3.4 A NOVEL FINANCIAL MARKET STRUCTURE FOR MITIGATING HURRICANE RISK

Daniel S. Wilks
Cornell University, Ithaca NY
dsw5@cornell.edu

1. INTRODUCTION

The risk of property damage and loss from hurricanes is a fact of life in coastal areas, which is being progressively magnified by the ongoing growth of coastal populations and property values (Cutter et al. 2007, Pielke et al. 2008). One very significant economic consequence 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 Process 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 unwilling to write new policies, greatly limiting coverage, or withdrawing from the windstorm insurance business altogether.

Conventional weather derivative contracts (e.g., CME Group 2008, Jewson and Caballero 2003, Zeng 2000) require market participants to find a willing counterparty, i.e., someone to take the opposite side of a contract. For example, the two participants in a conventional weather contract might be a ski resort operator wishing to protect against the adverse financial consequences of lower business volume in a low-snowfall winter, and a highway authority wishing to protect against the adverse financial consequences of larger operating costs in a high-snowfall winter. The former might contract to pay the latter if winter snowfall is above an agreed threshold, and the latter would pay the former if the winter snowfall were sufficiently low. In effect, both parties "bet" that adverse weather (from their individual perspectives) will occur, so that the negative impacts on their operations will be offset at least in part by the financial contract.

In order for this conventional bilateral market structure for hedging weather risk to work well, there must be comparable numbers of individuals (or, dollars at risk) who will be hurt by occurrence of an event (e.g., a heavy-snowfall winter) and its absence (a light-snowfall winter). Hurricane risk does not fit this model well, because there are many individuals and businesses who are hurt financially by hurricane landfalls, and few if any for whom the lack of landfalling hurricanes causes financial losses. Therefore a conventional bilateral hurricane market will only function well to the extent that speculators with very large financial resources (e.g., "hedge" funds) take financial positions that landfalling hurricanes will not occur.

This paper describes an approach to managing hurricane risk using a novel financial market structure, which allows participants to hedge against the risk that a selected coastal county or region on the United States Atlantic or Gulf coasts will be first hit by the next hurricane to make landfall in a calendar year. It differs from conventional bilateral markets in that it is one-sided, so that participants buy contracts from an Exchange, in effect "betting" that a hurricane will strike their area. The payments they receive in the event of a hurricane strike in their area are derived from the payments of market participants in other areas. The market structure is simple, and may offer an attractive alternative means to address the needs of both individuals and the insurance and reinsurance industries with respect to hedging potential financial losses from hurricanes.

2. MARKET STRUCTURE

2.1 *Overview*

The hurricane contracts described here are legally described as commodity options, and have been named Hurricane Risk Landfall Options, or HuRLOs. Purchase of these contracts allow market participants to hedge against the risk that one of seventy-eight 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. Figure 1 shows a map of the landfall areas (primarily, individual counties) for which HuRLOs may be purchased from an online Exchange. 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, buyers of Series 1 No Landfall contracts were paid in 2009, because there were no U.S. landfalling hurricanes in that year. 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.

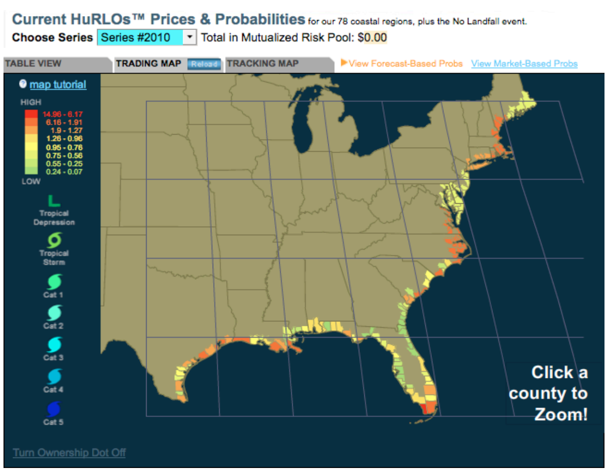

Figure 1. Screen from www.weatherrisksolutions.com, showing HuRLO landfall areas, color-coded with climatological first-landfall probabilities.

2.2 *Price Determination*

In a conventional bilateral market, price discovery is achieved through negotiations between buyers and sellers. This mechanism is not available in a one-sided market, so that an alternative, fair, risk-based pricing procedure is required. Define a vector of *pricing probabilities*, composed of the elements X^i_t, expressing the consensus market sentiment for each of the i = 1, . . . , I possible outcomes, where I = 78 explicit landfall areas + 1 No Landfall event = 79, at a given time t. The price P for a single contract in outcome i is proportional to its pricing probability,

$$P^i_t = \$1000 \ X^i_t \qquad (1)$$

The market is initialized with the climatological probabilities X^i_0 which, with the exception of the No Landfall probability (≈ 0.15), are indicated by the colors in Figure 1 for Series 1. Estimation of these initial probabilities is described in Section 2.3. Proceeds from all purchases in a given Series are collected into a Mutualized Risk Pool (MRP) for that Series, which will be shared among holders of contracts for the outcome which ultimately occurs.

Once the market begins to operate, the initial climatological probabilities are dynamically updated to reflect market activity, so that probabilities (and thus also prices) for outcomes being bought heavily will increase, whereas those for outcomes with little buying interest will decline. These adjustments are made using a recently introduced (Bequillard 2010) adaptive control algorithm, which is a novel variant of the Robbins-Monro stochastic approximation algorithm (Kushner and Yin 2003). Crucially, the adaptive control algorithm possesses the property that its adjustments to the pricing probabilities X^i_t converge to the consensus of market participants' judgments about the event probabilities, as revealed through the buying activity (Bequillard 2010). Thus, apart from a modest but inevitable lag in the response of the probability adjustments, prices computed using Equation (1) are fair and risk-based, in the sense that they reflect the market consensus for the outcome probabilities at any given time. In effect, 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. Examples of the probability updating process are given in Section 3.

2.3 *Climatological Probabilities*

As noted in Section 2.2, the vector of pricing probabilities $\boldsymbol{X}$ must be initialized in order for prices to be defined at the time the market is opened. Ideally the hurricane market is opened early in the year, so that climatological probabilities for the outcomes are appropriate starting points. However, the rather fine spatial resolution of the coastal county segments indicated in Figure 1 implies that raw climatological hurricane landfall relative frequencies are insufficiently smooth for this purpose.

The smoothed climatological probabilities in Figure 1 have been obtained through Monte-Carlo simulations based on the Hurdat data (Jarvinen et al. 1984) available from the National Hurricane Center (NHC) website. The procedure, proposed by Neumann (2006), is as follows. Each storm in the Hurdat data base, 1851-2008, is examined 1000 times, after a displacement of its track defined by an independent random draw from the circular uniform distribution with radius 50 nm. The subsequent courses of each of these perturbed tracks are then examined to find relative frequencies of crossings of the U.S. coastline segments indicated in Figure 1, at hurricane strength. It is recognized that storms earlier in the data base are less accurately portrayed, but the errors are least important near the U.S. coastline, which is the focus of the analysis. This method is similar to the HURRAN forecast method (Hope and Neumann 1970), which traces paths of analog historical storms displaced randomly from the current position of an existing storm, and the forecast method described in Wilks et al. (2009), which uses the same basic methodology initialized from a forecast future storm position.

The resulting probability estimates reflect the relative sizes of the counties, as well as the relationship of the local coastal geography to the climatological storm-path directions. For example, the most likely landfall location is Monroe County, Florida, which includes the Keys. These islands both present a large target, and are oriented nearly perpendicularly to the climatological average storm direction in this portion of the domain. In contrast, there is a distinct probability minimum on the northeast Florida and Georgia coastlines, reflecting the nearly directly northward climatological direction there.

2.4 *Market Termination and Contract Settlement*

Purchases of HuRLOs can continue until a possible hurricane landfall is imminent, or until 15 December, after which time the risk of U.S. hurricane landfall is nil. During the hurricane season, new purchase activity is suspended if and when a Hurricane Watch for one of the colored locations in Figure 1 has been issued by the NHC, meaning that hurricane conditions are possible within (approximately) 36 hours. At that point it is far from clear which coastline segment, if any, will receive a hurricane strike. Figure 2 shows the number of counties subtended by the NHC "cone of uncertainty" (considering cones falling fully on the coastline only) as a function of time ahead of hurricane landfall, for U.S. landfalling hurricanes 2002–2006. The variability in Figure 2 derives from the different sizes of the county coastlines, and the different angles of hurricane approach to the coast. The NHC cone itself provides approximately 90% probability coverage near the time of landfall (Wilks et al. 2009).

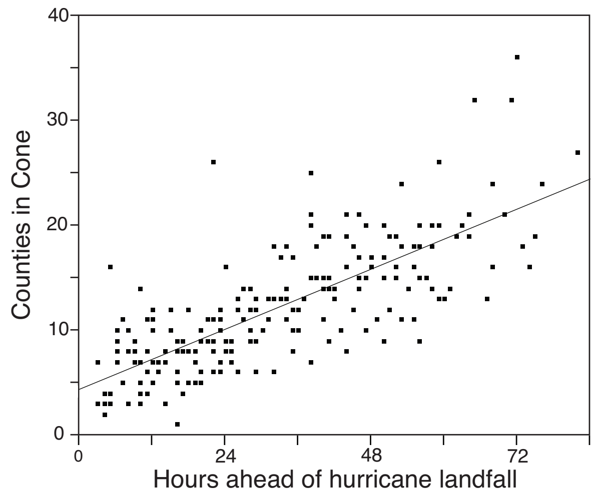

Figure 2. Numbers of coastal counties subtended by the NHC "cone of uncertainty" as a function of time before hurricane landfall, for US hurricanes, 2002–2006. Regression line is N = 4.3 + 0.24 h.

If the storm for which the hurricane watch was issued fails to make landfall on the U.S. as a hurricane, trading in the suspended Series is restarted. If a hurricane landfall occurs, its position for purposes of market settlement is determined by the first intersection of line segments connecting real-time NHC Advisory positions with the high-resolution representation of the coastline defined by the U.S. Census Bureau database available at www.census.gov/geo/www/cob/co2000.html (these data, rendered in map form, are also available as pdf files, at www.census.gov/geo/www/maps/stco_02.htm). Thus, holders of "in-the-money" contracts can be paid within a day or two of landfall.

Settlement amounts per contract held in the landfall county are determined simply as the total dollar amount in the MRP, divided by the total number of contracts for the landfall county that have been sold. Thus, purchasers of contracts for counties that were not hit fund the payouts for the county receiving the first strike. Because the contract prices are $1000 multiplied by the market probabilities at the time of purchase, and these probabilities have been updated continuously over the course of the market, these settlement amounts should be in

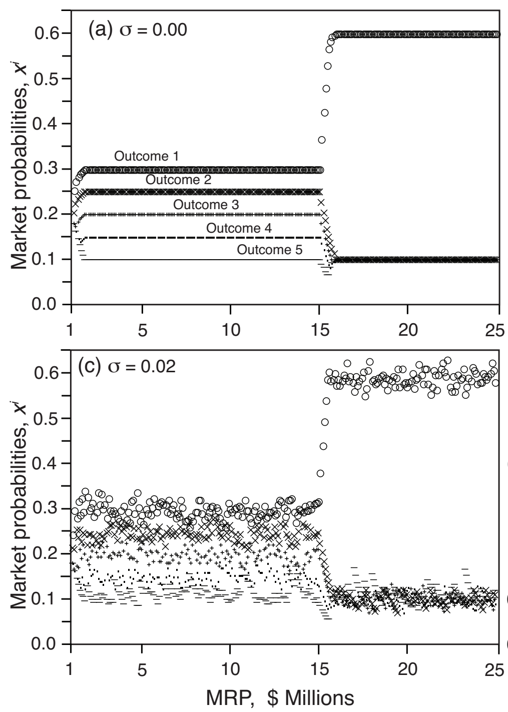

Figure 3. Adaptive convergence of market (pricing) probabilities in an idealized $I = 5$ outcome setting, when (a) The most favorably priced contract, only, is purchased at any given time. (b) – (d) are analogous results when determination of the most favorable price is obscured with Gaussian noise with the indicated standard deviation.

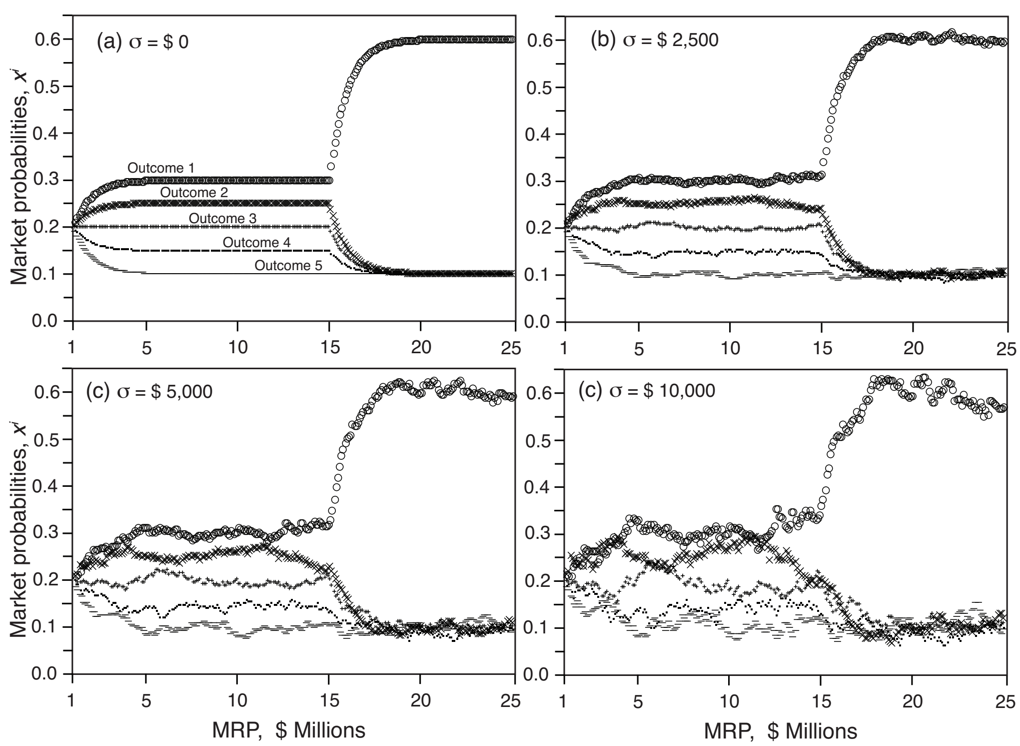

Figure 4. As Figure 3, but for a market where funds are allocated among the $I = 5$ outcomes according to the participant probabilities q^i.

the neighborhood of $1000 per contract, provided the market is well developed.

3. OPERATION OF THE ADAPTIVE CONTROL ALGORITHM

The capacity of the adaptive control algorithm to converge to market participants' beliefs about the outcome probabilities, as revealed by their buying activity, is illustrated in this section. Section 3a describes an idealized situation, with only I=5 outcomes. Section 3b shows a portion of a geographically explicit simulation involving Hurricane Charley (2004).

3.1. *Probability Convergence in an Idealized Setting*

Figures 3 and 4 illustrate the probability (and, through Equation (1), price) convergence of the adaptive control algorithm in a simplified setting. Only five outcomes, rather than 79, have been defined for this hypothetical market, for the sake of clarity in the diagrams. The five "climatological" probabilities assigned to the vector $\boldsymbol{X}$ to initialize the market are equal, at which time the MRP = $1,000,000.

An initially uniform distribution has been used here in order to illustrate the capacity of the adaptive control algorithm to respond promptly to participant sentiment. The market participants do not agree that all the outcomes are equally likely, and instead they invest money according to the following probabilities q^i for Outcomes 1 through 5: $q^1 = 0.30$, $q^2 = 0.25$, $q^3 = 0.20$, $q^4 = 0.15$, and $q^5 = 0.10$. This investment pattern continues until MRP = $15,000,000, at which time new information (corresponding perhaps to updated forecast information) becomes available, indicating $q^1 = 0.6$, and $q^2 = q^3 = q^4 = q^5 = 0.1$.

Figure 3a shows the evolution of the pricing probabilities as a function of total funds in the MRP, when only the most favorably priced contract, i.e. the contract for which $q^i - X^i_t$ is maximized, is purchased at any given time. Initially only contracts for outcomes 1 and 2 are purchased, which drives their prices up, and drives down the prices of the other outcomes. After a relatively modest additional inflow of investments to the MRP, the new equilibrium corresponding to the market consensus is achieved. When the market consensus changes at MRP= $15,000,000, adjustment of the pricing probabilities is similarly prompt.

Figures 3(b)–(d) illustrate that the adaptive control algorithm is robust to noise (i.e., uncertainty) in evaluation of which outcome is most favorably priced. Here the contract maximizing $q^i - X^i_t + \sigma z$ at any given time is purchased, where the standard deviations σ are indicated in the panel legends, and z is an independent standard Gaussian variate updated for each MRP increment of $100,000. The result is that the market probabilities (and thus prices) fluctuate rather than reaching stable equilibria, but these fluctuations are centered on the correct values, i.e., the q^i.

Figure 4 shows analogous results for a different model of investor behavior. Here all five outcomes are purchased within each $100,000 increment of MRP increase, with that $100,000 increment being allocated among the outcomes in proportion to the investors' probabilities q^i. Probability convergence in Figure 4a is somewhat less rapid than for the investor behavior modeled in Figure 3a, because (temporarily) overpriced contracts will be purchased. Panels (b) – (d) illustrate that the probability updating is robust to noise for this investor model also.

3.2 *Simulation for Hurricane Charley (2004)*

Figure 5a shows simulated prices for the Florida counties Lee and Charlotte during a simulated 2004 hurricane season. During this simulation, probabilities and prices for all I = 79 outcomes were simulated, but only two are shown in Figure 5a for clarity. On each simulated day through July, invested funds have been allocated to the various outcomes according to their climatological probabilities (dashed horizontal lines for the two counties in Figure 5a), perturbed on each day with noise. In August, when first Tropical Storm Bonnie, and then Hurricane Charley, threaten the U.S. coastline, the allocation probabilities (corresponding to participant probabilities q^i in Section 3.1) are based on disaggregations of the NHC Advisory forecasts, as described in Wilks et al. (2009).

As would be expected from the results in Section 3.1, the prices are stable through July, with some stochastic fluctuation around the climatological probabilities multiplied by $1000. Prices and probabilities are larger for Lee than the adjacent Charlotte County because of its longer coastline. Their probabilities spike with the approach of Hurricane Charley, rising to approximately 0.12 and 0.06 respectively at the time when simulated trading is suspended. These market probabilities are still fairly modest, because many counties might still plausibly receive the landfall of Hurricane Charley. Figure

5b shows that nearly all of the eastern coast of Florida is within the NHC "cone" at 36 h ahead of landfall, because of the oblique angle of approach of this storm.

(a)

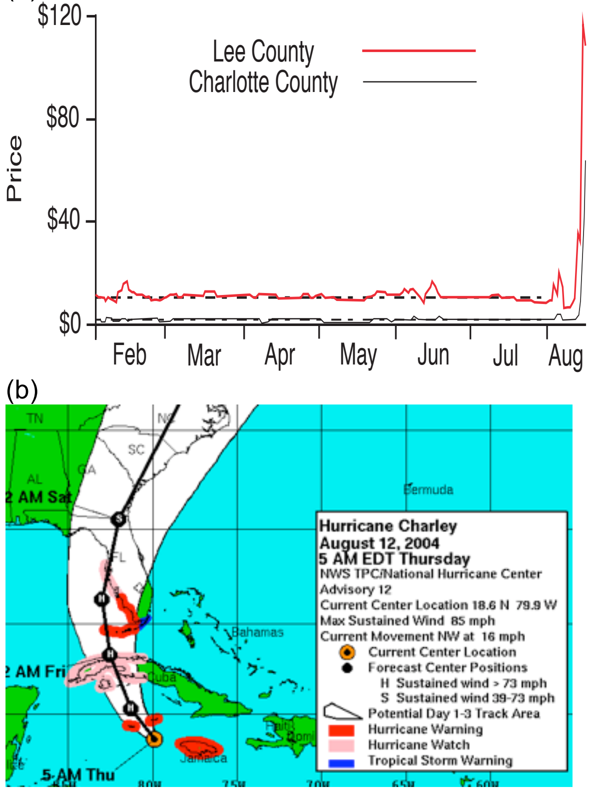

Figure 5. (a) Time course of simulated prices for Lee and Charlotte counties, Florida, until approximately 36 h before landfall of Hurricane Charley at Lee County. Horizontal dashed lines show the respective climatological probabilities. (b) NHC Advisory for Hurricane Charley, 36 h before landfall.

4. SUMMARY AND CONCLUSION

The novel, one-sided financial market structure described here appears to provide a promising mechanism for hedging weather risk in markets for events, such as hurricanes, for which natural counterparties may be few or nonexistent. The prices for the contracts vary through time in proportion to the probability that the next U.S. hurricane landfall will occur at a particular coastline segment, as assessed through an adaptive control algorithm that responds to different levels of buying for the different coastline segments. Proceeds of these sales are collected into a common ("mutualized risk") pool, and holders of contracts for the coastline segment eventually experiencing the landfall are paid from this pool in proportion to the number of contracts held. This market structure efficiently spreads hurricane risks across the entire Gulf and Atlantic coasts of the U.S., on the basis of the uncertainty regarding landfall location as quantified by probability assessments for the event; and without requiring a counterparty, i.e., someone other than the Exchange who is willing to sell the contract.

The landfall segments have been defined as counties in order that hedgers need only pay to hedge against local risks. However, the length scale of hurricane damage is typically somewhat larger than the length of the typical county coastline, so buying contracts for adjacent counties also will generally be advisable.

An additional aspect of the HuRLO markets that has not been mentioned previously is that the Exchange website also supports a secondary market, in which bilateral trades of previously purchased HuRLOs in a conventional bid-ask setting may be made. Neither short sales, nor margin sales, are allowed in either the primary or the secondary markets, so that many of the problems associated with conventional financial derivatives markets are avoided.

To date the market described here has operated only on a limited, preliminary basis. To the extent that it may attract broad participation in future years, it will be possible to compare the market probabilities with various forecast counterparts, as well as study other empirical properties of the market. Corresponding analyses derived from an economic laboratory simulation setting are described in Meyer et al. (2010).

5. REFERENCES

Bequillard, A., 2010: A robust market-making algorithm for nonbinary prediction markets. *SIAM Journal of Financial Mathematics*, submitted.

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

Cutter, S.L., L.A. Jonson, C. Finch, and M. Berry, 2007: The U.S. Hurricane Coasts: Increasingly Vulnerable?", *Environment*, **49:** 8-22.

Derrig, R.A., J.S. Fishman, M. Grace, and J. Schmit, 2008: Catastrophe Management in a

Changing World: The Case of Hurricanes", *Risk Management and Insurance Review*, **55**: 269-280.

Hope, J.R. and C.J. Neumann, 1970: An operational technique for relating movement of existing tropical cyclones to past tracks. *Monthly Weather Review*, **98**: 925-933.

Jarvinen, B.R., C.J. Neumann, and M.A.S. Davis, 1984: A tropical cyclone data tape for the North Atlantic basin, 1886-1983: Contents, limitations and uses. NOAA Tech. Memo. NWS NHC-22, 21 pp.

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

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

Meyer, R., Horowitz, M., Wilks, D.S., and Horowitz, K., 2010: A mutualized risk market with endogenous prices, with application to U.S. landfalling hurricanes. *Journal of Risk and Insurance*, submitted. (Also available as Wharton Risk Ctr. Working. Paper 2008-12-08, http://opim.wharton.upenn.edu/risk/library/WP2008-12-08_RM,MH,DW,KH_HuRLO.pdf).

Neumann, C.J., 2006: Hurricane strikes to coastal counties. Report to Weather Risk Solutions, LLC, November 2006, 13 pp.

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.

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*, **24**: 1052-1063.

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

Acknowledgements. This work has been supported by Kenneth Horowitz and Weather Risk Solutions, LLC.