A Mutualized Risk Market with Endogenous Prices, with Application to U.S. Landfalling Hurricanes

Robert Meyer
Michael Horowitz
Daniel Wilks
Alfredo Bequillard
Kenneth Horowitz

December 2008

Working Paper # 2008-12-08

Risk Management and Decision Processes Center
The Wharton School, University of Pennsylvania
3730 Walnut Street, Jon Huntsman Hall, Suite 500
Philadelphia, PA, 19104
USA
Phone: 215-898-4589
Fax: 215-573-2130
http://opim.wharton.upenn.edu/risk/

CITATION AND REPRODUCTION

THE WHARTON RISK MANAGEMENT AND DECISION PROCESSES CENTER

Since its creation 25 years ago, the mission of the Wharton Risk Management and Decision Processes Center has been to carry out a program of basic and applied research to promote effective corporate and public policies for low-probability events with potentially catastrophic consequences. The Risk Center has focused on natural and technological hazards through the integration of risk assessment, risk perception and risk financing with risk management strategies. After 9/11, research activities extended also to national and international security issues (e.g., terrorism risk insurance markets, protection of critical infrastructure, global security).

Building on the disciplines of economics, decision sciences, finance, insurance, marketing and psychology, the Center's research program has been oriented around descriptive and prescriptive analyses. Descriptive research focuses on how individuals and organizations interact and make decisions regarding the management of risk under existing institutional arrangements. Prescriptive analyses propose ways that individuals and organizations, both private and governmental, can make better decisions regarding risk. The Center supports and undertakes field and experimental studies of risk and uncertainty to better understand the linkage between descriptive and prescriptive approaches under various regulatory and market conditions. Risk Center research investigates the effectiveness of strategies such as risk communication, information sharing, incentive systems, insurance and regulation.

The Center is also concerned with training decision makers and actively engaging multiple viewpoints, including the expertise of top-level representatives in the world from industry, government, international organizations, interest groups and academics through its research and policy publications and through sponsored seminars, roundtables and forums.

More information is available at http://opim.wharton.upenn.edu/risk

A Mutualized Risk Market with Endogenous Prices, with Application to U.S. Landfalling Hurricanes

Robert Meyer

Michael Horowitz

Daniel Wilks

Alfredo Bequillard

Kenneth Horowitz*

Preliminary Draft: do not cite without permission from the authors.

* Robert Meyer is the Warren Johnson Professor of Marketing at the University of Miami and Co-Director of the Wharton Center for Risk Management and Decision Processes, University of Pennsylvania. Michael Horowitz is an undergraduate at the Wharton School of Business, University of Pennsylvania. Daniel Wilks is a Professor of Atmospheric Science at Cornell University. Alfredo Bequillard is a Financial Mathematician, formerly of Lehman Brothers Inc. Kenneth Horowitz is a principal at Weather Risk Solutions LLC.

A Mutualized Risk Market with Endogenous Prices, with Application to U.S. Landfalling Hurricanes

Abstract

This paper describes a new approach to commodity-option trading that allows market participants to hedge against the risk that a coastal county or region in the eastern United States will be hit by a hurricane during a given calendar year. Called *Hurricane Risk Landfall Options* or HuRLOs, the market differs from traditional weather derivatives in that market participants need not find a willing counterparty to take the opposite side of an option contract. Instead, the prices are set by an adaptive control algorithm that reflects previous purchasing decisions of other market participants. The empirical properties of this market are examined using data from an experimental market in which participants experience a hypothetical hurricane season during which they are allowed to buy HuRLOs in a primary market as well as sell and buy HuRLOs in a secondary market. The data show that aggregate market prices quickly converge to rational (efficient) levels among market participants after limited amounts of trading experience, and that the operation of the markets are free from such potential biases as investment procrastination or false-alarm effects. Some systematic anomalies are observed in the trading of individual HuRLOs, the most notable being that prices for the “No Landfalls” HuRLO display a "boomerang" bias, where it is overvalued by market participants immediately after a storm threat passes at specific locations.

1. Introduction

A major challenge facing residents of many of the world's coastal regions is how to manage the risk of property losses due to tropical storms and hurricanes. While storm threats have been an omnipresent fact of life in such areas for generations, the rapid growth of coastal populations and property values has served to magnify their financial impact (Cutter, et al., 2007; Pielke et al., 2008). In 2004 and 2005, for example, hurricane landfalls in the United States set world-wide records for both insured and uninsured losses, with seven of the storms (Katrina being the most devastating) being among the 20 most costly insurance catastrophes of all time (Wharton Risk and Decision Processes Center 2009). More recently, the destruction caused by Hurricanes Gustav and Ike in Louisiana and Texas has served as a vivid reminder that catastrophic losses may be much more the norm in the future than the exception.

One of the most publicized—and politically charged—consequences of these escalating losses has been the gradual degradation of the quality, and increase in the cost, of windstorm insurance available to coastal residents (Derrig, et al., 2008). The reasons for the insurance crisis are complex, but at the core lies a conflict between two seemingly incompatible forces: the natural reluctance of insurers to underwrite insurance policies for properties for which the probability of a catastrophic loss is ambiguous, and constraints on the prices that firms can charge residents to insure against these risks due both to regulatory controls and limits to affordability (Wharton Risk and Decision Processes Center 2009). The consequence is that in many states—particularly Florida—there has been a deterioration of the traditional private windstorm insurance market, with many major insurers declining to write new policies or greatly limiting coverage. In coastal Florida, for example, hurricane-event windstorm policies are typically written with substantial minimum percentage deductibles that effectively cause homeowners to be self-insuring against all but the most extreme hurricane events.

The purpose of this paper is to describe a new approach to managing hurricane risk using *Hurricane Risk Landfall Options* or HuRLOs (Weather Risk Solutions 2008a). HuRLOs are a newly-available class of commodity options that allow participants to hedge against the risk that a selected coastal county or region on the United States Atlantic and Gulf coast will be first hit by the next hurricane to make landfall in a calendar year. The market differs from other existing weather derivatives (e.g., CME Group 2008; Jewson and Caballero, 2003; Zeng, 2000) by not requiring market participants to find a willing counterparty to take the opposite side of an option contract. Instead, prices for options are set by an adaptive control algorithm that reflects previous purchasing decisions of other market participants. The premia collected from purchases are aggregated into a mutualized risk pool (MRP) that is then distributed among market participants at the time of a qualifying landfall event. At settlement participants who hold options for the county or region that is the site of landfall (or hold a "No Landfalls" option if none occurs) receive a pro-rata share of the MRP.

Because HuRLO markets are both simple in structure and can guarantee participants minimum payouts in the event of a landfall, they may offer an attractive alternative means to address the needs of the insurance and reinsurance industries with respect to hedging potential casualty losses from hurricanes. This mutualized risk market could also help consumers and businesses fill in gaps in their insurance coverage, including deductibles.

As a new instrument, however, the degree to which HuRLO markets can fulfill this theoretical potential is uncertain.   Unknown empirical properties of the market include such aspects as how rapidly the size of the mutualized risk pool might be expected to build over time as a season progresses, how buying behavior will respond to objective changes in the risk of hurricane landfalls, and, perhaps most importantly, the degree to which the market will act to correct local inefficiencies in the price of landfall options that arise as storms approach land.

In this paper we take a first step toward addressing these questions by reporting the empirical properties of HuRLO markets observed in an experimental market in which participants have the opportunity to purchase HuRLOs during the course of a simulated hurricane season. In the hypothetical season participants are exposed to a series of storms whose past and forecast future motions are designed to probe the sensitivity of the markets to events that might induce purchasing anomalies in real markets, such as storm threats to highly salient coastal counties that are frequently associated with hurricane landfalls (e.g., those in south Florida), plus a series of near-miss events.

Our primary finding is that participants make purchasing decisions that display, with experience, high levels of aggregate efficiency. Early in the simulated season participants act as if they over-value the ability to buy HuRLOs, something that produces inflated prices in coastal areas with high base-rate likelihoods of storm landfalls, as well as when storms first develop. These biases quickly vanish with trading experience, however, with aggregate market prices closely tracking objective likelihoods of storm landfalls. The overall effect is the production of a large MRP that would have provided substantial capital to participants who held options for the area that experienced the first storm landfall.

We organize our presentation in three sections. We first provide a detailed description of HuRLOs and the primary and secondary markets on which they trade, focusing on the primary market, and settlement procedures. We then offer initial evidence on the empirical features of HuRLO trading using data from an experimental market. We conclude with a discussion of the findings and their implications for the use of mutualized risk mechanisms for financial risk management in hurricane and other hazard contexts.

2. The HuRLO Markets

2a. Overview

HuRLOs are commodity options that allow market participants to hedge against the risk that one of seventy-eight coastal counties or regions on the United States Atlantic and Gulf coast will be first hit by the next hurricane to make U.S. landfall in a calendar year (Weather Risk Solutions 2008a). The first HuRLO markets were launched in October 2008 on an exempt board of trade operated by CME Alternative Marketplace Inc., a subsidiary of CME Group[1]. The initial launch offered two Series: one for the location of a first hurricane landfall (if any), and one for a second (Weather Risk Solutions,2008b). In addition to the seventy-eight geographical HuRLOs, participants could also purchase a "No Landfalls" option that would be exercised if no subsequent hurricane made landfall in the United States through the balance of the calendar year.

[1] HuRLOs trade on an electronic trading platform and are cleared by CME Clearing, a subsidiary of CME Group.

When a hurricane makes landfall first in one of the seventy-eight coastal counties or regions[2], that landfall triggers automatic exercise and settlement of the applicable options. The premiums collected from HuRLO purchases are aggregated into the MRP for the applicable HuRLO Series, to be allocated among holders of the HuRLOs for the coastal county or region where a hurricane makes first landfall, or holders of "No Landfalls" HuRLOs if no next hurricane makes landfall in the current calendar year. The settlement is in proportion to the number of options for the correct event that are held by each market participant.

Unlike traditional weather derivatives, market participants need not find a willing counterparty to take the opposite side of the contract. HuRLO prices in the primary market are based on an adaptive control algorithm and reflect the purchasing decisions of other market participants (described below in Section 2b). Market participants who have purchased HuRLOs may sell them to other market participants in a conventional (i.e. bilateral) secondary market. The secondary market thus provides a mechanism for participants holding positions to transfer them to other participants with differing beliefs, at agreed upon prices. Short sales are not permitted.

2b. How HuRLOs are Priced in the Primary Market

HuRLO prices in a given Series are based on the calculations of an adaptive control algorithm that tracks changes in market-based probabilities π^i_t for each outcome (coastal segment, or "No Landfalls"), *i*, based on the buying behavior of the market participants in each HuRLO Series up to time *t*. The price P^i_t for outcome *i* at any given

[2] Most of these 78 outcomes correspond to hurricane landfall in individual coastal counties, but multiple counties are aggregated into a few larger regions on the mid- and northern-Atlantic U.S. coastline because of lower hurricane frequencies there.

time t is proportional to the corresponding market probability, including a time-value-of-money adjustment that equalizes prices for purchases that may be made months apart:

$$P_t^i = \pi_t^i \, c \, \exp[rj/365] \quad . \qquad (1)$$

Here c is a constant defining the overall magnitude for prices and settlements in the HuRLO markets ("par" value, taken here to be $1,000), r is an annualized interest rate, and j is the number of days since the market opened for the current season.

The key element of price determination in Equation 1 is the HuRLO market (i.e., pricing) probability π^i_t for each outcome, which reflects the market-based probability for each of the 79 HuRLO outcomes in a particular Series. Following the purchase of each additional new HuRLO in a particular HuRLO Series, the prices for all 79 HuRLOs in that Series are recalculated to reflect a larger market probability for the outcome, i, that has just been purchased. Because the probabilities for all outcomes must sum to 1 in each Series, market probabilities for the 78 outcomes that were not purchased in the most recent transaction are decreased proportionally.

These probability (and, therefore, pricing) calculations are repeated after the purchase of each additional HuRLO. So, for example, if a block of 100 HuRLOs is ordered for a county i, the repricing calculations are iterated 100 times during the order-filling process. The result is that each of the 100 HuRLOs in the block costs slightly more than the previous HuRLO purchased in the block — the demand for HuRLOs in county i progressively increases the price of subsequent HuRLOs in this county while depressing prices for the options in the remaining 78 outcomes in the particular HuRLO Series. Prices move more quickly for a given level of buying activity when the market is comparatively thin (the MRP is relatively small), and move more slowly when the market is comparatively well developed (the MRP is relatively large).

The probability updating procedure used to determine pricing is a variant of the Robbins-Monro stochastic approximation algorithm (Kushner and Yin, 2003), although the specific equations used in the HuRLO markets represent a new member of this class of algorithms. The pricing probabilities converge to the consensus of market participants' probabilities for the outcomes, under various assumptions about investor behavior (Bequillard, 2008). That is, the adaptive control algorithm for updating the pricing probabilities automatically *learns* investors' probabilities for the outcomes in response to their collective actions in the market.

2c. Market Seeding

HuRLO markets are "seeded" with an initial stake in the MRP for each HuRLO Series. The seeding institution receives an equal number of HuRLOs in each of the 79 outcomes in each HuRLO Series. Because initially there are not yet market-based probabilities for the 79 outcomes, prices for the initial stakes are allocated according to historical (i.e., climatological) probabilities, π^i_0, reflecting the historical risks. The result is that the initial number of HuRLOs in each outcome is equal, and given by

$$N_0^i = \frac{MRP_0}{c} \quad (2)$$

For example, if the MRP for one HuRLO Series is seeded initially with MRP_0 = \$1,000,000, and c = \$1,000, then that HuRLO market begins with N^i_0 = 1000 HuRLOs in each outcome i. This allocation follows from Equation 1, with $j = 0$.

2d. Miscellaneous Special Rules

Minimum Settlement. If buying becomes very concentrated in one or a few HuRLO outcomes before the MRP has accumulated sufficient liquidity, settlements for HuRLOs in that outcome may become excessively diluted. Dilution is not allowed to exceed a percentage of par, which in 2008 was set at 50%. If an additional HuRLO purchase would drive the potential settlement below this level, the price for that HuRLO is increased sufficiently to cover the minimum settlement.

Minimum Pricing Probability. To ensure numerical stability, pricing probabilities are not allowed to fall below 0.0001. Accordingly, the minimum HuRLO price is 0.0001 c, or \$0.10 if c = \$1,000.

3. Empirical Properties: Evidence from an Experimental Market

3a. Motivation

In theory, HuRLOs provide a straightforward mechanism by which at-risk (and other) market participants can hedge against hurricane risk. Although market participants would be unlikely to have expertise in predicting whether and where hurricanes will make landfall, the trading platform offers tools that should, in principle, encourage market prices to efficiently mirror objective landfall odds. The platform, for example, displays hurricane Forecast Advisories and other data imported in real time from the U.S. National Hurricane Center (NHC), and market participants all have access to the purchasing decisions being made by other investors, such as the total number of HuRLOs purchased in each outcome.

Yet, there is still the possibility that market distortions may occur if participants believe that they hold private information about the behavior of hurricanes—or private beliefs about other participants' beliefs about hurricanes—that can be exploited to earn excess profits.   If biases were to arise in the HuRLO markets, what might they look like? While there has been no prior empirical work that directly answers this question, hypotheses might be drawn from past research on how individuals make decisions whether to invest in protective action in advance of hurricanes and other natural hazards (e.g., Meyer 2006). This work suggests that trading in HuRLOs could potentially be influenced by four psychological factors that could either distort valuations or suppress overall purchasing levels: procrastination biases, distorted beliefs about probabilities, speculative bubbles, and false-alarm effects.

Procrastination would be manifested by a tendency for participants to delay purchases of HuRLOs until storms actually threaten specific coastal areas, a bias that would act to diminish the overall size of the MRP and lend greater uncertainty to potential settlement values (O'Donoghue and Rabin 1999; 2001). The basis for this possibility is simple: at the outset of a storm season participants face a landscape where the probability of a landfall in any one coastal county or region is small, and objective odds remain unchanged until the first hurricane forms and begins to threaten land. A participant might thus see little downside in delaying the initial investment, either in the recognition that there would be little opportunity cost in delaying the decision if pre-storm prices are largely static, or in the hope that buying opportunities may emerge after seeing the purchases made by other participants.

A second source of concern is that market prices may be distorted by biased beliefs about landfall probabilities. While the market website provides participants with objective guidance on the likelihood of storm landfalls, participants face the challenge of translating this probabilistic information into discrete purchasing decisions for some or all of seventy-nine HuRLO outcomes in each available Series. Prior work on subjective perceptions of probability (e.g., Slovic 2000) suggests that these decisions could be influenced by two related biases: availability and information cascades (Kahneman and Tversky 1973; Bikhchandani, Hirschleifer, and Welsh 1992).

An availability bias would be the tendency for traders to be influenced by the mental ease with which a hurricane landfall could be imagined at a particular location (e.g., Chandler, et. al 1999; Folkes 1988; Kahneman and Tversky 1973). As an example, widespread news coverage of hurricane landfalls in Louisiana and Texas might cause participants to overvalue HuRLOs in those locations compared to those for which storm hits come less readily to mind, such as the coastal Northeast. Even if such availability biases do not arise, probability-related distortions could still occur if participants use allocation heuristics that focus purchases on only that subset of locations with the highest objective landfall odds—something that would have the aggregate effect of overly inflating prices for HuRLOs that have comparatively high objective probabilities, and yielding under-investment in outcomes with comparatively smaller chances of occurring.

In a related way, prices for specific HuRLOs might also be subject to information cascades that give rise to speculative bubbles (e.g., Bikhchandani, Hirschleifer, and Welsh 1992; Smith, *et. al* 1988). An information cascade arises when investors base valuations on the actions of other investors rather than objective information about the underlying value of the commodity. This could arise in HuRLO markets if participants believe that excessive prices being paid for certain landfall options reflect superior private information about hurricanes held by other participants — information beyond that held and distributed by the National Hurricane Center and other public forecasting services. Such biases might seem particularly at risk to arise when major hurricanes threaten highly-populated areas, when objective information about likely storm landfalls may be overwhelmed by less informed sources such as frenzied media coverage and rumors.

Finally, at the other extreme, HuRLO purchases might also be subject to the opposite effect of false-alarm or “cry wolf” biases, where the *absence* of a landfall that was previously thought to be likely could serve to suppress later purchases (e.g., Atwood and Major 1998; Breznitz 1984). Such an effect could manifest itself either in a reluctance to make early purchases of HuRLOs (inducing under-pricing in advance of landfalls) or, more seriously, a reluctance to participate in the market at all, thus suppressing MRPs for later Series.

3b. Method

We tested the degree to which HuRLO markets might exhibit such characteristics using data from an experimental market. Over the past thirty years experimental markets have emerged as a major tool used in both economics and behavioral finance to test the likely behavioral properties of new market instruments prior to launch, such as auction design (e.g., Plott 1997) and pricing mechanisms (e.g., Daniel, Hirshleifer, and Subrahmanyam 1998). Although experimental markets have the downside of simplifying the scale and features of real-world markets (e.g., levels of compensation are far lower), they have the advantage of allowing controlled stress tests of new instruments that would be impossible in field settings. Moreover, the literature suggests that financially-motivated experimental market participants often respond to market structures and incentives in a way that closely parallels traders in real markets, something that has spurred the increased use of laboratory settings as a means to test the empirical viability of new market products prior to launch (see, e.g., Kagel and Roth 1995; Smith 1988).

To undertake such a controlled test of the behavior of HuRLO markets, seventy-eight graduate and undergraduate students from a major Northeastern university school of business were recruited to participate in a session of simulated trading. Among the student participants 65% were finance majors and 46% had experience trading other securities. The test was conducted in the Northeast because it offered access to a participant pool that would be knowledgeable about hurricanes but have limited direct recent experience that might bias purchasing behavior in the simulation. For example, we wished to avoid purchasing biases that might arise if participants believed that the simulated hurricane was mimicking the path of a storm they had just experienced, or one with which they had great familiarity. The participant sample was thus designed to mimic a pool of financially-knowledgeable participants with neutral priors about the likely landfall location of hurricanes.

The experimental markets were conducted over a three-hour period in a large computer classroom, with each participant being paid $50 as compensation for participating. To provide a formal incentive for making decisions in a way that maximized profits, participants were told that at the end of the simulation they had a chance to win one of three additional $100 prizes based on their performance. These bonuses were awarded by means of a computerized lottery in which the chance of winning was directly proportional to realized earnings.

Experimental Procedure. Upon arriving in the computer lab, participants were told that they would be participating in an experimental market where their earnings were based on their ability to predict the landfall behavior of hurricanes. Each participant was endowed with a simulated trading account worth $750,000, an amount that pilot work suggested would be sufficient to allow participants to purchase a reasonably large number of HuRLOs over the course of the simulation, but not so large as to induce frivolous purchases or sales.

During the first 30 minutes, participants were given an introduction to the simulation, and allowed to experiment with the web-based trading interface. The interface was the same interface used in the actual HuRLO trading markets (http://www.weatherrisksolutions.com), consisting of three primary graphical elements, which we illustrate in Figure 1:

1. *A map page*, which displayed the location, current market-based and objective landfall probability forecasts (Wilks et al., 2009) for each HuRLO coastal county or region, current market prices, and storm locations and forecast movements during threats. The page also provides a short-cut window to the purchase interface;
2. *A primary-market trading page*, which provided a tabular summary of trading information relevant to each HuRLO, the participants' current holdings, and purchase shortcuts; and
3. *A secondary-market trading page* (not shown) which provided a tabular summary of HuRLOs that had been made available for sale by other participants.

After completing the warm-up phase the real market exercise began, proceeding for 2-and-a-half hours, interrupted by a 30-minute snack break to counter task fatigue. A post-experiment debriefing survey indicated that participants found the task to be an interesting one, without revealing difficulties either in navigating the interface or understanding the structure of the HuRLO markets.

The simulated hurricane season. The simulated hurricane season was structured in three time blocks. The first was a "preseason" block corresponding to the months of February through May, in which there was no hurricane activity. The passage of each month was compressed to four minutes of real time, during which participants were free to purchase and offer for sale HuRLOs at their discretion. The second block corresponded to the month of June, and contained the appearance of the first simulated tropical cyclone, named "Aisha". Beginning with the formation of the tropical depression that was to be Aisha, time in the simulation was slowed to 4-minute long "days," each containing the release of a new hypothetical NHC Forecast Advisory describing Aisha's strength, current movement, and anticipated future movement. Concurrently, newly-computed forecast probabilities for all of the 79 outcomes were displayed on the main trading-page table. The third block corresponded to the month of July, and contained the appearance of the second storm, named "Babar". During the course of Babar's existence simulated time was again incremented in 4-minute long days.

Storm information was conveyed in a manner meant to realistically simulate what would be typically provided to participants on the web-based interface from the NHC. Specifically, on each storm day the map screen would display a text advisory that described the storm's intensity, motion, and the location of any watches or warnings. Below that information, the storm was displayed on a hurricane tracking chart that illustrated its current location, most likely future movement, and a "cone of uncertainty" surrounding this forecast (see, e.g., Broad, *et. al* 2007).

As shown in Figure 2, Aisha and Babar were designed to mimic two common hurricane paths in the Atlantic Basin, as well as potentially provoke biases in purchasing behavior. Aisha formed in the southwest Atlantic and the first mock NHC Forecast Advisories hinted that it could be a threat to the outer banks of North Carolina and then, later, New England (the Cape Cod area). These threats, however, turned out to be false alarms, with the Aisha eventually making landfall in Nova Scotia as a tropical storm.

Babar, in contrast, was a major hurricane that initially appeared to be a serious threat to south Florida and the Keys, prompting hurricane watches in those areas. The storm passed south of these counties, however (a second false alarm), before eventually making landfall in southern Texas (Kenedy County). The landfall in Kenedy County, however, had a surprise element to it, as the day before the storm was bearing west-south-west on a heading that pointed to northern Mexico.

At the start of the simulation participants were not given any information about the number of storms they might see during the course of the session, and were explicitly told that there was no guarantee that a storm would make landfall; i.e., that the "No Landfalls" HuRLO was a legitimate possibility.

4. Results

4a. Overall behavior of the MRP

To provide an initial look at the overall functioning of the primary market, in Figure 3 we plot the cumulative growth of the MRP over time. In the figure the cumulative size of the pool is measured on the vertical axis, and the width of each colored-shaded bar segment reflects the number of primary-market purchases occurring in each trading period (wider bars indicate more participants making purchases). Contrary to the fear that speculators might delay purchasing until storms first appear, the data show that participants took active advantage of buying opportunities during the pre-season, with approximately 25% of the cumulative total MRP being built during this time. After this initial buying wave subsequent increases were driven by the appearance of the two storms and changes in the apparent immediacy of storm threats. For example, after Aisha is named and threatens New England (the second panel in Figure 3) there is a consistent growth in the pool, but the number and value of the purchases rapidly diminishes when the threat passes. We then see a similar—and more dramatic—surge at the end of the simulation when one of only a few outcomes became almost certain: Babar making landfall in one of the counties in south Texas, or there being no (U.S.) landfall at all.

A clearer view of how storm events affected purchasing behavior is provided in Figure 4, which plots period-to-period variation in the numbers of HuRLOs purchased (trading volume) as a function of time and storm news. During the pre-season we observe comparatively slow rates of purchasing in February, perhaps reflecting a desire by participants to "wait and see" the purchases of others. Purchasing then rapidly increases until showing a slight decrease in the final pre-season trading month (May), perhaps reflecting a desire among participants to keep funds in reserve for use during the actual hurricane season. During the storm events of June and July purchasing takes on a much different character: here trading volume is similar in overall volume to that observed in the pre-season, but closely tracks changing news about the likelihood that either Aisha or Babar will make landfall. The most notable anomaly is the high rate of purchasing that is observed with the first storm announcement (Aisha)—an exuberance that may have reflected exaggerated confidence among participants that early forecasts that the storm might affect North Carolina would be born out.

4b. The Efficiency of Investments

Of central interest in the experiment is not only the degree to which participants would actively purchase HuRLOs in the primary market but also the degree to which these purchases would act efficiently, such that market prices would reflect objective information about the likelihood of storm landfalls. As we noted earlier, a concern was that while participants might purchase in the primary market (which they did), the market value of HuRLOs over time might display such biased features as a tendency to overvalue HuRLOs for locations more typically associated with hurricane landfalls, or false-alarm effects, where purchases in landfalls that do not pay off suppress a desire to make subsequent purchases.

To examine whether such biases existed, in Figure 5 we plot time series of volume-weighted average objective probabilities compared to corresponding volume-weighted average market-based probabilities. As detailed above, the market-based probabilities are a function of trading volume and, in turn, are used to generate prices. The figure shows what might be seen as a surprising result: market participants appear to be extremely rapid learners, with aggregate market-based probabilities quickly converging to objective forecast-based likelihoods over time. More specifically, the data suggest that the primary market was subject to two separate bouts of naïve overvaluation that were not repeated: one at the very outset of the simulation when participants had the first opportunity to trade, and one that coincided with the first (but ultimately unrealized) U.S. hurricane landfall threat (Aisha).

To provide a high-resolution look at this learning effect, in Figure 6 we plot market-based versus objective forecast probabilities (Wilks et al., 2009) over time for two counties that have both strong historical associations with hurricane landfalls and high climatological base rates: Terrebonne Parish in coastal Louisiana and Monroe County, Florida, which includes the Keys. Consistent with an availability bias, at the start of the simulation we see a tendency for positions in these locations to be over-valued, with market prices being higher than those which would be rationally supported based on actuarial base rates. But in both cases this initial bias rapidly vanishes, with market-based probabilities closely tracking changes in the objective probabilities for the balance of the simulation.

Yet, it is important to emphasize that while market prices behaved efficiently on average, particularly as participant experience increased, examples of sustained inefficiency for individual HuRLOs remained. Two examples are provided in Figure 7, which plots market-based versus objective probabilities for the "No Landfalls" HuRLO, and Figure 8, which plots these values for Kenedy County, Texas, site of Hurricane Babar's ultimate landfall. Figure 7 suggests that participants over-valued the “No Landfalls” HuRLO at the outset of the simulation, and then again at three different later points in the simulation: when Aisha’s threat to North Carolina and New England vanishes in June (periods 11-12), at the outset of July before Babar's threat becomes apparent (periods (14-15), and finally when the storm’s brief threat to the Florida Keys passes (periods 20-21). One might characterize valuations of the "No Landfalls"

HuRLO, therefore, as displaying something of a “boomerang” bias, where participants over-correct for the easing of the threat to one area by overvaluing the likelihood that there will never be a landfall.

Kenedy County, Texas, offers an example of sustained underinvestment. Because of its lack of salience relative to higher base-rate counties in Louisiana, Florida, and North Carolina, participants act as if they undervalue this and other relatively low-probability counties for almost the entire duration of the simulation. Only at the very end---when landfall there became a plausible event—did market-based probabilities rise to meet normative values, but in the terminal period remain less than those prescribed by forecast-based probabilities.

4c. The Secondary Market

A final interest is the behavior of the secondary market during the course of the simulation. Analysis of the secondary market is informative because it provides evidence of the degree to which the excessive purchases observed in the primary market during the early stages were seen as “mistakes” by investors, versus conscious speculative investments; i.e., acquiring options at low prices with the goal of “flipping" them on the secondary market at higher prices.

In Figure 9 we plot the number of sell offers and buys in the secondary market, and the number of purchases in the primary market. The temporal pattern of sell offers—and the absence of buys—strongly suggests that the secondary market was primarily used by participants as a means to correct over-buying mistakes in the primary market, although usually unsuccessfully.

Specifically, the data shows a surge in sell offers during the second half of the Aisha's lifespan, when it would have been clear to participants that they had been overly exuberant in their desire to acquire HuRLOs both in the pre-season and, more overtly, when Aisha first developed and began to threaten the east coast. Consistent with this, there were few buyers willing to assume these possibly over-bought positions in the secondary market. It is only at the very end of the simulation—when Babar was about to make landfall in Texas—that we see evidence of the secondary market working efficiently. Here the high market prices of HuRLOs in Kenedy and neighboring counties drove buyers to the secondary market, where they found sellers who were willing to part with previously-acquired options at prices that they saw as profitable.

5. Discussion

In recent years we have witnessed a rapid growth in the development of new financial products designed to help firms and communities manage the risks of weather-related hazards. These include industry catastrophe bonds, loss warrants, and weather-related derivatives (Wharton Risk and Decision Processes Center 2009; CME group, 2007). While these products have played a useful role in the suite of risk-diversification tools available to insurers and re-insurers, significant challenges remain in overcoming the crisis of insurance coverage and costs faced by private homeowners and businesses in hurricane-prone areas.

The purpose of this paper was to describe HuRLOs – a new commodity option product that could help address these concerns. HuRLOs depart from traditional catastrophe bonds, loss warrants, and weather-related derivatives by offering market participants the opportunity to hedge or speculate on hurricane risk without needing to find a willing counterparty to take the opposite side of the contract. In the HuRLO markets, the risk of a hurricane landfall in one area is mutualized and underwritten by market participants who buy HuRLOs for other hurricane landfall areas or the "No Landfalls" HuRLO. Pricing in the primary market is based on an adaptive control algorithm and reflects the purchasing decisions of all market participants to date.

The likely empirical properties of the HuRLO markets were explored using data from a controlled experimental market. The study was motivated by the possibility that the markets, while easily understood, might nevertheless be subject to a number of biases, including a slow build up of the MRP due to hesitancy among participants to purchase HuRLOs early in the season before storms develop, and for HuRLO prices to be distorted by perceptual biases regarding objective landfall likelihoods. The experimental data, however, gave strong reason to believe that such biases, if they arise at all in a real-world implementation of these markets, will likely be transient. Participants exploited the opportunity to acquire HuRLOs at lower prices by buying prior to the formation of storms, and with limited trading experience made purchases at prices that, for the most part, were rationally consistent with objective probabilities of hurricane landfalls in different locations.

It should be emphasized that these findings were gathered from markets that likely differ from real-world markets in terms of both participants (primarily college students) and realism (here the hurricanes were purely hypothetical). The fact that high levels of efficiency were achieved so quickly despite the inexperience of participants speaks well for the ease with which the trading and pricing mechanism will likely be understood in real-world settings. Likewise, that this efficiency was achieved despite the participants having limited personal experience with hurricanes is also encouraging.

At the same time, the possibility remains that purchases in the face of real-world hurricanes might exhibit biases that we found little evidence for, such as the possibility that prices in certain landfall locations will be subject to speculative bubbles. While the laboratory simulation realistically simulated the kind of information residents typically receive from the U.S. National Hurricane Center, it made no attempt to simulate the frenzied media and social conditions that typically accompany major storm threats—conditions that could induce speculative bubbles and distorted behavior. An important next step in this empirical research program is thus to probe the performance of the HuRLO markets under such stresses.

Another aspect of the HuRLO market that could be the focus of future study is the degree to which the market valuations of landfall probabilities by themselves could be a valuable new source of information about the likelihood of storm impacts for use by residents and emergency planners in threatened areas. This possibility is motivated by the large volume of work documenting ability of prediction markets composed of heterogeneous traders to accurately forecast a variety of real-world events, ranging from election outcomes to new-product successes (e.g., Wolfers and Zitzewitz 2004). The possibility of using prediction markets to forecast hurricane movements was recently explored by Kelly, Letson, and Solis (2008), who found that a prediction market for hurricane landfall forecasts composed of meteorologists did a better job of predicting actual landfalls than any of the individual models used as basis for forecasts made by the NHC. It would be interesting to see whether similar predictive abilities would be observed in much larger HuRLO markets composed of traders with a diverse range of knowledge about storms.

Finally, although introduced in the context of hurricanes, it is important to emphasize that the mutualized risk pools used in the HuRLO market described herein is one that could be extended to a wide range of natural hazards such as earthquakes and floods. The one obvious boundary is that these markets would be inappropriate in settings where participants could affect the outcome, such as for wildfires or other hazards where humans can play a contributory role.

References

Atwood, L. E. and A. M. Major (1998). Exploring the "Cry Wolf" Hypothesis. *International Journal of Mass Emergencies and Disasters* 16:279-3.

Bequillard, A. (2008) A robust market-making algorithm for nonbinary prediction markets. Working paper, Lehman Bros., 15 pp.

Bikhchandani, Sushil, Hirschleifer, David, and Ivo Welsh (1992) "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy*, Vol. 100, No.5, pp. 992-1026,

Breznitz, S (1984). *Cry Wolf: The Psychology of False Alarms*. Hillsdale, NJ: Lawrence Erlbaum Associates Publishers

Broad, K., T. Leiserowitz, J. Weinkle, M. Steketee. (2007) Misinterpretations of the "cone of uncertainty" in Florida during the 2004 hurricane season. *Bulletin of the American Meteorological Society*. Vol 88, Issue 5 pp. 651-667.

Chandler, Carla C.; Greening, Leilani; Robison, Leslie J.; Stoppelbein, Laura (1999), "It can't happen to me . . . or can it? Conditional base rates affect subjective probability judgments", *Journal of Experimental Psychology: Applied.* 5(4), 361-378.

Cutter, Susan L., Jonson, Laurie A. Finch, Christina and Melissa Berry (2007). "The U.S. Hurricane Coasts: Increasingly Vulnerable?", *Environment*, 49 (7), 8-22.

CME Group (2008), *Alternative Investments: Hurricane Futures and Options*, http://www.cmegroup.com/trading/weather/files/hurricane_fact_card.pdf

Daniel, K., Hirshleifer, D, and A. Subrahmanyam (1998). "Investor Psychology and Security Market Under- and Overreactions". *Journal of Finance* **53** (6):

Derrig, Richard A., Fishman, Jay S., Grace, Martin, and Joan Schmit (2008), "Catastrophe Management in a Changing World: The Case of Hurricanes", *Risk Management and Insurance Review*, 55 (11), 269-280.

Folkes, Valerie (1988), "The Availability Heuristic and Perceived Risk", *Journal of Consumer Research*, 15(1), 13-23.

Jewson, S., and R. Caballero (2003) The use of weather forecasts in the pricing of weather derivatives. *Meteorological Applications*, 10: 377-389.

Kagel, J. H., and A. E. Roth (eds.) (1995), *Handbook of Experimental Economics*, Princeton, NJ: Princeton University Press.

Kahneman, Daniel, and Amos Tversky (1973), "The Psychology of Prediction", Psychological Review, 89(4), 237-251.

Kelly, David L. Letson, David, and Daniel Solis (2008), "Evolution of Subjective Hurricane Risk Perceptions: A Bayesian Approach," Working Paper, School of Business administration, University of Miami.

Kushner, H. J., and G. G. Yin (2003) *Stochastic Approximation and Recursive Algorithms and Applications.* Springer-Verlag, 2003.

Kunreuther, H., Meyer, R., Zeckhauser, R., et. al**, (2002)** High-Stakes Decision Making: Normative, Descriptive and Prescriptive Considerations, *Marketing Letters.* Vol. 13(3), 259-268.

Meyer, Robert (2006), "Why we Under Prepare for Hazards", in Ronald J. Daniels, Donald F. Kettl, and Howard Kunreuther (eds), *On Risk and Disaster: Lessons from Hurricane Katrina*, University of Pennsylvania Press, pp. 153-174.

O Donoghue, Ted, and Matthew Rabin (1999), "Doing It Now or Later," *American Economic Review*, 89(1), 103-124.

O Donoghue, Ted, and Matthew Rabin (2001), "Choice and Procrastination", *Quarterly Journal of Economics,* 116, 121-160.

Pielke, R.A., Jr., J. Gratz, C.W. Landsea, D. Collins, M. Saunders, and R. Musulin (2008) Normalized hurricane damages in the United States: 1900-2005. *Natural Hazards Review*, 9: 29-42.

Plott, Charles R. (1997), "Laboratory experimental testbeds: Application to the PCS auction", Journal of Economics and Management Strategy", 6(3), 605-38.

Slovic P.(2000) *Perception of Risk.* London: Earthscan; 2000.

Smith, Vernon L.; Suchanek, Gerry L.; Williams, Arlington W. (1988). "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets". *Econometrica* 56 (5): 1119–1151.

Weather Risk Solutions (2008a), "Hurricane Risk Landfall Options: Questions and Answers", http://www.weatherrisksolutions.com.faq.php.

Weather Risk Solutions (2008b), Press Release, "Weather Risk Solutions™ to Launch Hurricane Risk Landfall Options - HuRLOs™".

Wharton Risk and Decision Process Center (2009), *At War with the Weather: Managing Large-Scale Risks in a New Era of Catastrophes*. Cambridge, MA, MIT Press.

Wilks, D.S., C.J. Neumann, and M.B. Lawrence (2009) Statistical extension of the National Hurricane Center 5-day forecasts. *Weather and Forecasting*, submitted.

Wolfers, Justin, and Eric Zitzewitz (2004), "Prediction Markets", *Journal of Economic Perspectives*, 18(2), 107-26.

Zeng, L. (2000) Weather derivatives and weather insurance: concept, application, and analysis. *Bulletin of the American Meteorological Society*, 81: 2075-2082.

Figure 1: The HuRLO market trading interfaces. Top panel is the graphical interface that displays current holdings, market landfall probabilities, and price, and, in the event of a storm, storm location and expected movement. The lower panel shows the corresponding tabular interface.

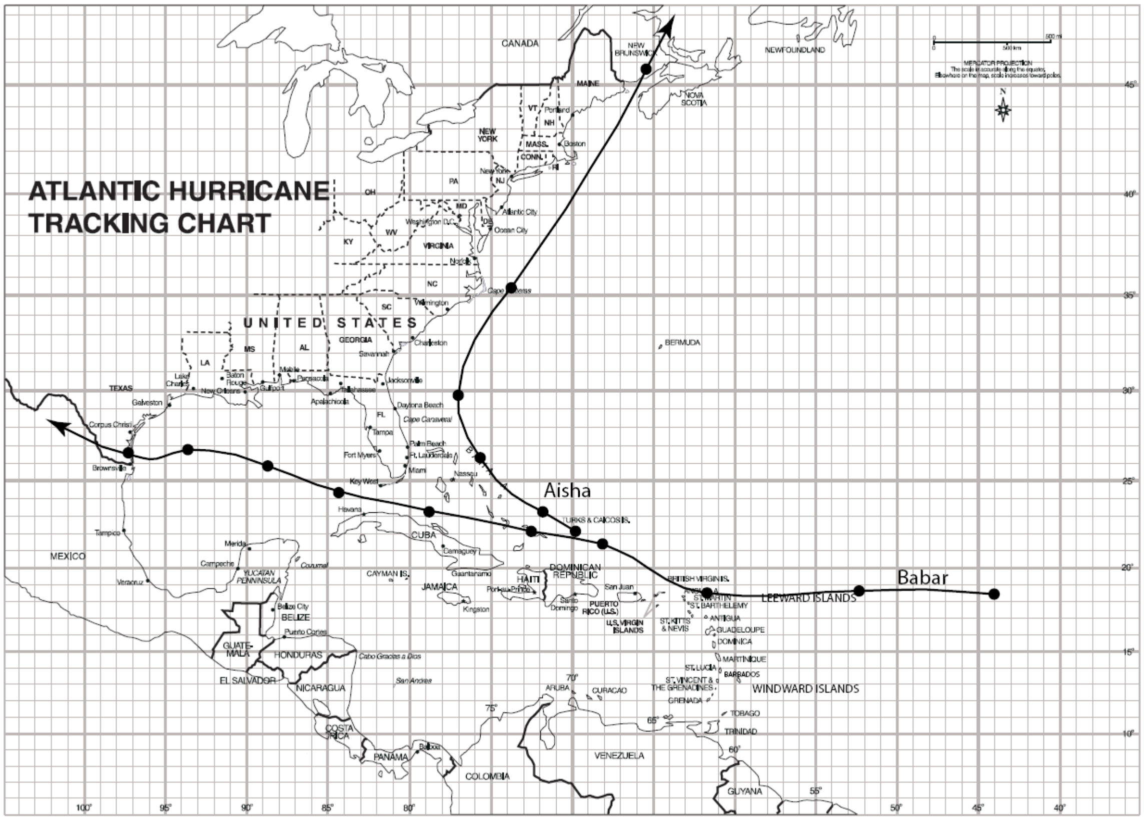

Figure 2: Tracks of the two simulated tropical cyclones

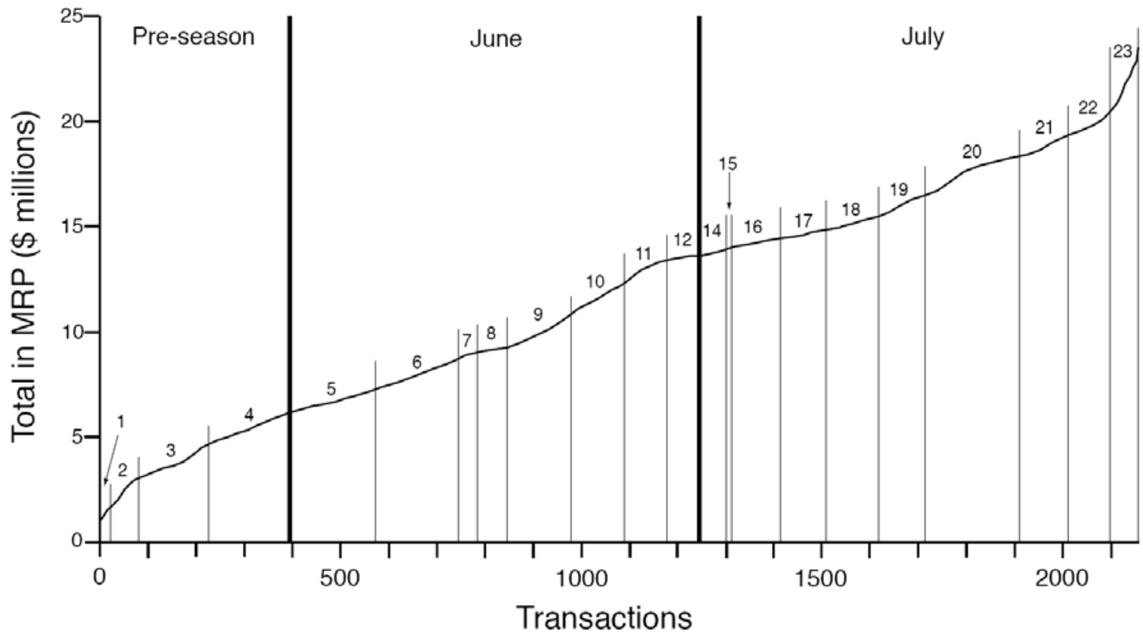

Figure 3: Growth of the MPR over time as a function of the number of options purchased. Each shade corresponds to a different 4-minute trading interval, corresponding to new months in the pre-season and individual days during storms 1 (Aisha) and 2 (Babar). The figure shows a rapid growth of the pool in the pre-season followed by surges of trading when storm landfalls seemed likely (particularly in the final trading day).

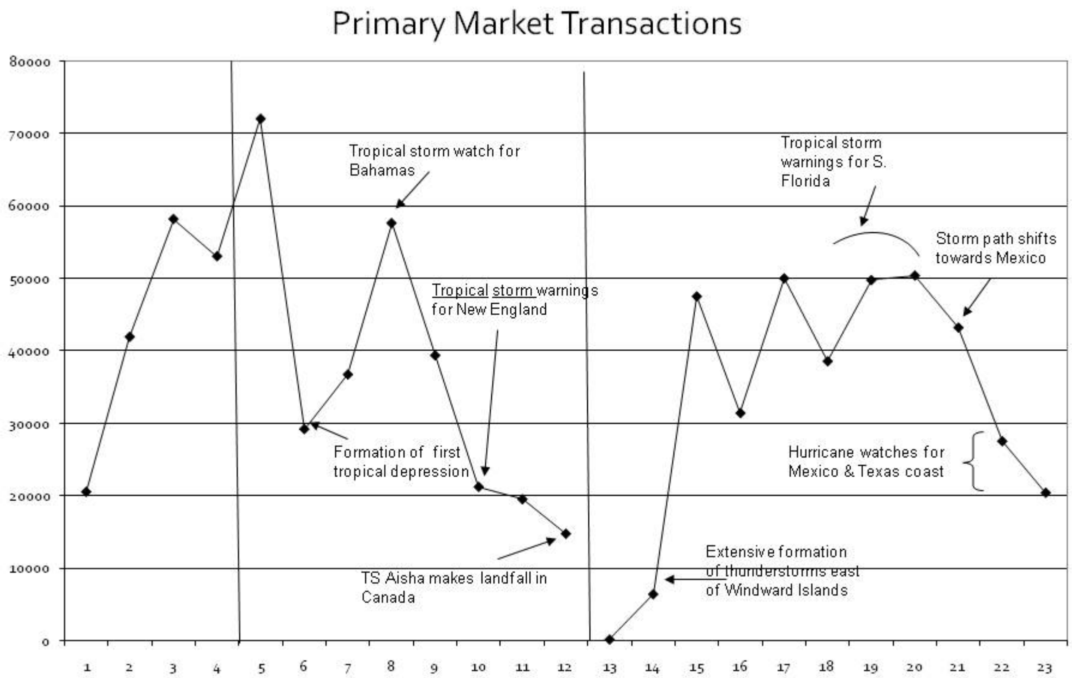

Figure 4: Plot of primary market trading volume over time as a function of storm events.

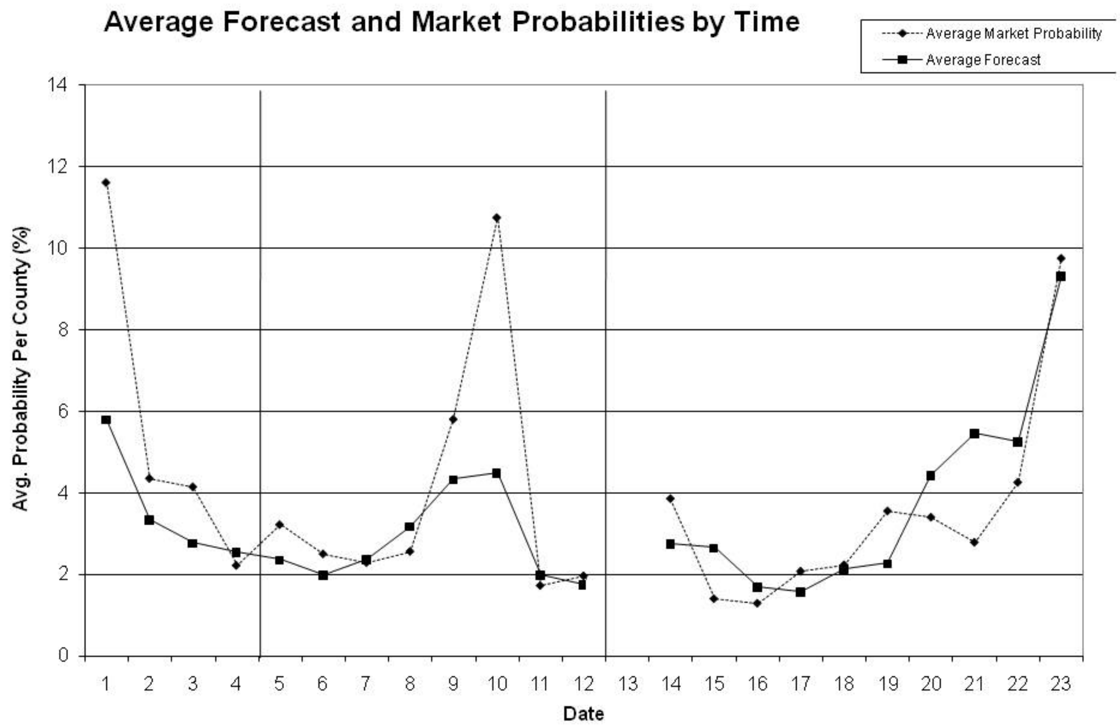

Figure 5: Plot of how volume-weighted Market Probabilities (corresponding to HuRLO prices through Equation 1) at the end of each period tracked changes in objective landfall probabilities over time. The figure shows a tendency to over-value options relative to the objective probabilities both at the start of the pre-season and given the first landfall threat (Aisha), but on-average convergence to objective probabilities thereafter.

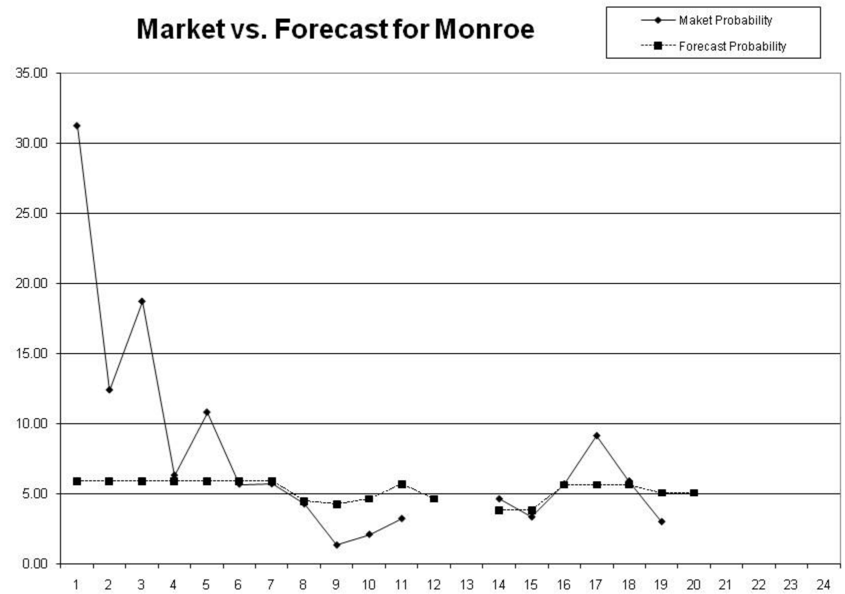

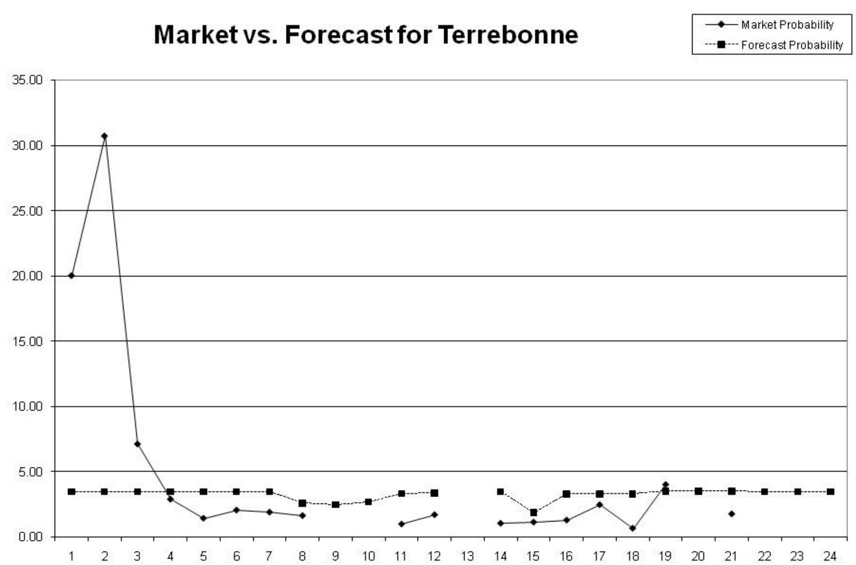

Figure 6: Plot of market versus objective probabilities at the end of each period in two heavily-traded options that show strong evidence of learning: Monroe County, FL (top) and Terrebonne, LA (bottom). Note that in the simulation Monroe was briefly placed under a hurricane watch due to the approach of Hurricane Babar (hence the heightened activity during the 16th-18th periods), whereas Terrebonne was never threatened. Gaps indicate periods in which no purchases were made.

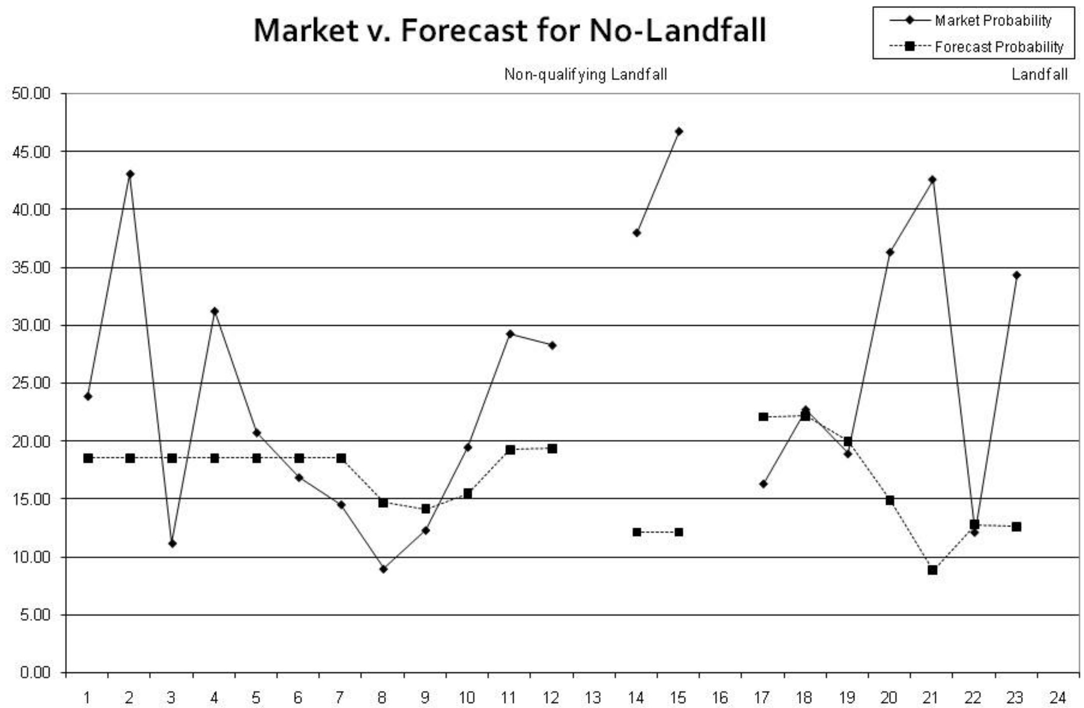

Figure 7: Plot of market versus objective probabilities at the end of each period for a the No-Landfalls option, which shows less evidence of learning: The figure shows excessive valuation of this option early in the simulation, then continued excessive valuation after each "near miss" event: when Aisha bypassed New England (periods 11-12) and when Babar by-passed South Florida (periods 20-21). Gaps indicate periods in which no purchases were made.

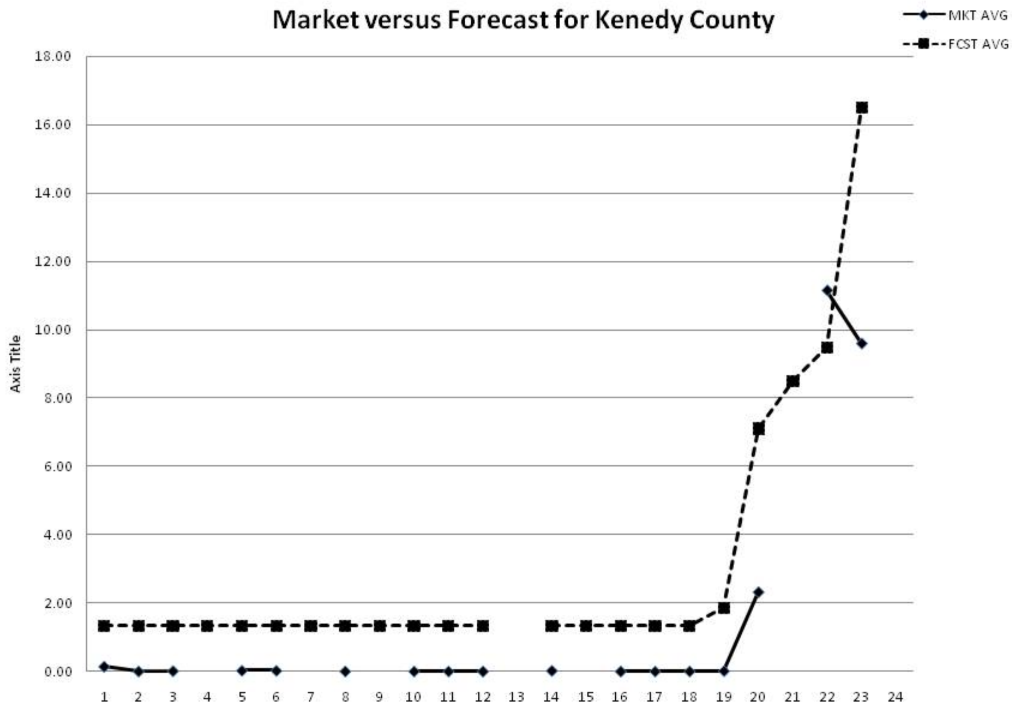

Figure 8: Plot of market versus objective probabilities at the end of each period for Kennedy County, Texas. The figure shows systematic under-investment through most of the simulation. Gaps indicate periods in which no purchases were made.

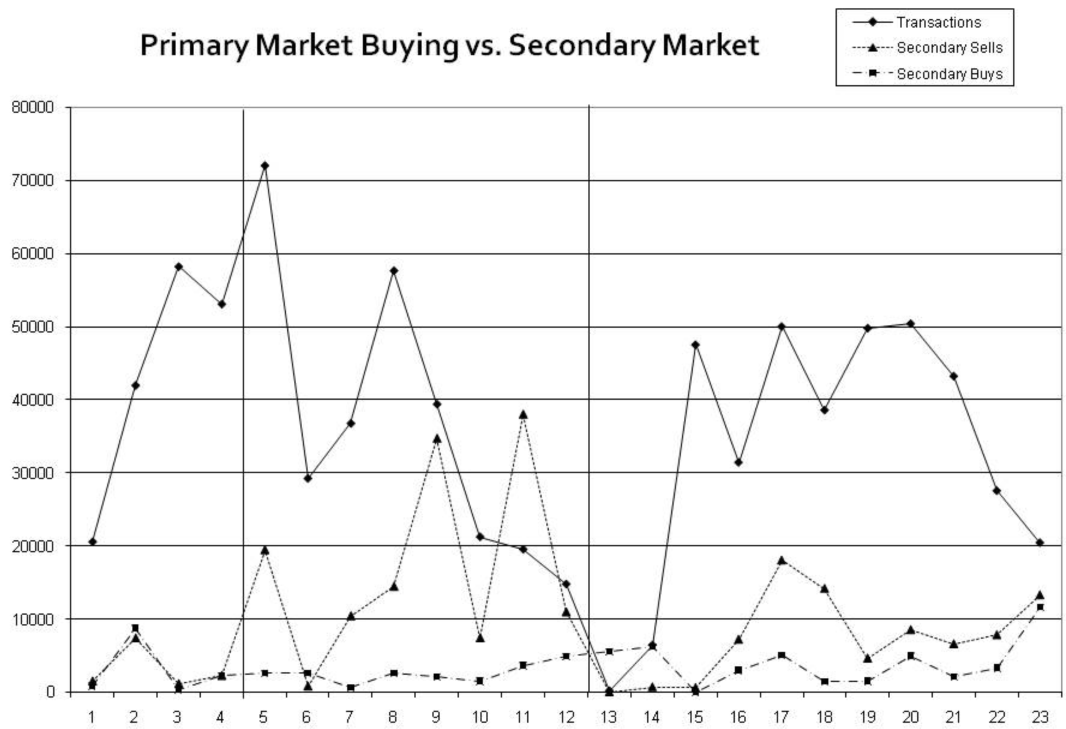

Figure 9: Plot of activity on the secondary market over time. The figure shows evidence of an active desire by participants to use the secondary market to sell many of the positions they secured during the exuberant buying period associated with the early stages of tropical storm Aisha. The figure also shows, however a reluctance of participants to acquire these positions.

.