The Pennsylvania State University

The Graduate School

College of Earth and Mineral Sciences

AN ANALYSIS OF THE ROBUSTNESS AND RELEVANCE OF METEOROLOGICAL TRIGGERS FOR CATASTROPHE BONDS

A Thesis in

Meteorology

by

Addison L. Sears-Collins

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science

May 2007

The thesis of Addison L Sears-Collins was reviewed and approved* by the following:

George S. Young
Professor of Meteorology and Geo-Environmental Engineering
Thesis Advisor

Arthur A. Small
Associate Professor of Meteorology

Keith J. Crocker
Professor of Insurance and Risk Management

William H. Brune
Professor of Meteorology
Head of the Department of Meteorology

*Signatures are on file in the Graduate School (i.e. Signatory Page).

ABSTRACT

Each year weather-related catastrophes account for an estimated United States dollars (USD) $40 billion in damage across the world, although only a fraction of this risk of loss is insured. Losses from hurricanes in the United States have increased over the past several years to the extent that many insurance companies have become increasingly reluctant to insure in certain locations along the coast. Several insurance companies have become insolvent as a result of the active hurricane seasons of 2004 and 2005. In order to cope with this hurricane risk, some insurance and reinsurance firms have shifted part of their risk to the capital markets in the form of catastrophe bonds.

Two problems are observed with catastrophe bonds based on parametric triggers (e.g. Saffir-Simpson scale rating of a hurricane at landfall). First, the trigger mechanisms are measured imprecisely, with the degree of imprecision depending on the choice of trigger mechanism, the available sensor systems, and the methods by which meteorologists analyze the resulting observations. Second, the trigger mechanisms might not relate well to the economic harm caused by the weather phenomena, suggesting that they were not selected on the basis of adequate understanding of relevant meteorology and its relationship to storm damage. Both problems are documented, and perhaps ameliorated in part, by a thorough study of the relevant meteorology and meteorological practices.

Development of a set of robust and relevant triggers for catastrophe bonds for hurricanes is the objective of this study. The real-time and post-landfall accuracy of measured hurricane parameters such as minimum central pressure and maximum sustained surface wind speed were analyzed. Linear regression and neural networks were then employed in order to determine the predictability of storm damage from these measurable hurricane parameters or combination thereof. The damage dataset consisted of normalized economic losses for hurricane landfalls along the United States Gulf and Atlantic coasts from 1900 to 2005.

The results reveal that single hurricane parameters and combinations of hurricane parameters can be poor indicators of the amount of storm damage. The results suggest that modeled-loss type catastrophe bonds may be a potentially superior alternative to parametric-type bonds, which are highly sensitive to the accuracy of the measurements of the underlying storm parameters and to the coastal bathymetry, topography, and economic exposure. A procedure for determining the robustness of a risk model for use in modeled-loss type catastrophe bonds is also presented.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

The idea to investigate hurricane measurement uncertainty as it pertains to the catastrophe bond market was proposed by George Young. I thank Arthur Small and Keith Crocker for serving on my thesis committee. I am grateful to Peter Black and William Brune who made my visit to the Hurricane Research Division possible. Thanks to Mark Powell, Chris Landsea, Eric Uhlhorn, Michael Black, John Kaplan, John Gamache, Matt Easton, Andrew Kleit, William Frank, and Jenni Evans for their insightful comments. Thanks are also to be expressed to Miriam Perkins of Applied Insurance Research, Gordon Woo of Risk Management Solutions, Brian Ohearne of Swiss Re, and Morton Lane of Lane Financial L.L.C. for information and data on current catastrophe bond offerings.

Chapter 1. INTRODUCTION

Losses from hurricanes in the United States have increased over the past several years to the extent that many insurance companies have become increasingly reluctant to insure in certain locations along the coast (McCarty 2006). Traditionally, insurance companies have purchased reinsurance from reinsurance firms in order to hedge against such catastrophic risks (McCarty 2006). Reinsurance contracts are to insurance companies what primary insurance is to individual policy holders. A primary insurance company pays a reinsurer a premium in exchange for protection. In the event of a large natural catastrophe, the reinsurer pays the insurance company a predetermined amount that is a function of the amount of losses the firm realized from the storm.

Until a series of damaging catastrophes in the early 1990s, the practice of using reinsurance contracts as a hedge against catastrophic risks worked fairly well (Bantwal and Kunreuther 2000). However, beginning in 1992, three catastrophes placed a tremendous strain on the reinsurance industry: Hurricane Andrew in 1992, the Midwest floods of 1993, and the Northridge earthquake of 1994. These events revealed the need for alternative forms of reinsurance.

Beginning in 1994, some reinsurers responded to the increased limitations on their capacity to finance large natural disasters by issuing catastrophe bonds. Since then, catastrophe bonds have been issued for risks associated with both natural and man-made disasters. Catastrophe bonds facilitate the transfer of risk from bond issuers—most often reinsurance companies—to risk tolerant institutions in the capital markets such as hedge funds and investment banks. These bonds also enable bond issuers to expand their insurance capacity.

A catastrophe bond works as follows:

An investor purchases a catastrophe bond from a bond issuer. The bond issuer must pay regular interest to the investor until the bond reaches maturity. If the bond reaches maturity without the occurrence of a pre-defined catastrophic event, the bond issuer is obliged to return the original principal payment to the investor. If a catastrophe occurs prior to the maturity date and exceeds the strike levels (a.k.a. triggers) of the bond (e.g. insured losses > 100 million dollars in the case of an indemnity-type bond[1], hurricane wind speed >150mph, in the case of a parametric-type catastrophe bond[2]), the investor no longer receives interest payments from the issuer and looses all or part of the principal. The issuer can use the retained principal to cover damages associated with the catastrophe.

There are a variety of catastrophe bond triggers currently in use. Many of the early catastrophe bonds were indemnity-type. With this type, if the realized insured losses of the bond issuer exceed a certain threshold, the bond triggers. The advantage of these bonds are that the bond is tied to the actual losses realized by the bond issuer. The disadvantage of these transactions is the risk of moral hazard. Moral hazard is a term commonly used in insurance to describe what occurs when coverage against a loss might increase the risk-taking behavior of the insured (i.e. the bond issuer). For example, an insurance company that issues an indemnity-based hurricane bond might decide to expand insurance coverage in high risk, low-lying, coastal areas because it knows that the catastrophe bond would cover its excess losses in the event of a large catastrophe. To

---

[1] An indemnity-type catastrophe bond is one in which the trigger is based on the actual losses realized by the bond issuer.

[2] A parametric-type catastrophe bond is one in which the trigger is based on a physical, measurable trigger such as maximum sustained surface wind speed or barometric pressure.

protect itself against this scenario, some investors prefer firms to issue indemnity-type bonds that have multiple damage thresholds, corresponding to different layers of risk. For example, if the realized losses of a firm trigger a primary threshold (i.e. exceed a pre-specified loss amount), there is a payout. If the realized losses of a firm trigger a secondary threshold, there is a greater payout. The payout amount increases as subsequent thresholds of the bond are triggered. However, even when a firm is protected, indemnity-type bonds can also be abused by inflating reported losses.

As an alternative that eliminates the latter problem but not moral hazard, some of the triggers used in catastrophe bonds are based on modeled losses. If the modeled losses to the bond issuer exceed a specific threshold, the bond triggers. Investors lose either all or part of the principal that they originally invested. The losses are simulated, determined after a storm makes landfall. A risk modeling firm feeds the landfall-measurable parameters of the hurricane into a model and then estimates the total amount of damage caused by the storm. If this value exceeds a predetermined threshold, the bond triggers. The advantage to the issuer is that the bond could potentially be more closely related to the storm damages than it would be if triggered off of a single parameter such as wind speed and pressure. The disadvantage is that this best case is achieved only if the data used to develop the model and drive the simulations accurately reflects the relevant physics and economics. The interdisciplinary problems involved are quite complex, rendering verification by investors more challenging than with single-parameter triggers. This is called asymmetry in information where the bond issuer knows more about the model than does the investor.

To lessen the problem of moral hazard risk and asymmetries in information associated with some of the early indemnity bond issues, parametric-type catastrophe bonds were developed. These bonds are based on physical triggers (Table 1). The advantage of this bond type is that the information on the trigger is readily available from public sources such as the National Hurricane Center. The disadvantage is that they have the problem of basis risk in that the triggers might not be correlated to the losses realized by the bond issuer. As is demonstrated in this study, there is also measurement error associated with the trigger used in a parametric-type catastrophe bond.

In this paper, we demonstrate that current hurricane catastrophe bond triggers such as whether or not the Saffir-Simpson rating of a hurricane exceeds a predefined threshold are subject to a substantial degree of measurement uncertainty. Therefore, catastrophe bonds based on these triggers are only as useful as the accuracy of the underlying measurements used to determine whether or not the trigger reaches the pre-defined strike levels. Further uncertainty arises because such hurricane event parameters are based on the time/space interpolation and extrapolation of a limited number of observations, the representativeness of which is not guaranteed. Second, this processing requires subjective decisions by staff at the National Hurricane Center (NHC) (personal communication, Richard Pasch 2006). The resulting uncertainty in the actual value of the trigger and hurricane parameters has the potential to cause economic problems and legal uncertainty. This study explores the methodological nature of these triggers and suggests triggers that might be most robustly (i.e. with accuracy) measured.

In addition to the uncertainty associated with the triggers, the correlation between measurable hurricane parameters such as wind speed or barometric pressure and the resulting damage is uncertain. Powell et al. (2004) suggests that, from an insurance perspective, the maximum sustained surface wind speed at landfall might not be the most reliable measure of risk or damage potential because it does not take the entire spatial wind field into account. Given a set of "robust" triggers and hurricane parameter data, we use the empirical modeling techniques of linear regression and neural networks (Bishop 1996) to determine the measurable parameter or combination of parameters of a hurricane that is most correlated with storm damage. If a particular trigger can be measured accurately and is highly correlated with economic damage, it would appear to be a good candidate for use in a hurricane catastrophe bond. The meteorological triggers that have been used for past catastrophe bond offerings (Table 1) were attempts to achieve this delicate balance between using a trigger that is accurately measured and using a trigger that correlates well with the losses of the bond issuer.

TABLE 1. Meteorological triggers for catastrophe bonds.

Meteorological Triggers for Catastrophe Bonds Summary (1997-2005*)		
Types of Triggers	**Bond Issues**	**Comments/Notes**
Barometric Pressure (Bond triggers if barometric pressure is lower than a predetermined level)	Prime Capital I Hurricane Ltd. 2000	Bond covers specific areas or "gates" in New York, NY and Miami, FL New York, NY: pressure< 955 millibars (0.54% prob.). Miami, FL: pressure< 932 (0.17% prob.) for gate A OR <936 millibars for gate B (0.76% prob.)
Wind Speed (Bond triggers if wind speed exceeds a predetermined level)	Pioneer Ltd. 2002 Arbor I Ltd. 2003 Arbor II Ltd. 2003 Palm Capital Ltd. 2003	All issues are sponsored by Swiss Re based on measured peak wind gust during U.S. landfalling hurricane. This summer, I will work to find out the actual wind speed trigger values for these four issues.
Saffir-Simpson Scale Hurricane Rating (Bond triggers if Saffir-Simpson Scale strength rating exceeds a predetermined level) + Indemnity Trigger	Residential Re I—1997 Residential Re II—1998 Residential Re III—1999 Residential Re IV—2000 Residential Re V—2001 Residential Re VI—2002	All of the Residential Re issues are a combination of indemnity losses with a physical index trigger: Indemnity trigger: The bonds trigger when USAA losses exceed a certain threshold. Physical trigger: Losses must be caused by a Category 3, 4 or 5 storm on the Saffir-Simpson index in District of Columbia or any one of the following states: Alabama, Connecticut, Delaware, Florida, Georgia, Louisiana, Maine, Maryland, Massachusetts, Mississippi, New Hampshire, New Jersey, New York, North Carolina, Pennsylvania, Rhode Island, South Carolina, Texas, Vermont and Virginia.
Modeled Losses (After the occurrence of a catastrophe, the weather observations are fed into a catastrophe model. This model is then run against a company's exposure database to produce simulated losses. If the simulated losses are above a specified threshold, the bond is triggered)	Zurich Re., Trinom Ltd. 2001 Helix 04 Ltd. 2004 Atlantic & Western Re Ltd. I & II., PXRE 2005 Champlain Ltd., Montpelier Re 2005	It is worth examining the error in the weather observations taken during a catastrophic hurricane. Specifically, how reliable are the National Weather Service's measurements of wind speed (anemometer), pressure (dropsonde), storm surge (estimation), and rainfall (rain gauge). These measurements are fed in to the catastrophe models to produce modeled storm losses. It is the modeled storm losses that then determine if a bond triggers or not.

Chapter 2. DATA AND METHODOLOGY

a. Historical hurricane measurements

The estimated minimum central pressure and maximum sustained surface wind speed values at landfall were obtained from the National Hurricane Center (NHC) North Atlantic hurricane database (HURDAT). Jarvinen et al. (1984) originally prepared this dataset, but it continues to undergo extensive updates and revisions (Landsea et al. 2004; Partagas and Diaz 1996). During the period from 1900-2005, a total of 135 landfalling hurricane events were identified along the coastal United States.

b. Hurricane economic impact

The damage dataset used in this study is from Pielke et al. (2007). This data consists of normalized hurricane damage estimates from 1900-2005. Pielke et al. (2007) uses two different approaches for normalizing storm damage for each landfalling hurricane over the 106-y period. The first approach follows the methodology of Pielke and Landsea (1998), adjusting past storm damage for changing societal conditions by accounting for inflation and changes in wealth and population updated to 2005. The second approach follows the methodology of Collins and Lowe (2001), adjusting storm damage for changes in inflation and wealth at the national level and changes in population and housing units at the coastal county level updated to 2005. The two approaches produce similar damage figures (Fig. 1) as described in detail in Pielke et al. (2007). This dataset provides a continuous record of normalized damage estimates of landfalling hurricanes during the 1900-2005 period. As described in Pielke and Landsea (1998), this methodology is sensitive to the accuracy of the reported damage figures.

Data that indicate that a given hurricane caused more or less damage could alter our results.

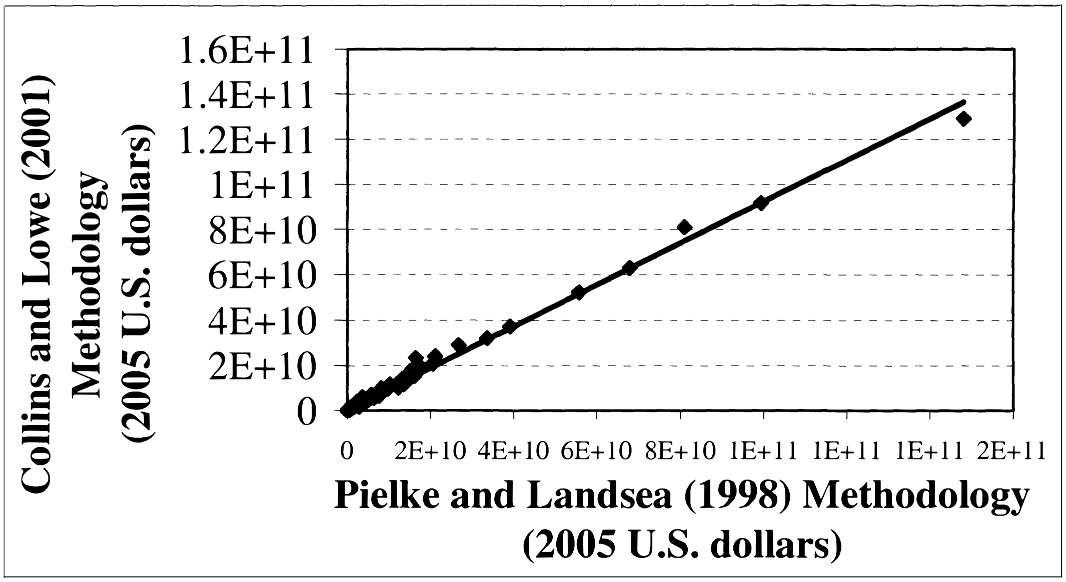

FIG. 1. Pielke and Landsea (1998) methodology vs. Collins and Lowe (2001) methodology for estimating storm damage for landfalling hurricane events from 1900-2005.

c. Catastrophe bond trigger data

Data on the common triggers for catastrophe bonds was obtained directly from offering circulars. An offering circular is a legal document produced by the bond issuer for each catastrophe bond issued, as required by the Securities Act of 1933. It includes the following information about the bond: the risk factors, the purpose of the offering, the credit rating, income tax considerations, a plan of distribution, a risk analysis report by a modeling firm (e.g. Equecat), and a detailed description of the conditions required for the bond to trigger. In theory, it contains any information that may be of use to a prospective investor, although as mentioned above asymmetric information situations can arise. Further information on the popular triggers was obtained through personal communication with professionals at Swiss Re, Munich Re, Lane Financial L.L.C, Guy Carpenter, and Applied Insurance Research (AIR).

d. Neural network architecture

Linear regression and neural networks (NNs) were used to extract patterns that might exist between hurricane parameters or combination of parameters (i.e. predictors) and the associated storm damage (i.e. predictand). A NN is a nonlinear regression technique that is modeled loosely on the small-scale structure of the human brain (Reed and Marks 1999). NNs recognize patterns and develop classifications through repetitive training. Unlike linear regression, NNs can develop models of phenomena in which the relationship between the predictors and predictand is nonlinear (Marzban and Stumpf 1996). Because of the inherent nonlinearity of atmospheric phenomena, NNs have gained wide acceptance as an empirical modeling tool in meteorology (Hsieh and Tang 1998).

The architecture of the NN in this study follows the standard construct of a feedforward network discussed in many data mining books (e.g. Reed and Marks 1999; Witten and Frank 2005). Developmental cases begin at an input layer and then pass through one or more hidden layers to an output layer. Each layer consists of a set of nodes. At each node, the weighted average of the predictors is computed where x_j is the jth predictor, w_{ij} is the weight for that predictor for node i, and u_i is the weighted average coming out of the ith node (Witten and Frank 2005):

$$u_i = \sum w_{ij} x_j \tag{2.1}$$

The weighted averages are "squashed" by a sigmoid function to ensure that all values are between 0 and 1:

$$f(u_i) = \frac{1}{1 + e^{-u_i}} \tag{2.2}$$

These values then become the predictors for a subsequent hidden layer. Thus, they serve as intermediate predictions. The predictions of one hidden layer become the input of the next until the weighted average of the predictions of the last hidden layer is squashed and returned as the final, output layer of values (i.e. the predictand).

After developmental cases are fed through the NN, the output is compared to known values of the predictand. The differences between this known predictand and the predicted output is the error. The error values are used to adjust weights in the NN through a process called backpropagation (Bishop 1996). These weights are analogous to the coefficients in a linear regression equation. A NN learns through iterative adjustment of these weights. This process is conducted via gradient descent so as to minimize the error of the output (Reed and Marks 1999). For each training case, the derivative of the error surface is computed with respect to the weights then the weights are adjusted in a manner proportional to the mean value of this error gradient averaged over all the training cases. Ideally, the weights are adjusted until the neural network converges on a global minimum. In practice, of course, there is some risk of this gradient descent method getting trapped in a local minima, so multiple runs with different initial weights are undertaken.

After training the NN on developmental data, it is necessary to test the newly created system on independent data in order to determine its true predictive capability (Reed and Marks 1999; Bishop 1996). A 10-fold cross-validation scheme was employed for the test phase of this study (Witten and Frank 2005). The advantage of using this cross-validation technique instead of other validation procedures is that more cases are used in the training phase (Witten and Frank 2005). Moreover, all of the cases are eventually used in the testing phase.

e. *Training the neural network*

Before training the network, we created two additional predictors from the maximum sustained surface wind speed values in HURDAT: the Saffir-Simpson scale rating (Simpson and Riehl 1981) and an interpolated Saffir-Simpson scale rating. The interpolated Saffir-Simpson scale rating was computed for each event as follows:

$$SS = c_o + \frac{(v_o - v_l)}{(v_u - v_l)} \tag{2.3}$$

where:

SS interpolated Saffir-Simpson Scale rating,

c_o Saffir-Simpson Scale rating of storm as reported by the NHC,

v_o observed 1-min maximum sustained surface wind speed,

v_l low wind speed threshold for Saffir-Simpson Scale rating of storm as reported by the NHC (Table 2),

v_u upper wind speed threshold for that category.

Note: The upper category 5 limit was set to 190mph because that was the fastest recorded maximum sustained wind speed at landfall in the dataset (Hurricane Camille, 1969).

TABLE 2. Saffir-Simpson Scale thresholds

Category	Wind speed (mph)
5	≥156
4	131–155
3	111–130
2	96–110
1	74–95

The final dataset consisted of a total of four predictors[3]: Saffir-Simpson rating at landfall, interpolated Saffir-Simpson rating at landfall, estimated central pressure at landfall (mb), and maximum 1-min sustained surface wind speed at landfall (m s^{-1}). These predictors comprised the input variables for the linear regression models and the first layer of the NNs. There were two sets of predictand, corresponding to the two different normalized storm damage figures from Pielke et al. (2007). These damage figures comprised the output variables for the linear regression models and the final layer of the NNs.

Since we identified 135 landfalling hurricane events during 1900-2005, the linear regression and NN models were trained on a total of 135 cases. Because of the limited period of record provided only 135 cases, it is possible that we did not have a sufficient number of representative cases to minimize overfitting. Overfitting results from having too little developmental data for the number of parameters tuned. Using a large number

[3] There were really only two independent predictors in total since three of the four predictors were a function of wind speed. But for the purposes of this paper, we used wind speed in three different forms to see if a different result was achieved.

of cases lowers the odds of obtaining coincidental relationships between the predictors and predictand, i.e. of overfitting.

The Waikato Environment for Knowledge Analysis (Weka) software was used to construct the linear regression and NN models. The goal was to develop a system that could predict storm damage from one or more of the four predictors. The predictive capability of a network was determined from the value of R-squared. There is no definitive rule on how to build an optimal NN. Therefore, the training and testing procedure involved altering the values for the learning rate and momentum, trying various combinations of predictors, and varying the number of hidden layers and nodes. The aforementioned learning rate controls the step size and how quickly the search converges to the minimum of the error surface. The momentum term makes the error adjustment process steadier by making the weight change more persistent. The number of hidden layers used ranged between 0 and 3, and the number of nodes was generally maintained between 2 and 6 so that the network was not too complex for the dataset.

Chapter 3. BACKGROUND AND LITERATURE REVIEW

a. Determining the trigger for a catastrophe bond

Developing an appropriate trigger for a catastrophe bond involves collaboration between the bond issuer and a catastrophe risk modeling firm. The three main modeling firms are Applied Insurance Research (AIR), Risk Management Solutions (RMS), and Equecat. For more information on the process of catastrophe risk modeling, refer to Dong et al. (1996), which describes the methodology in detail.

A bond issuer that wants to lower its exposure to a catastrophic hurricane will consult a modeling firm to determine the probability that a hurricane exceeding a certain intensity threshold will impact a pre-determined area of coastline. In a hypothetical scenario, Jones Reinsurance Company has $50,000,000 worth of reinsurance exposure in South Florida. The firm wants to decrease the chance that it will go bankrupt in the event that a hurricane of category 3 or greater makes landfall in South Florida. Such an event would likely cause Jones Reinsurance Company to declare bankruptcy because of the high number of property loss insurance claims. Thus, the firm wants to issue a catastrophe bond that would cover half of its risk exposure. AIR uses historical hurricane data to estimate the annual probability that a hurricane of category 3 or greater would make landfall in South Florida, assuming a category 3 storm would cause damage of greater than $25,000,000. The modeling firm then reports the results of this analysis to Jones Reinsurance. Jones Reinsurance is then able to issue a bond with an interest rate that is commensurate with its desired risk exposure.

Despite the error associated with the measured parameters of a hurricane, modeling firms and bond issuers often do not take this uncertainty into account during the pricing of a catastrophe bond (personal communication, Albert Selius, Swiss Re 2006). The only requirement is that the language in the offering circular, the publication that defines the terms of the catastrophe bond, must be unambiguous (personal communication, Keith Crocker 2006). The offering circular must specifically state the methodology for determining and measuring the trigger. For many bonds, the trigger is based on the measured parameters of the hurricane at landfall as published in the Tropical Cyclone Report, a report issued by the National Hurricane Center (NHC) after the passage of a hurricane. This report contains the statistics of a given storm and post-analysis best track estimates. Some offering circulars state that neither the risk modeling firm nor the National Weather Service (NWS) are responsible for errors that may exist in the hurricane parameters printed in the Tropical Cyclone Report (Grand Isle Limited 2007, Swiss Re internal publication).

b. Measurement methodology for the popular triggers

The process of determining storm intensity is a delicate combination of art and science. A lot of the final decision is based on the experience of the highly trained forecasters and scientists at the NHC (personal communication, Peter Black 2006). The NHC uses a combination of all of the available data from the various observational platforms in order to make a subjective decision on the intensity of a hurricane as it makes landfall and how much weight each instrument is given in the composite final decision of the storm parameters (personal communication, Richard Pasch 2006; Table 2). For example, land-based observing systems may fail as a hurricane makes landfall, in which case these systems cannot be used in the final analysis. Meteorologists at the NHC then have to analyze aircraft measurements made just prior to landfall and satellite imagery of the landfalling storm to make an educated guess on its landfall intensity.

The measurements from aircraft, ocean, and land-based observing systems are not without error, however. The NHC must take this error into account in order to make accurate and informed decisions on storm parameters. Modelling firms and bond issuers must also consider such instrumentation errors so that they can develop appropriate catastrophe bond triggers.

c. Trigger error

1) Maximum sustained surface wind speed measurement error

i. GPS dropwindsonde

The NHC maximum sustained surface wind speed (i.e. 1-min maximum sustained wind speed at a height of 10m) estimate at landfall for a given hurricane is subject to error from several sources (Table 2). For example, land-based observing systems often fail at high wind speeds as a result of power outages or structural failure (personal communication, Mike Black 2006). The data from the land-based observing systems that do not fail is often tainted by the presence of debris (Powell et al. 2004). These obstacles can result in substantial differences (as much as a factor of 2) in the measured wind speed between stations (Powell and Reinhold 1996). When this situation occurs, meteorologists at the NHC have to make an educated guess on the maximum sustained surface wind speed at landfall based on the available data from other instruments such as low-altitude hurricane reconnaissance aircraft, making measurements just before landfall and satellite imagery of the storm as it makes landfall. Over the years, the NHC has acquired such measurements from a number of different aircraft, including the National Oceanic and

Atmospheric Administration (NOAA) WP-3Ds, the Gulfstream IV jet, the Air Force C-130s, National Center for Atmospheric Research Electra, and a leased Lear-36.

One of the primary means by which these reconnaissance aircraft measure the maximum sustained surface wind speed in a hurricane before landfall is via the Global Positioning System (GPS) dropwindsonde (Hock and Franklin 1999). This instrument is regularly deployed in the eyewall from the WP-3D and C-130 aircraft and measures the ambient wind speed from flight level down to the surface. The accuracy of the wind speed measurement is 0.5 m s^{-1} (1.118 mph) over a range of 0–150 m s^{-1} (335.5 mph) (Hock and Franklin 1999).

The GPS dropwindsonde can be deployed in the eyewall as long as the center of the storm remains over the ocean, i.e. up until the moment when the storm makes landfall. Because the instrument provides a point measurement of wind speed, the data does not provide a complete representation of the pressure and wind field of the hurricane (Powell et al. 2004). Moreover, because of the turbulent flow in a hurricane, it is difficult to place a dropwindsonde in the location of the strongest surface winds. The actual error these effects introduce into the final NHC wind speed estimates has not been quantified to date and remains unclear (personal communication, Peter Black 2006). Resolving this question is challenging because both sampling and targeting issues are involved.

TABLE 3. Measurement methodology of the popular and potential triggers.

Measurement methodology of the popular and potential triggers	
Central pressure estimates and maximum sustained wind speed measurements.	i. Inertial and GPS navigation systems onboard the aircraft. 1. Flight level wind reduction factor applied to measurements at 700mb. ii. GPS Dropwindsonde iii. Stepped Frequency Microwaves Radiometer to measure wind speed. iv. Microwave Satellite Data and Imagery vi. NOAA near-polar orbiting satellites vii. Defense Meteorological Satellite Program satellite viii. Quick SCAT Scatterometer ix. Aqua x. Geostationary Satellite-based Dvorak estimates by the following agencies (used during pre-landfall period): 1. Tropical Analysis and Forecast Branch (TAFB) 2. NOAA Satellite Analysis Branch (SAB) 3. US Air Force Weather Agency (AFWA) xi. Tail Doppler Radar 1. Plane does a corkscrew pattern through the storm as it measures the wind profile. 2. Tail Doppler Radar will not allow you to get a direct measurement of surface wind speed (personal communication, John Gamache 2006) xii. Aircraft, ASOS and Official Surface Observing Systems, 1. e.g. NWS WSR-88D velocity data 2. "Observations include data from satellites, aircraft, airborne and ground-based radars, conventional land-based surface and upper-air observing sites, Coastal-Marine Network (C-MAN) stations, National Ocean Service (NOS) stations, ocean data buoys, and ships. Selected ship reports of winds of tropical storm force associated, and selected surface observations from land stations and from coastal and fixed ocean data buoys Data from many Automated Surface Observing System (ASOS) sites (measure wind speed over land), C-MAN stations, and buoys are used in many cases but sometimes incomplete due to power outages and other weather-induced failures prior to when peak winds and minimum pressures occurred" (Source: NHC's Tropical Cylcone Report for Hurricane Katrina).
Position estimations in real-time near landfall	NWS WSR-88D Doppler radars based on land

ii. Inertial and GPS navigation systems and the flight level wind reduction factor

Another primary means by which reconnaissance aircraft measure the maximum sustained surface wind speed in a hurricane before landfall is by inertial and GPS navigation systems onboard the aircraft. These aircraft generally fly radial flight-legs towards and away from the center of the storm at a flight level of ~700 hPa (Kossin et al. 2007). The wind at this altitude is measured by the inertial and GPS navigation system. This value is reduced by a reduction factor (R) of 0.90 to infer the 1-min maximum sustained wind speed at 10m. This reduction factor has a standard deviation of 0.19 (Franklin et al. 2003). Thus, the use of flight level winds to estimate 1-minute maximum sustained wind speed at 10m could easily result in errors of ±20%.

As a result, some studies have suggested that this method of extrapolating the maximum sustained surface wind speed from flight level wind speed using the R value is subject to a substantial degree of uncertainty (Dunion et al. 2003; Franklin et al. 2003; Powell et al 2003). Empirical evidence suggests that R varies within and between hurricanes, hence the large standard deviation mentioned above (Powell 2004; Franklin et al. 2003). R on the weaker left side of the hurricane center may be 4% higher than on the right side of the storm (Powell at al. 2004). Thus, it is necessary to know the overall structure of the boundary layer of a hurricane in order to determine the appropriate value for R and to reduce the likelihood of error in the resulting surface maximum sustained surface wind speed estimate. This information is often difficult to obtain (Uhlhorn et al. 2006). Powell et al. (2004) also found that increased surface roughness near the coast may result in values of R smaller than 0.90, leading to overestimations of the surface wind speed.

iii. Stepped Frequency Microwave Radiometer (SFMR)

In addition to the GPS dropwindsonde and flight level wind reduction, the SFMR has become an operationally important wind speed measuring device during the last several years. For over two decades, the National Oceanic and Atmospheric Administration (NOAA)/Hurricane Research Division (HRD) has used the SFMR to measure hurricane wind speed via research aircraft (Uhlhorn et al. 2006; Uhlhorn and Black 2003). The SFMR measures wind speed along the flight path by detecting passive microwave emissions from the sea surface. Ulhorn and Black (2003) and Black et al. (1995) have presented evidence that these emissions are strongly correlated with ambient wind speed.

In 2005, NOAA/Aircraft Operations Center (AOC) began to equip research aircraft with a next-generation SFMR (Uhlhorn et al. 2006). The accuracy of this instrument is 2.2 m s^{-1} ± 0.4% at 30 m s^{-1}. This accuracy is a factor of 2 greater than the original research SFMR that the HRD used prior to 2005 (Uhlhorn et al. 2006). Empirical evidence also reveals that the accuracy of the new system is comparable to the GFS dropwindsonde for surface wind estimates (Uhlhorn et al. 2006). One disadvantage of the SFMR is that it has a low bias at high wind speeds (Uhlhorn and Black 2003). It also loses accuracy in shallow water, so it is difficult to obtain a reliable measurement for the wind speed at landfall (personal communication, Mike Black 2006). The data from the SFMR is still taken into account, however, when the NHC makes the final decision on the maximum sustained surface wind speed at landfall.

iv. Tail Doppler Radar

Some reconnaissance aircraft also have Tail Doppler radar aboard, which can be operated in dual-Doppler mode. This radar measures the three-dimensional wind field of the inner core of a tropical cyclone. This data, however, requires a substantial amount of post-processing and is not available in real-time (Kossin et al. 2007).

v. Boundary layer wind streaks

No matter which sensors are used, boundary layer wind streaks located in the core of a landfalling hurricane may contribute to uncertainty in the estimates of maximum sustained surface wind speed at landfall. Wurman and Winslow (1998) presented Doppler On Wheels (DOW) mobile weather radar observations of these streaks in the inner core of Hurricane Fran (1996). These images revealed intense, sub-kilometer scale horizontal boundary layer rolls that triggered alternating bands of light (15-35 m s^{-1}) and strong (40 to 60 m s^{-1}) near-surface winds. These coherent, turbulent wind streaks might account for past observations of well-defined, small-scale linear swaths of hurricane damage (Wakimoto and Black 1994).

With respect to estimates of the 1-minute maximum sustained surface wind speed, the presence of boundary layer wind streaks in a hurricane can lead to overestimates or underestimates of the wind speed depending on whether the measurement device is located in a band of either strong or light surface winds, respectively. For example, a GPS dropwindsonde deployed into a band of light winds might not yield a measurement that truly reflects the upper limit of the wind damage capability of a given storm. This source of wind speed error combined with the other sources previously discussed has led risk modeling firms and reinsurers to incorporate central barometric pressure into catastrophe bond issues.

2) Minimum central barometric pressure measurement error

As with wind speed estimates, observations of the minimum central barometric pressure at landfall are subject to measurement uncertainty. The GPS dropwindsonde is the most common means of measuring the barometric pressure in hurricanes. Because dropwindsondes cannot be dropped over land, estimates of the landfall barometric pressure are derived from GPS dropwindsonde measurements taken just prior to landfall. The error in the pressure sensor on the GPS dropwindsonde is only ±1 hPa (Hock and Franklin 1999). There is also error, however, stemming from the difficulty of deploying the GPS dropwindsonde in the exact location of the minimum surface barometric pressure. Although this error is difficult to quantify, the lower reliability of land based observing systems during landfalling hurricane events means that the GPS dropwindsonde is the best device available for estimating central pressure. (personal communication, Peter Black, 2006).

Another targeting issue is that the dropwindsonde can blow into mesovortices embedded within the hurricane eye and eyewall, yielding an unrepresentative surface pressure reading. These, small-scale regions of vorticity can result in a lower measured pressure reading. These features have been well documented in the literature (Kossin and Eastin 2001; Black and Marks 1991; Marks and Black 1990; Bluestein and Marks 1987) but there frequency is not well known.

3) Error in the Historical Hurricane Database (HURDAT)

The catastrophe risk modeling firms that help the reinsurance company develop appropriate triggers for a catastrophe bond use historical hurricane data from the NHC Historical Hurricane Database (HURDAT). The uncertainty in this data is an additional source of error that must be considered in the development of a robust trigger for a catastrophe bond. The error of the barometric pressure and surface wind speed estimates in the HURDAT database is likely higher in earlier years because of the lower number of measurements taken for each storm and the evolution of the available sensor systems. Before the 1960s, there were often long periods of time, sometimes days, in which no measurements were taken of a particular hurricane. Also, the early reconnaissance aircraft rarely penetrated the eyewall, making it difficult to obtain measurements of intense hurricanes (personal communication, Chris Landsea 2006). Hurricane Carol (1954), for example, lasted ten days and made landfall on the east coast of the United States, yet only seven measurements of the storm were taken while it was of hurricane strength. For four days, the storm was spinning off the coast of Georgia without measurements being taken, so meteorologists would have had no idea about its intensity (personal communication, Chris Landsea 2006).

Chapter 4. RESULTS

Fig. 2 and Table 3 reveal that there is little to no linear relationship between the Saffir-Simpson Scale rating for a given landfalling hurricane and the associated storm damage. The results also indicate that the minimum central pressure and the maximum sustained surface wind speed at landfall are poor indicators of the amount of damage realized from a given hurricane (Table 3; Figs. 3 and 4). Using different combinations of predictors and performing linear regression against damage did not yield better results as evident from the low R^2 values in Table 3. Moreover, Fig. 4 shows that significance did not improve when the normalized damage figures were conditione on minimum central pressure values less than 960 hPa.

Neural networks did not perform substantially better than linear regression (Table 3). Numerous combinations of predictors and parameter settings were employed. Ten different random seeds were used, and 10 runs were generated in Weka for each random seed. The highest value for R^2 attained using a neural network was 0.125. Thus, the results illustrate that parametric triggers alone are “insufficient” for effective use in a catastrophe bond.

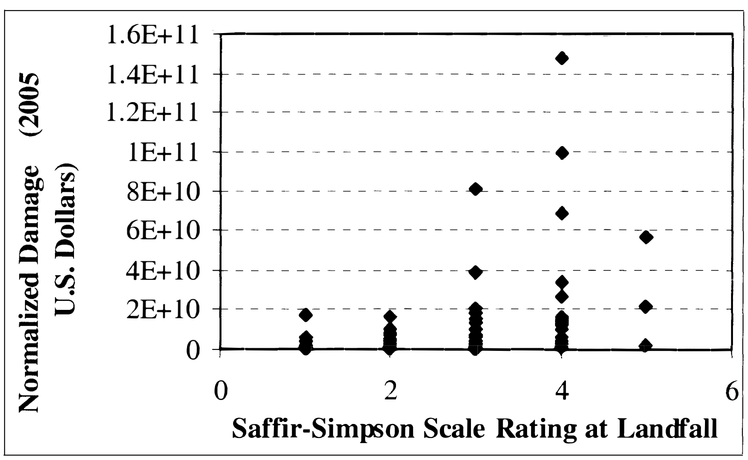

FIG. 2. Normalized hurricane damage vs. Saffir-Simpson scale rating at landfall

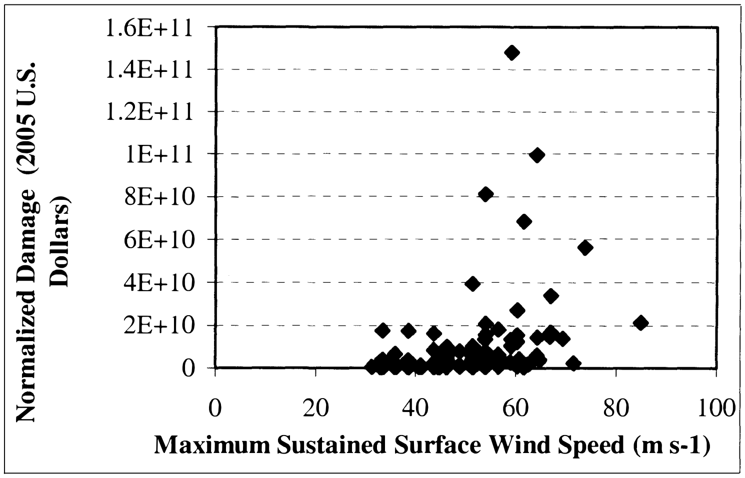

FIG. 3. Normalized hurricane damage vs. maximum sustained surface wind speed at landfall

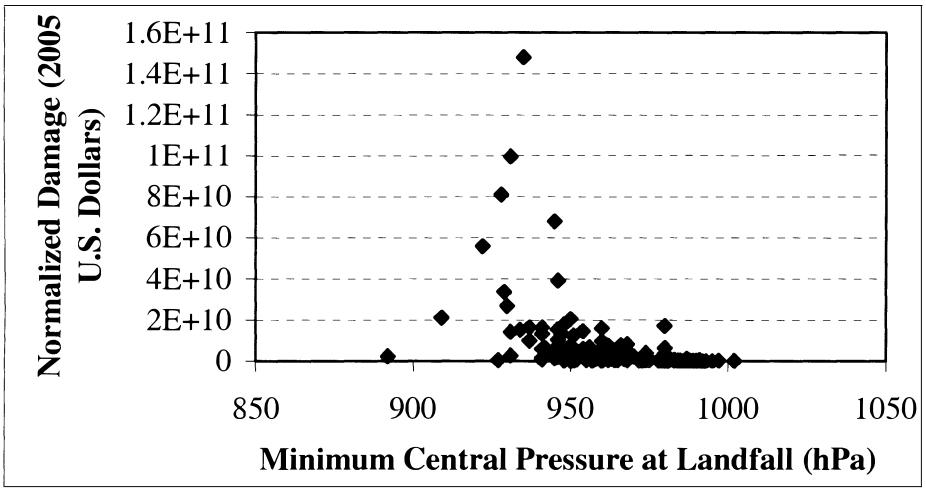

FIG. 4. Normalized hurricane damage vs. minimum central pressure at landfall

TABLE 4. Twenty-five representative models runs using the Weka software program

25 Representative Weka Runs					
Run	Algorithm	Parameter Settings	Attributes[4]	Correlation Coefficient	R^2
1	Linear Regression	Default	max_wind_speed_mph, pressure, saffir_simpson, interpol_saffir_simpson, normalized_pl05	0.31	0.096
2	Linear Regression	Default	max_wind_speed_mph, normalized_pl05	0.3239	0.105
3	Linear Regression	Default	pressure, normalized_pl05	0.3564	0.127
4	Linear Regression	Default	saffir_simpson, normalized_pl05	0.3342	0.112
5	Linear Regression	Default	interpol_saffir_simpson, normalized_pl05	0.315	0.099
6	Least Median Squares	Default	max_wind_speed_mph, pressure, saffir_simpson, interpol_saffir_simpson, normalized_pl05	0.4085	0.167
7	MultilayerPerceptron	Default	max_wind_speed_mph, pressure, saffir_simpson, interpol_saffir_simpson, normalized_pl05	0.3217	0.103
8	MultilayerPerceptron	Changed the epochs to 10000. Did not change anything else.	max_wind_speed_mph, pressure, saffir_simpson, interpol_saffir_simpson, normalized_pl05	0.0809	0.007
9	MultilayerPerceptron	Changed hidden nodes to 3,3,3. (i.e. Three layers with three nodes per layer.)	max_wind_speed_mph, pressure, saffir_simpson, interpol_saffir_simpson, normalized_pl05	0.3536	0.125
10	MultilayerPerceptron	Default	max_wind_speed_mph, pressure, normalized_pl05	0.2409	0.058

[4] Note: While there are redundancies in the attributes because more than one attribute is a direct function of wind speed, each attribute provides slightly different levels of information. For example, the dataset is grouped into fewer categories when grouped by saffir_simpson instead of max_wind_speed_mph because saffir_simpson can only be one of five values while max_wind_speed can be many more than five. This variability in the level of detail is what distinguishes each attribute that is a function of wind speed.

11	MultilayerPerceptron	Changed hidden layers parameter to 6,4,2.	max_wind_speed_mph, pressure, normalized_pl05	0.093	0.009
12	MultilayerPerceptron	Changed hidden layers parameter to 3.	max_wind_speed_mph, pressure, normalized_pl05	0.317	0.100
13	MultilayerPerceptron	Changed hidden nodes to 6,4,2.	max_wind_speed_mph, pressure, normalized_pl05	0.093	0.009
14	MultilayerPerceptron	Changed learning rate to 0.6 and momentum to 0.4	max_wind_speed_mph, pressure, normalized_pl05	0.2339	0.055
15	MultilayerPerceptron	Changed hidden layers parameter to 3,3,3.	year, max_wind_speed_mph, pressure, saffir_simpson, interpol_saffir_simpson, normalized_pl05	0.3252	0.106
16	MultilayerPerceptron	Changed hidden layers parameter to 3,3,3. Changed epochs to 1000	year, max_wind_speed_mph, pressure, saffir_simpson, interpol_saffir_simpson, normalized_pl05	0.2602	0.068
17	MultilayerPerceptron	Changed hidden layers parameter to 3,3,3. Changed momentum to 0.1	year, max_wind_speed_mph, pressure, saffir_simpson, interpol_saffir_simpson, normalized_pl05	0.3062	0.094
18	MultilayerPerceptron	Changed hidden layers parameter to 3,3,3. Changed learning rate to 0.2	year, max_wind_speed_mph, pressure, saffir_simpson, interpol_saffir_simpson, normalized_pl05	0.2955	0.087
19	MultilayerPerceptron	Changed hidden layers parameter to 2.	year, max_wind_speed_mph, pressure, saffir_simpson, interpol_saffir_simpson, normalized_pl05	0.1804	0.033
20	MultilayerPerceptron	Changed hidden layers parameter to 10.	year, max_wind_speed_mph, pressure, saffir_simpson, interpol_saffir_simpson, normalized_pl05	0.2375	0.056

Legend: *Multilayer Perceptron*=Neural network application in Weka, *Default*=Weka's default settings, *max_wind_speed_mph*=Maximum sustained surface wind speed at landfall, *pressure*=Central pressure at landfall, *saffir_simpson*=Saffir-Simpson scale rating at landfall, *interpol_saffir_simpson*=Interpolated Saffir-Simpson scale rating, *normalized_pl05*=Normalized damage using the Pielke and Landsea (1998) methodology.

Chapter 5. DISCUSSION

The analysis in this study suggests that single hurricane parameters such as minimum central pressure and the maximum sustained surface wind speed at landfall are poor predictors of the amount of storm damage. Combinations of hurricane parameters yielded similarly poor results. Thus, it will be necessary to include other parameters in order to achieve accurate hurricane damage models and more effective catastrophe bond triggers. Hurricane parameters that might considered for inclusion in catastrophe bond issues are flooding from rainfall, storm surge, radius of gale force winds, radius of maximum winds, topography, location of landfall, and angle of incidence. Societal parameters include population, wealth, and their distribution with altitude. An analysis of the relationship between these parameters and the storm damage is a future area of study, as is the distinction between wind damage and that due to flooding, a critical factor for the insurance industry.

The results presented here suggest that bond issuers should use modeled-loss[5] type bonds instead of bonds based on parametric triggers, which are subject to substantial measurement uncertainty. In order to test the robustness and validity of a given model, a bond issuer or risk modeling firm would use a similar procedure to the one employed in this study. The firm would first input parameters from past landfalling hurricanes into the model and generate corresponding damage estimates for each storm. These damage estimates would then be compared to the actual damage figures derived from either the Pielke and Landsea (1998) or Collins and Lowe (2001) methodologies in which storm damage was normalized for changes in inflation and wealth at the national level and changes in population and housing units at the coastal county level.. The firm could then

---

[5] See Table 1 for a definition of modeled-loss type bonds.

examine how well the model would have performed on past landfalling hurricanes. Models that produce damage estimates similar to the actual normalized damage figures would be considered reasonably robust.

For an alternative assessment of the impact between measurable hurricane parameters and the insurance and reinsurance industry, insured losses could be used in place of normalized storm damage in a future study. Since insurance companies cover damages resulting from wind but not floods, it is possible that the maximum sustained surface wind speed and minimum central pressure are better correlated with insured losses than normalized storm damage.

REFERENCES

Araya, Rodrigo, 2005: Catastrophic risk securitization: Moody's perspective. *Catastrophic Risks and Insurance*, **8**, 171-83.

Bantwal, V. and H. Kunreuther, 2000: A cat bond premium puzzle?. *Jrnl. of Psychology and Financial Markets*, **1**, 76-91.

Bishop, C. M., 1996: *Neural Networks for Pattern Recognition.* Clarendon Press, 482 pp.

Black, P., R. McIntosh, C. Swift, J. Carswell, K. St. Germain, I. Popstefanija, and M. Goodberlet, 1995: Ocean surface wind, stress, and rain rate measurements in tropical cyclones from concurrent airborne microwave scatterometer and radiometer measurements. *Reprints of the 27th Conference on Radar Meteorology*, Vail, CO, AMS, pp.623-625.

Black, P. G., and F. D. Marks, 1991: The structure of an eyewall meso-vortex in Hurricane Hugo (1989). Preprints, *19th Conf. on Hurricanes and Tropical Meteorology,* Miami, FL, Amer. Meteor. Soc., 579–582.

Bluestein, H. B., and F. D. Marks, 1987: On the structure of the eyewall of Hurricane Diana (1984): Comparison of radar and visual characteristics. *Mon. Wea. Rev.,* **115,** 2542–2552.

Collins, D. J., and S. P. Lowe, 2001: A macro validation dataset for U.S. hurricane models. Casualty Actuarial Society Forum, Casualty Actuarial Society, Arlington, VA, pp. 217-251, <http://www.casact.org/pubs/forum/01wforum/01wf217.pdf>.

Dong, W., H. Shah, and F. Wong, 1996: A rational approach to pricing catastrophe insurance. *Jrnl. of Risk and Uncertainty*, **12,** 201-218.

Dunion, J. P., C. W. Landsea, S. H. Houston, and M. D. Powell, 2003: A reanalysis of the surface winds for Hurricane Donna of 1960. *Mon. Wea. Rev.*, **131**, 1992–2011.

Hock, T. F., and J. L. Franklin, 1999: The NCAR GPS Dropwindsonde. *Bull. Amer. Meteor. Soc.*, 80, 407-420. Hock, T. F., and J. L. Franklin, 1999: The NCAR GPS Dropwindsonde. *Bull. Amer. Meteor. Soc.*, **80**, 407-420.

Hsieh, W.W., and B. Tang, 1998: Applying neural network models to prediction and data analysis in meteorology and oceanography. *Bull. Amer. Meteor. Soc.*, **79**, 1855-1870.

Franklin, J. L., M. L. Black, and K. Valde, 2003: GPS dropwindsonde wind profiles in hurricanes and their operational implications. *Wea. Forecasting*, **18**, 32–44.

Insurance Services Office, 1999: Financing Catastrophic Risk: Capital Market Solutions. New York, NY.

Jarrell, J. D., P. J. Hebert, and M. Mayfield, 1992. Hurricane Experience Levels of Coastal County Populations From Texas to Maine. NOAA Technical Memorandum NWS NHC-46 (August). Coral Gables, FL: NHC, 152 pp..

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 Technical Memorandum NWS NHC 22, Coral Gables, Florida, 21pp.

Kossin, J. P., J. A. Knaff, H. I. Berger, D. C. Herndon, T. A. Cram, C. S. Velden, R. J. Murnane, and J. D. Hawkins, 2007: Estimating hurricane wind structure in the absence of aircraft reconnaissance. *Wea. Forecasting,* in press.

--- and M. D. Eastin, 2001: Two distinct regimes in the kinematic and thermodynamic structure of the hurricane eye and eyewall. *J. Atmos. Sci.,* **58,** 1079–1090.

Landsea, C. W., C. Anderson, N. Charles, G. Clark, J. Dunion, J. Fernandez-Partagas, P. Hungerford, C. Neumann, and M. Zimmer, 2004: The Atlantic hurricane database re-analysis project: Documentation for the 1851-1910 alterations and additions to the HURDAT database. *Hurricanes and Typhoons: Past, Present and Future*, R. J. Murname and K.-B. Liu, Eds., Columbia University Press, 177-221.

Marks, F. D., P. G. Black, 1990: Close encounter with an intense mesoscale vortex within Hurricane Hugo (September 15, 1989). *Extended Abstracts, Fourth Conf. on Mesoscale Processes,* Boulder, CO, Amer. Meteor. Soc., 114–115.

Marzban, C., and G. J. Stumpf, 1996: A neural network for tornado prediction based on Doppler Radar-Derived attributes. *J. Appl Meteor.*, **35,** 617-626.

McCarty, Kevin, 2006: Is the housing market prepared for a natural disaster?: Testimony of the National Association of Insurance Commissioners before the Subcommittee on Housing and Community Opportunity of the House Committee on Financial Services. June 28, 2006. Rayburn House Office Building. <https://www.naic.org/documents/govt_rel_testimony_0606_mccarty.pdf>.

Munich Re, 2000: Jahresuberblick Naturkatastrophen 1999. Munich, Germany.

Partagas, J. F., and H. F. Diaz, 1996: Atlantic Hurricanes in the second half of the Nineteenth Century. *Bull . Amer. Meteor. Soc.,* **77**, 2899-2906.

Pielke, R. A., Jr., J. Gratz, C.W. Landsea, D. Collins, M. A. Saunders, and R. Musulin,

2007. Normalized hurricane damages in the United States: 1900-2005. Natural Hazards Review (submitted) <http://sciencepolicy.colorado.edu/publications/special/normalized_hurricane_damages.html>.

Pielke, R. A., Jr., and C. W. Landsea, 1998: Normalized hurricane damages in the United States, 1925-1995. *Wea. Forecasting*, **13**, 621-631.

Powell, M. D., J. D. Kepert, A. Boissonnade, and P. J. Vickery, 2004: Uncertainty in Hurricane Winds: What do new measurements and simulations tell us about Hurricane Andrew? *Preprints of the 26th Conference on Hurricanes and Tropical Meteorology,* 3-7 May 2004, Miami, Florida.

--- D. Bowman, D. Gilhousen, S. Murillo, N. Carrasco, 2004: Tropical cyclone winds at landfall: The ASOS–C-MAN Wind Exposure Documentation Project. *Bull. Amer. Meteor. Soc*, **85**, 845–851.

---, P. J. Vickery, and T. Reinhold, 2003: Reduced drag coefficient for high wind speeds in tropical cyclones. *Nature*, **422**, 279-283.

--- and T. Reinhold, 1996: Hurricane Andrew's landfall in South Florida. Part I: Standardizing measurements for documentation of surface wind fields. *Wea. Forecasting,* **11**, 304–328.

Reed, R.D., and R. J. Marks, 1999: *Neural smithing: Supervised learning in feedforward Artificial Neural Networks.* MIT Press, 352 pp.

Simpson, R. H. and H. Riehl, 1981. *The Hurricane and its Impact.* Baton Rouge, LA: LSU Press.

Uhlhorn, E. W. and P. G. Black, 2003: Verification of remotely sensed sea surface winds in hurricanes. *J. Atmos. Oceanic. Tech.*, **20**, 99-116.

---, P. G. Black, J. L. Franklin, M. Goodberlet, J. Carswell and A. S. Goldstein, 2006: Hurricane surface wind measurements from an operational stepped frequency microwave radiometer. *Mon. Wea. Rev.* (Submitted).

Wakimoto, R. M., and P. G. Black, 1994: Damage survey of Hurricane Andrew and its relationship to the eyewall. *Bull. Amer. Meteor. Soc.,* **75,** 189–200.

Witten, I. H., and E. Frank, 2005: *Data Mining: Practical Machine Learning Tools and Techniques.* Morgan Kaufmann Publishers, 525 pp.

Wurman, J. and J. Winslow, 1998: Intense sub-kilometer-scale boundary layer rolls observed in Hurricane Fran, *Science,* **280***,* 555-557.

Zimmerli, Peter, 2005: Insurance of atmospheric perils—challenges ahead. *Catastrophic Risks and Insurance,* **8**, 51-57.