Defense Resources Management Institute

WORKING PAPER SERIES

2006/05

Public Evacuation Decisions and Hurricane Track Uncertainty

Eva Regnier

Public officials with the authority to order hurricane evacuations face a difficult trade-off between risks to life and costly false alarms. Ideally, the decision should also be informed by measures of uncertainty and estimates of the value of waiting for updated, and more accurate, forecasts. Using a stochastic model of storm motion derived from historical tracks, this paper explores the relationship between lead time and track uncertainty for Atlantic hurricanes, and the implications for evacuation decisions and false alarm costs. Typical evacuation clearance times and forecast uncertainty imply that public officials requiring a 10% probability of failing to evacuate before a striking hurricane (a false negative) must accept that at least 75%, and for some locations over 90%, of evacuations will be false alarms. Reducing the decision lead time from 72 to 48 hours for major population centers could save on the order of $1B in evacuation costs annually.

Defense Resources Management Institute
School of International Graduate Studies
Naval Postgraduate School
Building 234
699 Dyer Road
Monterey, CA 93943-5138
831-656-2306
www.nps.navy.mil/drmi

The views herein expressed reflect the views of the author and are not the policy of DRMI, the Naval Postgraduate School, or the U.S. Government.

Defense Resources Management Institute
Naval Postgraduate School

The Defense Resources Management Institute (DRMI) is an educational institution sponsored and supervised by the Secretary of Defense and located at the Naval Postgraduate School in Monterey, California. Since 1965 the Institute has conducted professional education programs in analytical decision making and resources management for military officers of all services and for senior civilian officials of the United States and 149 other countries.

The mission of the Defense Resources Management Institute is to educate civilian and military government officials in the principles and concepts of defense resources management.

The unique course offerings create an interdisciplinary, interactive learning environment that facilitates the exchange of ideas among U.S. and international participants in Monterey, California, and locations worldwide. These programs are presented on a regularly scheduled basis at DRMI in Monterey and by specific arrangement at other locations in the United States and in other countries.

The success of DRMI reflects the breadth and depth of the in-house technical expertise that DRMI can draw upon. The DRMI Faculty, a community of scholars and practitioners, are experts in defense resources management and has published in major academic and technical journals. The faculty has extensive experience in providing technical assistance and educational programs worldwide over the past 40 years. Our educational strategy is one of collaboration, tailored to the specific environment of the participant.

The Defense Resources Management Institute specializes in four broad educational areas:

- Economic analysis of the public sector
- Operations research of public sector decision problems
- Public budgeting and fiscal management
- Defense organizations and management methods

For more information about our educational and technical assistance activities, please visit our website at **http://www.nps.navy.mil/drmi** or email us at drmiadmin@nps.edu.

Under second review in Management Science

Public Evacuation Decisions and Hurricane Track Uncertainty

Eva Regnier
Assistant Professor
Defense Resources Management Institute
Naval Postgraduate School
699 Dyer Road
Monterey, CA 93940

November 28, 2006

Abstract

Public officials with the authority to order hurricane evacuations face a difficult trade-off between risks to life and costly false alarms. The evacuation decision must be made on the basis of imperfect information, in the form of a forecast. Ideally, the decision should also be informed by measures of uncertainty and estimates of the value of waiting for updated, and more accurate, forecasts. Using a stochastic model of storm motion derived from historical tracks, this paper explores the relationship between lead time and track uncertainty for Atlantic hurricanes, and the implications for evacuation decisions and false alarm costs. Typical evacuation clearance times and forecast uncertainty imply that public officials requiring a 10% probability of failing to evacuate before a striking hurricane (a false negative) must accept that at least 75%, and for some locations over 90%, of evacuations will be false alarms. Reducing the decision lead time from 72 to 48 hours for major population centers could save on the order of $1B in evacuation costs annually. The savings differ by geographic location: for example, reducing the evacuation decision lead time from 72 to 48 hours lowers the false alarm rate from 76 to 69% at Miami, but only from 82 to 81% at New Orleans. At all lead times, New Orleans and Eastern Long Island would need to accept a higher rate of false alarm evacuations than Miami to achieve the same probability of a false negative.

1 Introduction

In recent decades, meteorologists have improved hurricane track forecast accuracy dramatically, and five-day forecasts are as accurate as three-day forecasts were fifteen years ago. At the same time, the stakes involved in balancing the costs of hurricane evacuation against the risks of inadequate preparation have grown with coastal population.

Pielke et al. (2006) estimate that, normalized for inflation and for growth in population and wealth, the long-term average economic cost of hurricanes is $10-11B per year in the U.S., updating their earlier estimate (Pielke and Landsea 1998). The damage caused by 2005's Hurricane Katrina, excluding deaths, is estimated by Knabb et al. (2005) at $81B and by Munich RE Group (2006) at $125B. Most hurricane damage can be prevented only by making infrastructure resistant to wind and flooding hazards. On the other hand, the costs of preparing for hurricanes—for example by evacuating people or moving assets such as ships and aircraft—can be reduced by improving forecasts, preparation processes, and decision processes.

The cost of evacuations is difficult to estimate, but a value of $1M per mile of coastline evacuated is commonly used (Whitehead 2003; Adams and Berri 1999).[1] The bus fire that killed 24 people evacuating to

[1] The actual cost depends on the severity of the storm and the population of the evacuated regions, among other factors. Whitehead (2003) used the results of intereviews conducted in North Carolina following Hurricane Bonnie to estimate the costs of mandatory evacuation orders, and arrived a estimates of $15 to $50M per storm in North Carolina, which is substantially less than $1M per mile. However, Whitehead included only direct costs and costs of evacuees' time, excluding economic disruption, lost wages, and public costs to implement the evacuation. For these and other reasons, his estimates likely understate the true costs of evacuations averaged over the vulnerable coastline. Adams and Berri (1999) argue that the costs may reach $50M per

escape Hurricane Rita in 2005 tragically demonstrated that false alarms can cost lives as well as dollars.

For a rough estimate of annual costs, assume all coastline placed under tropical storm warning is evacuated. The lead time for tropical storm warnings is 24 hours before the arrival of gale force winds; this is also approximately the evacuation lead time for small cities. Further assuming three tropical cyclone landfalls per year—the average during the relatively quiet 1976-2000 period (Powell and Aberson 2001)—and 460 miles of coastline evacuated per storm—the average length of coastline put under tropical storm warning (Jarrell and DeMaria 1999)—the total cost of evacuations comes to about $1.4B annually. On average, hurricanes affect about 150 miles of coastline per landfall, and therefore close to $1B of the evacuation cost is retrospectively unnecessary, i.e. the cost of false alarms.

For major population centers, the evacuation decision time can be 60 to 72 hours before hurricane landfall. At these lead times, the average track forecast error is 200-250 miles, and intensity (wind speed) forecast error is about 20 miles per hour. Public officials charged with ordering hurricane evacuations therefore face a difficult and dynamic problem in balancing the improving accuracy of updated landfall forecasts against the increasing risk associated with delaying an evacuation. The key to preventing future tragedies is not simply to order evacuations every time a hurricane threatens a population center. Rather the solution will lie in making the best possible use of available forecasts, taking into account their uncertainty, and balancing the lead time required for completing evacuation against the information that can be gained by waiting.

Decision support provided by the Federal Emergency Management Agency (FEMA) to public officials frames the evacuation decision as a one-time decision whose timing is determined by a point estimate of the time required to complete the evacuation. The value of waiting for updated forecasts is therefore neglected. Similar to ignoring option values in the analysis of investments, failing to account for the possibility of taking action after more information is received prevents a decision-maker from optimizing over all available alternatives. This framing also fails to provide a context for identifying time-scales at which improvements in evacuation processes to reduce cost or evacuation time would have the most value. Perhaps most important, the one-time decision framing fails to motivate planning for evacuations initiated with less than the designated lead time. This can happen when evacuation does not appear necessary at the designated lead time, but the storm takes an unexpected turn and later forecasts show the risk of a strike increasing. At this point, some preparations can still be made, such as local evacuation to high ground. However, little time remains for analysis once this situation arises, so it is very important to have considered short lead time strategies well in advance.

The purpose of this research is to make the connection between the timing of evacuation decisions and the time profile of information quality. Section 2 introduces the hurricane evacuation decision context for public officials on the Atlantic and Gulf coasts of the U.S. Section 3 describes the forecasts available to support these decisions; to the extent possible, measures of forecast accuracy are included. Earlier results indicate substantial geographic differences in track uncertainty (Regnier and Harr 2006); therefore, this paper presents results specific to four target locations—New Orleans, Miami, Norfolk, and Montauk, Long Island—highlighting the critical lead times that public officials must exploit to avoid future tragedies. Sections 4 through 6 use a Markov model of storm motion, developed in Regnier and Harr (2006), to quantify the time profiles of track uncertainty for storms affecting the four target locations. Section 4 measures track uncertainty in terms of strike probability. Section 5 examines the false alarm rate as a function of the probability-of-detection and lead time, and highlights potential reductions in false alarm rates that would be associated with reductions in the evacuation decision lead time. Section 6 translates false alarm rates

mile in some regions. The $1M per mile figure has been used for years without adjustment for inflation; because these estimates are very rough, I continue to use it without adjustment.

into annual expected number of false alarms, evacuation costs in dollars, and risk of a hurricane strike at an unprepared target. Section 7 concludes with a discussion of the results, including geographic differences and recommendations for exploiting the dynamic nature of hurricanes and hurricane forecasts.

2 The evacuation decision

The decision whether and when to initiate a public evacuation is generally made by an elected official. The locus of authority for this decision varies by state: for example, in South Carolina, only the governor can order an evacuation, whereas in North Carolina and Georgia the authority is at the local level. In some localities, there may also be degrees of evacuation that officials can choose, ranging from voluntary to mandatory, from limited to jurisdiction-wide. These options also vary by state—in Texas, even the governor lacks authority to order a mandatory evacuation.[2] The official must also decide which regions and which categories of people will be evacuated. Elected officials are usually guided in this decision by professional emergency managers who are more familiar with the decision context.

Emergency managers and elected officials generally receive National Hurricane Center (NHC) forecast information through a decision-support software called Hurrevac. Hurrevac structures the evacuation decision as a one-time decision, whose timing is set working back from the anticipated (deterministic) arrival of 39-mph (gale-force) winds, leaving enough time to complete an evacuation. Gale-force winds arrive approximately 10 hours before the storm's center (Powell and Aberson 2001).

Figure 1 is a screen shot from Hurrevac, showing the Katrina forecast on Friday, August 26th, 2005 at 11 p.m. The evacuation decision timing is shown in a pop-up in the top right; strike probabilities are displayed along the coast. Hurrevac fixes the evacuation decision time for various target locations and scenarios at somewhere between Friday night and Saturday morning, at least 42 hours before the (forecast) arrival of hazards (usually, gale-force winds).

The timing of the hurricane evacuation decision is determined based on a parameter called the evacuation clearance time, which is the time required from start to finish of an evacuation, rather than the time each evacuee spends evacuating. Evacuation clearance times are estimated based on simulations of the transportation network.

A typical evacuation clearance time for a small city is 24 hours; adding 10 hours' for the advance arrival of gale-force winds, the decision lead time will be at least 30 hours before the arrival of the storm's center. For large population centers, decision lead times are substantially longer. For Palm Beach County, the estimated evacuation clearance time exceeds 40 hours under some transportation scenarios (U.S. Army Corps of Engineers).[3] This is likely an underestimate: the population in that region has been expanding since the 1991 estimate, but the capacity of the major escape route has not grown apace. Evacuation decision lead times can easily reach 60 to 72 hours for many heavily populated areas. The New Orleans emergency management plan calls for ordering evacuation of vulnerable populations about 70 hours before the arrival of gale force winds (U.S. Army Corps of Engineers), for a total of approximately 80 hours before the arrival of the storm's center.

Another factor is uncertainty in the time-to-landfall: the mean absolute error in landfall timing for 1976-2000 was 8.8 hours for forecasts at lead times of 31-42 hours and 11.5 hours for forecasts at lead times of 43-54 hours (Powell and Aberson 2001). Extending the decision lead time with safety buffers not only decreases the accuracy of the estimate of time-to-landfall, but also decreases the accuracy of track and intensity forecasts available at the decision time.

[2] B. Long, FEMA Region IV, personal communication, 21 May, 2004.
[3] Lower Southeast Florida Hurricane Evacuation Study

[Figure 1 about here.]

Whether and when to initiate a costly evacuation are tactical decisions whose parameters—costs, time requirements, risk-reduction impacts—are determined by strategic decisions that are made long before a specific hurricane threatens. Strategic decisions that create the decision space for the tactical evacuation decision include:

- whether to invest in upgrades to the transportation infrastructure to speed evacuations;
- the delineation of evacuation zones that will create the increments of population whose evacuations can be staged; and
- the development of evacuation plans with multiple implementation timelines, e.g. the creation of 72-hour, 48-hour, and 36-hour evacuation plans options.

The cost-benefit tradeoffs for both strategic and tactical decisions are location-specific. They depend, of course, on available evacuation options vary with population, transportation infrastructure, and local topography. For example, in southern Louisiana, the coastal region is low-lying, so short-distance evacuation options are more limited. They also depend on the time-profile of forecast uncertainty. For example, the benefit of transportation upgrades depends on whether the potential reduction in evacuation time and decision lead time will delay the decision enough to substantially reduce false alarms—which may also differ by geographic location. At a tactical level, officials facing a 10% strike probability with 72 hours' lead time might appropriately decide to delay evacuation to wait for an updated forecast if they were located at Miami, but immediately take action if they were located at New Orleans. It is important for emergency managers and public officials, as well as owners of mobile assets such as ships and aircraft, to the quality of forecast information and how it changes with lead time.

3 Forecasts to support the evacuation decision

This section describes the forecast products available to support hurricane evacuation decisions. Tropical cyclones in the North Atlantic are divided into seven categories, in increasing order of intensity: tropical depressions, tropical storms, and category 1 through 5 hurricanes as defined on the Saffir-Simpson scale (http://www.nhc.noaa.gov/aboutsshs.shtml). Tropical storms and hurricanes are also called "named tropical storms". When a tropical cyclone forms, the NHC begins issuing storm-specific forecasts at six-hour intervals. The forecasts include track and intensity predictions, indicators of the uncertainty in the track, and a qualitative verbal discussion. The NHC also provides historical analyses of forecast error, which also form the basis of real-time measures of uncertainty about the track of current storms, including probabilistic forecasts. Although private weather services providers occasionally add their own analyses, almost all hurricane forecast information that the public receives is a repackaged form of NHC data.[4]

3.1 Deterministic forecast

The official NHC track predicts the position of a storm's center at 12-hour intervals to 48 hours in the future, and 24-hour intervals to 120 hours in the future. These positions are connected by a "skinny black line" to form a track (Figure 2). The official track is a subjective blend of the output of multiple deterministic physical models of the atmosphere and oceans, often called numerical weather prediction (NWP) models, and statistical models based on historical storm data.[5] Each model yields a storm track. The NHC forecasters

[4] J. Franklin, NOAA NHC personal communication, 27 Jan. 2006.

[5] For more information on the models the NHC forecasters currently use, see http://www.nhc.noaa.gov/aboutmodels.shtml.

consider the track generated by each of the models, as well as critical observations, which may include *in situ* aircraft data and remotely sensed satellite data. A considerable amount of judgment goes into the official track, and there is uncertainty about the track that varies in accordance with many factors, including location, intensity, and forward motion. The predictability of hurricane track and, correspondingly, the accuracy of forecasts, is a declining function of lead time (Figure 3).

[Figure 2 about here.]

[Figure 3 about here.]

Hurricane forecasters generally assess their performance based on the error in forecasts for the timing and location of landfall (Powell and Aberson 2001), or on track error (over land or ocean) for a given lead time (McAdie and Lawrence 2000). Landfall location is perhaps the most decision-critical variable that the NHC forecasts, as most valuable assets—both property and human lives—are onshore, and once over land, storms begin to lose their intensity. Between 1970 and 1998, 72-hour hurricane track forecast errors declined from a five-year average of about 470 mi to an average of about 255 mi (McAdie and Lawrence 2000), and longer-range forecasts have improved enough to allow for five-day forecasts that are as accurate as the three-day forecasts were fifteen years ago. Like track errors generally, landfall location errors have been declining in recent decades (Franklin et al. 2003). The number of landfalls per year is so small that trends are difficult to detect, which may explain earlier indications that landfall forecast accuracy improvements were negligible (Powell and Aberson 2001).

The official forecast also includes a forecast maximum sustained wind speed, also called intensity, at the same forecast intervals used for the track forecast; geographic extent of winds is also forecast out to 72 hours. Although the NHC verbal advisory (Section 3.3) sometimes says that intensity errors have averaged 20 kts per day, over the past 10 years the record shows intensity errors of about 10 knots for the 24-hour forecast, 15 for the 48-hour forecast and 19 for the 72-hour forecast.[6] The NHC also issues a wind-speed probability table which estimates the probability of sustained winds exceeding given thresholds; these forecasts are for maximum sustained winds of the storm, and are therefore not directly linked to the future location of the storm. They are calculated in a manner similar to the old strike probabilities described in Section 3.2.1. This product may be discontinued now that the wind-speed probability model, which generates location-specific probabilistic wind-speed forecasts (Section 3.2.2) is in place.

3.2 Probabilistic forecasts

In addition to historical error data, the NHC issues three types of uncertainty measures in real time: the error cone, probabilistic forecasts, and verbal messages. The error cone and probabilistic forecasts are based on historical forecast errors. The error cone—the white area shown in Figure 2—is simply the envelope formed by connecting circles drawn around each forecast point with radii corresponding to the mean track error for the given lead time over the previous four seasons.[7] Katrina passed along the western boundary of the error cone shown in Figure 2; the next forecast brought the track much closer to New Orleans, and brought New Orleans within the cone. Historical data show that, for about 75% of 5-day cones, the storm track verifies entirely within the cone.[8]

[6] http://www.nhc.noaa.gov/verification/pdfs/OFCL_10-yr_averages.pdf
[7] J. Franklin, NHC, personal communication, 27 Jan. 2006.
[8] J. Goerss, Naval Research Laboratory, Monterey, personal communication, 21 Feb. 2006.

3.2.1 Strike probabilities

From 1983 (Sheets 1985) through the 2005 hurricane season, the NHC's official probabilistic forecasts consisted of strike probabilities. The official strike probability model is a probability distribution about the track indicating the probability of the storm passing within 65 nautical miles (75 mi) of each location within 72 hours from the issuance of the forecast.[9] The strike probabilities were based on a bivariate normal distribution about the official forecast positions, where the standard deviations of cross-track and along-track errors were the sample standard deviations of basin-wide forecast errors for the 1970-1998 period at the corresponding lead time.[10]

A basic indicator of the uncertainty in a track forecast is the maximum—over all possible locations—strike probability at a given lead time. This measure reflects the confidence that the storm will pass near the forecast track, and is related to the "sharpness" (roughly the opposite of entropy, by its information theory meaning) of the probability distribution of the storm's future location. The maximum strike probability for the 12-hour forecast is 75-85%, but for the 72-hour forecast the track is highly uncertain, and the maximum strike probability—for the location directly in the path of the forecast track—is only 10-15% (Table 1).

[Table 1 about here.]

It is important to note that quantitative predictions of the timing of landfall are communicated as point estimates. Along-track error distributions in strike probability model include landfall timing error, but the distribution of the timing of landfall is folded into the 72-hour strike probability. The exception is that ranges may be given in the verbal discussions.

3.2.2 Wind-speed probabilities

The NHC replaced the strike probability model with an operational wind-speed probability product starting in the 2006 season. The new product expands the information content by providing the probability of winds exceeding three critical thresholds (35, 50, and 64 knots, equivalently 40, 58 and 74 mph) at each geographic location (Gross et al. 2004).[11] The exceedance probabilities for each threshold are the relative frequencies resulting from 1000 iterations of a Monte Carlo simulation. The simulated parameters are the location of the storm's center and maximum sustained wind speed. The distribution of the location of the storm's center is the official track location plus error, and the simulated wind speed is forecast plus error. The errors are sampled from historical error frequencies. The radius of winds of a given speed is calculated as a function of intensity through an empirical formula. The new product represents an important advance in forecasting decision-relevant parameters. By integrating uncertainty in both track and intensity, and by modeling variations in size to produce a forecast of an impact variable—wind speed—at each geographic location, this product takes a major step towards forecasting variables that are most decision-relevant.

Like the previous strike probability model, this model uses the same error distributions for storms throughout the basin: geographic differences in forecast errors are not modeled. However, factors other than lead time are also highly correlated with error. Slow-moving storms and more intense storms have smaller track errors than larger or weaker storms (Powell and Aberson 2001); larger storms are not necessarily the most intense. The dominant pattern of atmospheric steering forces, which change over the season, means that

[9] Strike probabilities were disseminated as a contour map and with numerical values for specific target locations. Strike probabilities were also are available in numeric form for designated locations, and for lead times less than 72 hours.

[10] J. Franklin, NHC, personal communication, May 11 2006.

[11] Examples of the new product are available at http://www.nhc.noaa.gov/pns1_2006_examples.shtml and http://rammb.cira.colostate.edu/projects/tc_wind_prob.

storms in certain regions are more predictable. Track errors tend to increase with latitude (Pike and Neumann 1987), and landfall position errors are 15-50% greater for storms traveling parallel to the coast than for storms traveling perpendicular to the coast (Powell and Aberson 2001). This paper models geographic differences in forecast uncertainty within the North Atlantic basin.

3.3 Verbal messages

Forecasters at the NHC have a wealth of information and experience that is not captured by the official track and strike probability models. For example, they can identify situations in which certain NWP models tend to be in error; they may know when observations are in doubt; and they can see agreement (or disagreement) among model tracks that give them greater (or lesser) confidence in their forecast than the historical average errors reflect. The NHC sometimes conveys this information to the public in verbal form. For example, the Katrina advisory discussion, 5 p.m. Friday, August 26, 2005 included the following:

> It is worth noting that the guidance spread has decreased [...] [T]his clustering increases the confidence in the forecast.

By contrast, the discussion for Hurricane Wilma on 5 p.m. October 19, 2005 read:

> Agreement among the track guidance models [...] has completely collapsed today. [...] Needless to say...confidence in the forecast track...especially the timing...has decreased considerably.

As discussed in Section 7.2, the tropical cyclone community is working to model forecast accuracy as a function of agreement among models as well as other early-storm indicators of forecast uncertainty.

3.4 Mapping forecasts into impacts

All weather forecasts are effectively decision-support tools: they are the distillation of a very large amount of information available to forecasters into a much smaller message intended to be valuable to end-users in decision-making. The forecasts issued by the NHC distill observed and modeled atmospheric parameters—temperature, pressure, etc. around the globe at many vertical layers—into storm-specific parameters—location, wind speed, speed of forward motion. A public official is interested in storm-specific parameters primarily as they will cause impacts at a stationary target location. The major impacts—flooding (due to storm surge and precipitation), winds, lightning, and tornados—are not forecast directly. Ultimately, public officials who make hurricane evacuation decisions are interested in the consequences produced by the interaction between these uncertain impacts and the public's actions which determines whether flooding will take lives. This paper takes the quantification of uncertainty one step beyond the storm-specific parameters, quantifying the uncertainty as it will affect evacuation decisions.

4 Strike probability

The center of Hurricane Katrina passed New Orleans shortly after noon Monday, August 29, 2005. The forecast issued at 11 p.m. Friday August 26, about 60 hours before the storm reached New Orleans predicted that Katrina would pass almost directly over the city (Figure 1). However, the official strike probability for New Orleans at that time was only 17%, and did not exceed 50% until less than 24 hours before the storm reached the city. The new wind-speed model, which was run experimentally in 2005, gave the following probabilities for hurricane-force winds at New Orleans: Friday morning (72 hours out): 6%; Saturday morning (48 hours): 28%; Sunday morning (24 hours): 60%.[12] The low strike probabilities were

[12] M. DeMaria, NOAA NESDIS, personal communication, 15 Sept. 2005.

surprising to many, even in the hurricane forecasting community, but as this paper will explore further, they are typical of the corresponding lead times.

Except where noted, the strike probabilities reported and used in calculations in the remainder of this paper are based on a statistical hurricane track model developed in Regnier and Harr (2006) (hereafter the RH model). The RH allows the calculation of strike probability for historical storms that occurred before the NHC added strike probabilities in 1985, and before the NHC issued track forecasts. Because the NHC strike probabilities and wind-speed probabilities are based on a track forecast, they cannot be calculated for any storm until a track is generated; i.e. they cannot be calculated for a hypothetical future storm. In addition, as discussed earlier, the NHC's strike probability models do not account for differences in uncertainty for storms threatening different targets, whereas the RH model does. Though the NHC and RH models are derived very differently—the NHC's reflect past forecast errors whereas the RH model is derived from historic tracks—their spread in future track locations are similar (Table 1).

Figure 4 shows official strike probabilities, RH strike probabilities, and probabilistic forecasts for the new wind-speed model. The wind-speed probabilities reflect the probability of hurricane-force winds reaching the city which, for a large and intense storm like Katrina, would occur if the center passed within a greater distance than the radii used to define a strike for either the NHC strike probability model or the RH model. The wind-speed probabilities are therefore expected to be higher than the strike probabilities given by the other two models. All three models show, however, that the probability of a hurricane strike at New Orleans remained less than 30% until 36 hours before the storm hit (about 24 hours before the arrival of gale-force winds). A more general comparison of NHC and RH model strike probabilities appears in the Technical Annex.

[Figure 4 about here.]

4.1 The RH Markov track model

The RH model is a discrete-time first-order Markov chain derived from 580 historical Atlantic tropical cyclone tracks—of which 263 were hurricanes— that comprise the HURDAT database (Jarvinen et al. 1984) for the period 1950-2005.[13] The state of a storm is defined by the location of its center and its direction of travel. The storm's location is the 1° latitude ×1° longitude cell containing the center within the region 0°N–70°N and 0°–100°W. Storm motion is modeled according to transitions among the states, which occur at 6-hour time steps, corresponding to the time increments recorded in hurricane verification tracks from the HURDAT data set, and the regular intervals for NHC-issued forecasts. Each observation for each storm is assigned to the appropriate state j in the Markov model. The transition probabilities q_{jk} are the historic relative frequencies of transitions from j to each state k, with $q_{jj} > 0$ allowed. For each target, cells within a strike zone, approximating a 75-mile radius around the target location, are assigned a strike probability of 1. For all other states, the strike probability, p_j, is defined as the probability that a storm in state j will ever pass through the target's strike zone, as calculated using the transition matrix of q_{jk}'s. Further details appear in the Technical Annex.

Results are given for four targets. New Orleans is still the most hurricane-vulnerable city in the U.S.,[14] with a large population living below sea level. The surrounding topography is also quite flat and provides few nearby flood-safe locations. Miami is at risk because of its large population, limited escape routes and limited nearby high ground. Although it is not usually considered hurricane vulnerable, and storms often

[13] In Regnier and Harr (2006), only 1950-2002 data were used.

[14] "Surprises in a New Tally of Areas Vulnerable to Hurricanes" by Cornelia Dean, New York Times October 10, 2006.

weaken before they reach New York, Long Island is a barrier island, and eastern Long Island, including Montauk, provides few escape options. Norfolk provides a contrast to the more vulnerable targets and is typical of the southern Atlantic coast states. The time profile of track uncertainty at Norfolk, especially at long lead times, is also of interest because the U.S. Navy faces the decision whether to sortie its Norfolk fleet from port in advance of a hurricane, with decision lead times of 72 to 120 hours, and at a cost in the tens of millions of dollars.

Table 2 summarizes the historical frequency of named tropical storms striking and threatening each target, with four additional targets: Galveston, Key West, Tampa Bay, and Wilmington, North Carolina. Threatening storms are defined as those that exceed a strike probability of 1% within 72 hours' minimum lead time of the target during their life cycle, as measured with the RH model. Storms that are below hurricane-strength at the time they threaten or strike the target are included; this tends to overstate the number of threats and strikes that induce evacuation at each location. On the other hand, hurricane preparations can occur for tropical storms forecast to intensify to hurricane-strength, even if they never do; these are excluded from the definition of threatening storms.

The Atlantic coast locations (Miami, Wilmington, and Norfolk) have a higher frequency of threats by both hurricanes and tropical storms generally than the Gulf coast locations (Galveston, New Orleans, Key West, and Tampa Bay), though in both Atlantic and Gulf coast locations, a similar fraction of threatening storms eventually strike the target. Montauk is unusual in that a much lower fraction of threatening storms actually strike, which might reduce residents' tendency to prepare.

[Table 2 about here.]

As an input to quantitative decisions under uncertainty, the essential information conveyed in a forecast is the probability of outcomes that affect alternatives and consequences. Strike probability and decisions determined by strike probability therefore form the basis for measuring information quality. Figure 5 shows the strike probability for historic hurricanes (1950-2005) at each of four target locations as a function of lead time. For threatening storms that never strike the target, lead time is measured from the time of closest approach to the target. Mean RH strike probabilities for striking storms are depicted in red, and threatening (but non-striking storms) are plotted in blue, with the mean $\pm$ one standard deviation (σ) plotted in dotted lines.

[Figure 5 about here.]

The mean strike probability for striking storms rises monotonically to 1 as lead time declines to zero. The striking storms represented include some whose forecasts did not consistently track directly over the target during the entire life cycle. Therefore, the strike probabilities in Figure 5 are expected to be lower than the NHC's maximum strike probabilities (Table 1) which apply to locations directly in the path of the official forecast. This relationship is observed, except for storms striking Norfolk at the 72- and 48-hour lead times. This exception is consistent with the observation that there is less uncertainty for storms striking Norfolk than storms striking most of the other studied locations. The NHC strike probabilities are a function of basin-wide forecast errors and therefore would tend to overestimate uncertainty for storms striking Norfolk.

The most notable feature of Figure 5 is that storms striking New Orleans and Montauk have low strike probabilities until they are very close to the coast. At these targets, it is difficult to distinguish striking from non-striking storms until 24 to 36 hours before landfall, because the mean strike probability for striking storms does not exceed the mean plus one standard deviation for non-striking storms until 24 hours before

the actual strike. Since threatening but non-striking storms outnumber striking storms 7:1 in the historical record at New Orleans and 12:1 for Montauk, the impact of false alarms is very important.

The plots for Norfolk and Miami are quite different. The mean strike probability for striking storms exceeds the mean plus one standard deviation for non-striking storms at 48 hours' lead time, and the strike probabilities are over 20% by this time. The false alarm rates for decisions made at this lead time will still be high, but not nearly as high as for Gulf locations.

Atlantic tropical cyclones commonly form in the open Atlantic, travel west, and then turn north and eventually east: the blue and green tracks in Figure 6 are typical. Note that strikes on the Atlantic coast are more common than strikes on the Gulf coast (compare Norfolk to New Orleans in Table 2)—i.e. the blue and green tracks are also typical of Atlantic tropical cyclones generally. These storms generally do not enter the Gulf, but, when they do, they leave little time for the coastal population to prepare. Storms striking the Gulf coast may form in the Gulf or the Carribean (like the pink track), providing less lead time before landfall. Table 2 gives the historical frequency of striking storms forming with less than 24, 48, or 72 hours' lead time. Generally, the frequencies are higher for Gulf coast locations (Tampa Bay and points west) than for Atlantic coast locations (Miami north through Montauk).

[Figure 6 about here.]

For non-striking storms, the strike probability hovers at or below 10%, and does not increase substantially as lead time declines because the storm may be far from the target at the time of its closest approach. For a given target location, a storm's strike probability may rise and fall over the life cycle. If evacuations are initiated whenever strike probability exceeds a threshold level, the false alarm rate can be high. Table 3 gives the percentage of storms that would trigger an evacuation if the decision were made based on a threshold strike probability. Only 1% of non-striking storms in the database exceeded a strike probability of 80%, and then only for Norfolk. For lower triggers, however, the rates can be much higher: over the course of their lifetimes, over 60% of non-striking storms exceed a 10% strike probability and at least 32% (49% for Norfolk) exceed the 20% threshold. For this reason, Regnier and Harr (2006) note that a one-time decision structure with fixed lead time may be less costly than constantly monitoring and reapplying a threshold trigger that is calibrated to a one-time decision.

[Table 3 about here.]

5 False alarm rate

It is possible to quantify the trade-off between risks of type I (false negative) and type II (false alarm) errors by making the following assumptions: binary weather outcomes (strike or no strike), fixed preparation lead time, and binary alternatives (evacuation or no evacuation). As discussed above, there are multiple opportunities to initiate evacuations, and lead time is an uncertain quantity. Moreover, weather outcomes vary tremendously. However, the simple 2×2 decision model allows for the calculation of error rates, a decision-relevant measure of uncertainty. These results can be used to identify critical lead times at which reductions or flexibility in evacuation time would be most valuable in each studied region. This is particularly relevant because alternatives are currently treated as binary in real-world decision processes.

The false alarm rate (hereafter, FAR), which reflects decision outcomes, can be used as a measure of forecast quality. The FAR is a function of both the forecast quality and the decision rule used to determine whether a given strike probability will trigger an evacuation. The choice of the strike probability threshold

used to trigger an evacuation represents a trade-off between the risk of failing to evacuate in advance of a striking storm and the possibility of a false alarm.

Figure 7 plots FAR vs. probability of detection (POD)—equal to 1 minus the probability of a type I error—for forecast lead times ranging from 12 to 72 hours, at the four target locations. A plot of POD vs. FAR is often called a receiver operating characteristics (ROC) curve, and may used both to aid decision-making and to measure accuracy in discriminating two categories of events (here, strikes and non-strikes) (Swets et al. 2000). The POD and FAR were calculated for all four targets using a single set of 10,000 simulated storms generated by the RH model. The Technical Annex provides more details on the generation of these curves. In Figure 7, the FAR is defined as the fraction of all storms triggering a retrospectively unnecessary evacuation, as percentage of all storms that threaten the target, with threats defined above. The choice of a denominator that is identical for all lead times allows for comparisons in the track uncertainty profiles at various lead times.

[Figure 7 about here.]

The curves in Figure 7 exhibit the typical concave ROC profile. For each target, the curves move to the right as lead time increases (toward higher FAR for a given POD), reflecting greater uncertainty. However, the profiles differ across the target locations. For Miami and New Orleans, long lead time forecasts are almost as good as short lead time forecasts for POD's up to 50% at Miami and up to 60% at New Orleans. For higher POD's they diverge dramatically. At Miami, 90% of non-striking storms can be screened out leaving only a 40% probability of type I error, even at 72 hours' lead time. The same could not be said at Norfolk until 12 hours before landfall. At Norfolk and Montauk, the curves representing different lead times are clearly distinct even for low POD's.

However, most public officials would require POD's higher than 60%. At the boundaries of the average NHC warning areas, the POD is approximately 95%. In light of the 2005 hurricane season, public officials are likely very concerned with avoiding inadequate preparation, and may require even higher POD's. For high POD's, the information quality is generally best at Miami, followed by Norfolk, New Orleans, and Montauk. This ordering is consistent with the strike probabilities (Figure 5).

To highlight the differences across the four targets, Table 4 shows the FARs incurred as a function of lead time for POD's of 90, 95, and 99%, and Figure 8 plots these values for POD of 95%. FAR is redefined with a more natural denominator as a % of evacuations, a measure of false alarm rate that is important to public officials concerned with educating the public about the necessity of evacuating although false alarms are likely.

The differences across the targets are consistent across lead times. Table 8 and Figure 8 show that the power to discriminate between striking and non-striking storms in making evacuation decisions is, at all lead times, ordered from best to worst as follows: Miami, Norfolk, New Orleans, and Montauk.

The most interesting result in Table 4 is how the FAR drops off as lead time declines. At New Orleans, reducing decision lead times from 72 to 42 hours has very little impact on the FAR, indicating that track uncertainty does not decline significantly. Between 42 and 36 hours, the FAR begins to drop off much more quickly as uncertainty is resolved. At Miami, on the other hand, track uncertainty falls dramatically between 72 and 42 hours, reflected in a substantial reduction in FAR.

For a 90% POD, officials at New Orleans would have to wait until 20 hours before landfall to get information that is as good as Miami's at 54 hours (i.e. FAR = 72%). Generally, the lags in information quality are 12 to 18 hours, but even these differences are very large relative to the lead times required for evacuation (Section 2).

[Table 4 about here.]

[Figure 8 about here.]

6 Cost of false alarms

If false alarms were free, the entire coastline would evacuate every time a tropical storm forms. In reality, this would be prohibitively expensive and disruptive—the average annual number of tropical cyclones and hurricanes is approximately ten, with six of these becoming hurricanes (Jarrell et al. 2001). Despite individuals' and officials' motivation to reduce risk, there will always be a trade-off between costly, and potentially deadly, false alarms and the risk of inadequate preparation for a hurricane strike. Table 5 summarizes this trade-off, varying two dimensions: POD and lead time. Emergency managers and officials charged with recommending or ordering evacuations implicitly, if not explicitly, choose the risk level (1−POD) they are willing to accept. Through investments in infrastructure and evacuation processes, they may be able to reduce their decision lead time. The long-run average annual number of Atlantic hurricanes (six), is used in the Table 5 calculations, though during the current active hurricane period, the rate has been higher: there were 15 hurricanes in 2005 and nine in 2004.[15]

[Table 5 about here.]

Atlantic coast locations are subject to a higher frequency of hurricane strikes (Table 2). On the other hand, as discussed previously, track uncertainty at long lead times is higher for Gulf coast locations. Although Miami experiences more threats and strikes than New Orleans, because of the differences in track uncertainty, the expected annual number of false alarms is similar at the two targets (Table 5). The expected annual numbers of evacuations, false alarms, and strikes at an unevacuated target are highest for Norfolk—which experiences the most hurricane strikes. The annual expected numbers of evacuations—at least 0.9 per year—are higher than most coastal residents experience. These values include all hurricanes, even those that weaken before they reach the target; for example, like the green storm in Figure 6, storms may pass over land and weaken before reaching more northerly locations.

Table 5 translates false alarms into evacuation costs per year for the region surrounding each target, assuming 460 miles of coastline evacuated and $1M per mile in evacuation costs (see Section 1 for explanation of these values). The shorter the decision lead time and the lower the POD, the lower these costs, which range from $287M (Miami, with 24-hour lead time) to $1.48B (Montauk, with 72-hour lead time) annually per target. Tropical storms that do not reach hurricane force are excluded; their inclusion would raise the total costs. Again, Montauk would not experience hurricane-force winds for most of these hypothetical storms.

This approach to estimating evacuation frequencies and costs differs from the Section 1 calculation of the total cost of false alarm evacuations. The Section 1 calculation is based on hurricane warnings, which are issued with 24 hours' lead time, whereas Table 5 analyzes decisions with lead times up to 72 hours; longer lead times will incur higher costs. The estimates in Table 5 include storms that threaten but do not make landfall, but exclude tropical cyclones that do not become hurricanes. Finally, this section treats each target separately, whereas the $1B in false alarms annually calculated in Section 1 applies to the entire U.S. The costs in Table 5 include partial duplication—evacuations for a single storm at both Miami and Norfolk would appear twice, for example. Nevertheless, the figures are fairly consistent: $1B annually for the entire

[15] Bill Gray's seasonal forecast of hurricane activity is perhaps the best known. His group's forecast, updated October 2, predicts six hurricanes in the Atlantic for 2006, which is the also the long-term average (http://tropical.atmos.colostate.edu/forecasts/2006/oct2006/).

U.S. coast, and costs on the order of $300-500M for 24-hour lead time evacuations of individual targets and surrounding areas—similar to $1M per mile for 460 warned miles of coastline.

A rough estimate for the value of reducing lead times along the entire coast from 72 to 48 hours is the difference between the sum, across the four targets, of the costs of false alarms for 72 -hour lead times minus the sum for the 48-hour lead time. The estimated values are are $427M, $542M, and $903M annually for 90, 95, and 99% POD respectively. The 72-hour decision lead time is representative of high-population, and therefore, high-cost evacuation areas.

False alarms may be even more harmful than their direct costs indicate because a high false alarm rate may reduce the public's willingness to prepare for future storms. Roulston and Smith (2004) have shown in a theoretical model that the optimal choice of an action threshold for ordering preparation is sensitive to a compliance rate that is a function of the false alarm rate in the forecast process. On the other hand, Dow and Cutter (1998) do not empirically observe this "crying wolf" effect in their empirical study and Baker (2002) reports that evacuation rates did not drop after two false alarms in 1985.

The risk of a hurricane striking an unprepared target (a type I error) is a critical factor that should not be lost in a discussion of FARs. Public officials must balance the cost of frequent FARs against the consequences of a type I error, which can be catastrophic. Using a decision rule based on required POD, the expected number of type I error events per year is estimated in Table 5. Shorter evcuation decision lead times and higher PODs generally lower the expected number of type I errors (equal to the number of strikes per year times 1− POD).

7 Discussion

7.1 Future hurricane preparation decision context

The hurricane threat and its costs are expected to increase for the forseeable future as the population on the Atlantic and Gulf coasts of the U.S. continues to grow. Prior to 1995, hurricane damages were increasing rapidly, but the rise was due to growing population and wealth, not to frequency or severity of hurricanes (Pielke and Landsea 1998; Pielke et al. 2006). Since 1995, hurricane frequency and intensity are also on the rise: the level of hurricane activity in the North Atlantic has more than doubled relative to the 1970–1994 period. There is an ongoing debate among tropical cyclone meteorologists and oceanographers as to whether this increase is a part of a multi-decadal cycle (Goldenberg et al. 2001), whether it is caused by rising sea surface temperatures, and therefore represents a long-term trend (Mann and Emanuel 2006; Trenberth and Shea 2006; Emanuel 2005; Webster et al. 2005), or whether the record is currently too sparse to show such a trend (Pielke, Jr. et al. 2005). There is widespread agreement that hurricanes, and, in particular, destructive hurricanes, can be expected to threaten and strike the U.S. at the 1995-2005 rate for at least the next 10 to 40 years.

As discussed in Section 3.1, forecast accuracy is constantly improving, and the tropical storms community has been developing forecast products that provide increasingly decision-relevant information. The new wind-speed probability model (Section 3.2.2) is a major advance, combining track and intensity errors to provide a probabilistic forecast of wind speed exceedances at each target location. Although both forecast accuracy and the relevance of products are improving, there will always be uncertainty at the lead times required for mass evacuation.

Further research on improving decision processes, and designing preparation and response plans to exploit forecast information should also be priorities. The National Hurricane Research Initiative Act introduced in the U.S. Senate in September 2006 calls for funding for "Research to improve the manner in which hurricane-related information is provided to, and utilized by, the public and government officials, including research

to assist officials of State or local government in determining the circumstances in which evacuations are required and in carrying out such evacuations."

Such research would provide tremendous value at much less marginal cost than infrastructure improvements or even forecast accuracy improvements. The OR/MS community is ideally positioned to make major contributions in improving the way forecast information is utilized, and working with meteorologists to target development of new products such that they can be integrated into decision processes (Regnier 2006).

7.2 The importance of forecasting

This paper's focus on decision processes and speeding evacuation is by no means intended to undercut the importance of forecasting. Modern observation systems, NWP models, and forecast methods have significantly reduced the possibility of tragedies such as the Galveston hurricane of 1900, which caused an estimated 8,000 deaths (Jarrell et al. 2001). New Orleans mayor Nagin reported that close to 80% of the population of New Orleans evacuated before Hurricane Katrina.[16] The death toll could have been far higher had the city lacked the days of advance warning provided by the NHC.

Moreover, the meteorology community is focusing resources on improving probabilistic forecasts. The new wind-speed probability model is the most recent official product to result from these efforts. In addition to retrospective analysis of error, described above, the hurricane meteorology community has begun to measure forecast uncertainty by looking at the sensitivity of model tracks to small perterbations in initial conditions to produce ensemble-based probabilistic forecasts (Toth 2001). Ensembles are essentially simulations run using NWP models, sometimes in reduced form (e.g. lower resolution), and perturbing initial conditions across a reasonable range to represent data uncertainty (Cheung 2001). The term ensemble forecast generally refers to a simulation using a single NWP model. A set of forecast tracks arising from multiple models can also be treated as an ensemble; variability within a multi-model ensemble can capture, to some degree, uncertainty in model formulation.

Goerss (2005), among others in the tropical cyclone meteorology community, is investigating statistically the relationships between early-storm indicators of track uncertainty, including agreement among numerical models, and developing reliable probabilistic forecasts that reflect differences in uncertainty and confidence in the official forecasts across storms, with the goal of producing an operational product.

Errors produced by NWP models can be ascribed in large part to the model specification, including the spatial and temporal resolution, and to errors in setting initial or boundary conditions. Recent work to improve NWP models includes increasing model resolution and improving data collection and assimilation methods, as well as coupling oceanographic and atmospheric models. The chaotic behavior of NWP models of the atmosphere is taken to imply that there is a fundamental limit to the ability to forecast meteorological events at the time-scales that are relevant to short-term decision-making, such as the conditions at a specific place on a given day (Wilks 2006; Leslie et al. 1998).

Improving hurricane forecasts at lead times of 48 to 72 hours is another way to reduce the false alarm rate: the savings that would result from improving 72-hour forecasts to current 48-hour accuracy would be equivalent to the savings from the same reduction in lead time. Investments are being made by increasing personnel at the NHC and by bringing new research results to operations through the Joint Hurricane Testbed Program that is sponsored by National Oceanic and Atmospheric Administration. The NHC has added four new hurricane specialists to its existing team of six who are the primary hurricane forecasting team for the entire North Atlantic. However, the rate of improvement in track forecast accuracy is about 2% per year (McAdie and Lawrence 2000), and accuracy is ultimately constrained by inherent limits in atmospheric

[16] "New Orleans braces for monster hurricane", www.cnn.com., Monday, August 29, 2005.

predictability. Reductions in decision lead times, therefore, have a greater potential to substantially reduce the human and economic costs of hurricanes and false alarms in the near term than improvements in forecast accuracy.

7.3 Flexibility and option value

This paper's analysis defines a framework for quantifying hurricane uncertainty in decision-relevant terms, as a function of lead time and geographic location. The translation of track uncertainty into decision space can be used to target improvements in planning to exploit the information available and to allow rational assessment of the trade-offs between early, precautionary preparation and later, more accurate forecasts.

Finding ways to delay irreversible costs and creating flexibility by preparing multiple evacuation plans that can be implemented at various lead times can also reduce false alarms. The rate at which track predictions improve as a storm advances provides a valuable input for the design of staged and calibrated response plans. Evacuation would be a possible outcome of a staged decision process, but less disruptive and less costly preparations would be initiated at early, high-uncertainty lead times. Such plans can reduce the costs of false alarms, while allowing appropriate preparations with short lead times.

As discussed above and illustrated in Figure 1, the evacuation decision is commonly framed as a one-time decision whose timing is determined by a point estimate of the time required to complete the evacuation. A two-alternative (evacuate or do not evacuate) structure with a fixed decision time manifests itself in many emergency managers' approaches. The U.S. Navy takes a similar approach when deciding whether to sortie its fleet from a port (Klein 2005).

Building flexibility into the decision process makes the evacuation decision dynamic. In a dynamic situation, optimal timing of evacuation will depend on quantifying forecast uncertainty at the time of each decision opportunity, the forecast quality that can be anticipated at later decision points, and the option value of waiting for updated hurricane track forecast information before taking costly, irreversible action. The decision to initiate each irreversible action is essentially an optimal-stopping problem; there are repeated opportunities to take this step.

The economic value of considering the evolution of uncertainty in deciding when to undertake the cost of evacuation is considerable (Regnier and Harr 2006). Despite Hurrevac's framing of the decision, delaying an evacuation is not equivalent to making no preparation for the storm. The option of sheltering in-place from winds or evacuating to high ground close to home is often available even after the arrival of gale-force winds. Moreover, there are many options for how to evacuate. In most regions, partial evacuations of smaller regions and local evacuations to higher ground within the city or county can be implemented with lead times much shorter than a complete evacuation would require. Although many people choose to evacuate long distances, for a perceived sense of safety or to stay with friends or relatives, with proper planning most people can evacuate tens of miles, not hundreds, sometimes within the county or to a neighboring county, reducing the length of each trip and evacuation clearance times overall because of lower traffic burdens.

The cost and risk-reduction impacts of evacution vary as a function of the lead time at which the evacuation is initiated—for example, reversing traffic flow is more expensive, but may speed evacuations. On the basis of a time-profile of forecast uncertainty, and an assessment of the reduction in uncertainty that can be anticipated, decision-makers can balance the costs of waiting to initiate protective action—which may include an increase in risk, or the possibility of a later, more costly preparation—against the value of waiting for more information in the form of more accurate forecasts—which includes the option to avoid incurring irreversible preparation costs if later information reveals that the target location will not be affected by the storm.

7.4 Strategic decisions

Although this research focuses on decisions made at time scales of days and hours in response to forecasts of a specific storm, the results can inform decisions with longer time scales. Strategic decisions made with a long-term objective, define the parameters for the storm-specific evacuation decision by affecting lead times required to implement various preparation actions, as well as the costs and effectiveness of the available actions.

Strategies for reducing evacuation decision lead times include: upgrading transportation infrastructure to speed evacuation and response, improving warning and communication systems and public education to speed response and reduce shadow evacuations, facilitating coordination across jurisdictional boundaries, and constructing local flood-protected shelters. To the extent that these improvements can reduce evacuation clearance times, they will also reduce false alarm rates. Transporting people from hospitals and nursing homes always involves risk, and road accidents are a further risk induced by evacuation. Moreover, as the population evacuating increases, these risks grow geometrically because transportation networks become overloaded which can leave evacuees stranded on the road and more vulnerable to wind and flooding hazards than they would be at home. Reducing false alarm rates will have the added effect of reducing the size of the population evacuating, thereby making evacuations safer for those who ultimately must move. In evaluating potential improvements in decision strategies and infrastructure, it is critical to understand how much additional forecast accuracy can be expected from a proposed reduction in lead time required to implement an evacuation.

7.5 Geographic differences

The differences among the four target locations analyzed are revealing and have implications for hurricane planning in each region. Hurricane Katrina made New Orleans's special vulnerability to flooding clear. Congressional testimony indicates that evacuation studies estimated a clearance time of 50 to 72 hours (though the storm intensity was not specified; more intense storms require more evacuation time).[17] This would seem to indicate that resources to speed New Orleans evacuations would be well-spent. However, the limits imposed by track uncertainty imply that unless evacuation decision time can be reduced to about 30 hours, the residents may have to accept that 75-90% of evacuations will be false alarms to achieve a 90-99% probability of detecting a striking storm.

On the other hand, in Miami, reducing lead time from 72 to 60 or 48 hours would substantially reduce false alarms. Specific recommendations for investments depend, of course, on the costs. Investments to expand the I-95 north of Miami, construct refuges of last resort close to the city, or develop a public system for rapidly transporting older and immobile residents, could have a substantial impact, reducing the cost of false alarms and making evacuations easier, and perhaps encouraging more inhabitants to evacuate.

7.6 A word on conservatism

Estimated evacuation clearance times are determined based on simulations of the transportation network, and may include parameters such as tourist density, or peak vs. non-peak baseline traffic. The parameters that are varied in establishing evacuation scenarios vary by locality. A key parameter that is always set to the same value is compliance rate: all evacuation clearance times are based on the assumption that 100% of the population evacuates in response to an order.[18] This means evacuation clearance times are biased towards being too long, which is intended to provide a safety margin. However, an early bias in decision

[17] The Testimony of Mr. Windell Curole, General Manager, South LaFourche Levee District before the Senate Committee on Commerce, Science and Transportation Subommittee on Disaster Prevention and Prediction, September 20, 2005, viewed at http://commerce.senate.gov/hearings/testimony.cfm?id=1607&wit_id=4594, December 8, 2005.

[18] C. Gregg, FEMA Region VI, personal communication 14 Feb. 2006.

time will increase the false alarm rate and may induce more risk to life than a decision made later in the storm.

Currently, Hurrevac decision times are determined solely based on transportation analyses. Geographic differences in track uncertainty identified here show that each locality's evacuation decision time should also be a function of the time profile of forecast uncertainty, even if the evacuation decision is not treated dynamically. Rational public officials facing a 10% strike probability with 72 hours' lead time might appropriately make different decisions if they were located at Miami instead of New Orleans. The official at Miami might delay an evacuation to wait for the next forecast, while a public official at New Orleans would go ahead and order the evacuation—because waiting would yield little information value. In fact, public officials who are aware of geographic differences in track uncertainty might be justified in adjusting their assessment of the uncertainty in a forecast up (for Gulf coast locations) or down (for Atlantic coast locations) relative to the official forecast.

A Technical Annex

A.1 Storm motion model

This paper's quantitative results are based on a Markov model (the RH model) of hurricane motion derived from 56 years of historical Atlantic tropical cyclone tracks (Regnier and Harr 2006). The RH model permits quantification of conditional probabilities of future evolution of the storm. The historical tracks come from the HURDAT data set (Jarvinen et al. 1984) which contains the tracks for all Atlantic hurricanes and tropical storms determined via post-analysis. Although HURDAT contains data for the years 1851-2005, only the 580 storms (of which 263 were hurricanes) from 1950 through 2005 were used in fitting the model because the record before 1950 is spotty. In fitting the model, tracks for both hurricanes and tropical storms that did not reach hurricane intensity were used.

A.1.1 States

RH use a discrete-time first-order Markov chain. The state of a storm is defined by the location of its center and its direction of travel. The hurricane location is the 1° latitude ×1° longitude cell containing the hurricane center within the region 0°N–70°N and 0°–100°W.

Atlantic hurricanes commonly form in the open Atlantic, travel west, and then some turn north and eventually east: the blue and green tracks in Figure 6 are typical. Storms striking the Gulf coast often form in the Gulf or the Carribean (see the pink track), providing less lead time before landfall. These storms still tend to follow the curving path, starting with a northwest tendency, transitioning to a northeast trajectory. The timing of the turns largely determines whether and where the storm will strike the coast. The northward recurvature typically happens in the region 10–25° N and 55–80° W. Therefore, for hurricanes in the region 10–25° N and 55–80° W, the direction of travel is also defined as a state variable. The direction of travel is categorized as "north" if its direction is primarily north, "west" if its direction is primarily west, and "other" otherwise (see Regnier and Harr (2006) for further details).

A.1.2 Strikes

A strike is defined to occur at a given geographic location (target) if the hurricane center passes through the 1° lat. × 1° long. cell containing the target, or in one of the adjacent cells to the north, south, east, or west of the stationary target, or in the diagonal cells to the southeast and southwest of the target, but not the cells to the northwest and northeast. This definition reflects the fact that in general the extent of hurricane-force winds is greater on the right-hand side of a hurricane, with respect to its direction of travel. For a given target, the number of states in the strike zone ranges from seven to 21 depending on the location

of the target. For targets whose strike zone is within the region 10–25° N and 55–80° W, where direction is also a state variable, there are 7×3 = 21 states in the strike zone. For targets whose strike zone does not overlap with the 10–25° N and 55–80° W region, there are only seven states in the strike zone. Between 0° and 70° latitude, the area of the strike zone ranges from approximately 12,000 to 33,000 sq miles, whereas the NHC definition of the strike area (NHC/TPC Glossary) is a circle with diameter = 125 nautical mi = 144 mi, which has an area of 16,252 miles.

A.1.3 Formation and transition probabilities

Storm motion is modeled according to transitions among the states, which occur at 6-hour time steps, corresponding to the time increments recorded in hurricane verification tracks from the HURDAT data set. Each observation for each storm was assigned to the appropriate state j in the Markov model. The probability distribution of hurricane formation across states is denoted as $\mathbf{r}$, where r_j is the fraction of storms in the database that were first observed in state j. The transition probabilities q_{jk} are the historic relative frequencies of transitions from j to each state k, with $q_{jj} > 0$ allowed. Both $\mathbf{r}$ and the transition matrix $\mathbf{Q}$ are independent of target location, and simply describe the formation and motion of a storm's center. Dissipation is modeled by assigning a state $j = 0$ to dissipation, so $q_{j0} > 0$ indicates that there is a strictly positive probability that the storm will dissipate before the next time-step.

A.2 Target-specific parameters

For each studied target location, and for each state, the state's lead time, τ_j was calculated as six hours times the minimum number of forward transitions (of nonzero probability) that separate the state from the strike zone. In some cases, the lead time is infinite, if there is zero probability (according to the model) for a storm to reach the target. For example, storms that are passing through Newfoundland have a zero probability of later striking Miami.

For each target, the instantaneous strike probability p_j is the probability that a storm in state j will eventually, with an unlimited number of transitions, pass through the target's strike zone. The p_j's are the solution to the simultaneous equations

$$p_j = \sum_{k \in S} q_{jk} p_k \quad \forall j \in S \text{ subject to the constraint } p_j = 1 \text{ if state } j \text{ is in the strike zone.} \tag{1}$$

Within the framework of the Markov model, all the information available at time t is contained in the state of the hurricane, and therefore the probability that a hurricane in state j will eventually strike, conditional on the information at time t, is p_j. As reflected in Equation 1, the value p_j depends on the strike probabilities in the next time step, which in turn reflect strike probabilities in the following time step. However, p_j compresses the future probabilities into a single, scalar value. A value of $p_j = 0.5$ could reflect either:

$p_k = 1$ or $p_k = 0$ $\forall k$ such that $q_{jk} > 0$. (i.e. that in the next 6 hours, all uncertainty will be resolved); or

$p_k = 0.5$ $\forall k$ such that $q_{jk} > 0$ (i.e. no reduction in uncertainty in the next 6 hours); or something in between.

The strike probabilities p_j are distinct from the climatological probability of strike conditional on the state of the storm as defined in the RH model. The climatological strike probabilities (i.e. the fraction of historical storms passing through state j that eventually pass through the strike zone) would not satisfy Equation 1. Either set of probabilities would reflect geographic differences in uncertainty and could have been used to generate the results in this paper. The advantages of using the climatological strike probabilities is that they might capture persistence of direction and potentially certain specific track patterns that reflect

common steering forces more reliably than the RH model. However, storm tracks generated with the RH model follow typical patterns quite well. The advantage of using the p_j's is that the number of historical storms passing through many of the states is quite small and therefore smoothing the strike probabilities by making adjacent states' strike probabilities mutually consistent may better measure strike probability conditional on each state j for future storms.

A.3 Validation

This model is very simple and is based purely on climatology, i.e. the historical record of weather events. It therefore has much lower forecasting skill, or relative accuracy, than the forecasting models. One of the useful features of this model is that it is fully stochastic: because of the Markov property, the probabilistic future of any storm is completely defined. This paper exploits the resulting ease of calculating the probabilities of many events of interest to examine geographic and temporal patterns of uncertainty.

A second advantage of the Markov model is that the probability distribution about the most likely track is very similar to the corresponding distribution in the NHC's strike probability forecasts. Table 1 compares the maximum strike probability at various lead times for the NHC forecasts and for the Markov model.

[Table 6 about here.]

In addition to the spread of the distribution of locations, which is similar for the two models (Table 1), the relationship between NHC and RH strike probabilities for a particular storm would depend on whether the official forecast track is close to the most likely track generated by the RH model. The speed of the storm matters as well: because the NHC strike probability model uses the same error radius for each lead time regardless of the storm's speed of forward motion, the strike probability and error cones for fast-moving storms are longer and the cross-track distributions narrower at each distance along the track.

Table 6 compares RH-model simulated storms with the historical storms, which were used to fit the RH model, in terms of their frequency of threatening and striking each of the four targets. The percentage of storms striking at each location are very similar, though historically a greater percentage of storms threaten the coast than the RH-simulated storms. This could reflect the fact that early in the data set, especially in the pre-satellite days, storms that stayed far off coast might not have been observed at all, even if they would have met the definition of a threat.

A.4 Strike probabilities

All Atlantic hurricanes between 1950 and 2005 were used to generate Figure 5; they were not required to be hurricane force at landfall. Storms were defined to threaten a target if they passed through a state with $\tau_j \leq 72$ hours and $p_j \geq 0.01$ at any time during their life cycle. For each set of storms (striking vs. non-striking) the p_j's were grouped within each lead time (6 hours to 72 hours, in 6-hour increments), and the mean, and mean $\pm$ one standard deviation calculated. Table 3 shows the percentage of all threatening storms that passed through at least one state with p_j exceeding each stated probability threshold at any time during their life cycle. By definition, all threatening storms exceed p_j of 1%.

A.5 False alarm rates

The ROC's at each target were generated on the basis of 10,000 storm tracks generated by Monte Carlo simulation using $\mathbf{r}$ as the probability distribution of formation and the transition matrix $\mathbf{Q}$. For a given target and lead time τ, each probability threshold (denoted ρ) for evacuation was considered. The storms whose strike probability exceeded the threshold at the given lead time (thus triggering an evacuation) were denoted $\Phi(\tau, \rho)$. The FAR is the fraction of the non-striking storms that trigger evacuation, while the probability of detection is the fraction of striking storms that will trigger an evacuation, at the given lead

time. Not all storms pass close enough to threaten any given target, and the number of storms passing within τ-hours' minimum lead time of the target is lower for smaller τ. In order to maintain the same denominator in the FAR calculation for the various lead times for each target, the set of "threatening" storms is all storms passing within 72 hours of the target and, within 72-hours' lead time, exceeding a strike probability of 1%. The FAR-POD pairs for all probability thresholds at the given lead time were then plotted in Figure 7 and the FAR values for PODs of 90, 95 and 99% are extracted in Table 4, with FAR redefined as in Equation 2.

$$
\begin{aligned}
POD(\tau,\rho) &= \frac{|K \cap \Phi(\tau,\rho)|}{|K|} \\
\text{Figure 4:} \qquad & FAR(\tau,\rho) = \frac{|\Phi(\tau,\rho)| - |K|}{|\Theta|} \\
\text{Table 3:} \qquad & FAR(\tau,\rho) = \frac{|\Phi(\tau,\rho)| - |K|}{|\Phi(\tau,\rho)|}
\end{aligned} \tag{2}
$$

Notation

S	the set of (10,000) simulated storms
K	the set of storms in S that strike the target
$\mathbf{s}$	a storm track, consisting of a sequence of states j
Θ	the set of threatening storms (differs by target)
	$= \{\mathbf{s} \in S : \exists j \in \mathbf{s}$ such that $p_j \geq 0.01$ and $\tau_j \leq 72\}$
$\Phi(\tau,\rho)$	storms that trigger a preparation, using threshold ρ and decision lead time τ

References

Adams, C. R. and D. J. Berri (1999). The economic cost of hurricane evacuations. First U.S. Weather Research Program Science Symposium, March 2931, 1999, http://box/mmm.ucar.edu/uswrp/abstracts/Adams_Christopher.html, accessed June 7, 2004.

Baker, E. J. (2002). Societal impacts of tropical cyclone forecasts and warnings. *World Meterological Organization Bulletin 51*, 229–235.

Cheung, K. K. (2001). A review of ensemble forecasting techniques with a focus on tropical cyclone forecasting. *Meteorological Applications 8*, 315–332.

Dow, K. and S. Cutter (1998). Crying wolf: Repeat responses to hurricane evacuation orders. *Coastal Management 26*, 237–252.

Emanuel, K. A. (2005). Increasing destructiveness of tropical cyclones over the past 30 years. *Nature 436*(4), 686–688.

Franklin, J. L. (2005). A verification of National Hurricane Center tropical cyclone forecasts in 2004. *Shore and Beach 73*(2/3), 25–28.

Franklin, J. L., C. J. McAdie, and M. B. Lawrence (2003). Trends in track forecasting for tropical cyclones threatening the United States, 1970–2001. *Bulletin of the American Meteorological Society 84*(9), 1197–1203.

Goerss, J. S. (2005). Quantifying tropical cyclone track forecast uncertainty and improving extended-range tropical cyclone track forecasts using an ensemble of dynamical models. Final report to Joint Hurricane Testbed.

Goldenberg, S. B., C. W. Landsea, A. M. Mestas-Nunez, and W. M. Gray (2001). The recent increase in Atlantic hurricane activity: Causes and implications. *Science 293*(5529), 474–479.

Gross, J. M., M. DeMaria, J. A. Knaff, and C. R. Sampson (2004). A new method for determining tropical cyclone wind forecast probabilities. *Proceedings of the 25th Conference on Hurricanes and Tropical Meteorology*.

Jarrell, J., M. Mayfield, E. Rappaport, and C. Landsea (2001). The deadliest, costliest, and most intense United States hurricanes from 1900 to 2000. NOAA Technical Memorandum NWS TPC-1. [Available online at http://www.aoml.noaa.gov/hrd/Landsea/deadly/index.html.].

Jarrell, J. D. and M. DeMaria (1999). An examination of strategies to reduce the size of hurricane warning areas. *Proceedings of the 23rd Conference on Hurricanes and Tropical Meteorology*.

Jarvinen, B. R., C. J. Neumann, and M. A. S. Davis (1984). A tropical cyclone data tape for the North Atlantic: Contents, limitations, and uses. NOAA Technical Memorandum NWS NHC 22.

Klein, P. M. (2005). The four Florida hurricanes of 2004 and their impact on the fleet. NRL Memorandum Report, Naval Research Laboratory, Washington, DC.

Knabb, R. D., J. R. Rhome, and D. P. Brown (2005). Tropical cyclone report: Hurricane Katrina, 23-30 August 2005. http://www.nhc.noaa.gov/pdf/TCR-AL122005_Katrina.pdf, updated August 10, 2006.

Leslie, L., R. Abbey Jr., and G. Holland (1998). Tropical cyclone track predictability. *Meteorology and Atmospheric Physics 65*, 223–231.

Mann, M. and K. A. Emanuel (2006). Atlantic hurricane trends linked to climate change. *Eos, Transactions of the American Geophysical Union 87*(24), 233–244.

McAdie, C. J. and M. B. Lawrence (2000). Improvements in tropical cyclone track forecasting in the Atlantic basin, 1978-98. *Bulletin of the American Meteorological Society 81*(5), 989–997.

Munich RE Group (2006). Annual Report 2005: Paving the way for opportunities. http://www.munichre.com/publications/302-04802_en.pdf.

Pielke, Jr., R. A., J. Gratz, C. W. Landsea, D. Collins, M. A. Saunders, and R. Musulin (2006). Normalized hurricane damages in the United States: 1900-2005. *Natural Hazards Review (submitted)*. http://sciencepolicy.colorado.edu/publications/special/nhd_paper.pdf.

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

Pielke, Jr., R., C. Landsea, M. Mayfield, J. Laver, and R. Pasch (2005). Hurricanes and global warming. *Bulletin of the American Meteorological Society 86*(11), 1571–1575.

Pike, A. C. and C. J. Neumann (1987). The variation of track forecast difficulty among tropical cyclone basins. *Weather and Forecasting 2*(9), 237–241.

Powell, M. D. and S. D. Aberson (2001). The accuracy of United States tropical cyclone landfall forecasts in the Atlantic basin: 1976-2000. *Bulletin of the American Meteorological Society 82*(12), 2749–2767.

Regnier, E. (2006). Doing something about the weather. *Omega: The International Journal of Management Science*. available online at www.sciencedirect.com.

Regnier, E. and P. Harr (2006). A dynamic decision model applied to hurricane landfall. *Weather and Forecasting*. in press.

Roulston, M. and L. Smith (2004). The boy who cried wolf revisited: The impact of false alarm intolerance on cost-loss scenarios. *Weather and Forecasting 19*(2), 391–397.

Sheets, R. C. (1985). The National Weather Service hurricane probability program. *Bulletin of the American Meteorological Society 66*(1), 4–13.

Swets, J. A., R. M. Dawes, and J. Monahan (2000). Better decisions through science. *Scientific American* (October, 2000).

Toth, Z. (2001). Ensemble forecasting in WRF. *Bulletin of the American Meteorological Society 82*(4), 695–697.

Trenberth, K. E. and D. J. Shea (2006). Atlantic hurricanes and natural variability in 2005. *Geophysical Research Letters 33*(L12704). doi:10.1029/2006GL026894.

U.S. Army Corps of Engineers, FEMA, N. F. Lower Southeast Florida hurricane evacuation study technical assessment: A summary for Broward County.

Webster, P., G. Holland, J. Curry, and H.-R. Chang (2005). Changes in tropical cyclone number, duration, and intensity in a warming environment. *Science 309*(5742), 1844–1846.

Whitehead, J. C. (2003). One million dollars per mile? The opportunity costs of hurricane evacuation. *Ocean and Coastal Management 46*(11-12), 1069–1083.

Wilks, D. S. (2006). *Statistical Methods in the Atmospheric Sciences* (2 ed.). Academic Press.

List of Figures

1 Screen shot from Hurrevac software, showing Hurricane Katrina forecast at 11 p.m. on August 26, 2005, with strike probabilities for coastal locations and, in the pop-up box, recommended evacuation timing for New Orleans and nearby parishes. Time progresses from right to left. . 24

2 Official NHC 3-day track forecast from the Hurricane Katrina graphics archive. 25

3 Official forecast errors from Franklin (2005), means over the 1993-2004 seasons. 26

4 Hurricane Katrina strike probabilities for New Orleans, plotted as a function of time before the storm's center passed the city. 27

5 Strike probabilities for striking and threatening (but non-striking) historical hurricanes, as a function of lead time. 28

6 Examples of historical storm tracks. Each storm started at its south end and moved north. The gray arrows indicate the average direction and speed of storms at each location. 29

7 Probability of detection as a function of false alarm rate (measured as % of threatening but non-striking storms) for lead times ranging from 12 to 72 hours, based on 10,000 simulated storm tracks. 30

8 False alarm rates to achieve a 95% probability of detection, measured as percentage of evacuations. 31

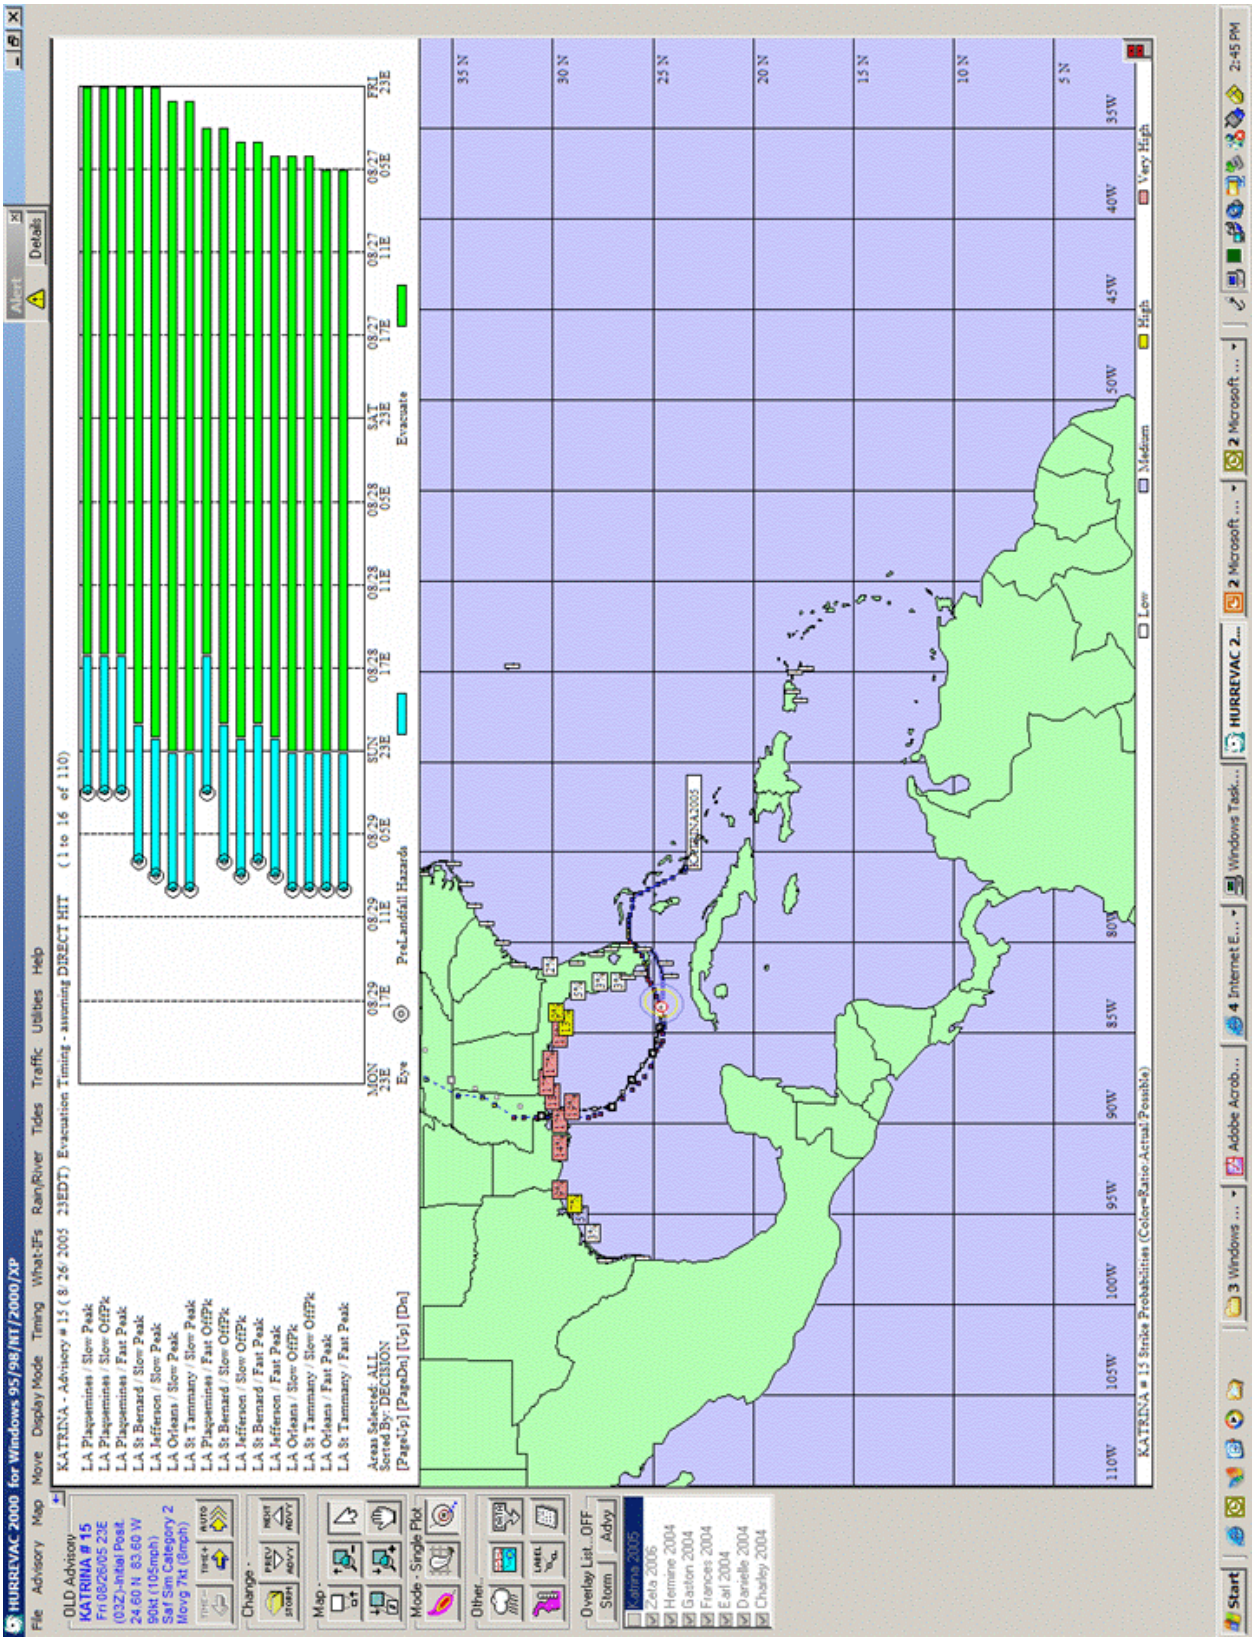

Figure 1: Screen shot from Hurrevac software, showing Hurricane Katrina forecast at 11 p.m. on August 26, 2005, with strike probabilities for coastal locations and, in the pop-up box, recommended evacuation timing for New Orleans and nearby parishes. Time progresses from right to left.

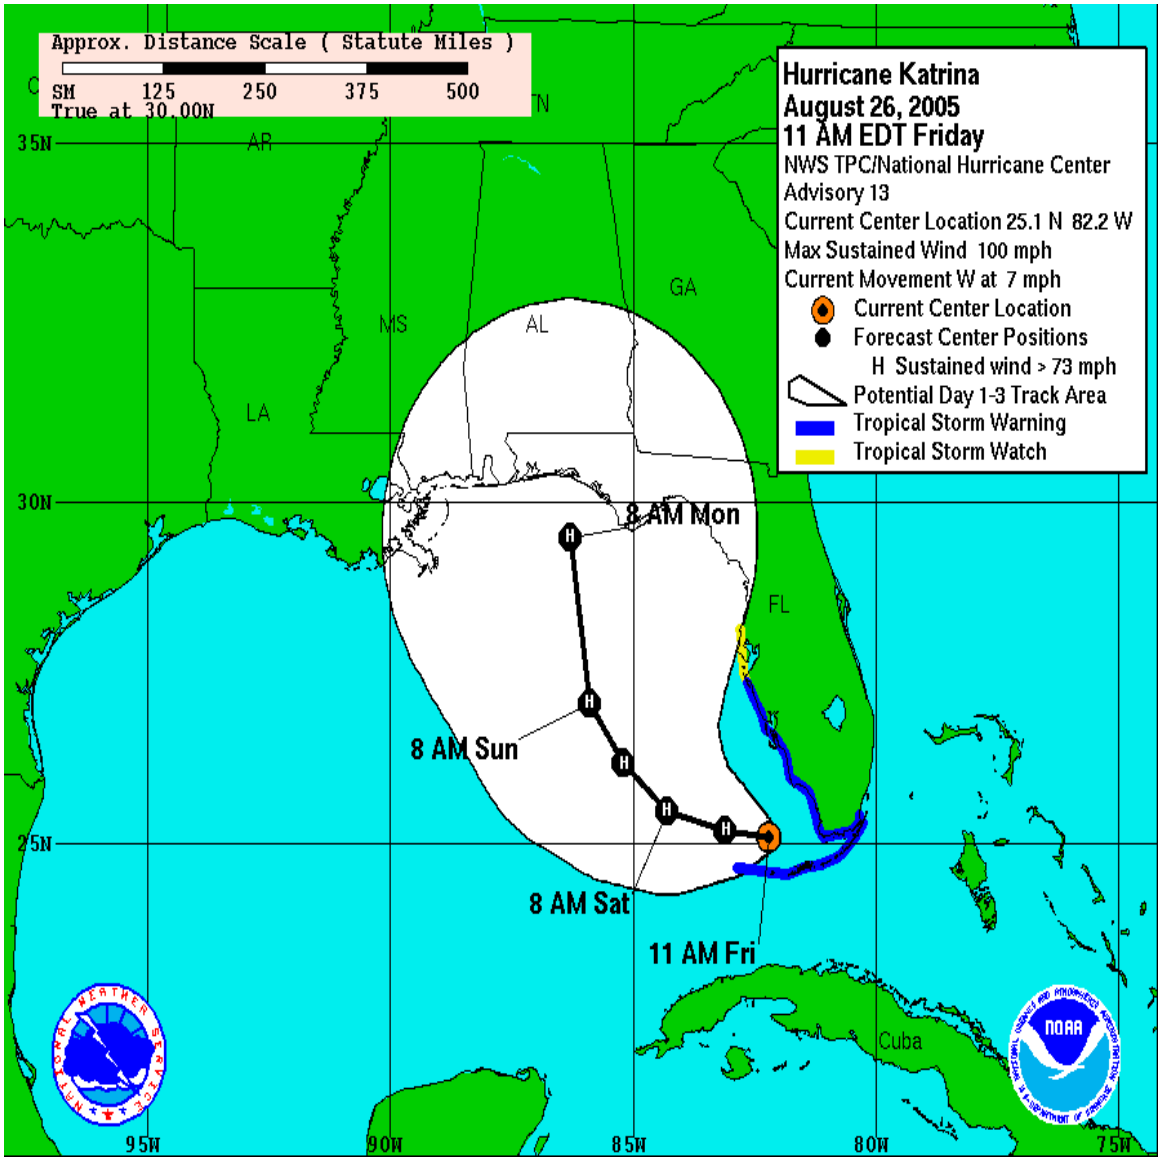

Figure 2: Official NHC 3-day track forecast from the Hurricane Katrina graphics archive.

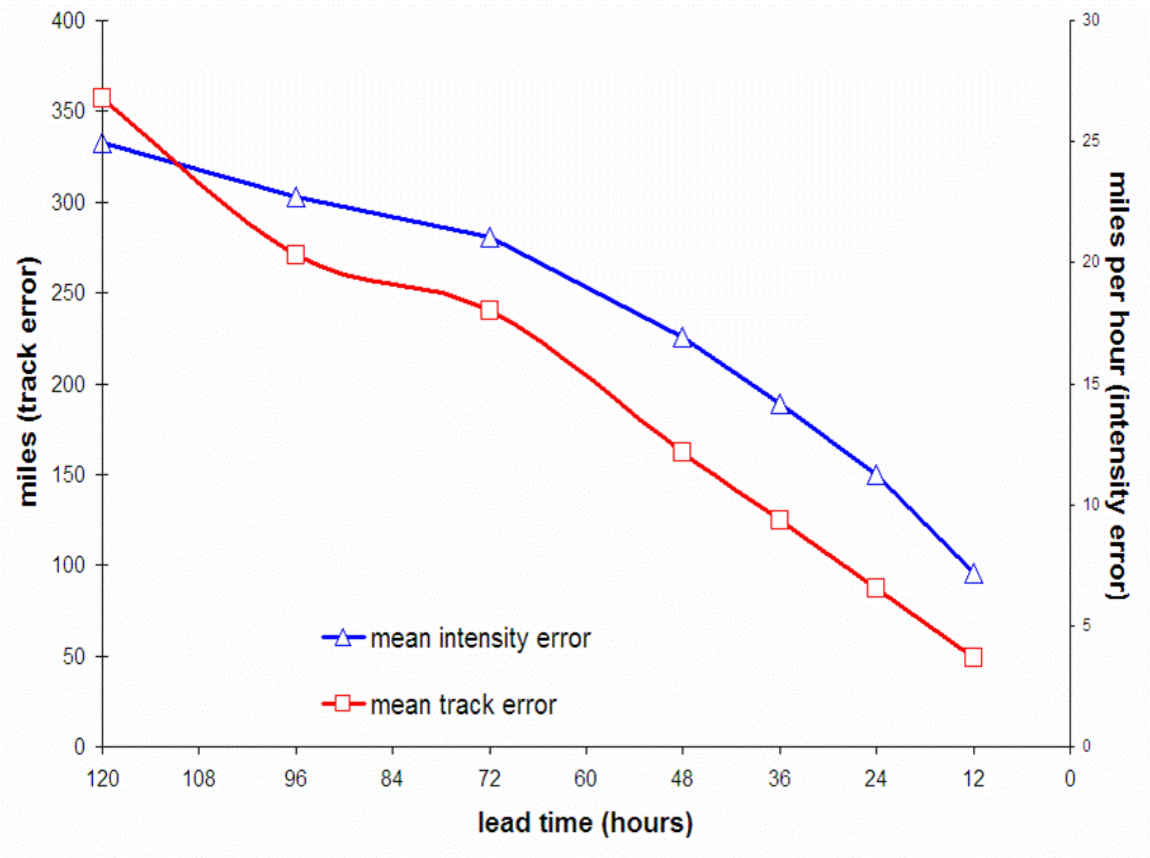

Figure 3: Official forecast errors from Franklin (2005), means over the 1993-2004 seasons.

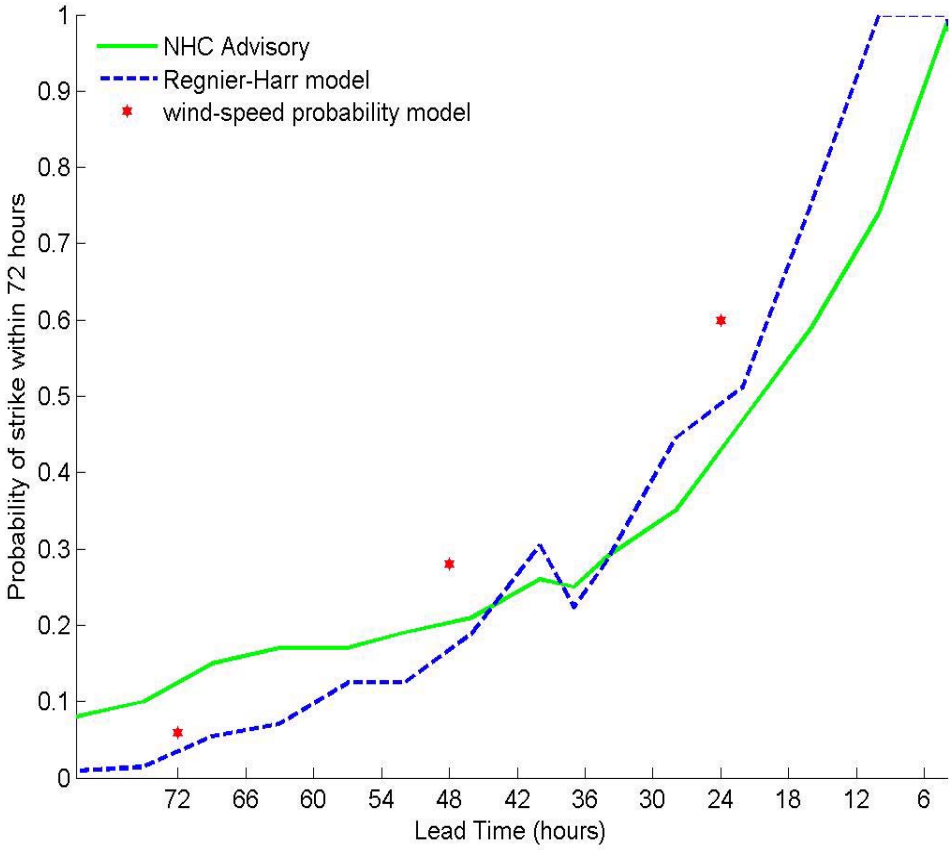

Figure 4: Hurricane Katrina strike probabilities for New Orleans, plotted as a function of time before the storm's center passed the city.

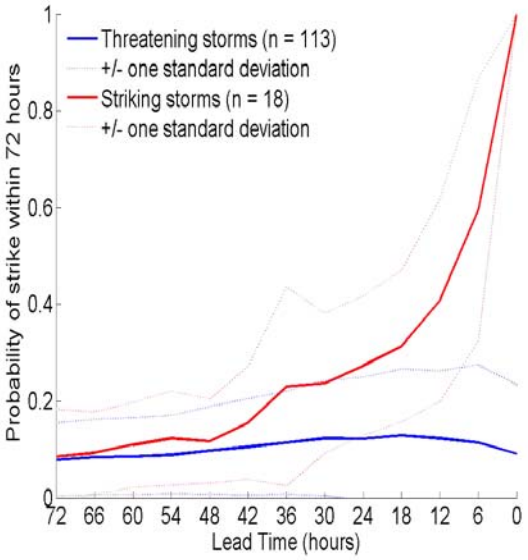

(a) New Orleans

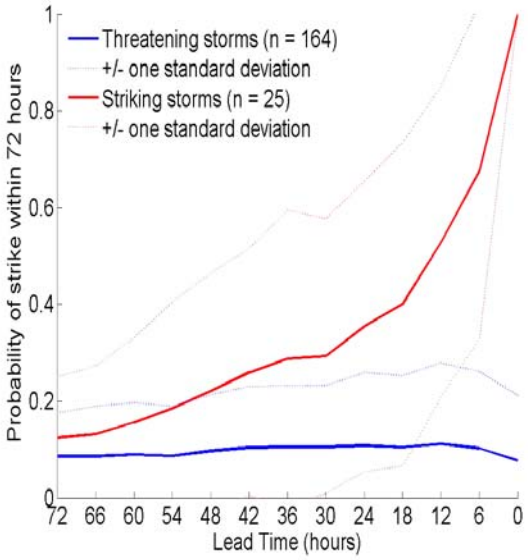

(b) Miami

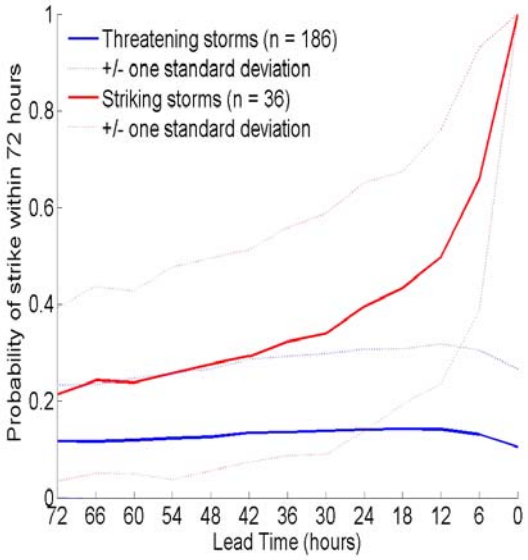

(c) Norfolk

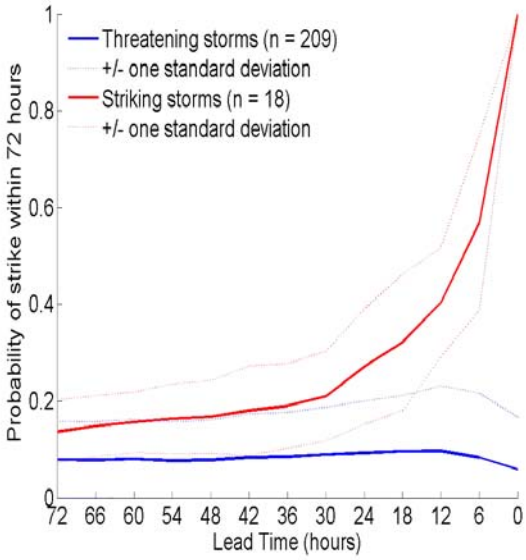

(d) Montauk

Figure 5: Strike probabilities for striking and threatening (but non-striking) historical hurricanes, as a function of lead time.

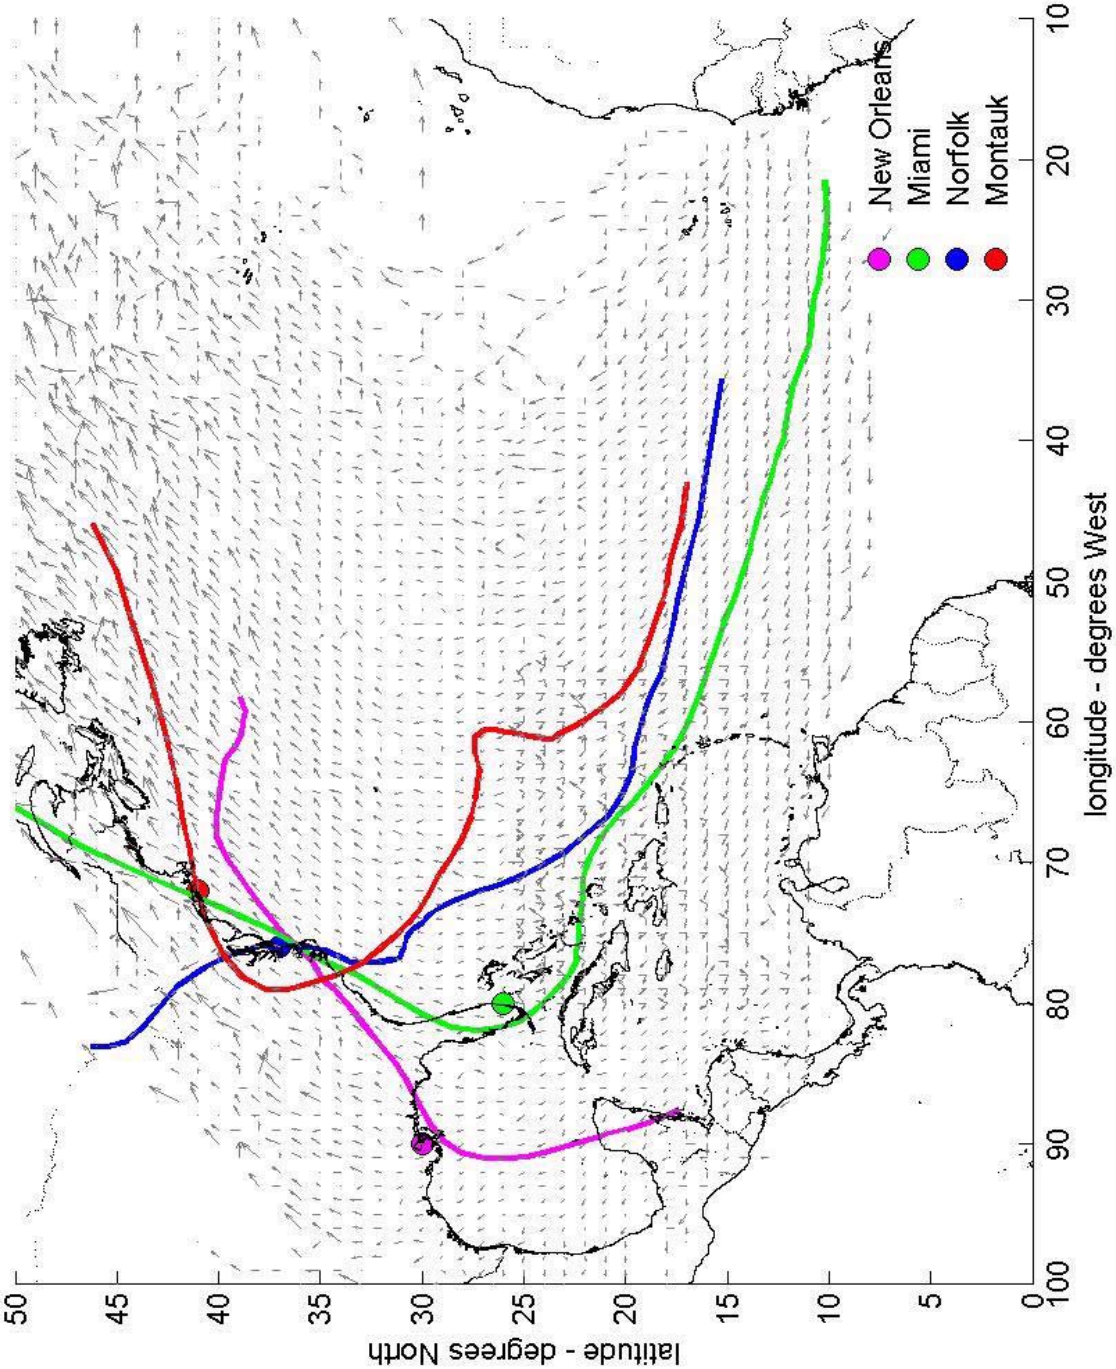

Figure 6: Examples of historical storm tracks. Each storm started at its south end and moved north. The gray arrows indicate the average direction and speed of storms at each location.

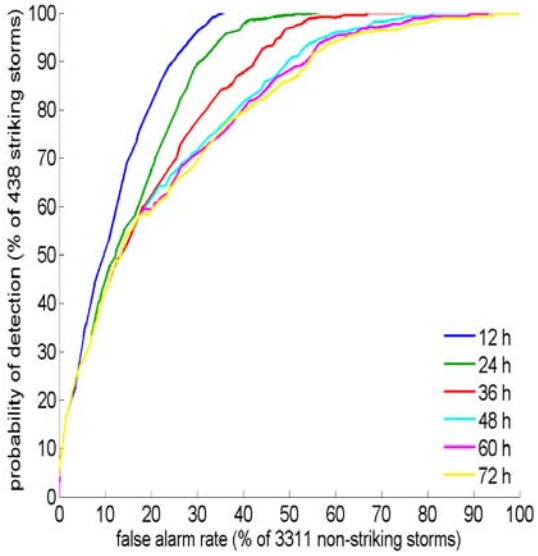

(a) New Orleans

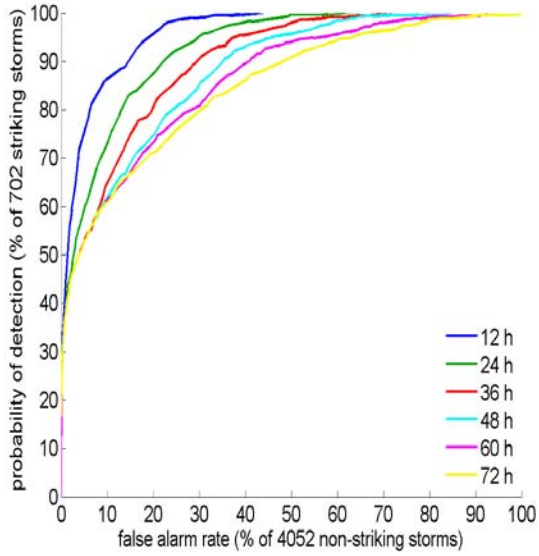

(b) Miami

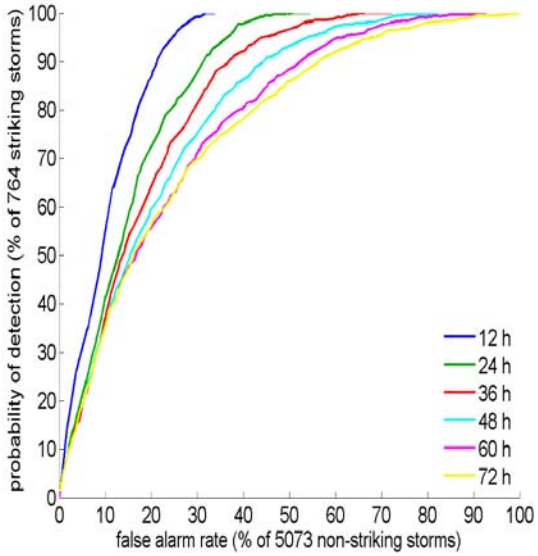

(c) Norfolk

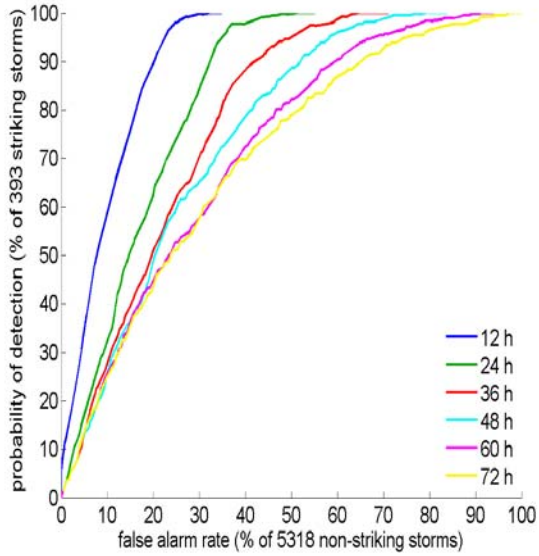

(d) Montauk

Figure 7: Probability of detection as a function of false alarm rate (measured as % of threatening but non-striking storms) for lead times ranging from 12 to 72 hours, based on 10,000 simulated storm tracks.

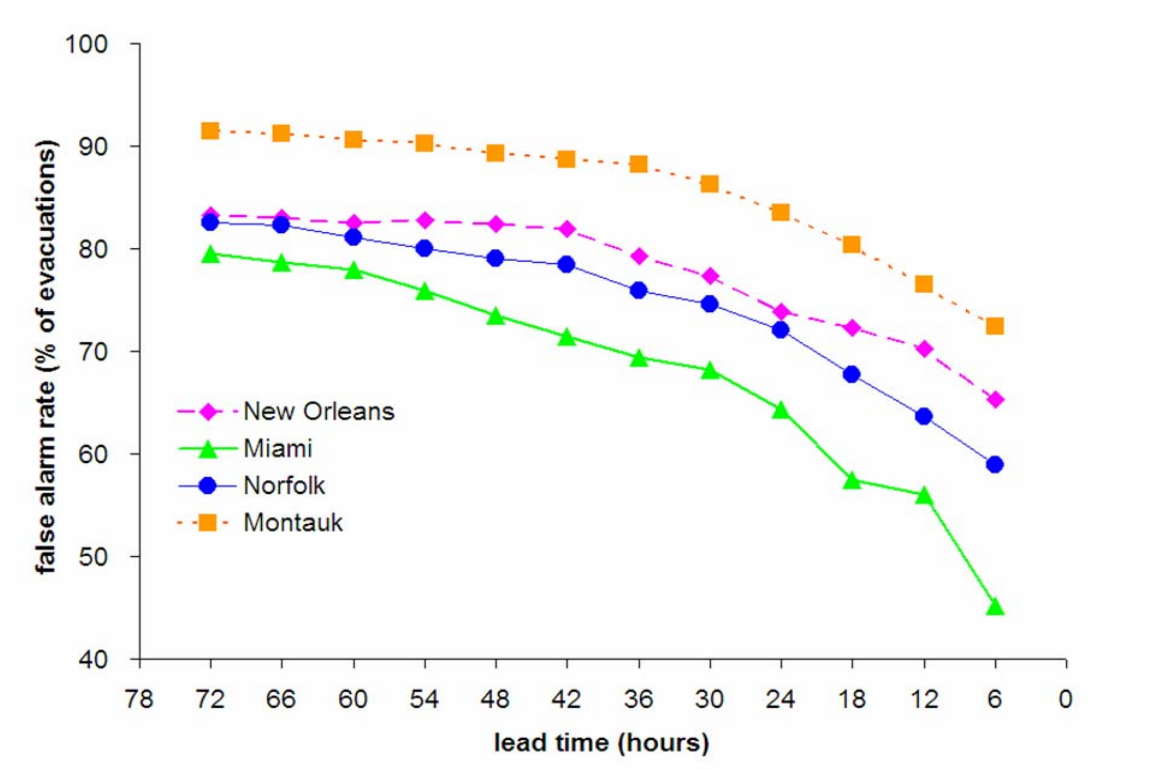

Figure 8: False alarm rates to achieve a 95% probability of detection, measured as percentage of evacuations.

List of Tables

1 Comparison of maximum strike probability as a function of lead time. NHC maximum strike probabilities are given at http://hurricanes.noaa.gov/prepare/products_probability.htm. . . . 33

2 Summary of historical tropical storms and hurricanes threatening and striking each target. . . 34

3 Percentage of historical hurricanes that exceeded the stated strike probability threshold at some point during their evolution. 35

4 False alarm rate (measured as % of evacuations) for each target as a function of probability of detection, calculated using 10,000 simulated storm tracks. 36

5 Summary of expected annual number of false alarms, evacuations, strikes at unevacuated targets (false negatives) and costs for hurricanes. Calculations based on 10,000 simulated storms and six hurricanes per year. 37

6 Comparison of frequency of simulated and historical storms threatening and striking each target. 38

	maximum strike probability (%)	
Forecast Lead Time (hours)	NHC strike probability model	RH Model
72	10-15	12
48	20-25	20
36	25-35	28
24	40-50	44
12	75-85	86

Table 1: Comparison of maximum strike probability as a function of lead time. NHC maximum strike probabilities are given at http://hurricanes.noaa.gov/prepare/products_probability.htm.

among all tropical storms, 1850-2005

target location	annual average threats	annual average strikes	strikes as % of threats	% of striking storms forming with less than 24 h	48 h	72 h
Galveston	3.7	0.5	13	67	78	81
New Orleans	4.0	0.5	13	48	76	83
Key West	4.6	0.6	13	52	73	82
Tampa Bay	4.9	0.6	11	58	77	87
Miami	5.1	0.7	13	58	71	74
Wilmington, NC	6.5	1.0	16	40	55	71
Norfolk	6.6	0.9	13	24	45	61
Montauk	6.6	0.4	7	24	52	64

among all hurricanes, 1850-2005

target location	annual average threats	annual average strikes	strikes as % of threats	% of striking storms forming with less than 24 h	48 h	72 h
Galveston	1.9	0.2	13	62	77	85
New Orleans	2.0	0.2	12	38	77	77
Key West	2.3	0.3	13	53	76	76
Tampa Bay	2.5	0.3	12	50	75	94
Miami	2.6	0.3	13	58	68	74
Wilmington, NC	3.3	0.5	16	34	52	69
Norfolk	3.3	0.4	13	30	48	61
Montauk	3.3	0.2	6	27	64	73

Table 2: Summary of historical tropical storms and hurricanes threatening and striking each target.

probability threshold (%)	among all threatening storms				among non-striking threatening storms			
	New Orleans	Miami	Norfolk	Montauk	New Orleans	Miami	Norfolk	Montauk
10	79	71	81	64	76	67	77	61
20	51	45	57	37	43	37	49	32
30	31	32	47	23	20	21	37	16
40	22	25	38	16	10	14	26	9
50	21	21	29	11	8	9	15	3
60	16	18	22	10	3	5	7	2
70	16	15	21	9	3	2	5	1
80	14	14	17	8	-	1	1	-
90	14	14	17	8	-	1	1	-
100	14	13	16	8	-	-	-	-

Table 3: Percentage of historical hurricanes that exceeded the stated strike probability threshold at some point during their evolution.

lead time (hours)	New Orleans			Miami			Norfolk			Montauk		
	probability of detection (%)											
	90	95	99	90	95	99	90	95	99	90	95	99
6	64	65	68	36	45	61	54	59	60	72	72	74
12	68	70	72	50	56	62	61	64	66	75	77	79
18	70	72	75	52	57	69	65	68	71	80	80	84
24	72	74	78	60	64	74	70	72	74	83	83	86
30	76	77	81	66	68	75	73	75	78	85	86	88
36	78	79	81	66	69	77	73	76	80	87	88	89
42	80	82	84	69	71	78	76	78	82	88	89	90
48	81	82	85	69	74	79	77	79	83	89	89	91
54	82	83	85	72	76	81	79	80	83	90	90	91
60	82	83	86	72	78	82	80	81	84	90	91	92
66	82	83	87	74	79	83	80	82	85	90	91	93
72	82	83	87	76	79	84	81	83	86	91	91	93

Table 4: False alarm rate (measured as % of evacuations) for each target as a function of probability of detection, calculated using 10,000 simulated storm tracks.

	lead-time (hours)	New Orleans			Miami			Norfolk			Montauk		
		probability of detection, %											
		90	95	99	90	95	99	90	95	99	90	95	99
false alarms	24	0.7	0.8	1.0	0.6	0.8	1.4	1.1	1.3	1.6	1.1	1.2	1.5
	48	1.1	1.3	1.7	1.0	1.3	1.8	1.5	1.9	2.5	1.8	2.0	2.6
	72	1.2	1.4	2.0	1.4	1.8	2.4	2.0	2.4	3.1	2.3	2.6	3.2
evacuations	24	0.9	1.1	1.3	1.0	1.2	1.8	1.5	1.7	2.0	1.3	1.5	1.7
	48	1.4	1.6	1.9	1.4	1.7	2.2	2.0	2.3	3.0	2.0	2.2	2.8
	72	1.5	1.7	2.3	1.8	2.2	2.8	2.4	2.8	3.6	2.5	2.8	3.5
strikes without evacuation	24	0.025	0.013	0.003	0.043	0.021	0.003	0.043	0.020	0.004	0.0252	0.009	0.0024
	48	0.025	0.013	0.003	0.040	0.021	0.004	0.047	0.022	0.004	0.0246	0.012	0.0012
	72	0.026	0.013	0.003	0.038	0.022	0.005	0.046	0.023	0.004	0.0222	0.012	0.0018
costs of false alarms ($M)	24	323	369	467	287	390	641	510	598	713	505	566	676
	48	520	599	769	470	601	836	710	873	1,158	814	920	1199
	72	563	646	922	655	819	1,114	909	1,093	1,426	1036	1180	1481

Table 5: Summary of expected annual number of false alarms, evacuations, strikes at unevacuated targets (false negatives) and costs for hurricanes. Calculations based on 10,000 simulated storms and six hurricanes per year.

target location	percent of tropical storms threatening historical	threatening simulated	striking historical	striking simulated
New Orleans	38	33	5	4
Miami	50	41	7	7
Norfolk	63	51	8	8
Montauk	64	53	4	4

Table 6: Comparison of frequency of simulated and historical storms threatening and striking each target.