# A Novel Financial Market for Mitigating Hurricane Risk. Part II: Empirical Validation

ROBERT J. MEYER

*The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania*

MICHAEL HOROWITZ

*Harvard Business School, Cambridge, Massachusetts*

DANIEL S. WILKS

*Department of Earth and Atmospheric Sciences, Cornell University, Ithaca, New York*

KENNETH A. HOROWITZ

*Weather Risk Solutions LLC, West Palm Beach, Florida*

(Manuscript received 17 May 2013, in final form 7 October 2013)

ABSTRACT

This paper explores the empirical features of a novel commodity option trading instrument described in the companion paper (Part I) that allows market participants to hedge against the risk that a coastal county or region in the eastern United States will experience a hurricane landfall. In this instrument investors can speculate on whether a landfall event will occur in any one of a number of coastal counties or regions, with option prices being determined by an adaptive control algorithm that reflects previous purchasing decisions of other market participants. In this paper, the authors report the results of an experiment designed to test the empirical robustness of this mechanism using data from traders buying landfall options over the course of a simulated hurricane season. In the experiment traders are given the opportunity to buy landfall options in the primary market as well as sell and buy options in a conventional bilateral secondary market. The data show that aggregate market prices quickly converge to rational (efficient) levels among market participants after limited amounts of trading experience. Some systematic anomalies are observed in the trading of options for individual outcomes, however, with the most notable being an initial tendency to overvalue landfall options that have the highest prior probabilities and for valuations of the ''No Landfall'' option to be inflated immediately after a storm threat passes without making landfall.

## 1. Introduction

The rising economic cost of tropical cyclones worldwide has brought increased calls for improved approaches to hurricane risk management (e.g., Kunreuther and Michel-Kerjan 2009). In a separate paper, Wilks and Horowitz (2014, hereafter Part I) describe one such new mechanism for managing hurricane risk that uses a novel financial market structure. Market participants buy ''landfall options'' that trigger settlements if a hurricane makes first landfall along a selected coastal county or region of the United States Atlantic or Gulf coasts in a calendar year. The structure differs from traditional weather-derivative markets[1] in that contracts are purchased from an exchange that requires only buyers, prices are automatically set based on purchasing volumes, and settlements are derived from the pool of payments made by all market participants. The market structure is simple and may offer an attractive alternative means to address the needs of individuals, businesses, and the insurance and reinsurance industries, with respect to hedging potential financial losses from hurricanes.

In this paper, we report findings about the empirical robustness of this new market structure using data on

*Corresponding author address:* Daniel S. Wilks, 1123 Bradfield Hall, Department of Earth and Atmospheric Sciences, Cornell University, Ithaca, NY 14853.
E-mail: dsw5@cornell.edu

[1] In a conventional weather-derivative market, buyers and sellers trade an asset whose value depends on whether or not a specific triggering event occurs, such as a freeze occurring before a certain date.

DOI: 10.1175/WCAS-D-13-00033.1

trading behavior observed in an experimental market. In this market laboratory traders were given the opportunity to buy landfall options in a primary market over the course of a simulated hurricane season, as well as sell and buy options in a conventional bilateral secondary market. Among our findings is strong evidence that prices in the aggregate quickly converge to those predicted if traders held unbiased beliefs about the probability of landfall in specific counties, albeit subject to a localized tendency to overvalue options with high likelihoods.

In section 2, we first briefly review the theoretical structure of the options market, and then describe the potential behavioral threats to the approach as a pragmatic tool for hurricane risk management. In section 3, we describe the experimental methodology, and in section 4 we report the research findings. We conclude, in section 5, with a discussion of the implications of the work for the general problem of financial risk management for low-probability, high-consequence events.

## 2. The HuRLO market and potential behavioral biases

The hurricane contracts described in more detail in Part I are commodity options and have been named Hurricane Risk Landfall Options (HuRLOs). Purchases of these contracts allow market participants to hedge against the risk that 1 of 78 coastal counties or regions on the United States Atlantic and Gulf coasts will be first hit by the next hurricane to make U.S. landfall in a calendar year (the official Atlantic "hurricane season" runs from June through November). The contracts are offered in multiple series, with Series 1 contracts pertaining to the first U.S. hurricane landfall in a given year, Series 2 pertaining to the second, and so on. In addition to the 78 explicit landfall areas, "No Landfall" HuRLOs are available in each series, which contracts pertain to the possibility that no (further) U.S. landfalling hurricanes will occur in the year to which the market pertains. For example, in years with a single U.S. landfalling hurricane, buyers of Series 2 (and higher) No Landfall contracts would be paid, but buyers of Series 1 No Landfall contracts would not.

When a hurricane makes landfall, first in 1 of the 78 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 a mutualized risk pool (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 Landfall 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 (Bequillard 2013, manuscript submitted to *Int. J. Theor. Appl. Finance*; Horowitz et al. 2013; Part I) that is a new variant of the Robbins–Monro stochastic approximation algorithm (Kushner and Yin 2003). These prices are proportional to probabilities for each of the outcomes, which converge to the consensus probabilities of market participants as revealed through their purchasing activity (Bequillard 2013, manuscript submitted to *Int. J. Theor. Appl. Finance*; Horowitz et al. 2013). 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.

In theory, HuRLOs provide a straightforward mechanism by which at-risk (and other) market participants can hedge against hurricane risk. Although the market participants would be unlikely to have expertise in predicting whether and where hurricanes will make landfall, they all have access to the purchasing decisions being made by other investors (such as the total number of HuRLOs purchased for each outcome) that, ideally, will rationally reflect the objective forecasts provided the National Hurricane Center (NHC).

Yet, there is still the possibility that market distortions may occur if participants believe that they hold private information about the behavior of a hurricane that can be exploited to earn excess profits. While there has been no prior empirical work that informs the kinds of biases that may arise when individual traders form beliefs about likely hurricane landfalls, hypotheses might nevertheless 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 2012). In particular, 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

[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.

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 essentially 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 prestorm 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. Although market participants would have access to objective guidance on the likelihood of storm landfalls provided by the NHC, participants face the challenge of translating this probabilistic information into discrete purchasing decisions for some or all of the 79 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 et al. 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 Georgia. Even if such availability biases do not arise, probability-related distortions could still occur if participants use allocation heuristics that focus purchases on 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 underinvestment in outcomes with comparatively smaller chances of occurring.

In a related way, prices for specific HuRLOs might also be subject to speculative bubbles where assets briefly trade at prices far in excess of their true value, only to quickly collapse (e.g., Bikhchandani et al. 1992; Smith et al. 1988). This could arise in HuRLO markets if participants believe that they see high prices being paid for certain landfall options and conclude that this is due to others having superior knowledge or information about where a hurricane is likely to make landfall—information beyond that held and distributed by the NHC 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). The locus of this effect could either be a reluctance to make early purchases of HuRLOs (inducing underpricing in advance of landfalls) or, more seriously, a reluctance to participate in the market at all, thus suppressing MRPs for later series.

## 3. Experimental market

### a. Description

We tested the degree to which HuRLO markets might exhibit the characteristic distortions outlined in section 2 using data from an experimental market. Over the past 30 yr, 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 et al. 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 ways that closely parallel 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 et al. 1988).

To undertake such a controlled test of the behavior of HuRLO markets, 78 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%

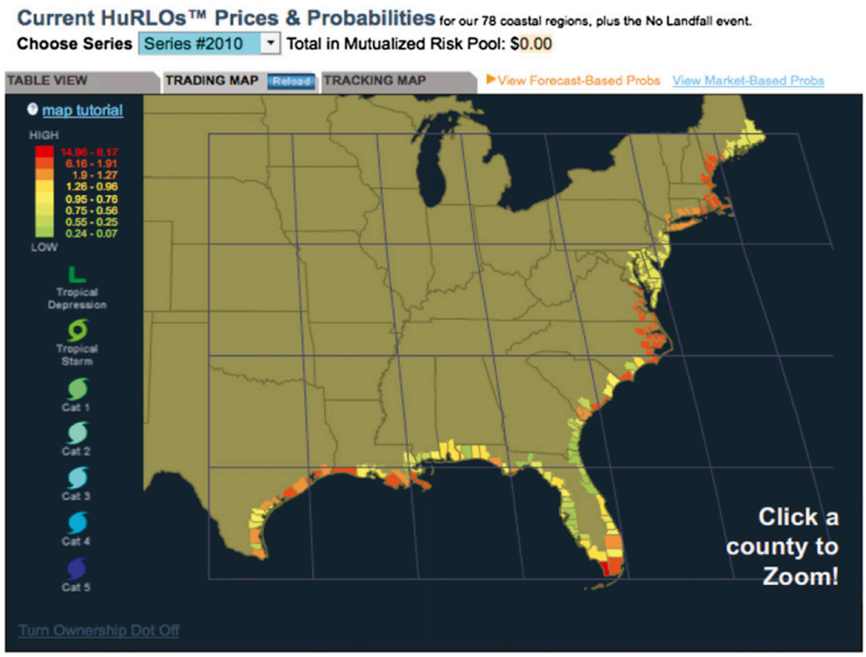


FIG. 1. Screen from the market interface (at www.weatherrisksolutions.com), showing HuRLO landfall areas, color-coded with climatological first landfall probabilities.

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 (the experimental markets were conducted before the 2012 landfall of Hurricane Sandy). 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 locations of hurricanes.

The experimental markets were conducted over a 3-h period in a large computer classroom, with each participant being paid $50 as compensation for participating. To provide an 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. We choose to incent participants using a probabilistic currency both because of the limited size of the available payoff pool and its precedent for use in experimental economics as a means for encouraging risk neutral trading (see, e.g., Berg et al. 1986; Selten et al. 1999). In the final section we will discuss the limitations posed by the compensation scheme, which was clearly much smaller than that which would be faced by real-world traders.

### b. Experimental procedure

Upon arriving in the computer laboratory, participants were told that they would be participating in an experimental market where their earnings would be based on their ability to predict the landfall behavior of hurricanes. Each participant was endowed with a simulated trading account worth $750 000 (U.S. dollars), 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 min, participants were given an introduction to the simulation and allowed to experiment with the web-based trading interface. The interface was developed by Weather Risk Solutions (www.weatherrisksolutions.com) for the possible future commercial implementation of the HuRLO market and consisted of three primary graphical elements:

1) a map page (Fig. 1), 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

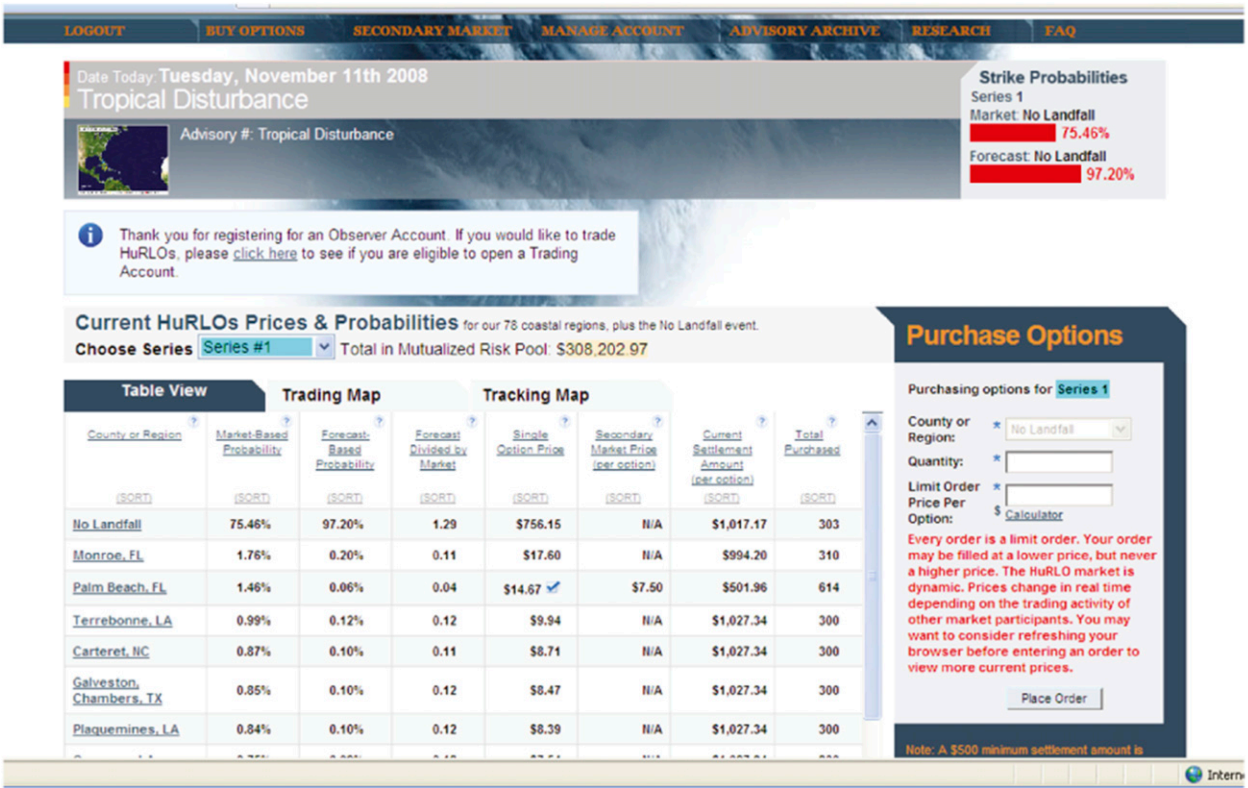


FIG. 2. Tabular market interface (from www.weatherrisksolutions.com), showing HuRLO landfall areas, respective prices, and related forecast and market information.

threats (the page also provides a shortcut window to the purchase interface);

2) a primary market trading page (Fig. 2), which provided a tabular summary of trading information relevant to each HuRLO, the participant's 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, and bid for purchase, by other participants.

After completing the short warm-up phase, a single session of the real market exercise began, proceeding for 2.5 h, interrupted by a 30-min snack break to counter task fatigue. A postexperiment 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.

### c. *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 4 min 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 become Aisha, time in the simulation was slowed to 4-min-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-min-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, based on real-time data 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

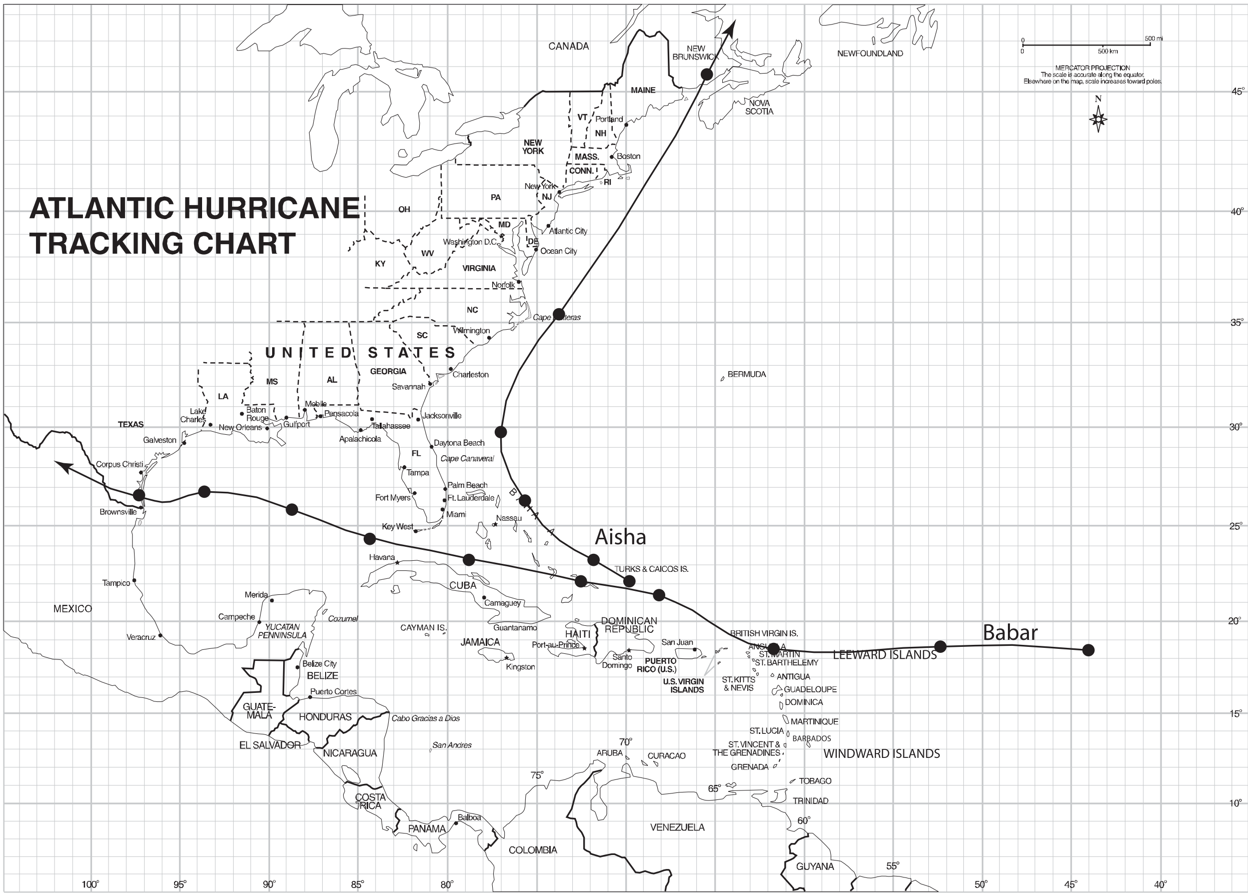


FIG. 3. Tracks of the two simulated tropical cyclones.

location of any watches or warnings. Below that information, the storm was displayed on a hurricane tracking chart that illustrated its current location, likely future movement, and the "cone of uncertainty" surrounding this forecast (Broad et al. 2007).

As shown in Fig. 3, 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 Aisha eventually making landfall in New Brunswick as a tropical storm.

Babar, in contrast, was a major hurricane that initially appeared to be a serious threat to south Florida and the Florida 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 landfall the storm was bearing west-southwest 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, that is, that the No Landfall HuRLO was a legitimate possibility.

## 4. Results

### *a. Overall behavior of the MRP*

To provide an initial look at the overall functioning of the primary market, in Fig. 4 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 segment reflects the number of primary market purchases occurring in each trading period (wider segments indicate more purchases). The data show that participants took active advantage of buying opportunities during the preseason, with approximately 25% of the cumulative total MRP being built during this

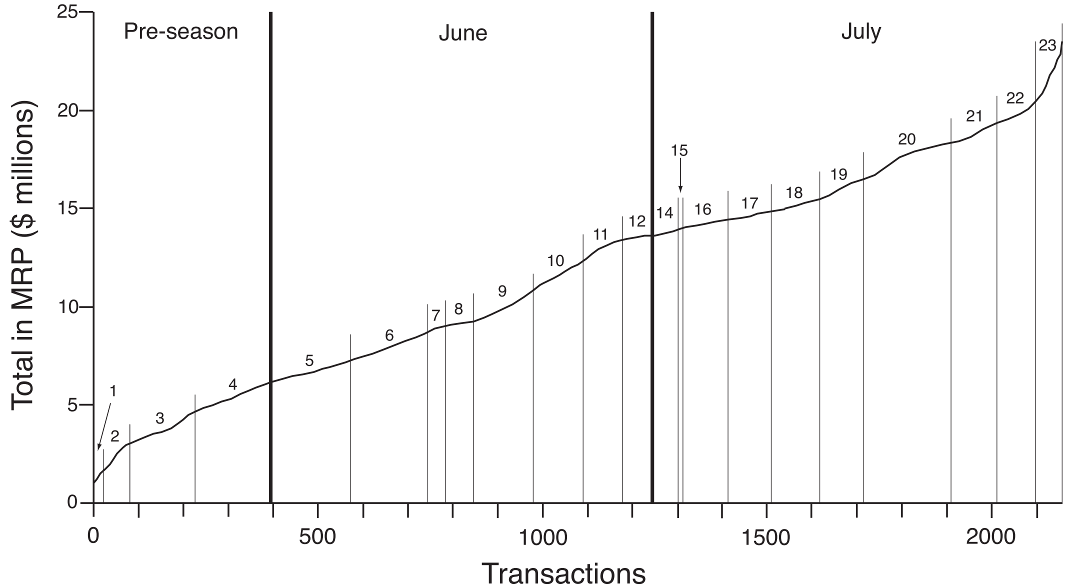


FIG. 4. Growth of the overall pari-mutuel pool over time as a function of numbers of transactions. Each numbered segment indicates a different 4-min trading interval that corresponds to new months in the preseason and days during the two storms. The figure shows a rapid growth of the pool in the preseason followed by surges of trading when storm landfalls seemed likely (particularly in the final trading day).

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 (beginning in Segment 8), there is a consistent growth in the MRP, but the number and value of the purchases rapidly diminishes when the threat passes (Segment 12). 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 (United States) landfall at all.

A different view of how storm events affected purchasing behavior is provided in Fig. 5, which plots

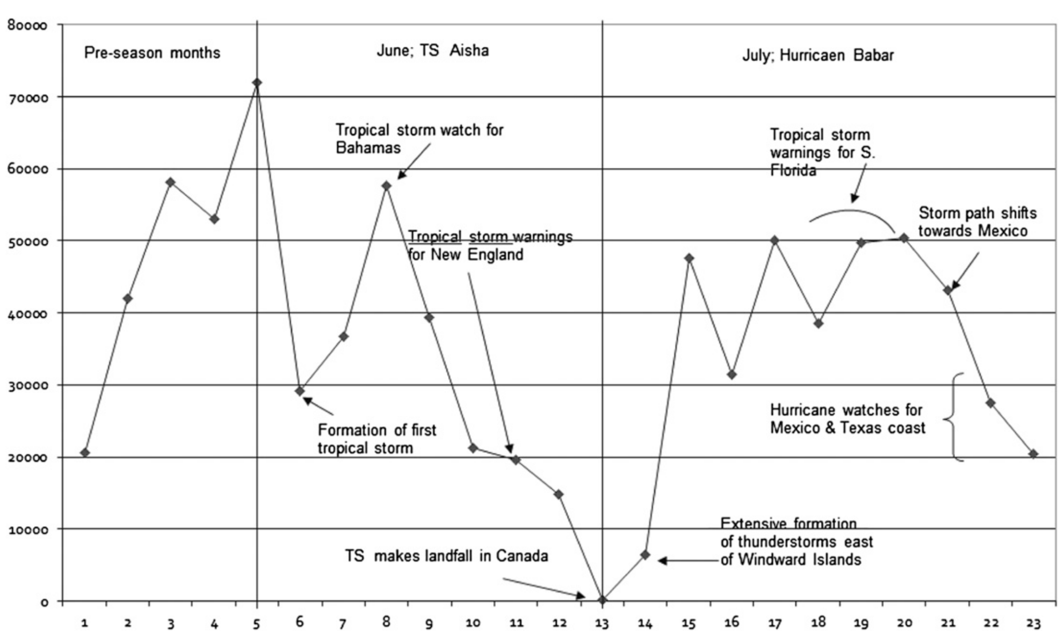


FIG. 5. Plot of primary market trading volume over time as a function of storm events.

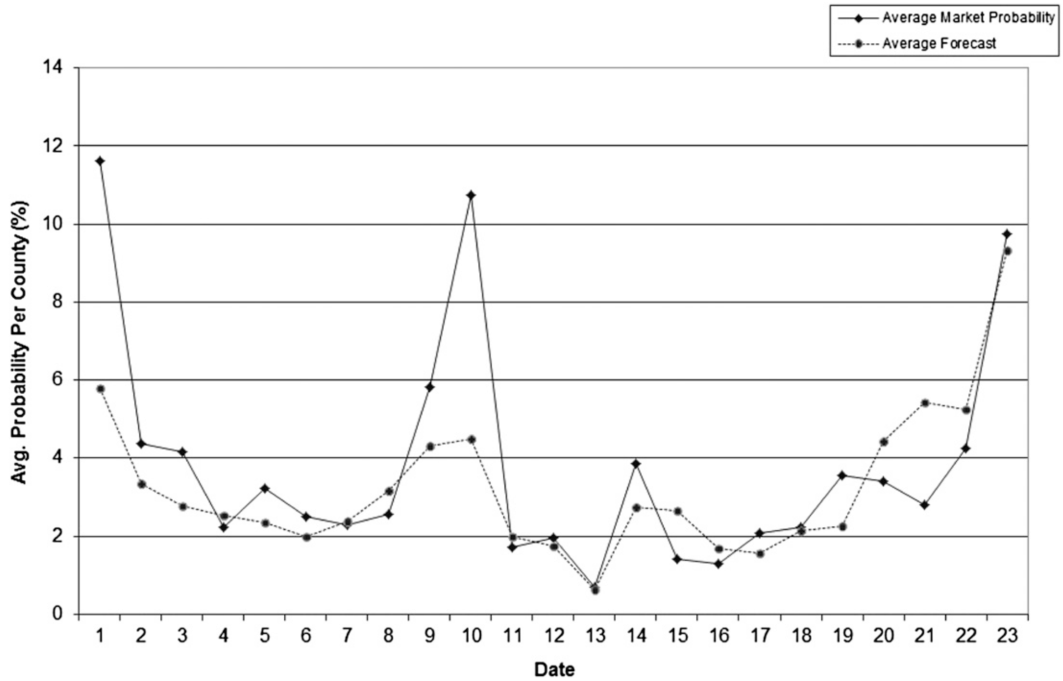


FIG. 6. Plot of how market probabilities (proportional to HuRLO prices) tracked changes in objective landfall odds over time, aggregated over all options. The figure shows a tendency to overvalue options both at the start of the preseason and given the first landfall threat (Aisha), but on average convergence to objective probabilities thereafter.

period-to-period variation in the numbers of HuRLOs purchased (trading volume) as a function of time and storm news. During the preseason 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 preseason 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 preseason, 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.

### b. *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 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 for landfalls that do not pay off suppress a desire to make subsequent purchases.

To examine whether such biases existed, in Fig. 6 we plot a time series of volume-weighted average objective probabilities (see Wilks et al. 2009) compared to corresponding volume-weighted average market-based probabilities. As detailed in Part I, the market-based probabilities are a function of trading volume (probabilities for heavily bought outcomes increase at the expense of those for which there is little buying interest) and, in turn, are used to generate prices. The figure suggests that market probabilities closely paralleled objective forecast-based likelihoods in almost all time segments, suggesting that, for the most part, prices were consistent with participants holding unbiased (rational) expectations about the value of different landfall options. There were, however, two notable exceptions: in the opening time segment when participants had the first opportunity to trade and one that coincided with the first (but ultimately unrealized) U.S. hurricane landfall threat (Aisha).

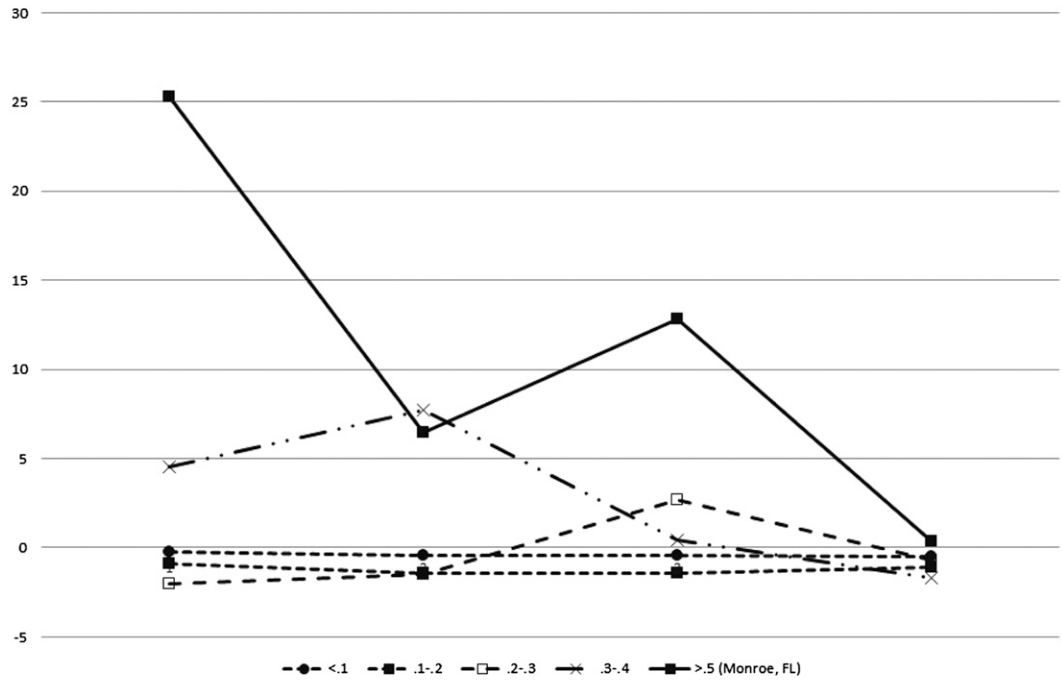


FIG. 7. Plot of the difference between market vs objective probabilities over four preseason trading rounds by preseason base landfall probability. Figure shows an initial overvaluation bias for the options with the highest priors that vanishes with time.

A county-level analysis of the of market probabilities in the first trading period provided deeper insights into the basis of the initial overvaluation bias. Rather than applying to all options, the bias shown in Fig. 6 appeared driven by a high willingness to pay for the two outcomes that had the highest preseason probabilities: Monroe County, Florida (including the Florida Keys) and the No Landfall option. The evolution of these prices within the first trading period are plotted in Fig. 7, which shows that the overbidding for these two options had different characterizations. The probabilities for Monroe County showed evidence of a short-term ''bubble,'' with market probabilities first rapidly increasing to well above market levels, but then collapsing near the end of the trading period. The increase in No Landfall probabilities, however, was less systematic, peaking near the end rather than the middle of the trading period. In both cases, however, it would be reasonable to conclude that prices were inflated by a tendency for participants to overvalue the outcomes that had the highest base likelihoods.

Whereas Fig. 6 suggests that for most subsequent periods aggregate market probabilities (and thus prices) closely tracked normative values, there were localized cases of sustained differences. Two examples are provided in Fig. 8, which plots market-based versus objective probabilities for the No Landfall HuRLO, and Fig. 9, which plots these values for Kenedy County, Texas, site of Hurricane Babar's ultimate landfall. Figure 8 suggests that not only did participants overvalue the No Landfall HuRLO at the outset of the simulation (see Fig. 7), but also 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 Landfall HuRLO, therefore, as displaying something of a ''boomerang'' bias, where participants overcorrect for the easing of the threat to one area by overvaluing the likelihood that there will never be a landfall.

Kenedy County, Texas (Fig. 9), offers an example of sustained underinvestment. Because of its apparently perceived 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 of the simulated season—when landfall there became a plausible event—did market-based probabilities rise to meet normative values, but in the terminal period these still remain less than those prescribed by forecast-based probabilities.

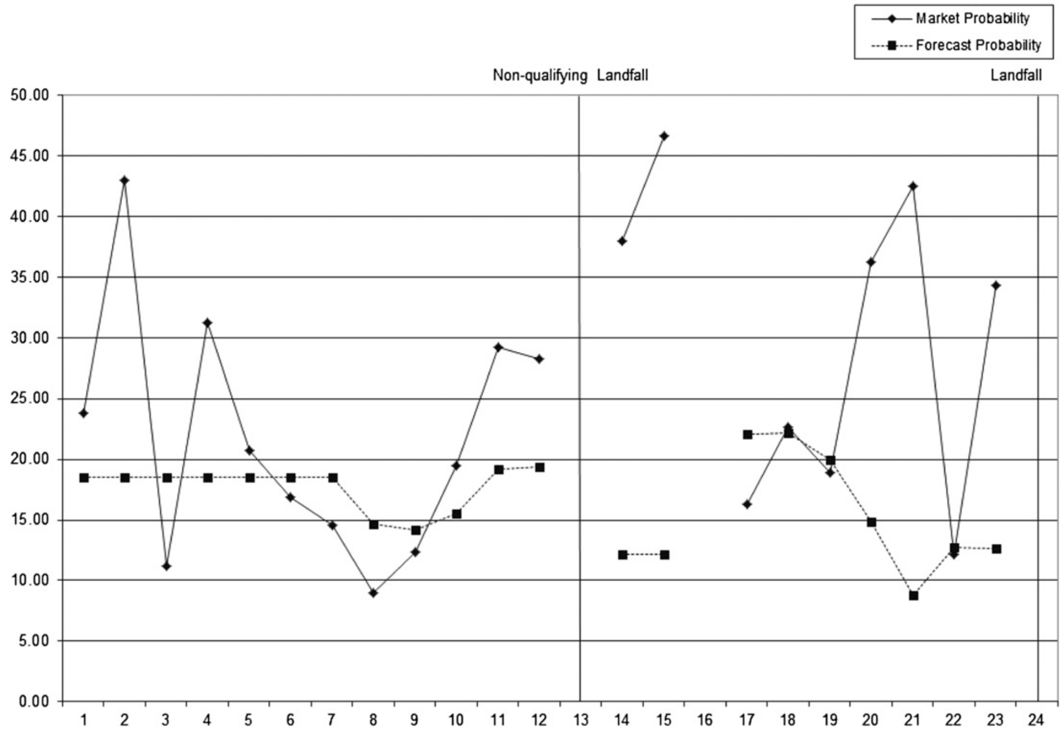


FIG. 8. Plot of market vs objective probabilities for the No Landfall option, which shows less evidence of learning. The figure shows excessive valuation of this option early in the simulation, then continued excessive valuation after each ''near miss'' event: when Aisha bypasses New England (periods 11–12) and when Babar bypasses south Florida (periods 20–21). Gaps indicate periods in which no purchases were made.

### c. *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; that is, acquiring options at low prices with the goal of ''flipping'' them on the secondary market at higher prices.

In Fig. 10, 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 overbuying mistakes in the primary market, although usually unsuccessfully.

Specifically, the data show a surge in sell offers during the second half of 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 preseason 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 overbought positions in the secondary market. It is only at the very end of the simulation—when Babar was about to make landfall in Texas—do we see evidence of the secondary market working efficiently. Here the high market prices of HuRLOs for Kenedy and neighboring counties drove buyers to the secondary market, where they found sellers who were willing to part with previously acquired options at a price 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 (Kunreuther and Michel-Kerjan 2009). While these products have played a useful role in the suite of risk diversification tools available to insurers and reinsurers, 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 has been to report experimental evidence on the empirical properties of a potential

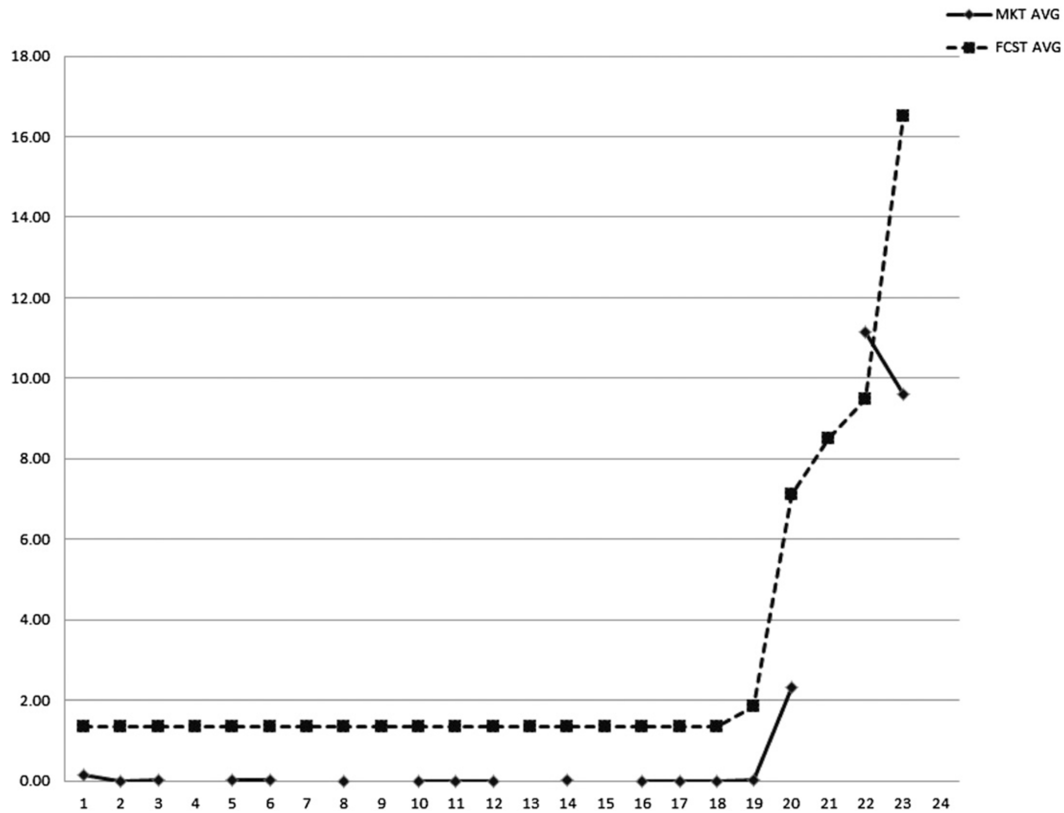


FIG. 9. Plot of market vs objective probabilities for Kenedy County, Texas. The figure shows systematic underinvestment through most of the simulation.

new commodity option product described in Part I 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 Landfall HuRLO. Pricing in the primary market is based on an adaptive control algorithm and reflects the purchasing decisions of all market participants to date.

Empirical properties of the HuRLO market structure 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 buildup 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, would 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 participants' (primarily college students) realism (here the hurricanes were purely hypothetical) and scale of incentives (relatively low stakes). The effect that relaxing these restrictions might have on the findings, however, is unclear. On the positive side, the fact that high levels of efficiency were achieved so quickly despite the inexperience of participants and the relatively low stakes (a small chance to win a $100 prize) would seem to bode well for efficiency in the real world, where traders would presumably have much greater experience, face much higher stakes, and have far more time to contemplate the wisdom of individual trades and purchases. On the negative side, the possibility remains that purchases in the face of real-world hurricanes might

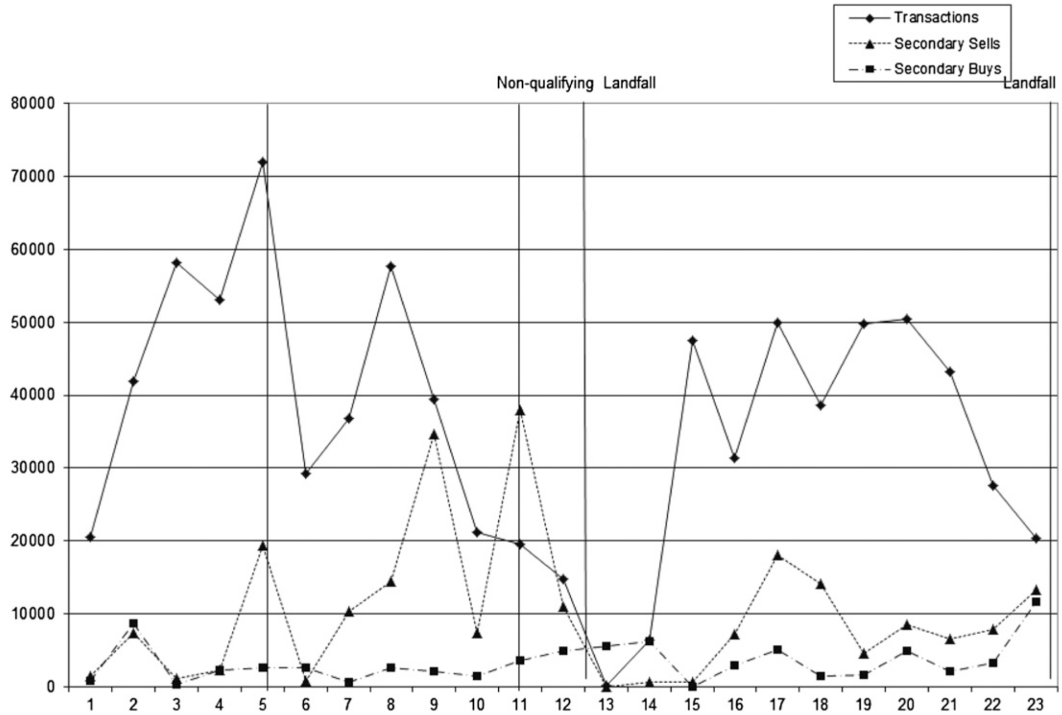


FIG. 10. Plot of activity on the secondary market over time. The figure shows evidence of an active desire by participants to use the secondary market to sell many of the positions they secured during the exuberant buying period associated with the early stages of tropical storm Aisha. The figure also shows, however, a reluctance of other participants to acquire those positions.

exhibit biases that we found little evidence for, such as the potential for prices in certain landfall locations to be subject to speculative bubbles. While the laboratory simulation realistically simulated the kind of information residents typically receive from the NHC, 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.

Likewise, another important feature of the HuRLO market that we did not attempt to study in the laboratory experiments was the effect that the ownership of assets in a particular coastal location might have on prices. Given a high concentration of wealth in a particular location, it is possible that this could act to locally distort prices above their actuarial value. Hence, the findings here should be seen as projecting the likely behavior of neutral speculators rather than home or business owners who are investing in HuRLOs to protect a specific locational asset.

Given this, an important next step in this empirical research program is thus to probe the performance of the HuRLO market not only given traders with specific locational interests, but also under a wider range of incentive schemes. Likewise, future work should also expand testing to include a broader range of participants, most notably home and business owners in hurricane-prone areas).

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 the 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 et al. (2012), 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 or 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.

*Acknowledgments.* This paper describes intellectual property owned by Weather Risk Solutions LLC, which is protected by United States and international patents. Funding for the experimental market was provided by Weather Risk Solutions LLC.

## REFERENCES

Atwood, L. E., and A. M. Major, 1998: Exploring the ''cry wolf'' hypothesis. *Int. J. Mass Emerg. Disasters,* **16,** 279–302.

Berg, J. E., L. A. Daley, J. W. Dickhaut, and J. R. O'Brien, 1986: Controlling preferences for lotteries on units of experimental exchange. *Quart. J. Econ.,* **101,** 201–306.

Bikhchandani, S., D. Hirschleifer, and I. Welsh, 1992: A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.,* **100,** 992–1026.

Breznitz, S., 1984: *Cry Wolf: The Psychology of False Alarms.* Lawrence Erlbaum Associates Publishers, 265 pp.

Broad, K., T. Leiserowitz, J. Weinkle, and M. Steketee, 2007: Misinterpretations of the ''cone of uncertainty'' in Florida during the 2004 hurricane season. *Bull. Amer. Meteor. Soc.,* **88,** 651–667.

Chandler, C. C., L. Greening, L. J. Robison, and L. Stoppelbein, 1999: It can't happen to me. . .or can it? Conditional base rates affect subjective probability judgments. *J. Exp. Psychol. Appl.,* **5,** 361–378.

Daniel, K., D. Hirshleifer, and A. Subrahmanyam, 1998: Investor psychology and security market under- and overreactions. *J. Finance,* **53,** 1839–1885.

Folkes, V., 1988: The availability heuristic and perceived risk. *J. Consum. Res.,* **15,** 13–23.

Horowitz, K. A., A. L. Bequillard, A. P. Nyren, P. E. Protter, and D. S. Wilks, 2013: Activity relating to ongoing financial events. U.S. Patent 20130024342A1, 116 pp. [Available online at http://www.google.com/patents/US20130024342.]

Kagel, J. H., and A. E. Roth, Eds., 1995: *Handbook of Experimental Economics.* Princeton University Press, 721 pp.

Kahneman, D., and A. Tversky, 1973: The psychology of prediction. *Psychol. Rev.,* **89,** 237–251.

Kelly, D. L., D. Letson, F. Nelson, D. S. Nolan, and D. Solis, 2012: Evolution of subjective hurricane risk perceptions: A Bayesian approach. *J. Econ. Behav. Organ.,* **81,** 644–663.

Kunreuther, H. C., and E. O. Michel-Kerjan, 2009: *At War with the Weather: Managing Large-Scale Risks in a New Era of Catastrophes.* MIT Press, 416 pp.

Kushner, H. J., and G. G. Yin, 2003: *Stochastic Approximation and Recursive Algorithms and Applications.* Springer, 474 pp.

Meyer, R. J., 2012: Failing to learn from experience about catastrophes: The case of hurricane preparedness. *J. Risk Uncertainty,* **45,** 25–50.

O'Donoghue, T., and M. Rabin, 1999: Doing it now or later. *Amer. Econ. Rev.,* **89,** 103–124.

——, and ——, 2001: Choice and procrastination. *Quart. J. Econ.,* **116,** 121–160.

Plott, C. R., 1997: Laboratory experimental testbeds: Application to the PCS auction. *J. Econ. Manage. Strategy,* **6,** 605–638.

Selten, R., A. Sadrieh, and K. Abbink, 1999: Money does not induce risk neutral behavior, but binary lotteries do even worse. *Theory Decis.,* **46,** 211–249.

Slovic, P., 2000, *The Perception of Risk.* Earthscan Risk in Society Series, Routledge, 512 pp.

Smith, V. L., G. L. Suchanek, and A. W. Williams, 1988: Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica,* **56,** 1119–1151.

Wilks, D. S., and K. A. Horowitz, 2014: A novel financial market for mitigating hurricane risk. Part I: Market structure and model results. *Wea. Climate Soc.,* **6,** 307–317, doi:10.1175/WCAS-D-13-00032.1.

——, C. J. Neumann, and M. B. Lawrence, 2009: Statistical extension of the National Hurricane Center 5-day forecasts. *Wea. Forecasting,* **24,** 1052–1063.

Wolfers, J., and E. Zitzewitz, 2004: Prediction markets. *J. Econ. Perspect.,* **18,** 107–126.