A STOCHASTIC DAILY MEAN TEMPERATURE MODEL FOR WEATHER DERIVATIVES

Jeffrey Viel[1, 2], Thomas Connor[3]

[1]National Weather Center Research Experiences for Undergraduates Program
and
[2]Plymouth State University

[3]Atmospheric and Environmental Research Inc.
Lexington, MA

ABSTRACT

Weather derivatives are usually priced by analyzing the climatologic data of an underlying weather index. This research proves when using temperature as an underlying weather index, climatology was not fully representative of future outcomes. Weather derivative contracts based on temperature are measured by degree days, which are a metric of energy consumption. Previous research attempted to develop techniques to model degree days, these techniques were based on invalid statistical assumptions and lacked robustness. Due to the fact that degree days are an aggregate monthly metric and path dependent, it was important to model the complete behavior of a time series by simulating the daily mean temperature.

This research provides an in depth statistical analysis of the daily mean temperature time series for eighteen cities from the Chicago Mercantile Exchange (CME). The residuals, which represented the difference between the observed data and trend, were used to develop two models to simulate a possible temperature time series for 2007. A distribution of ten thousand possible outcomes were created for each model, then analyzed against the climatologic data sets. Ultimately, this research exhibits that statistics extracted from the analysis of the residuals could be simulated to produce realistic outcomes of degree days for weather derivative contracts.

1. INTRODUCTION

Weather derivatives allow businesses and organizations to protect themselves against unfortunate weather fluctuations. Weather fluctuations can be categorized as either catastrophic or non-catastrophic events. Although weather derivatives are a form of insurance, they protect against non-catastrophic weather events. For example, a company may not suffer tangible damage (i.e. property damage) from a non-catastrophic weather event, but if a businesses' livelihood relies upon consistent weather, a long period of undesired conditions could severely damage its revenue. Insurance policies on the other hand, protect against catastrophic weather events such as hurricanes and tornadoes, in which home and property damage is more likely. Weather derivatives allow businesses to mitigate weather loss damagers by referencing weather observations. By comparison, insurance contracts require businesses to prove weather related damages.

The first weather derivative transaction in the United States occurred during the 1997-1998 winter season, collectively one of the most significant El Nino seasons (Mraoua et al, 2005). Since their introduction into the Chicago Mercantile Exchange (CME) in 1999, these financial derivative contracts have been on the rise ever since. The constellation of market makers and market participants today offer the weather market greater depth, breadth and financial security than ever. Its numbers include several of the strongest financial institutions on the globe.

[1,2] Jeffrey Viel
147 City Depot Road Charlton, MA 01507
jpviel@plymouth.edu
[3]Thomas Connor
AER 131 Hartwell Avenue Lexington, MA 02420
tconnor@aer.com

(WRMA, 2006)

Due to the fact that weather events and energy prices have always been highly correlated, the potential for weather derivatives was first discovered by energy companies. Today, interest has grown outside of the energy market and into the leisure and agricultural markets. Companies who set pricing for weather derivatives rely on the accuracy of the weather data.

Weather derivative contracts are typically structured as swaps, futures, and call-put options based upon different underlying weather indices (Alaton, 2003). All contracts traded on the CME are based off of the data collected from official weather stations identified by the CME. The CME has structured contracts that trade on accumulated degree days for time periods, ranging from months to seasonal strips, written for the eighteen stations traded on the exchange. The market participants collectively determine the strike location for the HDD and CDD contracts, for which the CME has prescribed a $20.00 per degree day tick size. The participants may buy and sell as many contracts as required to cover their temperature exposure.

When using temperature as the underlying index in weather derivative contracts, temperature is measured by degree days. Degree days are an aggregate monthly metric of energy consumption, and are calculated with the daily mean temperature. Equation 1 shows how the daily mean temperature is calculated by averaging the maximum and minimum daily temperature of a day. Equation 2 exhibits how each day's calculated daily mean temperature is then differenced against a reference temperature (18°C or 65°F). Equation 3 is lastly the summation of these degree days are than taken for each month.

$$T_i = \frac{T_{max} + T_{min}}{2} \qquad (1)$$

$$\begin{aligned} HDD_i &= 18°C - T_i \\ CDD_i &= T_i - 18°C \end{aligned} \qquad (2)$$

$$\begin{aligned} HDD_m &= \sum_{i=1}^{n} HDD_i \\ CDD_m &= \sum_{i=1}^{n} CDD_i \end{aligned} \qquad (3)$$

Heating and cooling degree days simply represent the number of degrees that the temperature deviates from a reference temperature. This is useful for energy companies because if the daily mean temperature is less than a reference temperature, than energy is more likely to be used for heating, and vice versa.

Several studies have been conducted on the pricing methodology behind weather derivative contracts. Alaton et al. developed a pricing model for weather derivatives with payouts depending on temperature for Stockholm, Sweden. Another study conducted by Mraoua et al. developed a pricing model for Casablanca, Morocco. Only one of these two studies researched the possible time dependencies in the residuals. Most importantly, the previous research did not include a statistical test for normality, and the authors admitted that the empirical frequency of small temperature differences were higher than predicted by the fitted normal distribution

Referencing previous research, this paper will focus on an improved daily mean temperature model for each of the eighteen cities selected using historical data from 1997 to 2006 from the Chicago Mercantile Exchange.

2. DATA AND METHODS

2.1 Data

Ten years of surface temperature data were collected (EarthSat, 2009). The surface data were composed of daily observations that were taken from locations for which CME contracts existed. These daily observations contained the maximum, minimum, and mean temperature for each day, along with the corresponding heating and cooling degree days.

2.2 Methodology

Using the Interactive Data Language (IDL), the daily mean temperatures were parsed. These daily mean temperatures were then plotted. A linear fit was first applied to the raw daily mean temperature time series and then removed, leaving a time series with no general linear trend. Figure 1 is a plot of the daily mean temperature time series from Logan International Airport in Boston, MA with the corresponding linear trend.

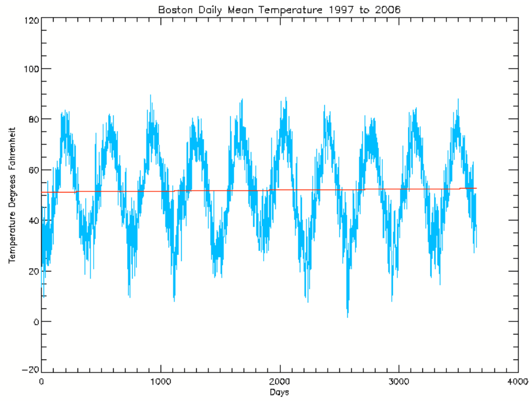

Figure 1: Boston Daily Mean Temperature 1997 to 2006 with linear trend

With the trend removed from the original time series of daily mean temperature, a discrete Fourier transformation was performed, separating the sequence into a summation of sin and cosine waves. After a spectrum was calculated with the frequencies, a power spectrum was plotted. With a frequency of one cycle per year, the largest signal was determined, which represented the seasonal signal. A mask was then applied to the power spectrum. The mask was a threshold used to set all the lower signals to zero, so only the seasonal signal would remain. Figure 2 is a plot of the power spectrum of the daily mean temperature for Logan International in Boston, MA with the determined mask threshold.

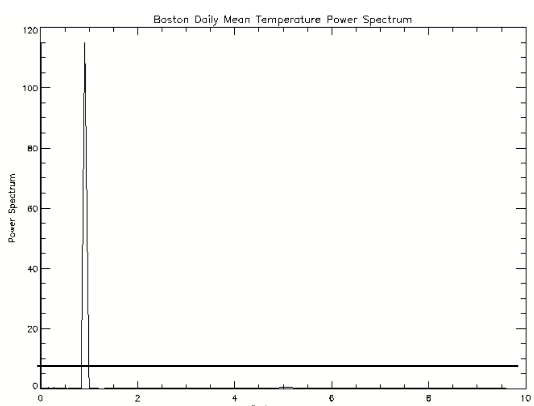

Figure 2: Boston Daily Mean Temperature Power Spectrum

Figure 3 is a plot of the product of an inverse Fourier Transformation from the masked time series. This signal with no noise represents the seasonal mean temperature signal.

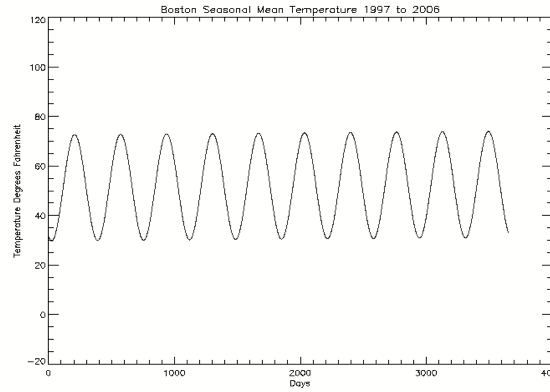

Figure 3: Boston Seasonal Mean Temperature

The intermediate residuals were then calculated by differencing the observed daily mean temperature from the corresponding seasonal mean time series. A pseudo variance was calculated by squaring the intermediate residuals. A discrete Fourier transformation was then performed on the pseudo variance, breaking the signal into a summation of sin and cosine waves. After a spectrum was calculated with these frequencies, a power spectrum was plotted. Figure 4 is a plot of the power spectrum. With a frequency of one cycle per year, the largest signal was determined, and then a mask was applied.

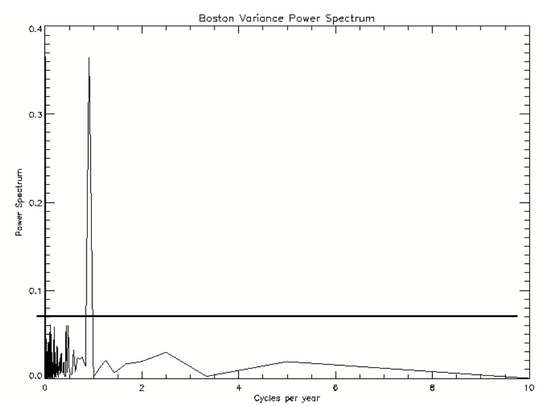

Figure 4: Boston Daily Mean Temperature Variance Power Spectrum

Figure 5 is the product, which represents an estimate of the daily mean temperature standard deviation. The raw residuals were then calculated by dividing the intermediate residuals by the estimated standard deviation of the daily mean temperature. Figure 6 is a plot of the residuals, which represents the difference between the observed index and the trend. This process was repeated for all eighteen cities.

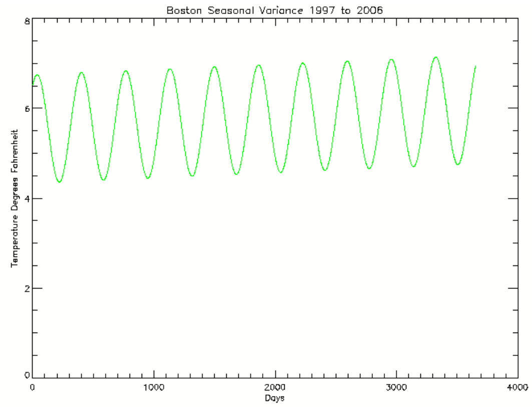

Figure 5: Boston Seasonal Variance

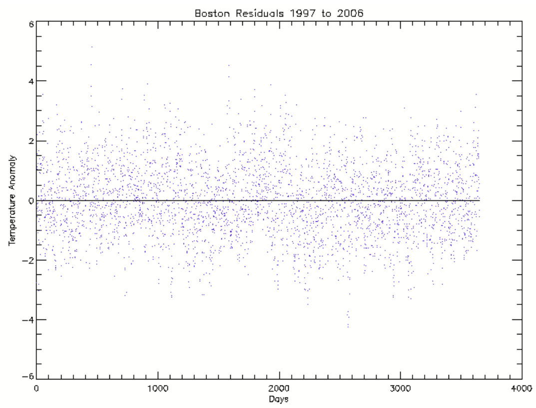

Figure 6: Boston Daily Temperature Residuals

2.3 Distribution Tests

In order to create the most realistic daily mean temperature model for each city, the distribution of the raw residuals had to be determined. Previous research relied on a visual representation of the data, but for this research a more in depth was in order. The statistical moments were calculated for each city using IDL. Table 1 is an example of the statistical moments for Boston. The residuals were then plotted in normalized histograms with an overlaid normal curve Figure 7, and against a normal distribution in a QQ plot Figure 8.

Table 1: Statistical Moments

City	Mean	Standard Deviation	Skewness	Kurtosis
Boston	0.00284524	1.57218	0.135301	0.0161614

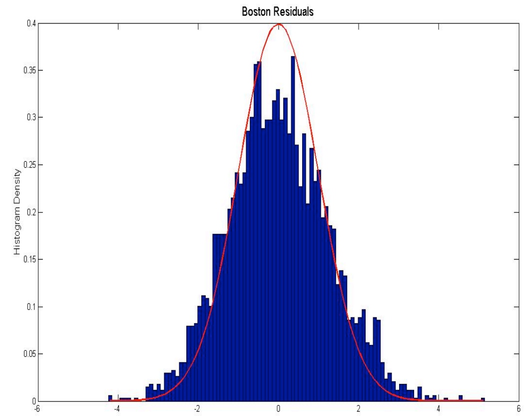

Figure 7: Boston Residual Histogram with Normal Curve

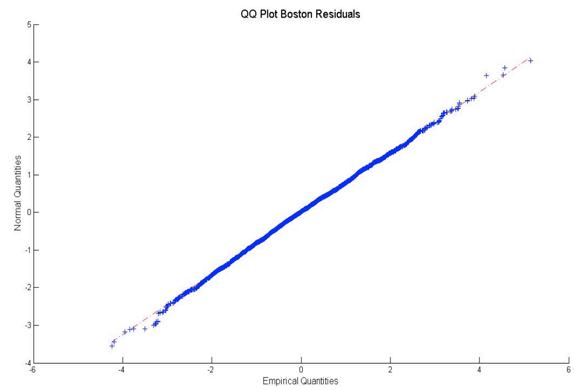

Figure 8: Boston Residual QQ Plot

After producing the statistical moments and the two other plots, the residuals were put through three tests for normality. The Lilliefor, Jarque-Bera, and Anderson-Darling tests were calculated at 95% confidence. The results for Boston, in which the null was a normal distribution and the alternative was a non-normal distribution, are given in Table 2. These steps were repeated for all eighteen cities.

Table 2: Tests of Normality

Test	Lilliefor	Jarque-Bera	Anderson-Darling
Boston	Reject the null	Reject the null	Reject the null

2.3 Models

After performing a thorough statistical analysis of the raw residuals, two models were created. Both models were derived from a basic autoregressive model with a lag of five days. The coefficients for each model were determined from an autocorrelation function. The autocorrelation function was calculated for each day of the year. The coefficients for each day of the ten year data set were then summed, and averaged. After determining the coefficients, the observed raw residuals were broken into the four climatologic seasons (i.e. DJF, MAM, JJA, and SON).

$$\varepsilon_t = \sum \alpha_{1\beta_t} + \alpha_{2\beta_{t-1}} + \alpha_{3\beta_{t-2}} + \alpha_{4\beta_{t-3}} + \alpha_{5\beta_{t-4}} \qquad (4)$$

Equation 4 was the equation used to produce each model. ε represents the residual being simulated at time step *t*, and α represents the autocorrelation coefficient. β represents the previous randomly selected residuals at each time step, respectively.

Lastly, in order to simulate one year's worth of daily mean temperature, the last four observed daily mean temperature values in December 2006 were used to initiate the simulation. As shown in Figure 9, observed residuals were selected randomly from the corresponding climatologic season with some memory of the previous residuals.

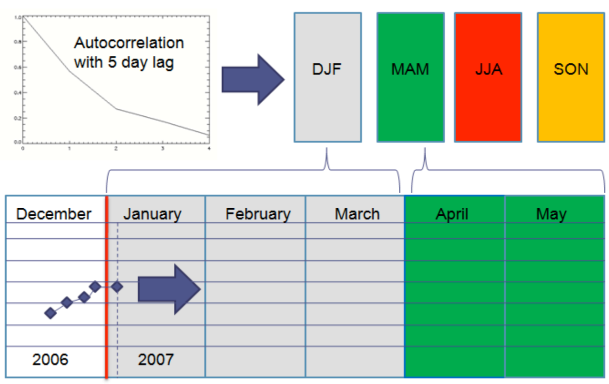

Figure 9: AR(5) Observed Residuals

As shown in Figure 10, the second model had many of the same characteristics as the first, except instead of selecting randomly from the corresponding observed climatologic season, the residuals were selected from a normal distribution.

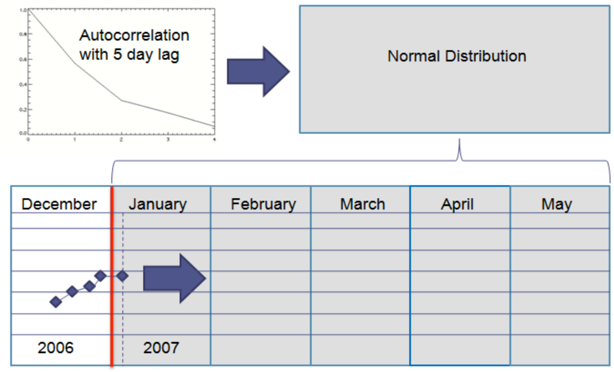

Figure 10: AR(5) Normal Distribution

Each model was then iterated ten thousand times, to build a large distribution of possible outcomes for 2007. Using a standard reference temperature of 18°C or 65°F degrees Fahrenheit, the heating and cooling degree days were calculated for each month.

Using IDL, the degree day data was plotted for each month in a histogram to evaluate the distribution of total degree days. Lastly, in order to properly compare the simulated data degree days against the historical data, a convolution was applied with a Gaussian curve to the histogram of the historic data. This process was repeated for all eighteen cities.

3. RESULTS

After reviewing the normalized histograms and QQ plots it was very difficult to visually determine whether a city was from a normal distribution. The majority of the normalized histograms had a great deal of action on the tails, and didn't appear to be from a normal distribution. The QQ plots had similar attributes, but additional information was still necessary. The results from the three tests of normality were most conclusive. Chicago was the only city, out of a group of eighteen with residuals from a normal distribution.

In order to properly evaluate each model, a city with residuals from a non-normal distribution was selected. Boston was selected as the non-normal city because it, in fact, possesses residuals that do not belong to a normal population according to the tests applied in section 2. Below, Figures 11-15 are samples of each model's simulation for Boston and Chicago.

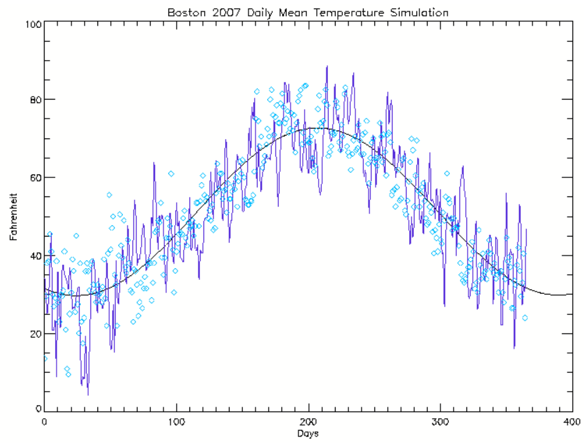

The black sinusoidal line represents the seasonal mean temperature for each city. The light blue diamonds represent the observed daily mean temperatures from the previous year. The purple line represents the simulated daily mean temperature data from each model.

Figure 11: AR(5) Boston Residuals

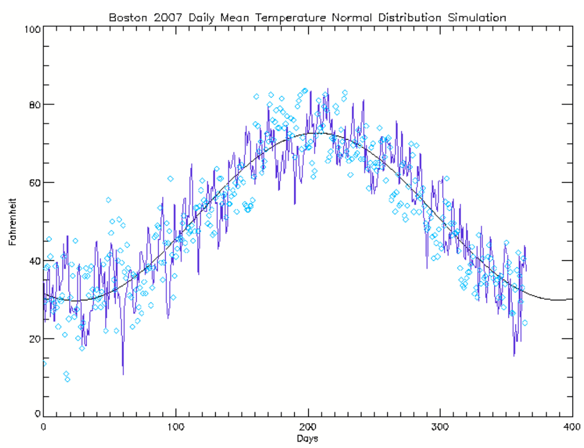

The black sinusoidal line represents the seasonal mean temperature for each city. The light blue diamonds represent the observed daily mean temperatures from the previous year. The purple line represents the simulated daily mean temperature data from each model.

Figure 12: AR(5) Boston Normal Distribution

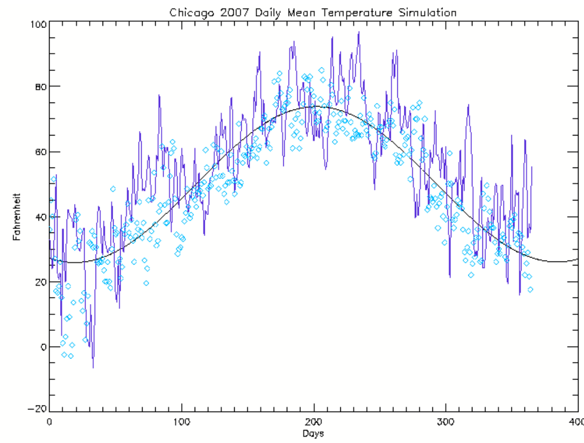

The black sinusoidal line represents the seasonal mean temperature for each city. The light blue diamonds represent the observed daily mean temperatures from the previous year. The purple line represents the simulated daily mean temperature data from each model.

Figure 13: AR(5) Chicago Residuals

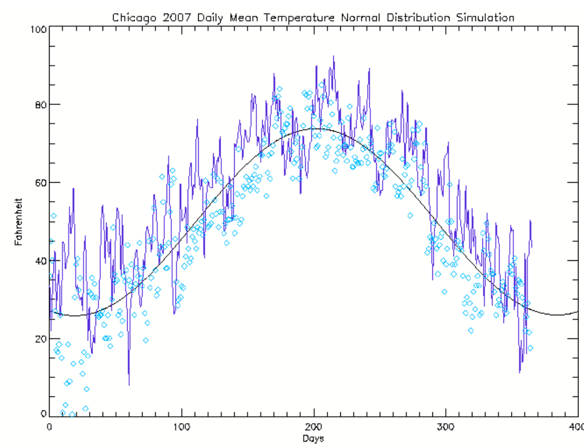

The black sinusoidal line represents the seasonal mean temperature for each city. The light blue diamonds represent the observed daily mean temperatures from the previous year. The purple line represents the simulated daily mean temperature data from each model.

Figure 14: Chicago AR (5) Normal Distribution

Intuitively, each city would have more realistic outcomes with the corresponding model that selected from their distribution. Although both models were run for each city, the remainder of this research assumed that climatologic data of Boston would be compared with the model that used the observed residuals, and Chicago would use the normal distribution.

Using the climatologic data, the heating degree days for January were calculated for each city. Figures 15 and 16 are histograms of the climatologic data of the heating degree days for Boston and Chicago.

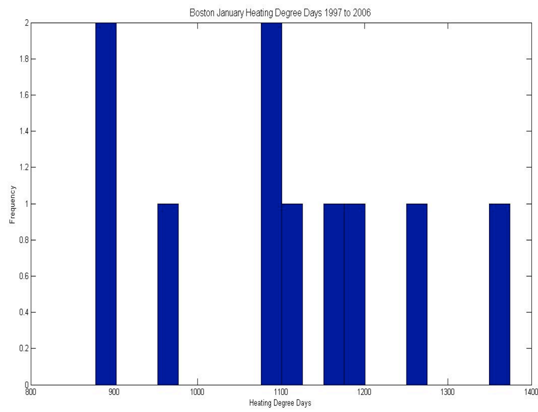

Figure 15: Boston January HDD 1997 to 2006

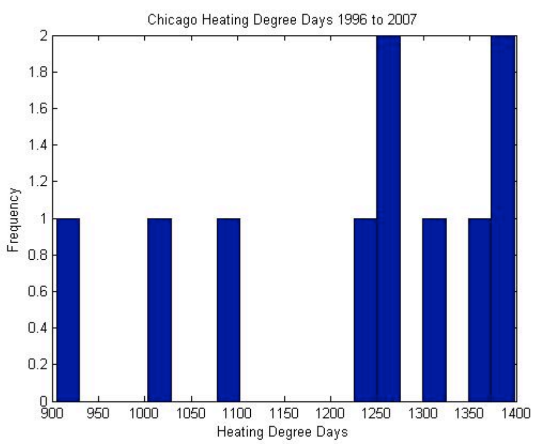

Figure 16: Chicago Heating Degree Days 1997 to 2006

As shown in Figures 15 and 16, the ten year data set had insufficient number of years to produce a histogram that resembled a normal curve. A convolution with a Gaussian curve was applied to the historic data histogram for Boston and Chicago.

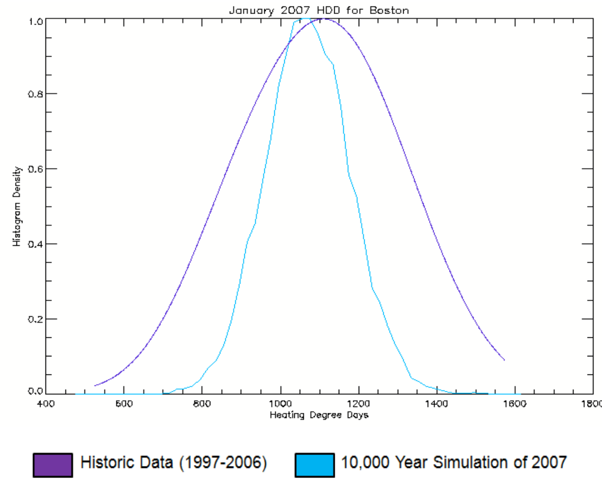

Figure 17: Boston Kernel Average and Simulated Heating Degree Days

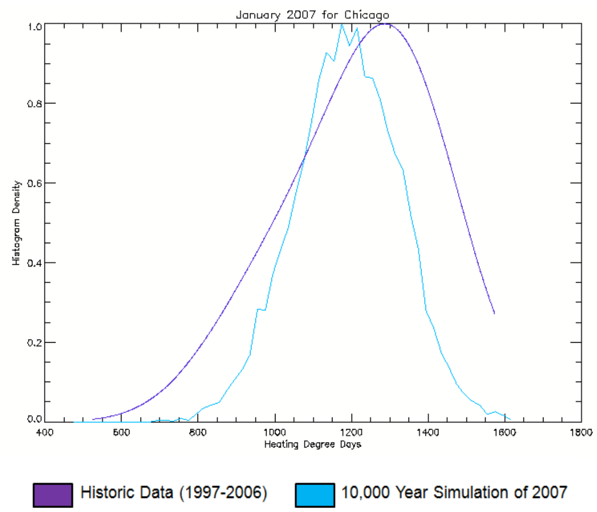

Figure 18: Chicago Kernel Average and Simulated Heating Degree Days

Due to the fact that the Gaussian curve had a large half width, the kernel estimation produced a much wider distribution than the simulated data. The mean and mode of each plot provided valuable information. In each case, the simulated data had a mode that was less than the historic data. The plots were constructed in order to assess where the historic data would place the maximum likelihood of HDD verses where the maximum likelihood occurs for the simulation.

.

4. CONCLUSION

Since weather derivatives are path dependent, taking the straight average of the climatologic heating and cooling degree days will not yield representative outcomes for the next year. This research attempts to simulate the best representation of next year's outcomes.

After removing the seasonality from the residuals and performing a thorough statistical analysis of the residuals, it was determined that the residuals cannot be assumed to be from a normal distribution for all cities. The normalized histograms and QQ Plots provided visual evidence that there was a great deal of action on the wings of the distributions.

The previous figures showed that the January 2007 heating degree day simulations, positioned the simulated distribution within the historic kernel density curve, but depicted a warmer winter than the historic data. This is very valuable because in weather derivative contracts, for every degree day off, the payout can increase very rapidly, and having this information in advance can prove very beneficial.

Future work, will involve additional metrics and techniques that will address the validity of the assumptions and statics used in these models. Additional, more robust models and pricing methods will be explored.

5. ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. ATM-0648566. Additional support was provided by AER internal research and development funds.

At AER, the author would like to thank mentors Thomas Connor, Mark Leidner, and Dr. Hillary "Ned" Snell, and William Gallery for technical and program assistance. At the Oklahoma University the author would like to especially thank Daphne LaDue for making the National Weather Center REU possible.

6. REFERENCES

Alaton, Peter, Djehiche, Boualem, and Stillberger, David, 2002: On Modelling and Pricing Weather Derivatives. Applied Mathematical Finance, 9, 1-20.

Bari, Driss, and Mraoua, Mohammed, 2005: Temperature stochastic modeling and weather derivatives pricing: empirical study with Moroccan data. Scandinavian Journal of Statistics. 4, 746 – 767.

Climatological Data (http://www.cmegroup.com/CmeWeb/ftp.wrap/pub/weather/usa/temperature/historical/daily)

Climetrix (http://www.climetrix.com/weathermarket/marketoverview/default.asp)

Earthsat, (http://www.earthsat.com/)

Jewson, Stephen, and Anders Brix, 2005: Weather Derivative Valuation The Meteorological, Statistical, Financial and Mathematical Foundations. Cambridge UP, 121-147.

Plymouth State University - Meteorology (vortex.plymouth.edu)

Weather Risk Management Association (http://www.wrma.org)