property casualty insurance industry. Many authors considering the value of climate forecasts make much of the need for flexibility in decision making processes in a given industry on timescales appropriate to the forecast horizon for there to be any potential gain from the information.[2] Since property insurance contracts are typically renewed annually, shorter term forecasts are of little value in the property insurance market.

Luo et al (1994) consider the effect of early season forecasts on corn growers' demand for crop insurance. They are primarily concerned with a form of adverse selection, where growers selectively insure when forecasts indicate lower than average productivity. Their result, that improved extended forecasts may increase the crop insurance program's average costs, imply that pricing of the insurance is not very flexible.

---

[2] See for example Anderson 1973.

Hoy (1998) considers the impact of the symmetric introduction of improved information about loss distributions to insurance markets. He finds that improved information can lower average costs by allowing insurers to more efficiently separate consumers by risk type. Hoy's gains in welfare depend on insurers having the power to impose a risk premium above expected losses, or facing a regulatory solvency constraint, similar to our results below. He does not consider extended term forecasts, however, merely improvements in the characterization of unconditional loss distributions. Ligon (1996) and Ligon and Thistle (1996) also consider gains from information in the context of adverse selection, and find consumers benefit from improved information regarding their loss process. They do not address the role of forecasts where the underlying loss processes are known. We do not consider adverse selection in this paper, although we do look at the effects of low probability biases among consumers who otherwise face the same risk process.

There are numerous empirical studies of the value of ENSO and long term weather forecasts,[3] but most of these use expected profit maximization as a decision rule, and value the information to producers. As such, they do not capture consumers' and producers' risk attitudes and the gain from risk reduction using these forecasts. There are some empirical studies on corn silage and gas storage that hint at the potential

---

[3]Most notably for agriculture. See, for example, Adams et al (1995) and Solow et al (1998).

insurance value of climate forecasts, but they do not address insurance per se.[4] Lave

(1963) and Babcock (1990) note that long term forecasts may actually reduce profits

in some farm sectors, especially those with inelastic demand. This is because the

information leads to productivity increases that raise output and lower prices in a

competitive market. In our insurance models, forecasts do not affect productivity in a

competitive market, except where they lower the cost of complying with solvency

regulations. Babcock notes that the conventional wisdom that information

improvements are supply increasing and producer welfare increasing ignores price

effects. In our model of an insurance market, information improvements are felt chiefly

as price effects. As we will see, forecasts reduce profits for insurers with the power to

price above the actuarially fair rate, but this is because such information, on average,

reduces the maximum price consumers are willing to pay for a policy, rather than by

increasing supply.

C.   Review of the Basic Model

Consider an individual with wealth X who is subject to a random loss L

described by the   cumulative distribution function F(L).   In Mossin's (1968)

formulation they can choose the quantity of insurance Q_κ, $0 \leq Q_\kappa \leq 1$, to purchase,

paying premium $Q_\kappa \lambda$ in return for compensation $Q_\kappa L$. The individual's problem is,

then, to choose the level of insurance Q_κ to consume that maximizes their expected

---

[4]See Nichols 1999 for gas storage and distribution, and McNew 1999 for
corn storage.

utility:

$$\max_{\{Q_\kappa\}} \int_0^\infty U\Big(X - Q_\kappa \lambda - \big(1 - Q_\kappa\big)L\Big)dF(L) \ ,$$

where $U(\bullet)$ denotes an individual's utility of final wealth. We will assume a tractable quadratic form for utility like that of Blazenko (1986). Expected utility for the consumer of insurance is then

$$X - Q_\kappa \lambda - (1 - Q_\kappa)\mu - \tfrac{\kappa}{2}(1 - Q_\kappa)^2 \sigma^2 \tag{1}$$

where $-\tfrac{\kappa}{2}$ is the marginal rate of substitution between the mean μ and variance σ^2 of loss L. The consumer's expected utility maximizing demand for insurance is then

$$Q_\kappa = 1 - \frac{\lambda - \mu}{\kappa\sigma^2}$$

Similarly, an insurer chooses Q_ι, the level of insurance to provide, to maximize his expected utility:

$$\max_{\{Q_\iota\}} \int_0^\infty U\Big(Y + \lambda Q_\iota - LQ_\iota\Big)dF(L) \ ,$$

where Y is the insurer's initial wealth. Insurer's expected utility is

$$Y + Q_\iota \lambda - Q_\iota \mu - \tfrac{\iota}{2}Q^2 \sigma^2 \ ,$$

where $-\tfrac{\iota}{2}$ is his marginal rate of substitution between the mean μ and variance σ^2 of loss L. The insurer is willing to supply insurance coverage Q_ι to maximize expected

utility:

$$Q_\iota = \frac{\lambda - \mu}{\iota\sigma^2}.$$

In equilibrium $Q_\kappa = Q_\iota$ and we derive the premium rate

$$\lambda = \mu + \left(\frac{\iota\kappa}{\iota + \kappa}\right)\sigma^2$$

and the insurer's expected profits are

$$\frac{\iota\kappa^2}{\left(\iota + \kappa\right)^2}\sigma^2. \tag{2}$$

As Blazenko (1986) notes, profits for risk-neutral investors in an insurance firm are maximized if $\iota = \kappa$. Thus, a risk-neutral monopolist will maximize profits by selecting managers who behave as though their coefficient of risk aversion is the same as that of their customers. Thus, premium and quantity[5] will be

$$\lambda = \mu + \tfrac{\kappa}{2}\sigma^2, \quad Q = \tfrac{1}{2} \tag{3}$$

and expected profits, $(\lambda - \mu)Q$, are $\tfrac{\kappa}{4}\sigma^2$. In a competitive insurance market, risk-neutral insurers will be constrained to offer full insurance on actuarially fair terms, with $\lambda = \mu$ and $Q = 1$, and expected profits will always be zero.

If insurers only offer full insurance contracts, then consumers will be willing to purchase them if their utility with the contracts is at least as great as without:

---

[5]Note that the profit maximizing level of insurance for the monopolist, 0.5, is a natural result of specifying a quadratic form for utility.

$$\int_0^\infty U(X - \lambda)dF(L) \geq \int_0^\infty U(X - L)dF(L),$$

which under our utility model is

$$X - \lambda \geq X - \mu - \frac{\kappa}{2}\sigma^2.$$

Thus, consumers will purchase full insurance whenever $\lambda \leq \mu + \frac{\kappa}{2}\sigma$. In a competitive market, premium is equal to the expected loss, as before, and a monopolistic insurer will still set $\lambda = \mu + \frac{\kappa}{2}\sigma$. With full insurance, however, the monopolistic insurer's profits will be higher than under the variable model if the consumer's utility function is concave. The fact that they are exactly twice as large in this example is simply an artifact of the quadratic functional form used here.

D.   Competitive Equilibrium with Common Beliefs

In order to consider the effects on insurance markets of introducing the use of forecast information, we shall suppose that in each time period covered by an insurance contract, there are two possible states of nature s, $s \in \{0,1\}$, which occur with the probabilities π and $(1 - \pi)$, respectively. Conditional on the state s, the loss L is drawn from the c.d.f. $F_s(\bullet)$ with mean μ_s and variance σ_s^2. We would naturally suspect that reality may be more complex than this model. However, our motivating example of a twelve to eighteen month climate forecast lends itself to this simple approach. This is because, from the point of view of a property insurance market, the primary value of such a forecast is to reveal whether an El Niño or La Niña phase of the El

Niño/Southern Oscillation cycle may be imminent. In many regions in the United States, one of these phases is often associated with increased property catastrophe risk. So, for example, the risk of property damage from hurricanes in Florida tends to be greater during a strong La Niña phase, while storm damage on the California coast is greater during a strong El Niño phase. Thus, in our example, one state (when s = 1) will correspond to a *catastrophe* state with elevated catastrophe risk. The other state (s = 2) will correspond to a *normal* state, with a smaller catastrophe risk.

We define a catastrophe risk as the risk of a loss event where the losses of individual consumers are correlated. If an insurance company pools many identical, independent risks, then the variance of the average loss per consumer is negligible. Pooling identical, perfectly correlated risks does not reduce the variance of the average loss. In our example, we are concerned only with the catastrophic, or correlated, part of a consumer's risk. It is these catastrophe risks about which we expect to learn from an extended climate forecast.

As a measure of value for the use of forecasts for the consumer we employ static comparisons between the unconditional mean of the consumer's conditional expected utility given a forecast, and the expected utility in the absence of a forecast. Similarly, for the insurer, we compare the unconditional mean of his conditional expected profit given a forecast to the expected profit in the absence of a forecast. In the absence of any forecast, the conditional mean and variance of the catastrophe loss in a given contract period is not known. Then the expected loss each period is

$$\mu = \pi\mu_1 + (1 - \pi)\mu_2 \tag{4}$$

with its variance equal to the mean of the conditional variance plus the variance of the conditional mean:

$$\sigma^2 = \pi\sigma_1^2 + (1 - \pi)\sigma_2^2 + (\pi - \pi^2)(\mu_1 - \mu_2)^2. \tag{5}$$

Suppose that it now becomes possible to make a "perfect" forecast prior to each contract period, so that the c.d.f. $F_s(\bullet)$ from which the catastrophe loss will be drawn that period is known. Then the contract terms will reflect the consumer's and insurer's knowledge of the conditional mean and variance of that loss process. As we saw above, in a competitive equilibrium, the quantity consumed is $Q = 1$ and the premium λ is equal to the expected loss. Thus, when the state is forecast to be $s = 1$, the premium is equal to μ_1, and when the state is forecast as $s = 2$, the premium is μ_2. The difference in variance has no effect on premium and quantity of insurance coverage in a competitive market. The consumer's expected utility of forecasts is then

$$U(\mu,\sigma^2) - \left(\pi U(\mu_1,\sigma_1^2) + (1-\pi)U(\mu_2,\sigma_2^2)\right)$$

or

$$(X-\mu) - \left(\pi(X-\mu_1) + (1-\pi)(X-\mu_2)\right) = 0.$$

Thus, there is no gain in welfare for the consumer from the use of a forecast in this model.

In a comparative statics exercise like this, the use of a forecast has no value for the insurer either, for whom expected profits are zero in a competitive equilibrium. We might suppose that an insurer could earn positive profits if they had a monopoly on

private forecast information. Once other market participants shared the same information, however, the opportunities for profit would be gone. We will discuss this point in greater detail below.

Note also that we assume the consumer behaves as though there is no possibility that claims will not be paid. Rees, Gravelle and Wambach (1999) find that when consumers are risk averse and can observe the insurers' risk of ruin (insolvency), firms in a competitive market will always hold enough capital to insure solvency.[6] If forecast states have different maximum possible losses, then the use of forecasts will produce savings for the consumer through reduced average capital costs. Without forecasts, the insurer must have enough capital on hand to cover the maximum possible loss for the worst case state of nature in every period. This will always be greater than or equal to the amount of capital necessary to cover the maximum possible loss conditional on the forecast state. If insurers' risk of insolvency is difficult to observe, some form of regulation may be useful to reduce the risk that consumer's claims may go unpaid. As we will discuss below, forecasts may reduce the costs of complying with solvency regulations in some cases.

---

[6]Rees et al's result also requires that there be a finite quantity of capital sufficient to insure solvency, that insurers do not face restrictions on the composition of their asset portfolios, and that full insurance only is offered. It is unlikely that all of these conditions are met as a rule. If the consumer has reason to believe that the state may intervene to cover unpaid claims in the event of the insurer's insolvency, Rees et al's result is also not likely to hold.

E. Monopoly with Common Beliefs

Recall that where the insurer has monopoly power, a risk neutral manager will maximize expected profits by setting the premium equal to the maximum value of an insurance policy to the consumer

$$\lambda = \mu + \frac{\kappa}{2}\sigma^2 .$$

If consumers are free to pick the quantity insured, the equilibrium quantity Q given our formulation of utility will be 0.5. Substituting these values for quantity and premium into equation 1 yields the consumer's expected utility in a period with no forecast

$$X - \mu - \frac{3}{8}\kappa\sigma^2 . \tag{6}$$

Their expected utility conditional on a forecast s is

$$X - \mu_s - \frac{3}{8}\kappa\sigma_s^2 \tag{7}$$

where μ_s and σ_s denote the mean and variance conditional on the forecast state s. Combining equations 4-7, we find the change in expected utility to the consumer from the ability to forecast the state of nature s is

$$\frac{3\kappa}{8}\left(\pi - \pi^2\right)\left(\mu_1 - \mu_2\right)^2 \geq 0$$

which is strictly greater than zero so long as consumers are risk averse (i.e., $\kappa > 0$) and the forecast states of nature are distinct (i.e., $\mu_1 \neq \mu_2$ and $0 < \pi < 1$).

If the insurer only offers a full-insurance policy, then the premium is the same,

but the consumer's expected utility in a period with no forecast is now lower:

$$X - \mu - \frac{1}{2}\kappa\sigma^2 \; ;$$

their expected utility conditional on a forecast s is

$$X - \mu_s - \frac{1}{2}\kappa\sigma_s^2$$

and the change in expected utility from their ability to forecast the state of nature s is greater

$$\frac{\kappa}{2}\left(\pi - \pi^2\right)\left(\mu_1 - \mu_2\right)^2 \geq 0 \quad .$$

Their inability to purchase an intermediate level of insurance reduces their utility in any given period proportional to the variance of their potential loss, so a reduction in the average variance has a greater payoff.

Conversely, the expected profits for the risk-neutral monopolist insurer are greater without a forecast. Recall from equation 2 that a monopolist's expected profits are also proportional to the variance of the loss. Since the variance without a forecast, σ^2, is greater than the mean of the conditional variance, $\pi\sigma_1^2 + (1 - \pi)\sigma_2^2$, by the amount $(\pi - \pi^2)(\mu_1 - \mu_2)^2$, the change in the monopolist's expected profit from the use of public forecasts is

$$-\frac{\kappa}{4}\left(\pi - \pi^2\right)\left(\mu_1 - \mu_2\right)^2 \leq 0 \; .$$

If he restricts the consumers choice to a full insurance contract, the decline in profits is twice as large. Monopoly profits in a market without private information are reduced

by the use of forecast information. Another way to summarize this result is simply to note that better information reduces the consumers perceived risk on average. Since the monopolist's profits derive from the consumers aversion to this risk, these profits are also reduced.

F.   Competitive Equilibrium with Low Probability Bias

As we saw above, in a competitive equilibrium with common beliefs, neither the insurer's profits nor the consumers utility are changed by forecasting the state of nature which determines the parameters of the risk process covered by the insurance contract. It seems reasonable to suppose, however, that insurers and their customers may have different beliefs or private information about the risk process. One possibility is that the parameters governing the risk processes, and any forecast information, are public information, but some consumers of insurance have a low probability bias. That is, they underweight the probability of the less likely state of nature. Kunreuther et. al. (1978) and McClelland, Schulze and Coursey (1993) find evidence in laboratory experiments and field surveys that some buyers of insurance behave as though their subjective probability of infrequent, high value risks is zero.

To demonstrate the effects of such a bias, we suppose that the mean loss for state 1 (the *catastrophe* state) is much greater than for state 2 (the *normal* state): $\mu_1 \gg \mu_2$. We further suppose, as seems likely, that $\sigma_1^2 \gg \sigma_2^2$. The probability of a catastrophe state occurring, π, is assumed to be low. We represent the low probability bias by assuming there is some fraction of the consumers for whom the subjective value of π

is zero.[7] For the remaining, "unbiased" fraction the belief about π is the same as that held by the insurer. Finally, we stipulate that the insurer may not price discriminate between different customers. Maximizing utility with respect to quantity yields a competitive insurance premium $\lambda = \hat{\mu}$, where $\hat{\mu}$ denotes the belief about the mean loss shared by the insurer and unbiased consumers. That is, in a competitive market, insurers are constrained to offer insurance at actuarially fair rates, based on their own assessment of the expected loss. In the absence of a forecast, this is the unconditional mean loss μ defined previously in equation 4. With a forecast of state s, $\hat{\mu}$ becomes the conditional mean μ_s, where $s \in \{0,1\}$, as before. Thus, the change in the insurers' expected profits in a competitive market from the ability to forecast is zero, as it was under common beliefs.

The quantity of insurance demanded at these prices by the unbiased consumer is 1, as before, and the addition of forecasts does not change their utility. The behavior of the unbiased consumer is the same under either a full-insurance or variable-insurance model.

If the biased consumer is constrained to choose between no and full insurance, they will not purchase insurance in the absence of forecasts if $\mu > \mu_2 + \frac{\kappa}{2} \sigma_2^2$. Since $\mu_1 > \mu$, they will not purchase insurance when a catastrophe state is forecast either. They will only purchase insurance when a normal state is forecast, since then premium will

---

[7]The results here are analogous to what we would obtain if we instead modeled one consumer whose subjective probability of a catastrophe state was allowed to vary between zero and π.

equal the biased consumers expected value of a loss, μ_2. From the point of view of the

biased consumer, the increase in expected utility due to the forecasts is the probability

weighted sum of his utility in the two forecast states minus his utility before the

forecasts were available, given his belief that the catastrophe state will not occur:

$$\pi\left(X-\mu_2-\tfrac{\kappa}{2}\sigma_2^2\right)+(1-\pi)(X-\mu_2)-\left(X-\mu_2-\tfrac{\kappa}{2}\sigma_2^2\right).$$

This simplifies very readily to

$$(1-\pi)\frac{\kappa}{2}\sigma_2^2 \geq 0. \tag{8}$$

Of course, an individual who discounts the probability of a rare event may, when

confronted with a prediction that it may be imminent, re-evaluate their assessment of

the risk. If the availability of extended forecasts convinces the homeowner to prepare

for a rare event they originally thought impossible, then their benefit from the forecast

will be greater. Should a rare, potentially costly event be forecast, however, and result

in small damages (recall that the variance of the loss in the disaster state is large), the

homeowner could as easily be confirmed in their original bias. Thus, we need not

assume that the persistence of an effective low-probability bias is incompatible with the

availability of accurate forecasts.

If we relax our assumption that $\mu > \mu_2 + \tfrac{\kappa}{2}\sigma_2^2$, then the biased consumer will

purchase insurance when no forecasts are available. If it still is true that $\mu_1 > \mu_2 + \tfrac{\kappa}{2}\sigma_2^2$,

then expected utility with forecasts is still greater than without them:

$$\pi\left(X-\mu_2-\tfrac{\kappa}{2}\sigma_2^2\right)+(1-\pi)(X-\mu_2) > (X-\mu)$$

or

$$\pi\left(\mu_1 - \mu_2 - \tfrac{\kappa}{2}\sigma_2^2\right) > 0 .$$

Finally, if $\mu_1 \leq \mu_2 + \tfrac{\kappa}{2}\sigma_2^2$, then the biased consumer fully insures at the actuarially fair rate in every period, and the introduction of forecasts has no effect on utility.

If the consumer is allowed to choose their optimal quantity of insurance, the result is similar, with the gain in expected utility from forecasts in the event that $\mu > \mu_2$ + $\kappa\sigma_2^2$ the same as in equation 8. For lesser values of μ the change in utility is generally less than under the full-insurance model.

G.   Monopoly with Low Probability Bias

We saw above that while consumers do not benefit from the introduction of forecasts in a competitive model with common risk beliefs, consumers with a low probability bias can benefit from the new information. Similarly, a monopolistic insurer--who's expected profits decline as a result of introducing accurate forecasts-- may also benefit from forecasts if some consumers hold low probability biases. As we noted above, information which reduces consumer's uncertainty about the risk they face reduces a monopolist insurer's profits. On the other hand, because of their subjective under-weighting of a low probability risks, some consumers may under-insure at prices the insurer is willing to accept. If forecast information induces them to purchase more insurance, the insurer's profits may be increased. Which effect dominates is a function of the share of the insurer's customers who hold these low probability biases.

Recall that under the full insurance model the monopolistic insurer maximizes profits by charging a premium that makes the consumer indifferent between holding

insurance and not holding insurance, which in our quadratic utility model is

$$\lambda^* = \mu + \tfrac{x}{2}\sigma^2.$$

Suppose that some percentage, call it m, of the insurer's customers act as though the state (s = 1) cannot occur. The most they will pay for a full insurance policy is

$$\lambda_2^* = \mu_2 + \tfrac{x}{2}\sigma_2^2,$$

which is less than λ^*. In the absence of a forecast, the insurer will only charge λ_2^* instead of λ^* if the profit from selling insurance policies to all the biased consumers is greater than the reduction in profits from charging the lower price to unbiased consumers:

$$m\left(\lambda_2^* - \mu\right) \geq (1-m)\left(\lambda^* - \lambda_2^*\right). \tag{9}$$

If this condition is violated, the insurer sells policies at the maximum premium the unbiased consumers will pay, and biased consumers will only purchase insurance when a normal state (s = 2) is forecast. The change in expected profit from forecasts for the monopolist will then be positive so long as

$$m > \frac{\left(\pi - \pi^2\right)\left(\mu_1 - \mu_2\right)^2}{\sigma^2 - \pi\sigma_1^2},$$

where the numerator is the variance of the conditional mean.

If instead the condition in (9) is satisfied for some m ≤ 1, then the change in the insurer's expected profits from the use of forecasts will be positive if

$$m < \pi\left(\frac{\lambda_1^* - \lambda_2^*}{\lambda_1^* - \mu_1}\right).$$

For larger percentages of biased consumers, the change in expected profits from forecasts is negative because the expected loss from forecasting catastrophe states is greater than the expected gain from forecasting normal states.

Finally, if m is so large that the insurer's expected profits when a catastrophe state is forecast are greater if he charges the premium λ_2^* (assuming $\mu_1 \leq \lambda_2^*$), then the insurer will charge this premium in every period, and the availability of forecast information will have no effect on his profits.

H. Private Information

In the preceding discussion we assumed that forecast information was publicly available. If an individual insurer has private forecast information which other market participants do not share, he may be able to generate increased profits. Of particular interest is whether his use of private information to maximize profits will produce clear signals which would allow other market participants to infer his risk beliefs. We will first consider the case of a monopolist who offers only full insurance contracts.

Suppose that the consumer has some prior experience with the risk they are insuring, so that the unconditional mean μ and variance σ^2 of the potential loss are public information. Then, given our utility function, a consumer will be willing to buy a full insurance policy if the premium is less than or equal to $\mu + \frac{s}{2}\sigma^2$. If the monopolist, conditional on private information, expects that the mean and variance of

next period's loss will be lower than average, his expected profits will be greatest if he can conceal this information and charge $\lambda = \mu + \frac{k}{2} \sigma^2$. Expected profits are then $\mu - \mu_2 + \frac{k}{2} \sigma^2$, where μ_2 is the conditional mean ($\mu_2 < \mu$).

If the insurer's private forecast anticipates a higher than average mean and variance of losses, profits will be greater if he can credibly transmit this information to the consumer:

$$\frac{k}{2} \sigma_1^2 > \mu - \mu_1 + \frac{k}{2} \sigma^2$$

where μ_1 and σ_1^2 are the conditional mean and variance ($\mu_1 > \mu$, $\sigma_1^2 > \sigma^2$). In fact, if $\mu_1 > \mu + \frac{k}{2} \sigma^2$, he will not be willing to provide insurance unless the consumer raises their estimate of the expected loss. Thus, in a market where the insurer has market power, private forecast information, and the ability to offer only full insurance, we would expect that his private information will only be made public when a forecast indicates a higher than average mean and variability of loss.

If the monopolistic insurer allows the consumer to choose intermediate quantities of insurance, then the profit maximizing premium when a lower than normal mean loss is forecast is

$$\lambda = \frac{\mu + \mu_2}{2} + \frac{k}{2} \sigma^2.$$

Note that this is lower than the premium that would be charged in the absence of a forecast. Myopic profit maximization in this case could send a signal that the conditional expected loss was lower than average. Similarly, if the conditional mean

and variance of the loss is forecast to be greater than average, the insurers profits are greater if he can share this information with the consumer.

In an otherwise competitive insurance market, if one insurer has access to forecast information, they may be able to earn positive profits. If the insurer expects a lower than average loss, they can earn positive expected profits by charging the market price (the unconditional mean). If a small reduction in price increases their market share, profits will be greater. A reduction in price does not necessarily signal the insurer's risk beliefs to other market participants if the premium is calculated as the expected loss plus a fixed cost: $\mu Q + c$. As Karl Borch (1990) notes, it will not in general be possible for others to identify which part of the fixed cost c is due to administrative costs and which part is due to any risk premium. In this case, the risk premium can be interpreted as profit for a risk neutral insurer. If the change in premium is not too large, the insurer may be able to capture greater market share and earn positive expected profits without signaling his private risk information. Similar to the case of the monopolist, the competitive insurer will seek to withdraw coverage or make their information public if a forecast indicates higher than average expected losses.

If the forecast information is widely held among insurers in a competitive market, but is initially unknown to the consumers, we would expect that insurers will still be constrained to offer insurance at its actuarially fair value. The result for consumers would be the same as in section D.

I.   Regulation

If competitive constraints drive premiums down toward the expected loss, then insurers operating in markets with substantial correlated risks may face a significant risk of ruin (bankruptcy) in the event of a catastrophe state. If it is difficult for consumers to monitor an insurers risk of ruin directly, regulation of this risk may be necessary to maintain public confidence that claims will be paid in the event of a catastrophe. A common regulatory device is to establish a fixed reserve ratio for insurers. If an insurer covers a ratio Q of property at risk with total value X, then he is required to have on hand some minimal capital RXQ, where R is a fixed reserve ratio. The risk reserve must typically be held in low risk, relatively liquid assets, so the insurer incurs some opportunity cost for each dollar in the reserve RXQ. Let ξ denote the return on capital that the insurer would otherwise earn over and above the return to capital held in the low risk, liquid asset. Then the competitive insurer will only be willing to insure if the premium is at least as great as the expected loss plus his opportunity cost of complying with the regulation; the competitive premium will be λ = μ + ξRX.

If full insurance policies only are offered, then consumers purchase insurance so long as the cost of the regulation is lower than $\frac{\kappa}{2} \sigma^2$, as in the case of a monopolistic insurer. Consumers' utility is not affected by the use of forecast information. This is a natural result of full insurance--the insured consumer's face no risk--and the fact that changes in the conditional mean do not affect the cost of regulatory compliance in this

model. Consumers will not choose full insurance if they have the option, but will instead choose to retain part of the risk themselves. Then the use of forecast information will increase their expected utility so long as either $\sigma_1^2 \neq \sigma_2^2$ or $\mu_1 \neq \mu_2$.

The benefits to fully insured consumers of the introduction of accurate forecast information are zero in this model because the insurers capital costs do not change with the risk he assumes, and of the inflexible regulatory constraint. The regulatory constraint does not provide opportunities to use improved information to reduce the costs of reducing the risk to consumers from a potential default by the insurer. A regulation fixing a minimum risk of ruin, rather than a fixed reserve ration, gets around this difficulty. In practice, such a regulation is rarely employed due to the difficulty in calculating an insurer's risk of ruin.[8]

Under competition, insurers in our model are constrained to price insurance at the expected value of the risk. And yet, we sometimes observe greater risk premiums in capital markets. We can readily imagine that insurers in especially risky lines of business might face higher capital expenses which they pass on to their customers. Then we would expect climate forecasts to reduce costs under both the variable or the full insurance models. In general, however, competition in these markets will drive premiums down toward the expected value of the risk, and thus to reduce the potential value of forecast information.

In the fixed risk of ruin regulatory scheme, firms set the size of their risk reserve

---

[8]Borch 1990. -

to satisfy

$$\mathbf{Pr}\,\mathrm{ob}\Big(LQ > Y_{min}Q\Big) = \psi \ ,$$

where QY_{min} represents the minimum secure assets the insurer must have on hand to

ensure that the risk of non-payment of any claims does not exceed the probability ψ.

The competitive premium will be $\lambda = \mu Q + \xi F^{-1}(1 - \psi)$, where $F^{-1}(1 - \psi)$ is the same as

the inverse survival function of the probability ψ. Under the full insurance model, the

change in the consumers expected utility from the use of forecasts is now

$$\xi\Big(F^{-1}\Big(1- \psi\,\big|\mu,\sigma^{2}\Big) - \pi F^{-1}\Big(1- \psi\,\big|\mu_1,\sigma_1^2\Big) - (1-\pi)F^{-1}\Big(1- \psi\,\big|\mu_2,\sigma_2^2\Big)\Big) \ ,$$

where $F^{-1}(1 - \psi|\mu,\sigma^2)$ is the inverse survival function of ψ conditional on mean μ and

variance σ^2. The exact change in expected utility depends on the functional form

specified for the loss distribution $F(L)$, but in general if $\sigma^2 > \pi\sigma_1^2 + (1 - \pi)\sigma_2^2$ we expect

that

$$F^{-1}(1 - \psi|\mu,\sigma^2) > \pi\,F^{-1}(1 - \psi|\mu_1,\sigma_1^2) + (1 - \pi)F^{-1}(1 - \psi|\mu_2,\sigma_2^2).$$

These models indicate that flexible regulatory constraints that allow insurers to

adjust costs to reflect transient changes in expected risk yield the greatest benefits to

consumers from the introduction of extended climate forecasts. Note that we model the

insurer as putting his own capital at risk. If instead we assume that the insurer has none

and borrows from a competitive capital market, the reinsurer (lender) will require

payment $\mu Q + \xi YQ$ in a competitive market, where YQ is the amount of reinsurance

cover provided. The amount of Y is determined by the regulatory constraint, so that in

the examples above QY is either $F^{-1}(1 - \psi)$ or RXQ. ξ is now the market rate of return,

and the results of this analysis are otherwise unchanged.

J.   Conclusion

Better extended-term climate forecasts can be of considerable benefit to consumers in a competitive insurance market. For these benefits to be realized, the improved information must be widely shared among insurers, and must reduce the cost of regulations intended to limit the risk of insurer default. Inflexible solvency regulations, such as fixed reserve ratios, result in reduced gains from forecasts under the variable insurance model and zero gains under the full insurance model. In the absence of the ability to accurately forecast parameters of interest to the property catastrophe insurance business, it would not be surprising if regulatory structures were established lacking such flexibility. Their persistence in an era when it may be possible to accurately forecast catastrophe risks would limit the value of the improved information to consumers.

Similarly, Rees, Gravelle and Wambach find that if consumers are fully informed regarding the insurer's risk of insolvency, then under certain conditions competitive insurers will always provide enough capital to insure solvency. Such a market is analogous to the case of a flexible regulatory regime like the fixed risk of ruin case we consider above. In both cases insurers have the flexibility to exploit forecast information to reduce the capital costs associated with limiting their risk of insolvency. Where it is difficult to calculate an insurer's risk of ruin, neither case is likely to occur. Without solvency regulations or fully informed consumers, competition in insurance

markets drives the premium down toward the expected loss and thus the effect of forecast information on the quantity or average cost of insurance is reduced or eliminated.

An insurer in a competitive market may be able to generate positive expected profits from privately held forecast information when a forecast indicates that expected losses will be lower than average if this information can be concealed. If other market participants cannot discern between differences in an insurer's administrative costs and risk premium, then the informed insurer may also be able to increase market share at the expense of less-well-informed competitors without signaling the content of his private information.

Monopolistic insurers, likewise, may increase their expected profits through the use of private forecast information. Their ability to maximize profits without signaling their private information to other market participants is greater if they can limit consumers to a choice between a full insurance contract or no contract. When consumers can choose to retain some share of the risk they face, myopic profit maximization by the insurer may send clear price signals regarding their private information. In either the competitive or the monopolistic case, insurers with private information will wish to reveal that information when the forecast mean and variance of the consumer's risk process are greater than for the unconditional expectation, and to conceal their information when the forecast mean and variance are lower.

When some consumers hold a bias which leads them to act as though the

probability of a rare catastrophe state is zero, the use of forecast information may provide benefits to the monopolistic insurer that outweigh the reduction in profit that would otherwise result. This situation arises when the profits from selling insurance to biased consumers when the forecast state accords with their beliefs are greater than the reduction in profits from selling insurance to unbiased consumers. Profits from unbiased consumers are reduced because a forecast may reduce the uncertainty against which they are willing to insure, while profits from biased consumers may be increased if they would otherwise always under-insure in the absence of a forecast.

In each case considered here, under competitive constraints we see that the value to consumers of property catastrophe insurance of a perfect annual climate forecast derives from the indirect effects of a reduction in average risk. That is, it is through the specification of additional constraints on information, regulation, or competition that we generate value to the consumer from forecasts, rather than from equilibrium conditions in the basic model for a competitive insurance market.

K. Bibliography

Adams, Richard M., Kelly J. Bryant, Bruce A. McCarl, David M. Legler, James J. Obrien, Andrew R. Solow (1995) "Value of improved long range weather information," *Contemporary Economic Policy*. July, 13(3):10-18.

Anderson, Lee G. (1973) "The Economics of Extended-Term Weather Forecasting," *Monthly Weather Review*. February, 101 (2):115-125.

Babcock, Bruce A. (1990) "The Value of Weather Information in Market Equilibrium," *American Journal of Agricultural Economics*. February,72:63-72.

Blazenko, George (1986) "The Economics of Reinsurance," *The Journal of Risk and Insurance*, June, 53.2:258-277.

Borch, Karl H. (1990) *Economics of Insurance*, North Holland: New York.

Kunreuther, Howard, Ralph Ginsberg, Louis Miller, Philip Sagi, Paul Slovic, Bradley Borkan and Norman Katz (1978) *Disaster Insurance Protection: Public Policy Lessons*, John Whiley and Sons: New York.

Hoy, Michael. (1988) "Risk Management and the Value of Symmetric Information in Insurance Markets," *Economica*, 55.219:355-6.

Lave, Lester B. (1963) "The Value of Better Weather Information to the Raisin Industry," *Econometrica*. January-April, 31(1-2):151-164.

Ligon, James A. (1996) "Consumer Risk Perceptions and Information in Insurance Markets with Adverse Selection," *Geneva Papers on Risk and Insurance Theory*, 21.2:191-210.

Ligon, James A. and Paul D. Thistle (1996) "Information Asymmetry and Informational Incentives in Monopolistic Insurance Markets," *Journal of Risk and Insurance*, 63.3:434-457.

Luo Haiping, Jerry R. Skees and Mary A. Marchant (1994) "Weather Information and the Potential for Intertemporal Adverse Selection in Crop Insurance," *Review*

of Agricultural Economics, 16.3:441-51.

McClelland, Gary H., William D. Schulze and Don Coursey (1993) "Insurance for Low-Probability Hazards: A Bimodal Response to Unlikely Events," *Journal of Risk and Uncertainty*, 7:95-116.

McNew, Kevin. (1999) "The Value of El Niño Forecasts in Agricultural Commodity Markets: THe Case of U.S. Corn Storage," in Rodney F. Weiher, ed. *Improving El Niño Forecasting: The Potential Economic Benefits*. 41-42.

Nichols, Richard A. (1999) "Case Study of 1997-98 El Niño for Minnegasco" in Rodney F. Weiher, ed. *Improving El Niño Forecasting: The Potential Economic Benefits*. 27-28.

Rees, Ray, Hugh Gravelle and Achim Wambach (1999) "Regulation of Insurance Markets," *The Geneva Papers on Risk and Insurance Theory*, 24:55-68.

Solow, Andrew R., Richard F. Adams, Kelly J. Bryant, David M. Legler, James J. O'Brien, Bruce A. McCarl, William Nayda and Rodney Weiher (1998) "The Value of Improved ENSO Prediction to U.S. Agriculture," *Climatic Change*. 39:47-60.