Comparison of Local and Basinwide Methods for Risk Assessment of Tropical Cyclone Landfall

Timothy M. Hall

NASA Goddard Institute for Space Studies, New York, New York

Stephen Jewson

Risk Management Solutions, London, United Kingdom

ABSTRACT

Two statistical methods for predicting the number of tropical cyclones (TCs) making landfall on sections of the North American coastline are compared. The first method—the "local model"—is derived exclusively from historical landfalls on the particular coastline section. The second method—the "track model"—involves statistical modeling of TC tracks from genesis to lysis, and is based on historical observations of such tracks. Identical scoring schemes are used for each model, derived from the out-of-sample likelihood of a Bayesian analysis of the Poisson landfall number distribution. The track model makes better landfall rate predictions on most coastal regions, when coastline sections at a scale of several hundred kilometers or smaller are considered. The reduction in sampling error due to the use of the much larger dataset more than offsets any bias in the track model. When larger coast sections are considered, there are more historical landfalls, and the local model scores better. This is the first clear justification for the use of track models for the assessment of TC landfall risk on regional and smaller scales.

1. Introduction

Powerful tropical cyclones (TCs) are among the most devastating of natural phenomena, and there is intensive effort to predict their landfall rate along coastlines. Such predictions are needed, for example, by insurance companies to set insurance rates and by governments to establish building codes. The time horizon of interest (seasonal to decadal) is well beyond the scope of numerical weather prediction models, while free-running climate models have insufficient resolution to resolve TCs. Instead, researchers have relied on a variety of statistical models. The most direct approach is to estimate future landfall rates on a segment of coastline from the historical landfalls on that segment (e.g., Elsner and Bossak 2001; Tartaglione et al. 2003). This method, however, is hampered by a dearth of data in regions that are small or experience low activity.

Data limitations can be ameliorated, in principle, by using the additional information contained in TCs that come close to the coast segment in question but do not make landfall or make landfall elsewhere. One such approach is to develop basinwide TC track models, which use historical data across the ocean basin to simulate entire (or partial) TC tracks (e.g., Vickery et al. 2000; James and Mason 2005; Emmanuel et al. 2006). Many such simulations can be performed, resulting in many more landfall events on the coast section in question than the historical record. Sampling error is thereby reduced. Tropical cyclone track modeling is the approach generally taken by the insurance industry.

Track models come with a price, however: they may suffer from bias due to inappropriate or missing statistical descriptions of physical processes. Accumulated along the simulated TC life cycle and projected onto the coast segment, such biases could easily offset the reduction in sampling error. As far as we know, the track models have never been rigorously compared to other models. How do their landfall rates compare to local analyses? What is the balance between sampling error and bias? These are the questions we address in this paper.

We compare TC landfall predictions along the North American Atlantic coast produced by a "local model,"

which only uses landfall data on the coastline segment in question, and a "track model," which simulates the trajectory of TCs from genesis through lysis. After presenting the local and track models, we discuss the analysis of landfall probability and devise and apply a scoring system that allows direct comparison between the models. We find that bias in the track model is more than compensated for on most regional-scale coast segments by the reduction of sampling error compared to the local model. This is the first time, to our knowledge, that the use of basinwide statistical track models has been rigorously justified for use in TC landfall risk assessment.

2. Local model

The local model makes predictions of future TC landfall rates on a segment of coastline using only historical landfall events on that segment. Consider the case of i historical TC landfalls in m yr on some segment of coastline. To the extent that TCs behave independently, TC landfall can be considered as a Poisson process (e.g., Bove et al. 1998). (We have verified empirically the Poisson character of landfall in the TC track model described in section 3.) The most straightforward way to model landfall in a subsequent year is to draw from a Poisson distribution, $f(n) = e^{-\lambda}\lambda^n/n!$, for the number n of TC landfalls in a year. The rate (mean landfall count per year) is $\lambda = i/m$.

This approach is unrealistic for small i. In the most severe case, one may have a coastline segment with no historical landfall events. Naïvely, $i = 0$ implies $\lambda = 0$, and zero probability of landfall is predicted, even though TC landfall cannot be ruled out as meteorologically impossible. The problem is that the true underlying Poisson rate is not known. The simple Poisson model does not account for the fact that $i = 0$ landfalls in m yr is perfectly consistent with underlying rates that are nonzero.

A better approach is to sum over all possible Poisson rates that could have resulted in i landfalls in m years (e.g., Epstein 1985). That is, we compute the probability of n landfalls in a year, given i observed landfalls in m yr, as

$$f(n|i) = \int f(n|\lambda)f(\lambda|i)\,d\lambda. \qquad (1)$$

Inside the integral $f(n|\lambda)$ is the probability of n landfalls in a single year, given a rate λ. This is simply the Poisson density:

$$f(n|\lambda) = \frac{e^{-\lambda}\lambda^n}{n!}. \qquad (2)$$

The second factor $f(\lambda|i)$ in the integral is the probability of a Poisson rate λ, given i TC landfalls in m yr. It can be factorized using Bayes's theorem:

$$f(\lambda|i) \propto f(i|\lambda)f(\lambda). \qquad (3)$$

The first term on the rhs of (3) is again the Poisson density, here for i landfalls in m years, given an annual rate λ. The second term on the rhs of (3) is the "prior" distribution, which summarizes any previous knowledge of the rate. We assume our prior knowledge to be uninformed, and choose a uniform value for the probability of λ; that is, $f(\lambda) = c$, and, therefore, $f(\lambda|i) = cf(i|\lambda)$. The constant c is determined by the normalization requirement:

$$\int f(\lambda|i)\,d\lambda = c\int f(i|\lambda)\,d\lambda = 1. \qquad (4)$$

Substituting the Poisson density $f(i|\lambda) = e^{-\lambda m}(\lambda m)^i/i!$, where m is the number of historical years, into (4), one finds $c = m$. The (posterior) distribution is then

$$f(\lambda|i)\frac{m}{i!}e^{-\lambda m}(m\lambda)^i. \qquad (5)$$

Note that while (5) has the form of a Poisson distribution, relative to (2) the roles of i and λ have been reversed. The random variable is now λ, and as a function of the λ expression (5) is a Gamma distribution. We can now perform the integration in expression (1) to obtain

$$f(n|i) = \frac{(i+n)!}{i!n!}\left(\frac{m}{m+1}\right)^{i+1}\left(\frac{1}{m+1}\right)^n, \qquad (6)$$

which is an example of the negative binomial distribution. Note that a different derivation of (6) can be found in Elsner and Bossak (2001).

Expression (6) constitutes the local model for the landfall. It predicts the probability of n landfalls on a coastline segment, given i observed landfalls in m years.

3. Basinwide track model

Hall and Jewson (2007) describe a statistical model of TC tracks in the North Atlantic from genesis to lysis based on the National Hurricane Center's (NHC) North Atlantic hurricane dataset (HURDAT; Jarvinen et al. 1984). Hall and Jewson (2007) used HURDAT data from 1950 to 2003. Here, we extend the data period to 1950–2005, which encompasses 595 TCs. Observations prior to 1950 are less reliable, as they precede the era of routine aircraft reconnaissance.

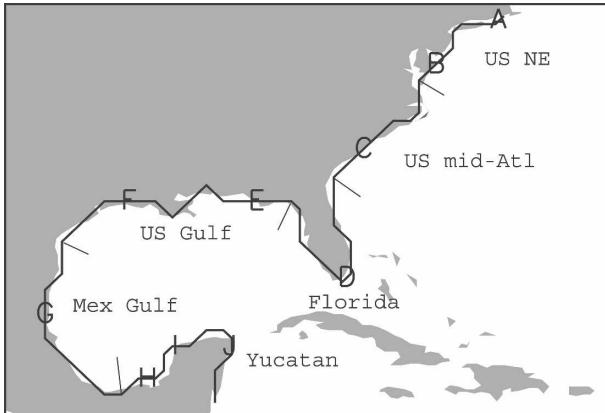The track model consists of three components: 1) genesis, 2) propagation, and 3) lysis (death). The num-

FIG. 1. Eastern North American coastline divided into 39 segments. Select segments are labeled with letters for reference, as are two additional segments, 4 and 17. Also shown are regions, as labeled and delineated by the short line segments.

ber of TCs in a simulation year is determined by random resampling of the historical annual TC number. Genesis sites are simulated by sampling a pdf composed of a sum of Guassian kernels around historical genesis sites. For propagation, we compute mean latitude and longitude 6-hourly displacements and their variances, by averaging of "nearby" historical displacements. Standardized displacement anomalies are modeled as a lag-one autoregressive model, with latitude and longitude treated independently. The autocorrelation coefficients are computed from "nearby" historical anomalies. Finally, TCs suffer lysis with a probability determined by averaging nearby historical lysis rates.

In each of these model components it is necessary to choose length scales. In the case of the means, variances, and autocorrelation, these length scales weight the averaging of historical data, thereby defining what is "near" to a current simulation point. In the case of a genesis site, a length scale constitutes the kernel bandwidth, which is uniform over the domain. The length

scales are selected to maximize an average log likelihood of the observations, given the statistical models (e.g., Aldrich 1997). The average log likelihood is determined, in turn, by a jackknife ("drop one out") out-of-sample calculation (Quenouille 1949; Tukey 1958).

Hall and Jewson (2007) describe the details of the model formulation. They have also performed extensive diagnoses of the track model, comparing ensembles of 1950–2003 simulations to the historical tracks in terms of track density across the basin, rates at which tracks crossed various latitude and longitude lines, and landfall rates. In many regions and diagnostics the simulated tracks were statistically indistinguishable from the historical tracks (as determined by Z score tests across the simulation ensemble). In other cases the track model displayed biases, in particular an underestimate of landfall rate on the mid-Atlantic coast and the northern Gulf coast.

4. Landfall probabilities

Following Hall and Jewson (2007), we have divided the North American coastline from Maine to the Yucatan Peninsula into 39 segments of different lengths, as shown in Fig. 1. A "landfall event" occurs on a segment when a TC trajectory intersects the coastline segment heading sea to land. Note that with this definition a single TC can make multiple landfalls.

Figure 2a shows the 595 historical TC tracks from the 56-yr period 1950–2005, while Fig. 2b shows 3500 of the 10 660 tracks from a 1000-yr track-model simulation based on this historical data. Figure 3 shows the historical landfall rates on the 39 segments in the 56-yr period 1950–2005, expressed as landfalls per year per 100 km of segmented coastline. Also shown are the landfall rates predicted by the track model in a 1000-yr simulation based on the 1950–2005 HURDAT data. To compare to the historical landfall rates, we have broken the simulation into 17 segments of 56 yr each and computed
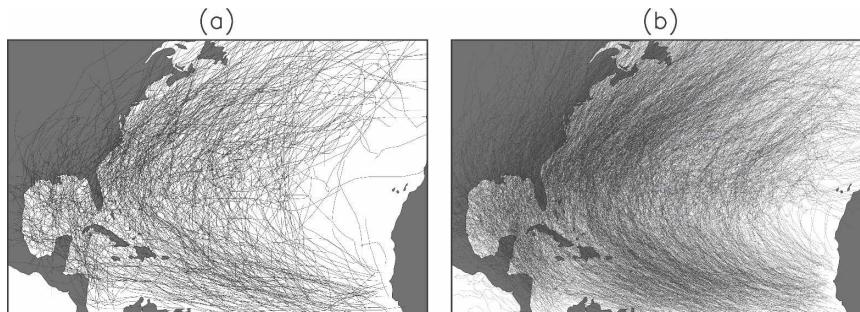

FIG. 2. (a) Historical TC tracks from the 56-yr period 1950–2005 and (b) simulated TC tracks from a 1000-yr simulation of the track model based on the same historical data (only every third TC is plotted, for clarity).
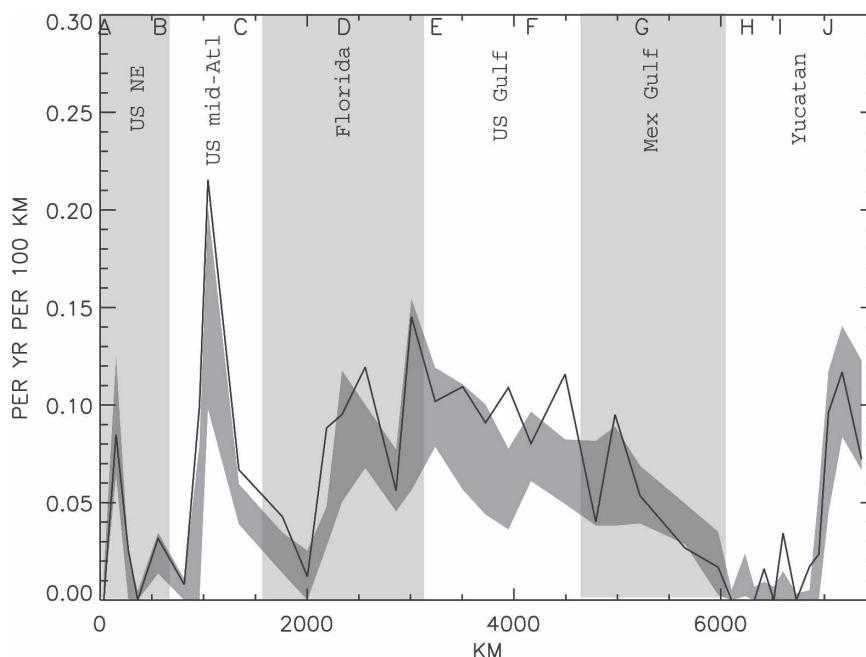
FIG. 3. Historical landfall rates in the 56-yr period 1950–2005 are expressed as landfalls per 100 kilometers of segmented coastline per year (solid). Landfall rates in a 1000-yr track-model simulation are expressed in the same units (dashed). Letters along the top axis indicate reference points, and shading indicates regions defined in Fig. 1.

the mean and standard deviation of the landfall rate at each coast segment across the 17 periods. The shaded region in the figure represents the mean ± one standard deviation. The simulated landfall mimics the geographic variation of the historical rates well, but there is an overall underestimate of the rate. Among the historical TCs, the rate of landfall over the entire segmented coastline defined in Fig. 1 is 4.7 yr^{-1}, while among the simulated TCs, the rate is 4.2 yr^{-1}. The model's underestimate of the landfall rate over the entire coastline is significant. Over the 17 periods of 56 yr each, the model predicts a 56-yr landfall count on the total segmented coastline of 225 ± 23, as compared with 265 total landfalls for the historical 56-yr period. The model bias on individual locations varies.

To elucidate better the different landfall probabilities of the local and track models, we examine in detail the two coast segments: segments 1 (southeast New England) and 14 (southwest Florida). In the period 1950–2005, there were no historical landfalls on segment 1. What, then, is the probability of making n landfalls in a subsequent year? The local model predicts $f(n|i)$ of expression (6) with $i = 0$ and $m = 56$ yr. In the 1000-yr track simulation, there are three landfalls on segment 1, and the probability for n landfalls in a subsequent year is $f(n|i)$ evaluated at $i = 3$ and $m = 1000$. The results are listed in Table 1. Both track and local models predict by

far the highest probability at zero landfalls. However, the local model predicts roughly 4 times the probability of (0.017 versus 0.004) of one or more landfalls, despite having a lower rate. This is because the historical record does not constrain well a Poisson rate that is so low. For the track model the Poisson rate is better constrained, and higher landfall numbers have lower probability.

On segment 14, there are 25 landfalls in the 56 historical years and 320 landfalls in the simulated 1000 yr. Here, the Poisson rates for both the track model and the local model are well constrained. Both model predictions have the probability peaked at zero landfalls,

TABLE 1. Local- and track-model probabilities of $n = 0$, 1, and 3 landfalls on two coast segments. The mean rates, expressed as landfall number over number of years, are also shown.

Coast segment 1 (SE New England)			
	Rate	$n = 0$	$n = 1$	$n = 2$
Local	0/56	0.982	0.017	0.000
Track	3/1000	0.996	0.004	0.000
Coast segment 14 (SW Florida)			
	Rate	$n = 0$	$n = 1$	$n = 2$
Local	25/56	0.631	0.287	0.068
Track	320/1000	0.726	0.232	0.037

but the track model predicts a greater probability for zero landfalls (0.726 versus 0.631) and lower probability for one or more landfalls. This is due to a significant negative bias in the track model's rate on this segment.

5. Scoring the models

To decide which of the models is genuinely better, we need to evaluate the predicted probabilities with actual landfall counts from historical years not included in model construction. The local and track models are scored using identical out-of-sample log-likelihood evaluation. Following a Quenouille–Tukey jackknife procedure (Quenouille 1949; Tukey 1958), we choose a year, j, in the 56-yr range 1950–2005 (the "out of sample" year) for which a model prediction is to be made, and consider a section, C, of coastline. For the local model the distribution $f(n|i)$ of expression (1) is calculated using the $i = i_{his}$ historical landfalls on C in the $m = 55$ yr excluding j (the "in sample" years). If the observed number of landfalls on C in year j is n_{obs}, then the model's likelihood is $f(n_{obs}|i_{his})$. We obtain a total score, $S_{loc}(C)$, for the local model on C by averaging the log likelihoods over all the out-of-sample years j.

For the track model we pick an out-of-sample historical year j and construct the model from the 55 in-sample years of HURDAT full-basin data in the 1950–2005 range excluding year j. We then simulate TCs over a large number of model years (1000) and count the landfalls in the coastline segments. The distribution $f(n|i)$ is computed as for the local model, but now $i = i_{sim}$ is the number of simulation TC landfalls that occur in the $m = 1000$ simulation years. The likelihood is $f(n_{obs}|i_{sim})$. We repeat this process for the all of the 56 out-of-sample years, each time reconstructing the track model and performing 1000 yr of TC simulations. The total score, $S_{tra}(C)$, of the track model is again the average of the log likelihoods.

The scoring procedures for the two models are identical. Only the sources and sizes of the landfall data differ. Any advantage to the track model will be a consequence of the fact that there are many more landfalls ($i_{sim} \gg i_{his}$). The negative binomial distribution (6) tends toward a more narrow Poisson distribution for large i, providing the possibility of a higher likelihood for the observed landfall and a higher score for the track model. Any disadvantage of the track model will be a consequence of the fact that its Poisson rate is wrong, causing the observed landfall count to fall well outside the distribution's peak and be erroneously assigned a low likelihood. In summary, the track model is more precise but less accurate, while the local model is more accurate but less precise. Our scoring scheme will capture the relative merits of precision and accuracy.

6. Results

Figure 4 shows the score difference $S_{tra} - S_{loc}$ on each coastline segment, plotted as a function of distance along the coast starting from the northeast. In 29 of the 39 coast segments, $S_{tra} > S_{loc}$. The significance of the track model "winning" 29 of 39 segments is high. If the models were equally likely to win, the probability of the local model winning 10 or fewer times is only 0.002. We discount this possibility and conclude that the track model is genuinely better overall at predicting local landfall.

On coast segments where the track model has a higher score either 1) the track model matches the historical landfall rate closely or 2) the track model has a modest bias and there are few historical landfalls. In case 1, the track model's landfall rate is accurate and only a small decrease in sampling error compared to the local model is enough for a higher likelihood. In case 2, the track model is inaccurate, but this is more than compensated for by its greater precision compared to the poorly sampled local model. On coast segments where the track model has a lower score, the opposite is true. Either 1) the track model suffers a large bias or 2) the track model suffers a modest bias and there are relatively many historical landfalls. The filled symbols in Fig. 4 indicate coast segments where the track model's landfall rate differs from the historical rate by at least one standard deviations, as computed across the seventeen 56-yr periods (i.e., its Z value magnitude is greater than 1). The track model loses in 9 of the 15 locations where this is true. In four of the six locations with $|Z| > 1$ that the track model wins, there are one or fewer historical landfalls, so that the model's great advantage in precision compensates for its poor accuracy. By contrast, in the one location (segment 19, along the U.S. Gulf Coast) that the track model loses despite having $|Z| < 1$, there are 15 historical landfalls, providing enough precision to the local model that its modest advantage in accuracy results in a win.

When many segments of coast are considered together, there are more landfalls, the sampling error of the local model is reduced, and the overall bias of the track model dominates the likelihood comparison. Performing an identical out-of-sample scoring comparison with all 39 segments of the North American Atlantic coast taken together results in a higher likelihood for the local model than the track model. We can also compare the models on the subsets of the coast indicated in Fig. 1. The track model has a higher likelihood for the U.S. Northeast (segments 1–5), the U.S. mid-Atlantic coast (segments 6–9), the Mexican Gulf Coast (segments 23–27), and the Yucatan Peninsula (segments
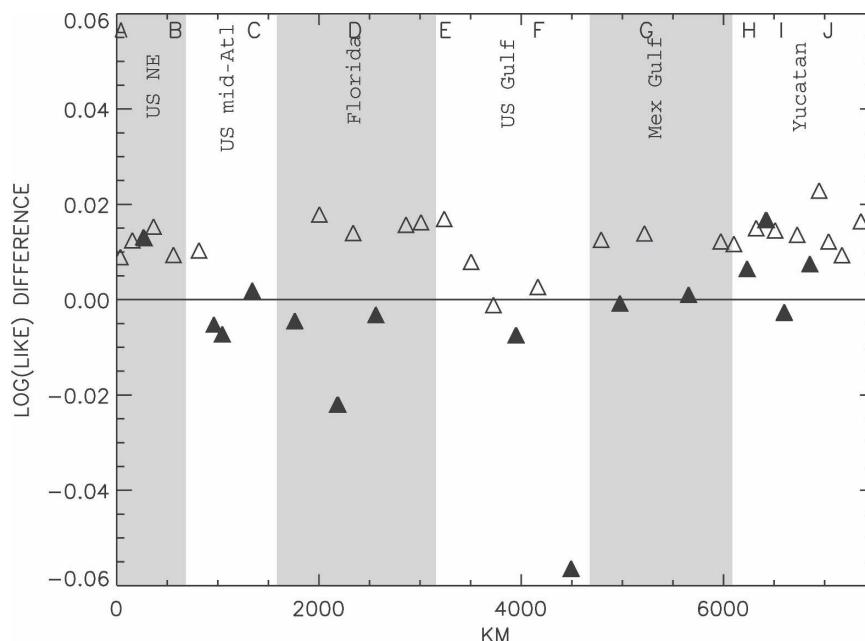
FIG. 4. The score difference, $S_{tra} - S_{loc}$, between the track and local models on each coast segment as a function of distance along the segmented coast, measured from the northeast. The track model beats the local model in 29 of 39 segments. Filled symbols indicate locations where the track model mean rate differs from the 56-yr historical mean by more than one standard deviation, as measured across 17 of 56-yr simulations. Letters along the top axis indicate reference points and shading indicates regions defined in Fig. 1.

28–39). The local model has a higher likelihood on the Florida peninsula (segments 10–16) and on the U.S. Gulf Coast (segments 17–22). These regional comparisons are summarized in Table 2, which also lists the historical and simulated landfall rates. The track model wins regions where the biases are small and the landfalls are relatively infrequent. The track model loses regions where the biases are large and the landfalls are relatively frequent.

7. Conclusions

We have compared two statistical models of TC landfall: 1) a "local model," which is built solely on historical landfall events on the coastline segment of interest; and 2) a "track model," which simulates entire TC tracks from genesis to lysis using historical data (HURDAT) over the full North Atlantic Basin. Both types of models have been used in the literature for predicting TC landfall rates (Elsner and Bossak 2001; Vickery et al. 2000). Track models have been preferred by the insurance industry because they make more complete use of historical data, but to our knowledge there has not been any rigorous demonstration that the consequent reduction in sampling error outweighs the potential increase in bias.

Our results justify, for the first time, the use of track models over local models for landfall risk assessment on regional and smaller scales. Over much of the North Atlantic coastline the track model of Hall and Jewson (2007), despite displaying significant bias, is genuinely better at predicting landfall rates than a local model, based on a jackknife out-of-sample evaluation of the log likelihood of observed landfalls. The track model has higher likelihood than the local model on coast sections with relatively few landfalls, because its greatly reduced sampling error more than compensates for its reduced accuracy. The regions where the local model has higher likelihood tend to be regions with many historical landfalls, reducing the local sampling error, or regions of particularly large track-model bias. When the entire coastline (from the U.S. Northeast through the Mexican Yucatan Peninsula) is taken together, the increased number of historical landfalls reduces the sampling error of the local model, and it has a higher log-likelihood score than the track model. On intermediate-sized regions the results are mixed, with the local model winning some (Florida peninsula, U.S. northern Gulf Coast) and the track model winning others (U.S. Northeast, U.S. mid-Atlantic, Mexican Gulf Coast, Yucatan Peninsula), depending on the relative magnitudes of sampling error and biases.

TABLE 2. Historical and simulated landfall rates (counts per year) on the six larger coastal regions in Fig. 1. Also listed is the model (local or track) with the higher log-likelihood score.

	U.S. Northeast	U.S. mid-Atlantic	FL	U.S. Gulf Coast	Mexican Gulf Coast	Yucatan Peninsula
Historical rate	0.25	0.63	1.11	1.64	0.57	0.54
Simulated rate	0.24 ± 0.05	0.46 ± 0.08	0.88 ± 0.09	1.37 ± 0.15	0.64 ± 0.14	0.52 ± 0.10
Higher likelihood	Track	Track	Local	Local	Track	Track

We note that the sampling error incurred by the local model could be reduced by judicious use of the less reliable landfall records from the nineteenth and early twentieth centuries (Elsner and Bossak 2001). On the other hand, biases in the track model may well be reduced with additional model development. The Hall and Jewson (2007) model is relatively simple, for example, taking no account of the dependence of TC tracks on date of year or TC intensity. It may be possible to draw on the strengths of both local and track models with a hybrid model, in which a track model is somehow optimally "tuned" to local landfall rates, at least in regions where landfall events are sufficiently frequent. Qualitatively, local information would dominate in active regions, while the track model would dominate in regions of rare landfall. The scoring scheme described here could be used to evaluate rigorously a hybrid model, or any other landfall prediction scheme.

So far we have only tested landfall for all named TCs taken together. It would be interesting to perform a separate analysis for the landfall of intense hurricanes (e.g., category 3 and higher). We expect that the track model would be better in this case, too, as the many fewer intense TCs will cause higher sampling errors for the local model. Similarly, the track model is likely to perform well compared to a local model for other regions, such as the Indian and Pacific Oceans, where the data are more sparse. To complete the evaluation of the track model in this way, its intensity component needs to be tested, not just the landfall. We are at present developing a statistical intensity component to the Hall and Jewson (2007) model, and will perform rigorous likelihood comparisons to local models when the intensity model is ready.

We have ignored the effects of long-term climate variability and change in our analysis, focusing instead on comparing model predictions of landfall rate assuming stationary distributions. In reality climate cycles such as ENSO have significant influence TC landfall (Bove et al. 1998), and there is evidence and growing concern that anthropogenic climate change is leading to increased TC durations and intensities (Webster et al. 2005). The statistical models described here can be combined with models that attempt to predict varia-

tions in the overall numbers of TCs on interannual time scales. Alternatively, or in addition, the model construction can be conditioned on the phase of a climate cycle, such as ENSO, or on certain past climatic conditions that are suspected to be more common under global warming, such as high sea surface temperature. Any such conditioning involves a reduction in data, and the relative benefit of the track model relative to a local model is enhanced.

REFERENCES

Aldrich, J., 1997: R. A. Fisher and the making of maximum likelihood 1912–1922. *Stat. Sci.,* **12,** 162–176.

Bove, M. C., J. B. Elsner, C. W. Landsea, X. Niu, and J. J. O'Brien, 1998: Effect of El Niño of U.S. landfalling hurricanes, revisited. *Bull. Amer. Meteor. Soc.,* **79,** 2477–2482.

Elsner, J. B., and B. H. Bossak, 2001: Bayesian analysis of U.S. hurricane climate. *J. Climate,* **14,** 4341–4350.

Emanuel, K. A., S. Ravela, E. Vivant, and C. Risi, 2006: A statistical deterministic approach to hurricane risk assessment. *Bull. Amer. Meteor. Soc.,* **87,** 299–314.

Epstein, E. S., 1985: *Statistical Inference and Prediction in Climatology: A Bayesian Approach. Meteor. Monogr.,* No. 42, Amer. Meteor. Soc., 199 pp.

Hall, T. M., and S. Jewson, 2007: Statistical modeling of North Atlantic tropical cyclone tracks. *Tellus,* **59A,** 486–498.

James, M. K., and L. B. Mason, 2005: Synthetic tropical cyclone database. *J. Waterway, Port, Coastal, Ocean Eng.,* **131,** 181–192.

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

Quenouille, M., 1949: Approximate tests of correlation in time series. *J. Roy. Stat. Soc.,* **11B,** 18–84.

Tartaglione, C. A., S. R. Smith, and J. J. O'Brien, 2003: ENSO impact on hurricane landfall probabilities for the Caribbean. *J. Climate,* **16,** 2925–2931.

Tukey, J. M., 1958: Bias and confidence in not quite large samples. *Ann. Math. Stat.,* **29,** 614.

Vickery, P. J., P. Skerlj, and L. Twisdale, 2000: Simulation of hurricane risk in the US using an empirical track model. *J. Structural Eng.,* **126,** 1222–1237.

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