

Evaluation of Catastrophe Models Using a Normalized Historical Record

Why It Is Needed and How To Do It

Roger A. Pielke, Jr.*
Christopher W. Landsea**
Rade T. Musulin***
Mary Downton****

Abstract

This article explains the role of catastrophe simulation models in managing exposures to hurricane losses. The credibility and accuracy of these models have been called into question by consumer groups and regulators. This paper suggests a practical method for establishing a baseline to test the skill of these models in estimating catastrophe losses. In establishing a baseline for model evaluation, performance can be measured quantitatively. Given the vast exposure of society to the impacts of catastrophic events, it is imperative that the financial markets of the world develop effective strategies to manage risk. It follows that such strategies will be more effective to the degree that they are based on reliable and publicly understood estimates of future risk. Finally, the transpar-

* Environmental and Societal Impacts Group, National Center for Atmospheric Research, Boulder, CO.

** Hurricane Research Division, AOML/NOAA Miami, FL.

*** Florida Farm Bureau Insurance Companies, Gainesville, FL.

**** Environmental and Societal Impacts Group, National Center for Atmospheric Research, Boulder, CO.

ent nature of the proposed methodology, in contrast to the complex and often proprietary structure of the computer models, highlights the societal factors underlying the sharp increase in catastrophic loss exposure.

Introduction

Worldwide, the costs of natural disasters have increased dramatically in recent decades (Swiss Re, 1997). Many in the insurance and reinsurance industries have noted this increase and simultaneously become aware that their exposure to catastrophic losses has been underestimated (Campbell, 1997). One result of this greater attention to the risks of catastrophic loss has been the development over the last ten years of a consulting industry that provides predictive information based on sophisticated computer models on the probabilities and consequences of catastrophic events. These consultants — including firms such as Applied Insurance Research, EQECAT, Risk Management Solutions, and Tillinghast, as well as internal experts within large insurance firms — provide quantitative estimates of risk and exposure that inform a range of financial decisions.

Generally, the models have indicated that both the rate and capitalization levels of the insurance industry have been inadequate in catastrophe-prone areas. If policy makers respond to the models, significant additional costs will be borne by consumers living in these areas, both in terms of insurance premiums and construction costs to comply with revised building codes. Predictably, the potential for consumer disruption based on computer modeling has led to considerable controversy in the heavily regulated property insurance market.

Many insurance professionals have embraced modeling as a way for the industry to correct glaring errors in prior methodologies used to estimate catastrophic exposure (Chernick, 1998). They note that the methods being used to simulate hurricane losses are similar to those accepted with little controversy in many other fields, including aircraft design, weather prediction, governmental budgeting, and financial management (Musulin, 1997). In contrast to most other industries in American society, however, residential property insurance is heavily regulated at the state level by elected officials or their appointees. This subjects the catastrophe modeling process to close public scrutiny. And unlike many other fields in which this tech-

nology has been employed, the long-term loss predictions produced by catastrophe models remain to be evaluated in a quantitative manner against a transparent, objective baseline.

Catastrophe models play a powerful role in decision making within the financial community, and the large stakes involved create a need for public understanding of their performance. These models have been extensively scrutinized in recent years by insurers, regulators, and financial rating agencies.

According to Moody's Investment Service (1997), the models are evaluated in the following manner:

- historical data to be input to the models are examined,
- parameters are stressed to assess sensitivity in outputs,
- catastrophe models are compared against one another,
- the resumes of modelers are examined.

A similar set of criteria is used by the State of Florida by its commission established in 1995 to evaluate catastrophe models in order to identify those which are "acceptable" for setting residential insurance rates (FCHLPM, 1997). The Commission also compares the meteorological characteristics of the simulated storms with the historical storms and performs many other tests on components of the models.

These efforts, while important, do not provide decision makers with a full set of information that would provide a means of evaluating the skill, quality, or value¹ of the model predictions, leading to wide criticism of the models as "black boxes" (Whitehead, 1997). Notably absent from these analyses has been a systematic reconciliation of modeled loss estimates with the long-term historical record. As a result, the insurance commissioners of both Texas and Florida have recommended that catastrophe models play no formal role in their state's insurance ratemaking (Nelson, 1998; Mah, 1997), and debate rages over the appropriate insurance premiums in locations exposed to the impacts of extreme events (Cordle, 1998).

Decision makers have long faced challenges in the evaluation of predictions that they use even as demand for predictive information

1. The terms "skill," "quality," and "value" have specific technical definitions in the forecast value literature. See Murphy (1993).

grows (Sarewitz, et al., in press). Because decisions are by their nature forward looking, decision makers seek knowledge of the future. Similarly, scientists and other experts have taken advantage of growing computer power and enormous databases to produce a growing range of predictive information. This growth in the predictive enterprise can become problematic if predictions are misunderstood, misleading, or misused. One of the central questions that decision makers ought to ask of their predictive information is: What is a good prediction and how do I know it when I see it?

To answer this question, this paper presents a transparent methodology that would allow for a quantitative evaluation of catastrophe models. The paper proceeds in four sections. The first presents the evaluation methodology. The next discusses possible applications. The next suggests how the methodology might be implemented in practice. Finally we present a preliminary comparison of loss estimates using this methodology to those generated from one of the leading computer models.

Our objective in proposing the methodology is threefold. First, in establishing a baseline for model evaluation, performance can be measured quantitatively. Second, given the vast exposure of society to the impacts of catastrophic events, it is imperative that the financial markets of the world develop effective strategies to manage risk. It follows that such strategies will be more effective to the degree that they are based on reliable and publicly understood estimates of future risk. Finally, the transparent nature of the proposed methodology, in contrast to the complex and often proprietary structure of the computer models, highlights the societal factors underlying the sharp increase in catastrophic loss exposure.

It is important to note that the proposed normalization methodology is based upon the establishment of a transparent objective baseline that the models should seek to outperform. It is not a substitute for modeling. It is, however, a necessary element in the rigorous evaluation of the predictive ability of catastrophe models.

Methodology

An evaluation methodology for catastrophe models is perhaps best introduced via the analogy of weather forecasts. How does one know

if a weather forecast is a good one? The answer is a bit trickier than one might think.

Consider the case of early tornado forecasts. In the 1880s a weather forecaster began issuing daily tornado forecasts in which he would predict for the day "tornado" or "no tornado." After a period of issuing forecasts, the forecaster found his forecasts to be 96.6% correct — a performance that would merit a solid "A" in any grade school. But others who looked at the forecaster's performance discovered that simply issuing a standing forecast of no tornados would result in an accuracy of 98.2%!

This finding suggested that in spite of the high degree of correct forecasts, the forecaster was providing predictions with little skill, and in fact could result in costs rather than benefits². In a similar manner, simply comparing the output of catastrophe models with recent events does not provide enough information with which to evaluate their performance. A more sophisticated approach is needed.

One way to evaluate weather forecasts is to compare the prediction with some baseline or standard. Climatology, i.e., historical weather information aggregated over time and space, provides such a baseline because it provides the best estimate of the future occurrence of weather events, absent any other information. A weather forecast is considered skillful if it improves upon a prediction based on climatology. For instance, the average high temperature over the past 100 years in London on September 6 might be say 10 degrees Celsius (i.e., the climatological mean for that date). Absent any other information, the best prediction of the temperature on next September 6 is thus 10 degrees. Any weather forecast for that particular day would be considered skillful if it were to improve upon climatology in comparison to the actual temperature recorded on that date.

The analogue to climatology in the insurance world is historical data adjusted to reflect current conditions. Using historical data to forecast future expected costs is the basis of actuarial science. In many lines of insurance, such as automobile, basic actuarial adjustments for inflation trends, claim reporting patterns, etc., have been carefully developed and are generally accepted by both insurers and regulators. In lines subject to severe catastrophic loss, however,

2. On the history of tornado forecasts and their evaluation see Murphy (1996 and 1997) and Galway (1985).

many actuaries have argued that historically accepted techniques for projecting future loss costs from prior loss data are fundamentally flawed (Musulin, 1997; Chernick, 1998), leaving little alternative to the use of computer modeling. While this may be true in the context of complex actuarial calculations, the historical record may be of great value in providing a benchmark against which to evaluate model output.

The fundamental problem in using historical catastrophic loss data is the infrequent nature of catastrophic events, necessitating a long experience period. The use of data from different points in time is difficult because underlying societal conditions have changed. Such changes are difficult to quantify from insurance industry databases, but can be measured by data available from other sources. A major limitation of prior actuarial methodologies was their reliance on insurance industry data alone to forecast future loss exposure.

For example, it is meaningless to compare the losses associated with the devastating 1926 Great Miami hurricane with those related to Hurricane Andrew in 1992. Even after adjusting for inflation, Miami was a far different place in 1992 than in 1926. Even the Miami of 1999 is not directly comparable to that of 1992. Computer models have sought to overcome this problem by simulating historical meteorological events with current societal conditions. But if past catastrophe losses could be effectively adjusted to reflect the change in societal conditions over time, a baseline would be developed against which their predictive skill could be evaluated, exactly analogous to the case of weather forecasts and climatology.

Normalization methods provide a means to place historical losses on an actuarial basis. The following example is based on normalized hurricane losses from Pielke and Landsea (1998). In principle, a similar methodology could be applied to any time series of catastrophes (cf. Litzenberger et al., 1996).

Normalized Hurricane Losses in the United States

To illustrate how to normalize past hurricane losses to present values, assume that losses are proportional to three factors: inflation, wealth, and population. Of course, it is possible that these factors

would be replaced and/or complemented by others, which represent changes to the insurance industry (e.g., changes in deductibles, policy types, etc.) or society in general. The result of normalizing the data will be to produce the estimated impact of any storm as if it had made landfall today.³

A generalized normalization method is determined as follows:

NL _{present} =	a storm's losses normalized to present values.
y =	year of storm's impact.
c =	county(ies) of storm's maximum intensity at landfall.
L _y =	storm's loss in year y, in current dollars (i.e., not adjusted for inflation).
I _y =	inflation factor, determined by the ratio of the present implicit price deflator for Gross Domestic Product to that of year y.
W _y =	wealth factor, determined by the ratio of the inflation adjusted present fixed reproducible tangible wealth expressed as per capita to that of year y.
P _{y, c} =	population factor, determined by the ratio of the change in the population of the coastal county(ies) most affected by the storm from year y to present.

A general formula is thus:

$$NL_{\text{present}} = L_y * I_y * W_y * P_{y, c}$$

For example, the following illustration converts the actual losses in the 1938 New England hurricane to 1995 values. This storm made

3. Inflation is accounted for using the implicit price deflator for Gross National Product, as reported in the Economic Report of the President. Wealth is measured using an economic statistic kept by the U.S. Bureau of Economic Analysis called "Fixed Reproducible Tangible Wealth" and includes equipment and structures owned by private business, owner occupied housing, nonprofit institutions, durable goods owned by consumers, as well as government owned equipment and structures. Wealth is accounted for in the normalization using a ratio (inflation adjusted) of today's wealth to that of past years (end of year gross stock). Because the measure of wealth is based on national figures, we have adjusted it back to per capita by removing from it the relative change in the entire U.S. population. Wealth data are available from 1925, consequently the normalization begins with that year. The final factor is population change based on data from the U.S. Census for each of the 168 coastal counties that lie along the coast from Texas to Maine.

landfall as a Category 3 hurricane through the states of New York, Connecticut, Rhode Island and Massachusetts causing an estimated \$306 million damage. The population of the 11 coastal counties impacted at that time was 2.336 million, while the 1995 estimated population had increased to 4.860 million, a factor of 2.08. The inflation and wealth factors are 11.75 and 2.224, respectively, between 1938 and 1995. Thus, the normalized damage that would be attributed to the 1938 New England hurricane if it struck in 1995 is the following:

$$\$306 \text{ million (1938)} \times 11.75 \times 2.224 \times 2.080 = \$16,629 \text{ million (1995)}$$

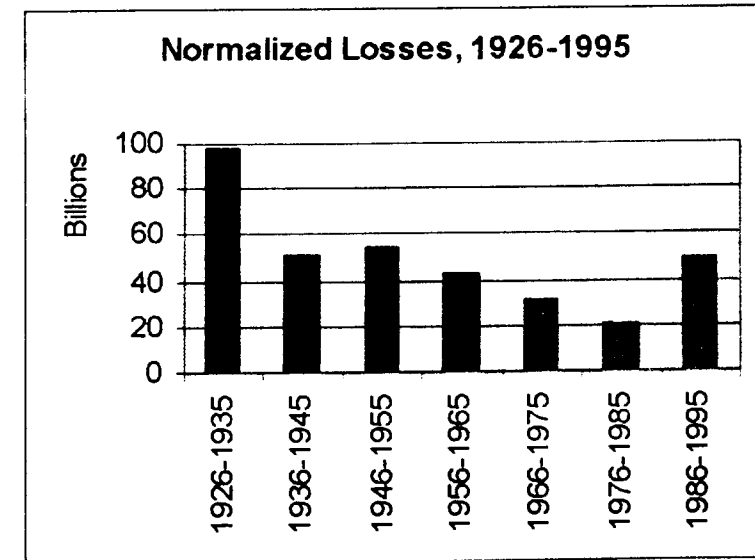
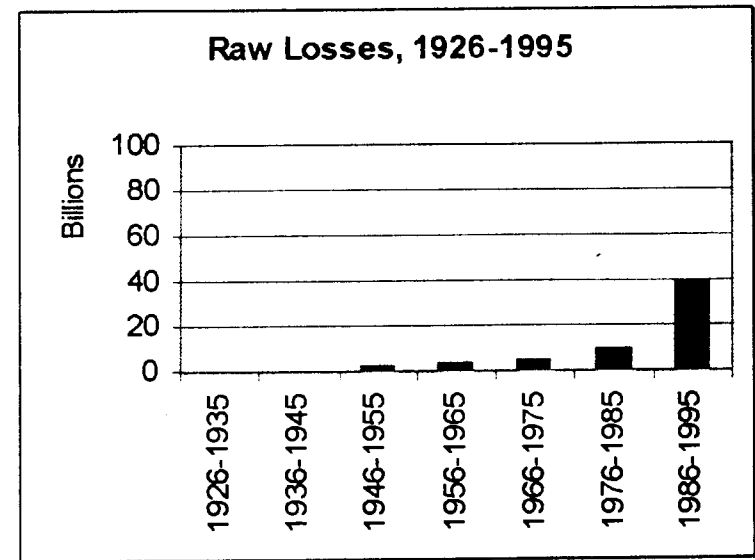
The non-normalized and normalized trend data on annual hurricane impacts in the United States from 1925 to 1995 are shown in Figure 1. It shows the estimated losses associated with ten year periods ending in 1995, as if each year's storms had made landfall in 1995. Clearly, the normalized information shows a much different time series than the data adjusted only for inflation.

Additional adjustments to the historical data would be required before detailed comparisons to model output could be made. For example, storm damage is generated by structures, rather than people, and the ratio of population to housing units has changed over time, as housing units have been growing faster than population. Another possible adjustment could account for the increase in the number of secondary residences and vacation homes in recent decades, which one would expect to be disproportionately located in high risk coastal areas. The census bureau database provides a wealth of information on these types of demographic trends.

Absent other information, the normalized record is the "best estimate" of expected losses, much like climatology is the best estimate of future weather. This "catastrophe climatology" provides a baseline against which one can evaluate the performance of a particular model or compare models against one another. Catastrophe models would show skill (and thus be more likely to have economic value to decision makers) if they can improve upon the loss estimates provided by the simple estimate.

In theory, the simple estimate should prove to be an easy target to beat, since catastrophe models have the advantage of contextual information. But the point of creating a "catastrophe climatology" is not to improve upon the catastrophe models, but instead to pro-

FIGURE 1



vide a baseline against which their performance can be measured in order to assess their value in decision making, and to stimulate future improvements in their predictions. A systematic effort to reconcile the available historical data with the output of computer models would also have great value in explaining the apparently large differences in loss potential predicted by models and that evident in the historical record, increasing public understanding of the modeling process and identifying areas where it needs to be improved.

Applications of the Methodology

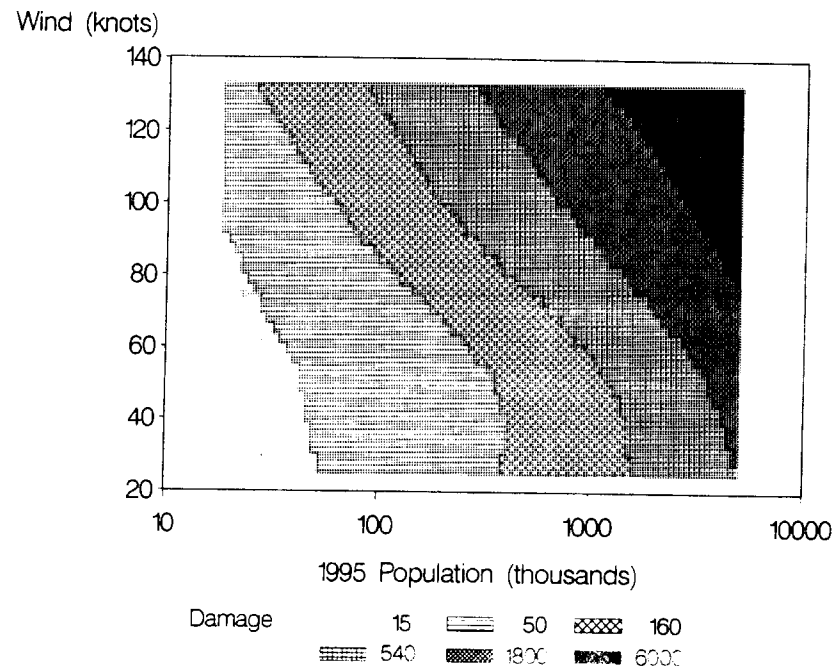
The development of a "catastrophe climatology" could be used to evaluate model performance in terms of loss estimates for particular storms, loss exceedence probabilities, probable maximum losses, etc., and to compare models against one another.

Single Storm Loss Estimates

In the aftermath of a hurricane's impact, there is immediate demand for information about the total losses. Catastrophe modeling firms provide such estimates, which could also be derived from a normalized record of past losses. For example, assume that a Category 4 hurricane makes landfall in New Orleans. One way to estimate losses is to look at the normalized damages of past storms, which make landfall in regions of similar population (or insurance coverage). Using the Pielke-Landsea (1998) database for illustrative purposes, Figure 2 shows how total normalized losses (on a logarithmic scale) compare with a rough estimate of storm wind speed at landfall. Figure 3 shows the distribution of events on this same scale. Using this diagram to estimate the damages of a landfalling storm, one would need only know the wind speed and affected population. Of course, in practice, one would want to be more sophisticated, but nonetheless, the general approach would remain the same.

To see how the simple methodology could be used to evaluate the performance of the models, consider the following example. Assume that for Company A the distribution of past analogue events suggests that a particular storm's impact will result in \$100 million in losses. A catastrophe model would show skill if its estimate of the losses (including associated uncertainty) were shown to be closer

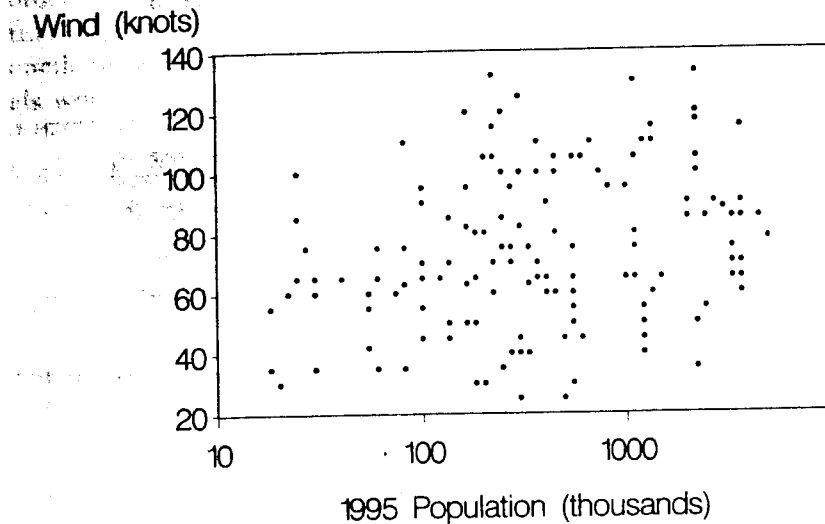
FIGURE 2
Normalized U.S. Storm Damages, 1925–1995
(millions of 1995 dollars)



than the estimate generated from the normalized historical record to the actual loss amount. Further assume that for Company A, Catastrophe Model B predicted losses of \$200 million, Cat Model C, \$175 million, and Cat Model D \$50 million. If the actual losses borne by Company A were later determined to be \$150 million, then only Cat Model C could be said to have useful skill, because its prediction improved upon that of the "catastrophe climatology." Cat Model B, missing the mark by an equal amount as simple forecast, showed no skill and Cat Model D showed negative skill in that it performed worse than the simple forecast.

Of course, in a real setting one would want to compare probability distributions of losses so as to consider uncertainty in the estimate. But, again, the general idea would be the same. It would be possible to review model performance retrospectively as well as in

FIGURE 3

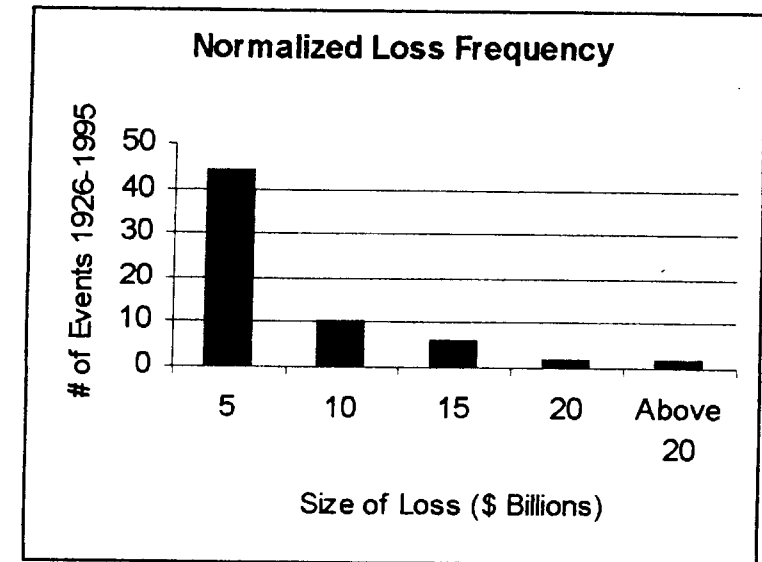


real time to provide their users with a sense of the relative value that they add to decision making.

A very rough measure of expected insured losses is easily computed on the back of an envelope. We have obtained from Property Claims Services, Inc. their record of insured losses due to hurricanes, which dates to 1949 (52 storms) and adjusted these numbers using the Pielke/Landsea methodology to 1995 values. The adjusted insured losses for Andrew (1992) and Hugo (1989) in 1995 dollars (i.e., had these storms made landfall in 1995) are \$19.4 billion and \$5.6 billion, respectively. These should be considered rough calculations, but nonetheless illustrate how a simple and transparent methodology can be used to generate baseline information that would allow for the evaluation of catastrophe models.

Loss Exceedence, Probable Maximum Loss

In a similar fashion to predicting and evaluating losses to a portfolio from a single storm, the historical database can be used to generate normalized estimates of loss exceedence probabilities and probable maximum losses. To assess loss exceedence probabilities, one can

FIGURE 4
Normalized Loss Frequency

simply graph normalized losses in terms of a frequency distribution. Figure 4 shows such a graph based on the Pielke/Landsea database of total societal losses. A graph based on losses to individuals' portfolios would have many more data points. Such a graph can be compared with model-based loss exceedence probabilities in order to determine whether or not the models represent an improvement over those generated from the historical record. Similarly, the worst-case event (in terms of losses) in the normalized record provides some guidance and independent calibration to model-based estimates of large losses. Consider that the worst case in the Pielke/Landsea database was the 1926 Miami storm, which resulted in more than \$80 billion (1998) in total normalized losses, or almost three times those of Andrew⁴.

4. One would expect the 1926 Miami hurricane to result in greater losses than Andrew because of its larger size and landfall over downtown Miami.

Implementation of the Methodology

The methodologies and results presented above represent a simplified way to extract useful information from the historical record. Further refinement of the methodology is essential before detailed comparisons are made to the output of computer models. Nonetheless, a wealth of information can be extracted from the historical record if past loss data can be suitably normalized.

To turn this methodology into something of practical value will require the support and participation of the insurance and reinsurance industries. Given the fact that these firms are in many cases betting their corporate existence on both the accuracy of the models and an ability to convince regulatory authorities to allow their use in ratemaking, it would seem that it would be in the community's interest to better understand the quality of services that they purchase. Further, the catastrophe modeling firms themselves should also encourage the evaluation of their products, to stimulate competition in the industry, to price their products commensurate with performance, as well as to establish an empirical basis for penetrating regulated insurance markets.

The following three steps suggest one way that the methodology might be implemented at low cost to wide benefit.

1. Determine the appropriate components of an insurance-based normalization

Experts from the insurance industry should meet to discuss and debate how a normalization methodology might be implemented to reflect the changing nature of insured losses. The Pielke/Landsea normalization was designed to adjust total losses by rather blunt factors. It may be possible for insured losses to be adjusted with greater precision, both in terms of data resolution and those factors most important to the growth in insured losses. The central output from this step would be an authoritative, consensus statement by industry experts as to what the most appropriate "catastrophe climatology" would be for the insurance industry.

While it is impossible to check the normalization process directly (since identical events rarely affect the same geographic location at different times), one can test how well the normalization

process can explain actual loss events. For example, losses from a storm in 1999 should be consistent with those from a past storm of similar meteorological intensity once adjustments are made for different societal conditions. Another way to evaluate the normalized database would be to assess whether or not the societal adjustments reveal the underlying climatic trends, which are objectively known. For instance, Pielke and Landsea (1999) show that their normalized damage dataset accurately reflects the climatic shifts associated with El Niño and La Niña. This provides confidence that the societal adjustments are reasonable. The resulting baseline would establish the foundation for the following steps.

2. Establish a center for the collection and storage of loss data

The power of the methodology lies in the cumulative body of experience of catastrophe losses, which is much more than simply the number of past hurricane events. It is the number of past events times the number of relevant portfolios. Thus, while any one company might not have a wealth of data on losses (or predictions), industry-wide there exists a large body of experience. Therefore, to best evaluate the information provided by catastrophe models, it would be in the industry's interest to share data on losses and on the predictions that they purchase. This information would be of use even well after the event. Given the potentially sensitive nature of a company's loss information, consideration should be given by the industry to empowering an industry association (such as the Reinsurance Association of America), an information-oriented group (like the Institute for Building and Home Safety), an existing organization in the financial services industry (such as Moody's), or an association of U.S. state insurance commissioners to oversee the collection and storage of historical loss data.

3. Evaluate model performance

The data gathered in Step 2 would then be compared against the "catastrophe climatology" generated in Step 1. The result would be a quantitative assessment of the skill of catastrophe models. There will likely be considerable need for technical analysis of the model's performance—including considerations of uncertainty—and it is

likely that such an evaluation will reveal strengths and weaknesses in the models. This process could lead to refinement of the models themselves. Further, if modelers are aware that they will be measured against a transparent baseline, this will provide incentives to use whatever information is available to improve model performance. For instance, it is well-established that the phase of the El Niño cycle is correlated with hurricane activity and related to hurricane damages (Pielke and Landsea, 1999). It is likely that such climatological information could play a role in improving catastrophe model performance. The evaluation of catastrophe models would need to be done on an ongoing basis as society changes and more events and predictions take place.

An important product of this effort would be a transparent reconciliation of historical data and model forecasts. For example, it should be possible to identify a number of variables that can account for the difference in actual losses in a given event and a model's estimate of the losses that would be generated by the same storm today. The list of variables might include changes in population, number of housing units, inflation, wealth, construction practices, or the ratio of commercial to residential structures. Identifying and measuring the relative contribution of the social factors driving the increase in catastrophic loss potential would serve several public policy objectives, including building a greater degree of understanding of the validity of model forecasts and prioritizing possible loss mitigation strategies.

The application of normalization is not limited to calculating aggregate event losses or testing computer models. This technique also provides a means of adjusting an individual insurer's historical data to reflect current conditions in a way that overcomes many of the limitations of traditionally accepted methods such as the "excess wind procedure." While normalization of past loss experience cannot provide the complex data generated by computer models, it can allow actuaries to "salvage" much more information from the historical record.

A Comparison of Normalized Losses to Those Generated by a Computer Model

We compared the results of the normalization methodology to those generated by a leading computer model. Applied Insurance Research

(AIR) furnished a dataset consisting of its estimates of the insured losses that would have been generated by actual hurricane loss events had they occurred in 1998, reflecting 1998 population, housing stock, insurance coverages, etc. The authors calculated total losses, both insured and uninsured, resulting from the same storms, using the methodology outlined earlier in the paper with the addition of one variable:

$H_{y,s}$ = housing factor, determined by the ratio of the change in the housing units in state "s" to the change in its population from year y to present.

The general formula used to normalize losses was,
 $NL_{\text{present}} = L_y \cdot I_y \cdot W_y \cdot P_{y,c} \cdot H_{y,s}$

Modeled insured losses were 56% of the normalized total losses for the period 1952–1995. Comparisons of PCS insured losses to total losses shows that about 50% of total losses have been insured in recent events. The ratio of modeled insured losses to normalized total losses has been roughly stable over several decades and is similar to the ratio of PCS insured to total losses. This is consistent with a conclusion reached by Pielke and Landsea (1998) that the results generated by computer models appear reasonable in the aggregate.

Conclusion

Properly used, predictive information has an important role to play in decision making. But misused or misunderstood predictions can lead to bad decisions. Using predictions without understanding their skill leaves open the possibility of their misuse. Catastrophe models play an important role in providing decision makers with information that allows for the securitization of risk, more accurate rate-making, improved insurer solvency monitoring, and evaluation of loss mitigation alternatives. But at the same time, they produce predictions that have not been adequately reconciled with the long-term historical record. This has severely inhibited the effective application of advanced modeling technology in the insurance industry's ratemaking process. This paper provides an additional methodology that can evaluate the performance of catastrophe models

against a transparent, objective baseline. This can serve to improve model performance, decision making by the insurance and reinsurance industries, and ultimately, how society responds to extreme events.

References

- Campbell, J., 1997. "Before the Next Cataclysm," *Regional Review*, 7:6-11, Boston: Federal Reserve Bank of Boston.
- Chernick, D. R., 1998. "Pricing the Hurricane Peril—Change is Overdue," *Casualty Actuarial Society Forum*, Winter: 23-54.
- Cordle, I. P., 1998. "A Storm Over Hurricane Insurance: Why We Pay More Than Other Areas," *The Miami Herald*, June 7: 1A.
- Florida Commission on Hurricane Loss Projection Methodology, 1997. *Report of Activities as of May 29, 1997*, State Board of Administration of Florida, Tallahassee, FL.
- Galway, J. G., and J.P. Finley, 1985. "The First Severe Storms Forecaster," *Bulletin of the American Meteorological Society*, 66:1389-1395.
- Litzenberger, R. H., D. R. Beaglehole, and C. E. Reynolds, 1996. *Fixed Income Research: Assessing Catastrophe-Resinsurance-Linked Securities as a New Asset Class*, New York: Goldman Sachs.
- Mah, C. H., 1998. Commissioner's Bulletin No. B-0015-98, Letter to all insurance companies, corporations, mutuals, associations, or other insurers writing property insurance in the State of Texas, Texas Department of Insurance, Austin, TX.
- Moodys Investors Service. 1997. *Moody's Approach to the Rating of Catastrophe-Linked Notes*. Report Number 26481. New York: Moodys Investors Service.
- Murphy, A. H., 1996. "The Finley Affair: A Signal Event in History of Forecast Verification," *Weather and Forecasting*, 11:3-20.
- Murphy, A. H., 1993. "What Is a Good Forecast? An Essay on the Nature of Goodness in Weather Forecasting," *Weather and Forecasting*, 8:281-293.
- Musulini, R. T., 1997. "Issues in the Regulatory Acceptance of Computer Modeling for Property Insurance Ratemaking," *Journal of Insurance Regulation*, Spring: 342-359.
- Nelson, B., 1997. Letter to the Florida State Legislature, August 21, Florida Department of Insurance, Tallahassee, FL.
- Pielke, Jr., R. A. and C. W. Landsea, 1999. "La Niña, El Niño, and Atlantic Hurricane Damages in the United States," *Bulletin of the American Meteorological Society*, 80: 2027-2033.
- Pielke, Jr., R. A. and C. W. Landsea, 1998. "Normalized Hurricane Damages in the United States: 1925-1995," *Weather and Forecasting*, 13:351-361.
- Sarewitz, D., R. A., Pielke, Jr., and R. Byerly, 2000 (in press). *Predictions: Decision Making and the Future of Nature*, Washington, DC: Island Press.
- Swiss Re, 1997. "Natural Catastrophes and Major Losses in 1996," *Sigma*, 3:1-38.
- Whitehead, S., 1997. "Risky Business: Proprietary Modeling and Insurance Ratemaking," *Journal of Insurance Regulation*, 15: 372-381.