

# A Quantile Regression Study of Climate Change in Chicago, 1960-2010

Julien Leider

Department of Mathematics, Statistics, and Computer Science,  
University of Illinois at Chicago

June 16, 2012

## Abstract

This study uses quantile regression combined with time series methods to analyze change in temperatures in Chicago during the period 1960-2010. It builds on previous work in applying quantile regression methods to climate data by Timofeev and Sterin (2010) and work by the Chicago Climate Task Force on analyzing climate change in Chicago. Data from the Chicago O'Hare Airport weather station archived by the National Climatic Data Center are used to look at changes in weekly average temperatures. The method described by Xiao et al. (2003) is used to remove autocorrelation in the data, together with the rank-score method with IID assumption to calculate confidence intervals, and nonparametric local linear quantile regression to estimate temperature trends. The results of this analysis indicate that the decade 1960-1969 was significantly colder than later decades around the middle of the yearly seasonal cycle at both the median and 95th percentile of the temperature distribution. This analysis does not find a statistically significant trend over the later decades, 1970-2010.

## 1 Introduction

Climate change, including change in key climate characteristics such as temperature and frequency of extreme weather events, has become an important area of research, particularly as it is hypothesized that this change may be induced by human activities. Evidence has accumulated that climate is changing at the continental and global level with shifts such as increased temperature and greater weather volatility, and that these changes can indeed be traced to human activities since the start of the Industrial Revolution [14, Chapter 9, Executive Summary]. The effects of climate change have not been verified for smaller regions due to greater natural variability over smaller

---

The author, Julien Leider, may be reached by email at [jleide2@uic.edu](mailto:jleide2@uic.edu). This research was supervised by Dr. Jing Wang and Dr. Dibyen Majumdar from the Department of Mathematics, Statistics, and Computer Science at the University of Illinois at Chicago. Dr. Wang may be reached by email at [wangjing@math.uic.edu](mailto:wangjing@math.uic.edu).

sample sizes [14]. However, research has been conducted on possible effects of climate change at the local level; for instance, Timofeev and Sterin research local effects in Russia in [15].

The topic of this paper is the analysis of climate change in Chicago over the years 1960-2010 inclusive. The Chicago Climate Task Force has previously conducted research on this and published its results as the Chicago Climate Action Plan [2]. They downscaled global climate models to consider the possible future effects of global climate change at the local level. They also analyzed historical data both to calibrate this downscaling and to observe changes that have already occurred. For that purpose, they relied on monthly, daily, and hourly climate records from 14 National Weather Service stations providing at least 40 years of continuous coverage up to 1990 [2, Chapter 2, p. 7]. In particular, in order to analyze changes within the city of Chicago, they used observations from the Chicago Midway Airport, the Chicago O'Hare Airport, and the University of Chicago. Their analysis of historical data includes inspection of monthly mean temperatures as well as certain observed extreme weather events.

This paper takes the analysis of Chicago climate data further by applying quantile regression techniques to local temperature data using the statistical package R [12] and the R `quantreg` [6] package. The use of quantile regression allows us to systematically evaluate patterns in extreme temperatures over time. It also allows us to estimate central tendency more robustly using the median. This paper combines quantile regression with time series methods in order to appropriately handle the autocorrelation structure of the data. Data archived by the National Climatic Data Center [8] from the Chicago O'Hare Airport, one of the sites used by the Chicago Climate Task Force, are used for this purpose.

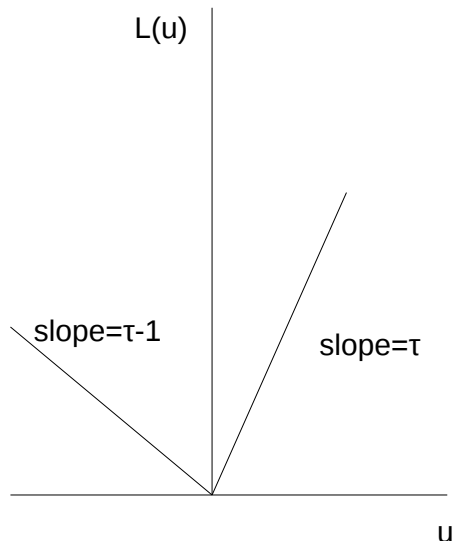
Quantile regression methods were successfully applied in order to carry out a cross-decade comparison of seasonal temperature patterns. This analysis indicates that the decade 1960-1969 was significantly colder than the later decades 1970-2010 both at very high temperatures (the 95th percentile) and at average temperatures (the median), but does not show a statistically significant trend over the period 1970-2010.

## 2 Quantile Regression

Quantile regression was originally proposed in 1978 by Roger Koenker and Gilbert Bassett as a robust alternative to least-squares regression [7] and was described at length by Koenker in [5]. However, the basic idea of quantile regression dates back to work by Ruder Boscovich in the 18th century which was continued by Pierre-Simon Laplace and Francis Edgeworth in the 19th century [5, p. 2-5]. Due to the computational difficulty of quantile regression, it was not until the advent of linear programming together with increased computing power that the method was rediscovered and put to use in the 20th century. Quantile regression offers a number of advantages over least-squares methods. While ordinary least squares regression typically assumes that the error terms are IID, normally distributed, and homoscedastic, quantile regression does not require these restrictive assumptions. Furthermore, since quantile regression estimates quantiles of the conditional distribution rather than the mean, it is more resistant to outliers than least-squares methods.

In contrast to the least-squares loss function  $L(u) = u^2$ , quantile regression makes

Figure 1: The quantile regression loss function.



use of the asymmetric loss function (shown in Figure 1)

$$L(u) = |u| \cdot (\tau \cdot I(u \geq 0) + (1 - \tau) \cdot I(u < 0)) = u \cdot (\tau - I(u < 0))$$

where  $\tau$  corresponds to the quantile to be estimated [5]. Note that if  $\tau = 0.5$ , i.e., the median is being estimated, then this loss function becomes simply

$$L(u) = |u|$$

and the sum of the absolute values of the residuals is minimized to perform regression.

To see that this loss function leads to the quantiles, note that we seek to minimize

$$\begin{aligned} E[L(X - \hat{x})] &= \int_{-\infty}^{\infty} L(x - \hat{x}) dF(x) \\ &= (\tau - 1) \int_{-\infty}^{\hat{x}} (x - \hat{x}) dF(x) + \tau \int_{\hat{x}}^{\infty} (x - \hat{x}) dF(x) \end{aligned}$$

where  $F(X)$  is the cumulative distribution function (CDF) of the random variable  $X$ . The derivative of this expectation with respect to  $\hat{x}$ , when  $f(x)$  is the probability density function (PDF) of  $X$ , is

$$\frac{d}{d\hat{x}} [(\tau - 1) \int_{-\infty}^{\hat{x}} (x - \hat{x}) dF(x) + \tau \int_{\hat{x}}^{\infty} (x - \hat{x}) dF(x)]$$

$$\begin{aligned}
&= \frac{d}{d\hat{x}} [(\tau - 1) \left( \int_{-\infty}^{\hat{x}} x dF(x) - \hat{x} \int_{-\infty}^{\hat{x}} dF(x) \right) - \tau \left( \int_{\infty}^{\hat{x}} x dF(x) - \hat{x} \int_{\infty}^{\hat{x}} dF(x) \right)] \\
&= (\tau - 1) (\hat{x} f(\hat{x}) - \hat{x} f(\hat{x}) - 1 \cdot \int_{-\infty}^{\hat{x}} dF(x)) - \tau (\hat{x} f(\hat{x}) - \hat{x} f(\hat{x}) - 1 \cdot \int_{\infty}^{\hat{x}} dF(x)) \\
&= (\tau - 1) (-F(\hat{x})) - \tau (1 - F(\hat{x})) \\
&= F(\hat{x}) - \tau
\end{aligned}$$

To find the minimum, this derivative is set equal to zero; then  $F(\hat{x}) = \tau$ . Note that the second derivative of this expectation is just  $f(\hat{x})$ , which is a nonnegative function, so we have indeed found a minimum. Thus, minimizing the value of this loss function applied to the residuals does lead to estimating quantiles of the response variable.

In practice, the CDF  $F(x)$  is unknown. Instead, the empirical CDF  $F_n(x) = \sum_{i=1}^n I(x_i \leq x)$ , based on sample observations, is used. To calculate the sample quantiles, we minimize the same expectation; in this case,  $\int_{-\infty}^{\infty} L(x - \hat{x}) dF_n(x) = \frac{1}{n} \sum_{i=1}^n L(x_i - \hat{x})$ . Since  $\frac{1}{n}$  is a constant for any given sample of size  $n$ , minimizing that expression is equivalent to minimizing  $\sum_{i=1}^n L(x_i - \hat{x})$ .

To fit a model  $y_i = \beta^T x_i + \epsilon_i$ , we estimate  $\beta$  using

$$\hat{\beta} = \operatorname{argmin}_{\beta \in \mathbb{R}^d} \sum_i L(y_i - \beta^T x_i)$$

where  $d$  is the number of parameters in our model, so that  $\beta$  and  $x_i$  are vectors of length  $d$ . This computation cannot be carried out analytically, in contrast to the computation of least squares regression. Instead, this can be reformulated as a problem in linear programming [5]. Let  $u_i$  and  $v_i$  be slack variables corresponding to the positive and negative parts of the residuals  $y_i - \beta x_i$ , respectively; thus, there are  $n$  of each. Note that  $u_i$  is zero when  $v_i$  is non-zero, and vice versa. Also, note that when  $a$  is positive,  $L(a) = a \cdot \tau$ , and when  $a$  is negative,  $L(a) = a \cdot (\tau - 1)$ . Thus, by separating the residuals into their positive and negative parts, we are able to obtain a linear objective function, which allows us to use techniques of linear programming to find a solution. The solution is given by

$$\begin{aligned}
\hat{\beta} &= \operatorname{argmin}_{(\beta, u, v) \in \mathbb{R}^d \times \mathbb{R}_+^{2n}} \{L(1_n^T u + 1_n^T v) | X\beta + u - v = y\} \\
&= \operatorname{argmin}_{(\beta, u, v) \in \mathbb{R}^d \times \mathbb{R}_+^{2n}} \{\tau 1_n^T u + (\tau - 1) 1_n^T v | X\beta + u - v = y\}
\end{aligned}$$

where  $1_n$  corresponds to the  $n$ -element vector containing all ones and  $X$  is an  $n \times d$  matrix containing the  $n$  observations of the explanatory variables. This corresponds to a linear programming problem where we seek to minimize an objective function giving hyperplanes over a polyhedral constraint set in  $(2n + 1)$ -space.

For our purposes, quantile regression offers a number of advantages over other approaches. Most important, extreme weather is a particularly troubling potential aspect of climate change and an increase in temperature variability is of at least as great interest as a change in mean temperature over time. Thus, it would be undesirable to assume at the outset that the data are homoscedastic for purposes of least squares regression. Furthermore, quantile regression does not require assuming that the error terms are normal, which leads to more robust results. These factors, together with the ability to estimate quantiles rather than the mean, makes quantile regression significantly more robust than least squares regression and more suited to our purposes.

In addition, this analysis uses nonparametric methods exclusively. In particular, local linear regression with a normal kernel is used to estimate the temperature trend, relying on code in Koenker's `quantreg` package [6]. The description in Appendix A of [5] of available nonparametric quantile regression methods in the `quantreg` package was relied on for choosing the most appropriate method. Three nonparametric methods are discussed there: local polynomial regression, splines, and penalty methods. In conducting this analysis, both local linear regression and splines were tried on the data. There were significant boundary effects using splines, so local linear regression is used throughout instead. The use of nonparametric methods avoids difficulties with selecting an appropriate model, which is desirable as we did not have an expectation at the outset as to the most appropriate model for the data.

Since our data occur in the form of a time series, there is a potential autocorrelation structure which requires modifications to the traditional methods of quantile regression, which implicitly assume that the data are independent and identically distributed (IID). We used the autocorrelation function (ACF) and partial autocorrelation function (PACF) to assess the autocorrelation of the data. Let  $x$  be a time series variable of interest, where  $\bar{x}$  is the mean of  $x$  and  $n$  is the total number of observations of  $x$ . Let  $\rho(h)$  be the ACF, and let  $\gamma(h)$  be the autocovariance function. Then the sample ACF is given by

$$\hat{\rho}(h) = \frac{\hat{\gamma}(h)}{\hat{\gamma}(0)}, -n < h < n$$

where

$$\hat{\gamma}(h) = \frac{1}{n} \sum_{t=1}^{n-|h|} (x_{t+|h|} - \bar{x})(x_t - \bar{x}), -n < h < n$$

[1, p. 19]. Similarly, the sample PACF  $\hat{\alpha}(h)$  is given by

$$\hat{\alpha}(0) = 1$$

and

$$\hat{\alpha}(h) = \hat{\phi}_{hh}, h \geq 1$$

where  $\hat{\phi}_{hh}$  is the last component of

$$\hat{\phi}_h = \hat{\Gamma}_h^{-1} \hat{\gamma}_h$$

while

$$\Gamma_h = [\gamma(i-j)]_{i,j=1}^h$$

and

$$\gamma_h = [\gamma(1), \gamma(2), \dots, \gamma(h)]'$$

[1, p. 95].

We then used the method in [17] to transform the data, as described in more detail in section 4. As part of that process, an autoregressive moving average (ARMA) model is fit to the data. The ACF and PACF were used in determining the order of this model, since the maximum lag at which the ACF is non-zero is associated with the order of the moving average (MA) portion of the ARMA model, while the maximum lag at which the PACF is non-zero is associated with the order of the autoregressive (AR) portion of the model [13, p. 108]. Using transformed data that are more nearly IID allowed us to reliably calculate confidence intervals, as will also be described in section 4.

### 3 Data Used

The primary criteria for data selection were that data be available continuously for at least 40 years during the last half-century, that the data come from a weather station located in Chicago, Illinois, and that the data come from a reliable, National Weather Service certified source. The website of the National Climatic Data Center (NCDC) was used since the NCDC is the “world’s largest archive of weather data” and is the official archive for the National Weather Service and other branches of the United States government [9]. The NCDC’s data set of hourly global surface data (DS3505) was used to locate potential data sources [8]. Since a search of other sources, including the archives of the weather website Weather Underground [16] and a list of National Weather Service weather stations in Illinois [10], did not identify any appropriate sources in Chicago not listed in the NCDC’s data set, the NCDC data set was taken to be exhaustive. In this set, only the O’Hare and Midway weather stations offered data over a suitable time period; both included more than 50 years of data. However, the Midway data included gaps of over a year in some cases, while the O’Hare data included gaps of a few days at most. Thus, the data from O’Hare (WBAN number 94846) are used as the most complete and continuous available data set in Chicago. We consider the 51 years from 1960-2010 inclusive, although the O’Hare data go back further to 1946, since this still provides us with a sufficiently large time period to draw meaningful comparisons while keeping the amount of data manageable. For each day, the temperature observation at 1200 UTC was used so as to have a suitable basis for computing weekly or monthly averages: this way, it was possible to average over comparable data points that were not themselves the result of averaging. Python [11] code written by the author was used to process the data files.

### 4 Data Analysis

The main focus of this analysis was to apply quantile regression methods to the data. However, quantile regression methods, particularly those for calculating confidence intervals, assume independent and identically distributed (IID) data [4]. Since the data used for this analysis are a time series, it is necessary to take account of their autocorrelation structure. In particular, in order to apply quantile regression to the original time series, it would be important to first deseasonalize the data, since otherwise the data would obviously be non-IID. Thus, the initial approach of this analysis was to try to detrend and deseasonalize the data using time series methods, so as to then apply quantile regression to the resulting stationary residuals. The ACF and PACF were used to analyze the autocorrelation structure of the data and to assess the success of our attempts to deseasonalize and detrend them. If the data were IID, both of these functions would be expected to equal zero (within confidence bounds) at all non-zero lags [1].

Deseasonalizing the data emerged as the main obstacle. A number of methods were tried, including differencing over the period of the data [13, p. 60-62], periodic regression [13, p. 72-73], kernel smoothing combined with differencing, and a moving average filter combined with linear regression [1, p. 31-32]. None of these methods produced satisfactory results, as can be seen from the plots in Figure 2. The ACF for the results of regression combined with a moving average filter is less periodic than

the original ACF but does not die off at any lag. The ACF for the results of periodic regression exhibits a similar trend. The PACF plot for the results of differencing shows spikes at yearly intervals. The PACF plot for the results of differencing combined with kernel smoothing does not die off at any lag.

Using weekly and monthly averaged data looked like a potential solution, but was unhelpful. This method offered the possibility of removing some of the noise in the data, making it more manageable for plotting purposes, and eliminating any gaps in the data, since we never had more than a few consecutive days of missing values. When averaging, we averaged over those values that were present in any given week or month. This approach did not solve the problem; if anything, decreasing the granularity of the data made attempts at deseasonalizing them even less successful, as was revealed in the data plots and the plots of the ACF and PACF.

In order to overcome the difficulty of deseasonalizing the data, the data were instead divided into five decades and the temperature pattern over a single seasonal cycle was analyzed for each decade. Data points within the same decade at the same time of year were treated as replicates for that particular point in the seasonal cycle. Since there is then only one seasonal cycle to analyze per decade, this eliminates the need to deseasonalize the data, as the seasonal pattern no longer contributes to the autocorrelation structure of the data. This did not completely eliminate all autocorrelation in the data, but it made it manageable as long as weekly averaged data were used and transformations were performed where appropriate to make the data IID. Since there were data for 51 years, the last “decade” contains data for 11 rather than 10 years. This approach is similar to that used in [15], where the authors look at temperature during each of the four seasons separately over a period of roughly 50 years.

Upon reaching this point, the next step in the analysis was to try applying the transformation described in [17] in order to obtain completely IID data. The first step in this transformation is to obtain a fit to the data using local polynomial regression; we used local linear quantile regression at the median for this purpose. This required selecting an appropriate bandwidth. The authors of [17] use the bandwidth  $1.06s_X T^{-1/5}$  with a Gaussian kernel when investigating their method numerically, where  $s_X$  is the standard deviation of the independent variable (in this case, time), and  $T$  is the total number of observations. This analysis uses the code in the R `quantreg` package [6] to perform local linear quantile regression; since this code uses a Gaussian kernel, that was a good candidate bandwidth. We compared the results of using that bandwidth to those obtained using three simple rule of thumb bandwidths:  $\frac{X_{max}-X_{min}}{20}$ ,  $\frac{X_{max}-X_{min}}{10}$ , and  $\frac{X_{max}-X_{min}}{5}$ , where, in each case,  $X_{max} - X_{min}$  equals the range of the independent variable, which in this case is time. As shown in Figure 3 for decade 1, the bandwidth in [17] and the bandwidth  $\frac{X_{max}-X_{min}}{10}$  produce nearly identical results (similar results are obtained for the other decades). On the other hand, the two other bandwidths do not work well with the method for removing autocorrelation in [17] since, as it turns out, they lead to residuals that are too autocorrelated, so that the transformation is ineffectual. Thus, the bandwidth in [17] is used both for removing autocorrelation and for all subsequent analysis.


The second step in the transformation in [17] is to fit an AR or ARMA model to the residuals of the fit found in the first step. For an  $AR(k)$  model, this gives

$$u_t = a_1 u_{t-1} + \dots + a_k u_{t-k} + \epsilon_t$$


where  $\epsilon_i$  is the  $i$ th residual of this  $AR$  model and  $u_i$  is the  $i$ th residual from the preceding

Figure 2: ACF/PACF plots with 90% confidence intervals.


**Detrending with regression**



**Detrending with differencing**



**Detrending with differencing and kernel smoothing**



**Detrending with periodic regression**


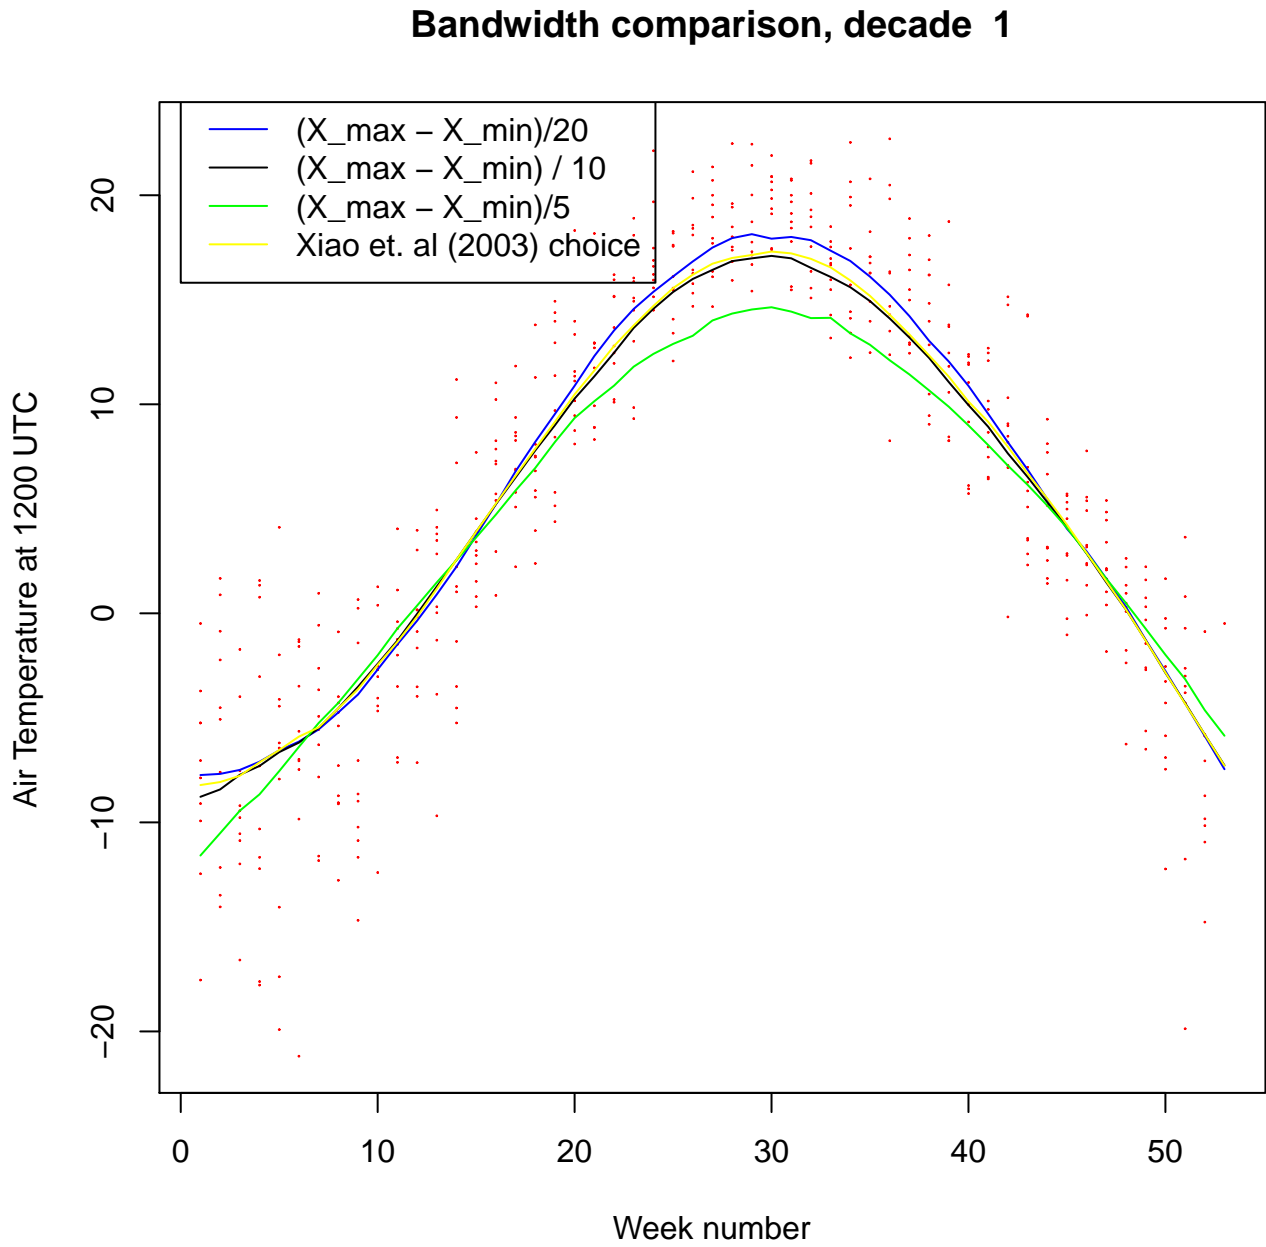




Figure 3: Bandwidth comparison.



nonparametric fit. None of the sets of residuals from the preceding fit appeared IID, so an ARMA model was fit for each decade. To compare model fits, Akaike’s Information Criterion (AIC) was used to select models leading to a minimum AIC, as suggested in [13, p. 52]. This involved trying to fit ARMA( $p, q$ ) models with  $p$  and  $q$  between 0 and the largest non-zero (outside of confidence bounds) lag of the ACF (for the MA portion) or PACF (for the AR portion). For each decade, this resulted in choosing an AR model ( $q = 0$ ); the orders selected for decades 1-5 were 3, 7, 3, 9, and 1, respectively. Then we applied the transformation suggested in [17], which is given by

$$\hat{Y}_t = Y_t - \sum_{j=1}^{\tau} \hat{a}_j(Y_{t-j} - \hat{m}(X_{t-j}))$$

where  $Y_t$  corresponds to the original series,  $\hat{Y}_t$  is the transformed series, and  $\hat{m}(X_t)$  is the value at  $X_t$  of the fit obtained using local linear quantile regression. This transformation improved the autocorrelation structure for the first decade, bringing it closer to the IID assumption; the original and transformed data for this decade are shown in Figure 4. It made the other decades worse, so the original version of those data was used for the analysis. However, those data are close enough to IID, based on their ACF and PACF, that it is safe to construct confidence intervals for the results.

It was also necessary to select an appropriate method for constructing confidence intervals. After experimenting with each of the methods offered in the R `quantreg` package [6], including direct estimation of the variance-covariance matrix under either an IID or non-IID assumption, estimation of the variance-covariance matrix using kernel smoothing, the Markov Chain Marginal Bootstrap (MCMB) method, four other bootstrap methods, the rank-score method under an IID error assumption, and the rank-score method under broader assumptions, as well as the modified MCMB algorithm provided in the package `rqmcm2` [3] and described in [4], it turned out that only the two rank-score methods provided reasonable results. The other methods gave confidence bounds that were either unlimited at certain points or varied wildly. Since the authors of [4] recommend the rank-score method with IID error assumption and find that it performs at least as well as the other rank-score method but at lower computational cost, that is used throughout for constructing confidence intervals. We use a 90% confidence level throughout.

## 5 Results

The results of quantile regression for the median and two extreme quantiles, the 95th and 5th percentiles, are shown in Figures 5-7. The most extreme quantiles that still produced reasonable confidence interval results were used. For the 95th percentile, there is a warming pattern both in the early and middle portions of the year which at several points is significant (confidence interval bounds do not overlap) across multiple decades. There is a similar pattern for the median. The confidence interval results for the 5th percentile are significantly broader, particularly in summer, and there are not any clear cross-decade patterns. There is also a spike in the lower confidence band for the third decade which inspection of the data indicates is not related to outliers in the original data. It is likely that difficulties obtaining tight confidence bounds for the 5th percentile explain the difficulties finding clear cross-decade patterns there.

Figure 4: Original and transformed data for the first decade.

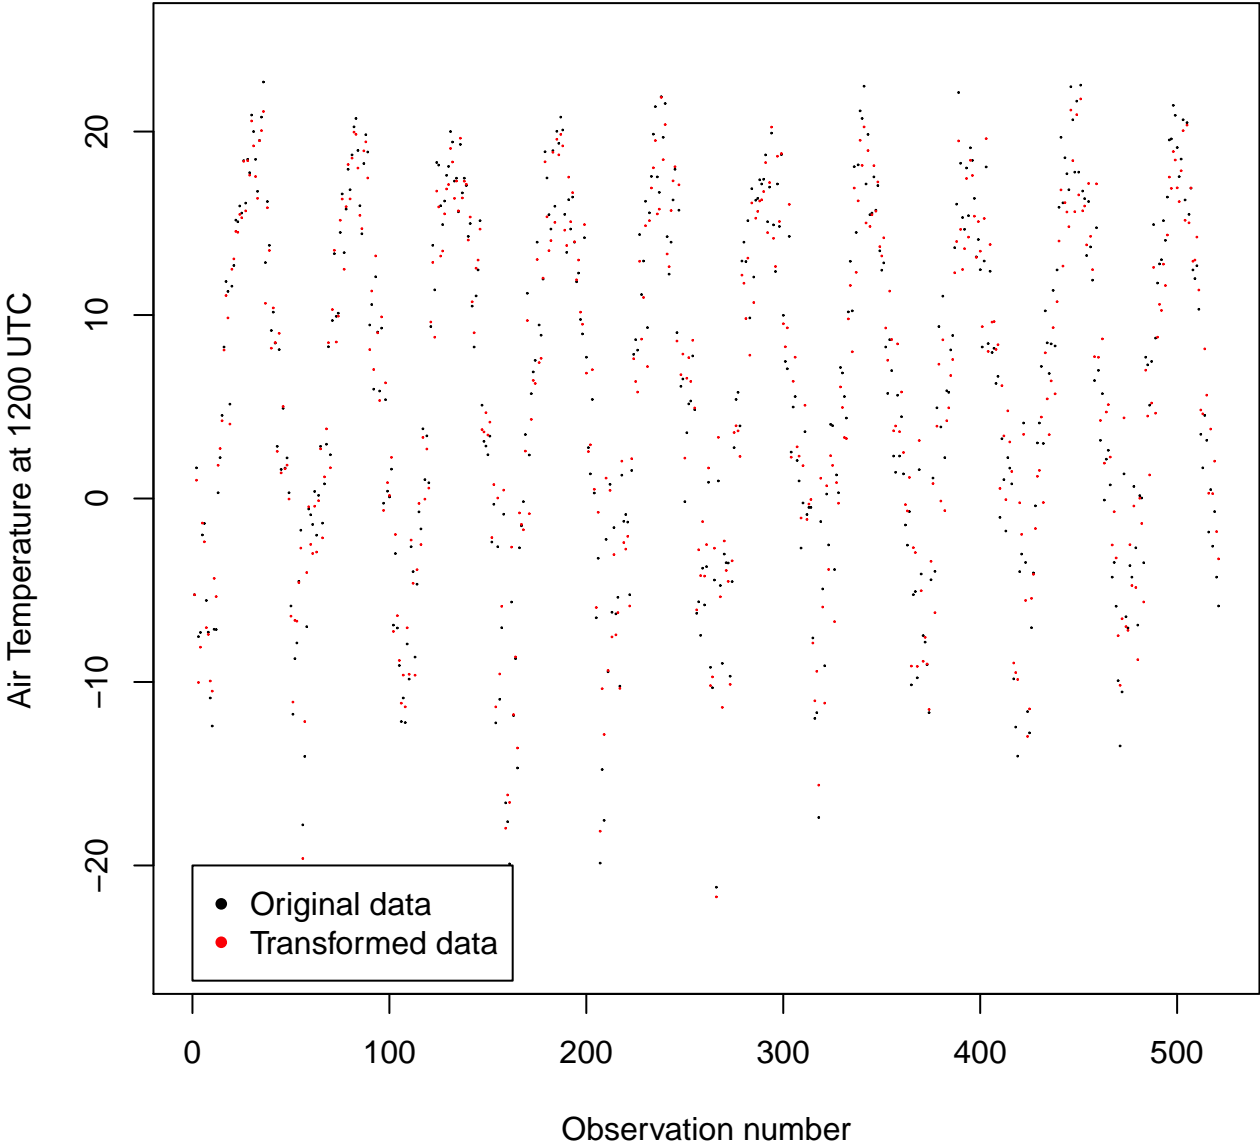


Figure 5: Decade-by-decade comparison, 95th percentile with 90% confidence intervals.

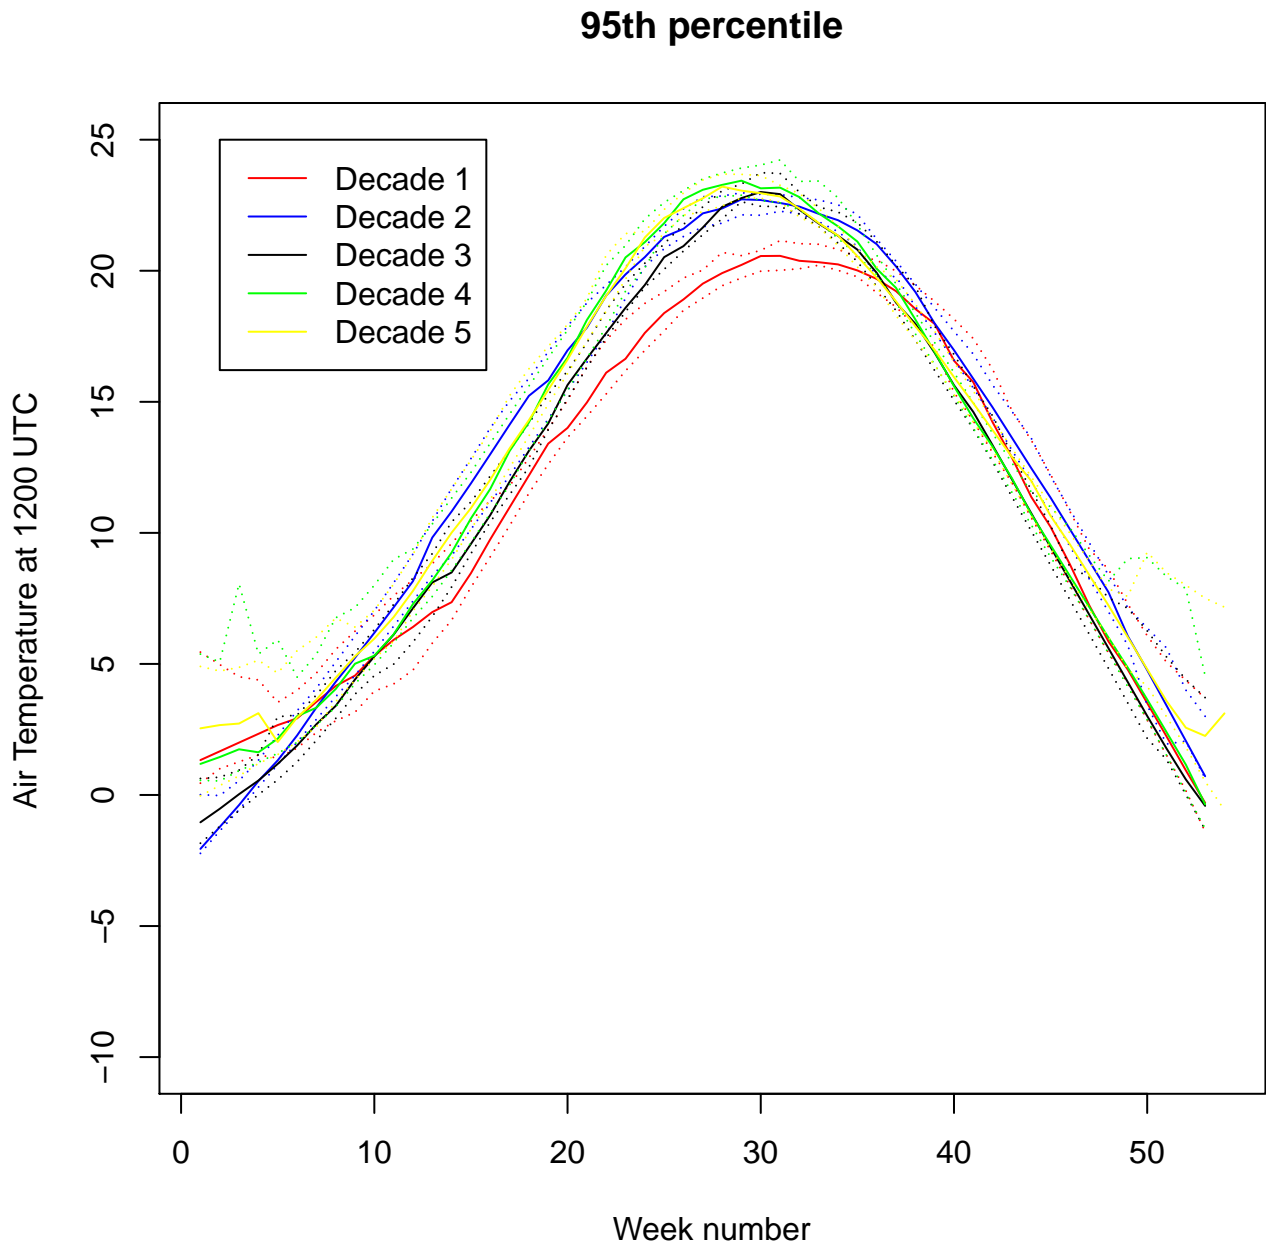


Figure 6: Decade-by-decade comparison, the median with 90% confidence intervals.




Figure 7: Decade-by-decade comparison, 5th percentile with 90% confidence intervals.

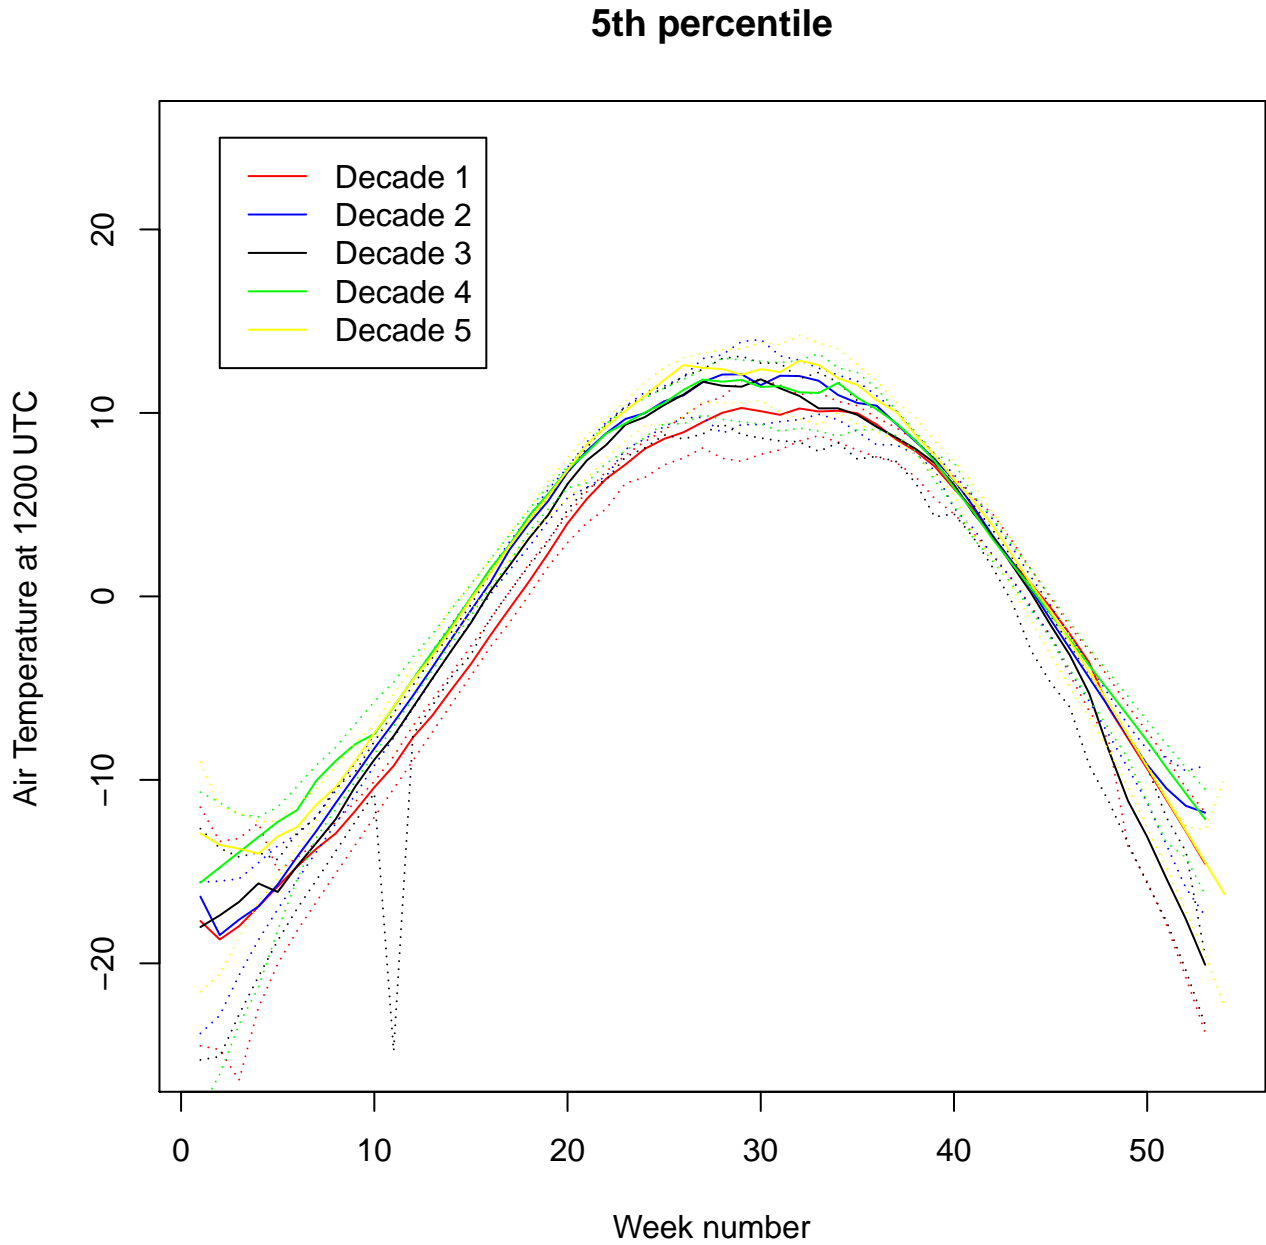


Figure 8: Pairwise decade comparison—the first and fifth decades with 90% confidence intervals.

### Decade 1 versus decade 5 at the 95th, 50th, and 5th percentiles

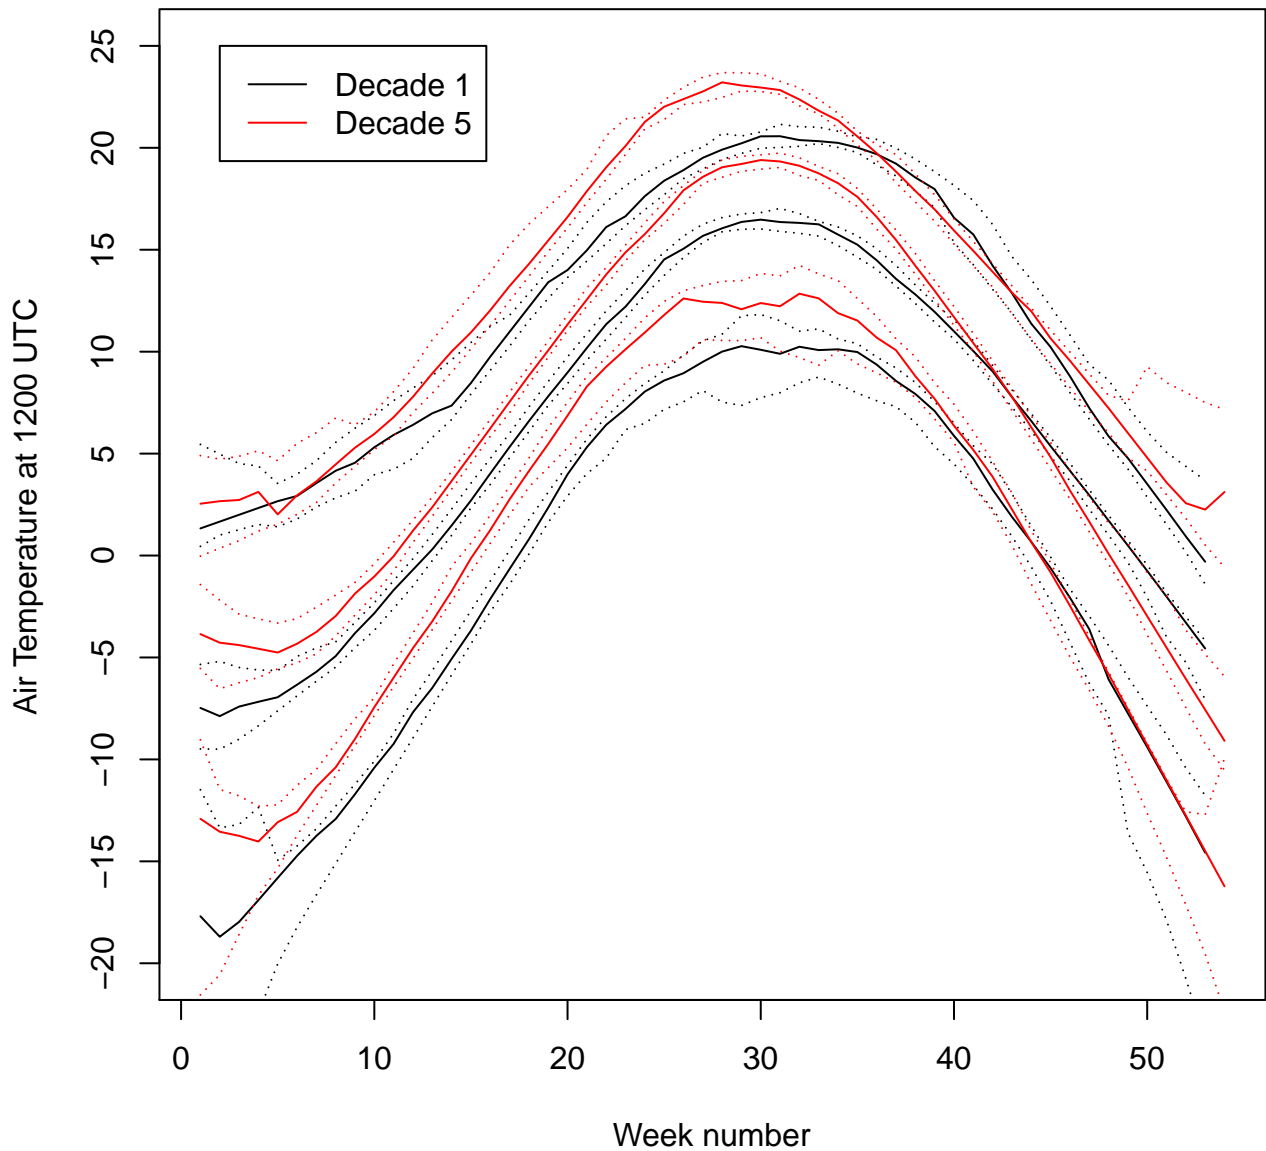



Figure 9: Pairwise decade comparison—the second and fifth decades with 90% confidence intervals.

### Decade 2 versus decade 5 at the 95th, 50th, and 5th percentiles





The results of two pairwise decade comparisons are shown in Figures 8 and 9. Comparing the first and fifth decades, there is significant warming at both the median and the 95th percentile. The comparison of the first and fourth decades, not shown, mirrors these results. On the other hand, comparing the second decade with the fourth (not shown) or fifth, there is no significant change. Thus, there appears to be some evidence the the first decade, including the years 1960 to 1969 inclusive, was unusually cold.

## 6 Conclusion

In this paper, a combination of time series and quantile regression methods was successfully used to analyze Chicago temperature data. The method in [17] proved useful for removing some of the autocorrelation in one decade, while the rank-score method with IID assumption provided a working method for computing confidence intervals. Using nonparametric local linear quantile regression with modified code from the R `quantreg` package, it was then possible to carry out a cross-decade comparison of seasonal temperatures. This analysis found evidence of warming from the decade 1960-1969 to later decades, but did not find a time trend over the rest of the study period.

## 7 Acknowledgements

The author would like to thank Dr. Jing Wang and Dr. Dibyen Majumdar at the Department of Mathematics, Statistics, and Computer Science, University of Illinois at Chicago for their supervision of this work. The author is also indebted to Rebecca Lowery, librarian at the Richard J. Daley Library at the University of Illinois at Chicago, for her valuable assistance in locating data sources. Finally, the author would like to thank the referee for valuable suggestions that have improved this paper.

## References

- [1] P. J. BROCKWELL AND R. A. DAVIS, *Introduction to Time Series and Forecasting*, Springer, New York, 2 ed., 2002.
- [2] CHICAGO CLIMATE TASK FORCE, *Chicago Climate Action Plan*. <http://www.chicagoclimateaction.org/>.
- [3] M. KOCHERGINSKY AND X. HE, *rqmcmb2: Markov Chain Marginal Bootstrap for Quantile Regression*, 2007. R package version 1.0.2-1.
- [4] M. KOCHERGINSKY, X. HE, AND Y. MU, *Practical confidence intervals for regression quantiles*, *J. Comput. Graph. Statist.*, 14 (2005), pp. 41–55.
- [5] R. KOENKER, *Quantile Regression*, Cambridge University Press, Cambridge, 2005.
- [6] R. KOENKER, *quantreg: Quantile Regression*, 2011. R package version 4.71.
- [7] R. KOENKER AND J. GILBERT BASSETT, *Regression quantiles*, *Econometrica*, 46 (1978), pp. 33–50.
- [8] NATIONAL CLIMATIC DATA CENTER, *Surface data, hourly global*. <http://gis.ncdc.noaa.gov/map/isd/>.
- [9] NATIONAL CLIMATIC DATA CENTER, *What is NCDC?*, September 2010. <http://www.ncdc.noaa.gov/oa/about/whatisncdc.html>.
- [10] NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION’S NATIONAL WEATHER SERVICE, *XML Feeds of Current Weather Conditions*, April 2007. [http://www.weather.gov/xml/current\\_obs/seek.php?state=il&Find=Find](http://www.weather.gov/xml/current_obs/seek.php?state=il&Find=Find).
- [11] PYTHON SOFTWARE FOUNDATION, *Python v2.7.2 documentation*, 2012. <http://docs.python.org/index.html>.
- [12] R DEVELOPMENT CORE TEAM, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2010. ISBN 3-900051-07-0.
- [13] R. H. SHUMWAY AND D. S. STOFFER, *Time Series Analysis and Its Applications: With R Examples*, Springer, New York, 3 ed., 2011.
- [14] S. SOLOMON, D. QIN, M. MANNING, Z. CHEN, M. MARQUIS, K. AVERYT, M. TIGNOR, AND H. MILLER, eds., *Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, 2007*, Cambridge University Press, Cambridge, 2007.
- [15] A. A. TIMOFEEV AND A. M. STERIN, *Using the quantile regression method to analyze changes in climate characteristics*, *Russian Meteorology and Hydrology*, 35 (2010), pp. 310–319.
- [16] WEATHER UNDERGROUND, *Weather forecast and reports*, 2012. <http://www.wunderground.com/>.
- [17] Z. XIAO, O. B. LINTON, R. J. CARROLL, AND E. MAMMEN, *More efficient local polynomial estimation in nonparametric regression with autocorrelated errors*, *J. Amer. Statist. Assoc.*, 98 (2003), pp. 980–992.