

JIRSS (Year)

Vol. 17, No. 02, pp 37-55

DOI: 10.29252/jirss.17.2.4

Stochastic Models for Pricing Weather Derivatives using Constant Risk Premium

Jeffrey S. Pai¹ and Nalini Ravishanker²

¹Warren Centre for Actuarial Studies and Research, University of Manitoba, Canada.

²Department of Statistics, University of Connecticut, USA.

Received: 11/09/2017, Revision received: 02/05/2018, Published online: 10/08/2018

Abstract. Pricing weather derivatives is becoming increasingly useful, especially in developing economies. We describe a statistical model based approach for pricing weather derivatives by modeling and forecasting daily average temperature data which exhibits long-range dependence. We pre-process the temperature data by filtering for seasonality and volatility and fit autoregressive fractionally integrated moving average (ARFIMA) models, employing the preconditioned conjugate gradient (PCG) algorithm for fast computation of the likelihood function. We illustrate our approach using daily temperature data from 1970 to 2008 for cities traded on the Chicago Mercantile Exchange (CME), which we employ for pricing degree days futures contracts. We compare the statistical approach with traditional burn analysis using a simple additive risk loading principle for pricing, where the risk premium is estimated by the method of least squares using data on observed prices and the corresponding estimate of prices from the best model we fit to the temperature data.

Keywords. ARFIMA model, Burn analysis, Daily temperature, Prediction.

MSC: 62P05 .

Jeffrey S. Pai (Jeffrey.Pai@umanitoba.ca)

Corresponding Author: Nalini Ravishanker (nalini.ravishanker@uconn.edu)

1 Introduction

Weather conditions have a significant impact on virtually every sector of the economy in almost every country globally. Estimates from the US Department of Commerce indicate that about one fifth of the US economy is directly affected by the weather and that roughly one third of the economy is weather sensitive, every state being impacted. While the economy of any country can be badly hurt by severe weather events, it can also be helped by less severe events, for example, a heat wave could increase sales of air conditioners and tickets for water parks. Weather related catastrophes include floods, storms (hurricane or typhoon), dry and cold events (droughts, forest fires, heat wave, heavy frost, snow storms, and ice storms), tornadoes, hailstorms, or snow/mud avalanches. Between 1998 and 2008, there were 293 catastrophic events, each event causing \$25 million or more of insured property losses and affecting a significant number of policyholders and insurers, resulting in a total insured property damage sum of \$199 billion (Insurance Services Office, see www.iso.com). Traditionally, insurance has served as the main tool for protection against unexpected weather conditions. Weather derivatives, developed from the financial services industry, are gaining popularity as important alternate weather risk management tools (Berlage, 2013).

Weather derivatives are financial contracts with payouts that depend on weather variables such as temperature, humidity, rain or snowfall. The first weather derivative was dealt between Consolidated Edison Co. and Aquila Energy in 1996. Over-the-counter trading of weather derivatives began in 1997, by energy traders Aquila, Enron, and Koch Industries. Exchange-traded weather futures contracts and corresponding options were first handled by the Chicago Mercantile Exchange (CME) in 1999. Currently, CME weather contracts are available in more than 45 cities in the United States, Europe, Canada, Australia and Asia. They mostly trade in temperatures, and sometimes in snowfall, rainfall, frost, and hurricanes. According to the Weather Risk Management Association (WRMA), weather contracts traded by CME reached \$8 billion within a few years of inception, with a record high of \$45.2 billion in 2006. The trading volumes decreased after the demise of the US energy company Enron (which pioneered weather derivatives) in 2001, but the gap left by energy companies was soon covered by investment banks and hedge funds. The weather derivatives market, together with most of the global financial markets, was again hurt during the financial crisis in 2008. However, it has been growing again since 2010, with increased count and diversity of end users of weather derivatives and the associated weather risk management tools. The CME weather derivatives market grew 20% just between 2010 and 2011, with about 730,000 contracts traded worldwide. According to the latest survey

conducted by the Weather Risk Management Association (WRMA, see www.wrma.org), the customized weather derivatives market grew approximately 30% between April 1, 2009 and March 31, 2010. According to this source, the total notional value for OTC traded weather risk contracts rose to \$2.4 billion, while the overall market grew to \$11.8 billion. There is considerable interest in emerging market economies as well. For instance, the Associated Chambers of Commerce and Industry of India (ASSOCHAM) has been pushing for the introduction of regulation that would allow for the use of weather derivatives, which will be an enormously successful market given the size of the Indian agricultural sector. Chokshi (2012) discussed the emergence of weather derivatives as a hedging tool, and studied the feasibility of weather derivative contracts in the context of risk management issues in agriculture, also see Below (2008) and Zhou *et al.* (2016).

The Black-Scholes model (Black and Scholes (1973)) is the most successful traditional arbitrage-free approach for pricing derivatives, and is applicable in a market which operates continuously without friction. However, pricing weather derivatives are complex mainly because (a) the weather index is itself a meteorological observation which is not directly traded on the financial market (Dischel (1998), Cao and Wei (2004), Davis (2001) and Leggio and Lien (2002)), and (b) there is little liquidity in the weather derivatives market. Traditional arbitrage-free pricing models that are usually used with financial derivatives are not appropriate for pricing weather derivatives, and it is not possible to construct a portfolio with perfect duplication in pricing. Cao and Wei (2004) claimed that the market price of risk is a key factor in the valuation of a weather derivative payoff, and that any reasonable pricing model must incorporate accurate modeling of the underlying and the assessment of the market price of risk. Their main finding is that the market price of risk associated with the temperature variable is significant. Platen and West (2005) claimed that the risk premium tends to decrease when the insurance companies are in competition. Geman (1999) suggested that the meteorological index (say temperature) be replaced by an energy price which has a high correlation with the meteorological index.

The Historical Burn Analysis is a traditional actuarial approach (Brix *et al.* (2002), Jewson (2004), Platen and West (2005) and Hamisultane (2008)) which is attractive due to its ease of implementation and applicability to any type of weather derivative, thus avoiding the issues in (a) and (b) discussed earlier. It is based on the law of large numbers, and has been widely used by the insurance industry for many years. For example, in pricing the HDD futures in Boston for next January, one simply calculates the average payoff of the same futures based on the past years. The actuarial approach can be improved by using a statistical analysis based on historical temperature data,

yielding a forecast formula to simulate future temperature and thus calculate the average payoff of the derivatives. Unlike the Historical Burn analysis which is directly based on the CDD/HDD, this blended statistical approach relies on time series modeling of pre-processed daily temperature data followed by a simulation of the CDD/HDD index with the corresponding payoffs. For earlier work, in this area, see Jewson and Caballero (2002), Campbell and Diebold (2005), Caporin and Pres (2009), Hamisultane (2008) and Schiller *et al.* (2012). Jewson *et al.* (2005) discussed valuations covering meteorological, statistical, and financial issues that arise in the pricing and risk management of weather derivatives. Chincarini (2011) focused on the market efficiency by examining the accuracy of national weather forecasts, simple weather models, as well as the actual market prices of weather futures.

All weather contracts are based on actual observations of weather variables at specific weather stations. There is evidence that the daily average temperature exhibit evidence of long-range dependence (Caballero *et al.* (2002)). This paper describes an approach for pricing weather derivatives based on model fitting and prediction of daily average temperature using an autoregressive fractionally integrated moving average ARFIMA(p,d,q) model. The data description is provided in Section 2. In Section 3, we describe a statistical modeling framework for daily average temperature using the class of autoregressive fractionally integrated moving average (ARFIMA) models, while in Section 4, we present an application to pricing temperature derivatives using the data for Boston as illustration. Analyses for other cities may be done in a similar fashion.

2 Data Description

All weather contracts are based on actual observations of weather variables at specific weather stations. For our analysis, we use average daily temperature from January 1, 1970 to November 8, 2008 measured at 17 different locations in the US. The data were provided by MDA EarthSat (www.mdacorporation.com). The underlying index of a weather derivative defines a measure of weather which governs when and how payouts on the contract will occur. The most common indices in the market are Heating Degree Days (HDDs) and Cooling Degree Days (CDDs). Let T_{\max} and T_{\min} respectively denote the maximum and minimum daily temperature, and let the average daily temperatures be defined as $T_{\text{ave}} = \frac{T_{\max} + T_{\min}}{2}$. Let X denote the average temperature for a given day. The daily HDD and CDD values are defined as

$$\text{HDD} = \max(0, 65 - X).$$

$$\text{CDD} = \max(0, X - 65).$$

The cumulative HDD is $\text{cHDD} = \sum_t \text{HDD}_t$, and the cumulative CDD is $\text{cCDD} = \sum_t \text{CDD}_t$. To determine a weekly, monthly, or seasonal HDD or CDD index value, we simply add the daily HDD or CDD values for the appropriate time period. Consider the following examples:

- HDD seasonal contract from November 1 to March 31 (151 days): $\text{cHDD} = \sum_{t=1}^{151} \text{HDD}_t$;
- CDD seasonal contract from May 1 to September 30 (153 days): $\text{cCDD} = \sum_{t=1}^{153} \text{CDD}_t$;
- HDD monthly contracts for October, December, January and March (31 days), November and April (30 days), February (28 days): $\text{cHDD} = \sum_t \text{HDD}_t$, where the summation is over the number of days in the selected month;
- CDD monthly contracts for April, June, and September (30 days), May, July, August, and October (31 days): $\text{cCDD} = \sum_t \text{CDD}_t$, where the summation is over the number of days in the selected month.

That is, suppose we are looking at pricing monthly contracts based on HDD, for October, November, December, January, February, March and April. If we look at pricing seasonal contracts, we would look at the period November - March. For CDD, the corresponding periods are April, May, June, July, August, September, and October for monthly contracts, and May - September for seasonal contracts.

3 Long Memory Models for Daily Temperatures

Stochastic modeling of daily average temperatures has been discussed in the literature, for example, see Alaton *et al.* (2002), Caballero *et al.* (2002), Campbell and Diebold (2005), and Pai and Ravishanker (2009). Campbell and Diebold (2005) modeled daily average temperature using time trend, seasonality, a long autoregression (AR), and generalized conditional heteroscedastic (GARCH) with time varying seasonal deter-

ministic regressors:

$$\begin{aligned}
 T_t &= \text{Trend}_t + \text{Seas}_t + \sum_{l=1}^L \rho_{t-l} T_{t-l} + \sigma_t \varepsilon_t; \\
 \text{Trend}_t &= \sum_{m=0}^M \beta_m t^m; \\
 \text{Seas}_t &= \sum_{p=1}^P [\sigma_{c,p} \cos(2\pi p d(t)/365) + \sigma_{s,p} \sin(2\pi p d(t)/365)]; \\
 \sigma_t^2 &= \sum_{q=1}^Q [\gamma_{c,q} \cos(2\pi q d(t)/365) + \gamma_{s,q} \sin(2\pi q d(t)/365)] + \sum_{r=1}^R \alpha_r (\sigma_{t-r} \varepsilon_{t-r})^2 \\
 &\quad + \sum_{s=1}^S \beta_s \sigma_{t-s}^2
 \end{aligned}$$

where $d(t)$ is a repeating step function which cycles through $1, \dots, 365$ (dropping Feb. 29), $\varepsilon_t \sim iid(0, 1)$, $L = 25$, $M = 1$, $P = 3$, $Q = 3$, $R = 1$, and $S = 1$.

Alaton *et al.* (2002) adopted a similar model with trend and a sine function to incorporate the seasonal effect but defined a piecewise constant function on σ_t , $t = 1, \dots, 12$. Caballero *et al.* (2002) assumed that the daily temperature had long-memory, and fit univariate autoregressive fractionally integrated moving average (ARFIMA) models to the daily average temperature at each location. They filtered out the long memory effect at each location using the estimated d for that location, and then fit a suitable vector autoregressive moving average (VARMA) model to the resulting time series. Pai and Ravishanker (2009) pre-processed the daily average temperature data by first filtering for seasonality and volatility, and then modeled the pre-processed data via vector ARFIMA processes.

In this section, we briefly summarize fitting ARFIMA(p, d, q) models to pre-processed daily average temperatures. Prior to fitting univariate and vector ARFIMA(p, d, q) models to the time series data from these stations (see the first row of Figure 1 for Boston, Dallas, and Houston for 2004-2006), we pre-processed the data for each station separately, by filtering for seasonality and volatility, using the entire series. Let z_t denote the raw temperatures, which exhibits seasonal behavior and some time-dependent variation. We form the 365-day centered moving averages, $\tilde{z}_t = \sum_{i=-182}^{182} z_{t+i}/365$, yielding 35 years of data from January 1, 1971 to December 31, 2005.

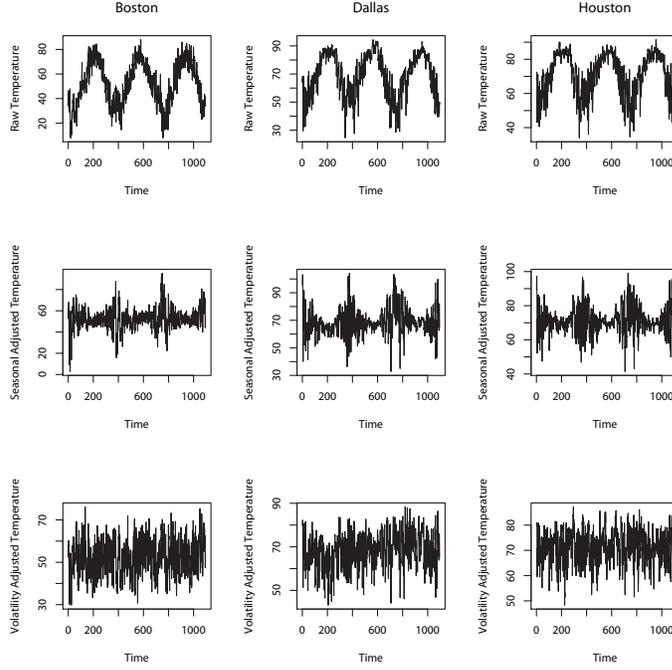


Figure 1: Raw and Pre-processed Daily Average Temperatures Data for 2004-2006 (Boston, Dallas, and Houston).

Dividing z_t by \tilde{z}_t removes the long-term secular trend and cyclical component, if any; let $s_t = z_t/\tilde{z}_t$. We compute seasonal indices (see the first row of Figure 1) $\bar{s}_k, k = 1, \dots, 365$ as $\bar{s}_k = \sum_{i=1}^{35} s_{k+(i-1)*365}/35$, for $k = 1, \dots, 365$, and $\bar{s}_k = \bar{s}_k \frac{365}{\sum_{j=1}^{365} \bar{s}_j}$, for $k = 1, \dots, 365$. Then, obtain the seasonal adjusted series as $z_1^a = z_1/\bar{s}_1, \dots, z_{365}^a = z_{365}/\bar{s}_{365}, z_{366}^a = z_{366}/\bar{s}_1$, and so on; these reveal strong time-dependent heteroscedasticity (see the second row of Figure 1). To eliminate the heteroscedasticity, we first calculate the sample standard deviation of $\{z_t^a | t = k + (i-1)*365, i = 1, \dots, 35\}, k = 1, \dots, 365$. Let \bar{z}^a be the sample mean of z_t^a (see the second row of Figure 1), we obtain an adjusted series as $z_t^v = \frac{z_t^a - \bar{z}^a}{\bar{\sigma}_k} + \bar{z}^a$ where $\bar{\sigma}_k = \sigma_k \frac{365}{\sum_{j=1}^{365} \sigma_j}$, for $k = 1, \dots, 365$. Note that for $k > 365, \bar{\sigma}_k = \bar{\sigma}_{k-365}$. The seasonal indices for the mean have a one-year cycle that peaks in July while the temperature volatility bottoms during the summer months (see Figure 2). We fit univariate and vector (p, d, q) models to the pre-processed data for 1995-2004, for $n = 3650$, (see the

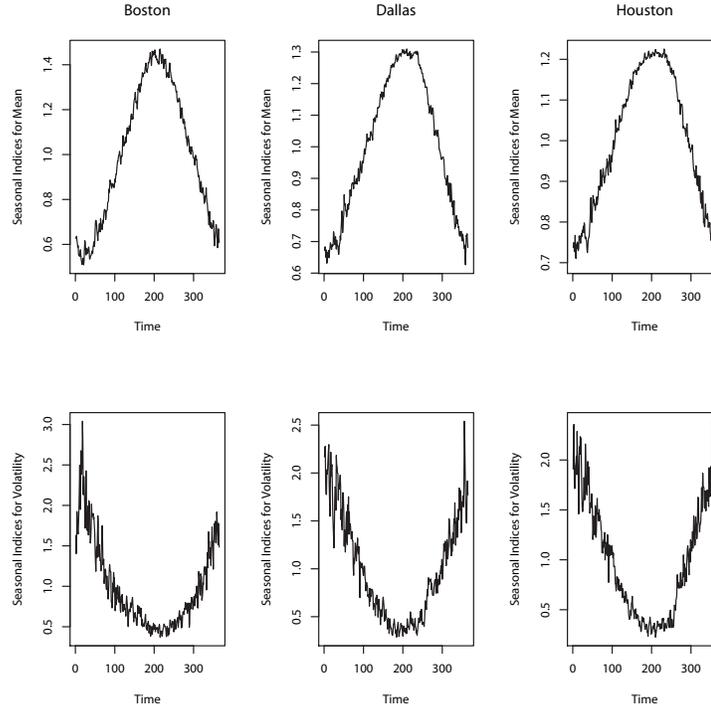


Figure 2: Seasonal Indices for Boston, Dallas, and Houston.

third row of Figure 1).

A time series $\{z_t\}$ is generated by an ARFIMA(p, d, q) process with the mean μ if

$$\phi(B)(z_t - \mu) = \theta(B)a_t, \quad (3.1)$$

where $\phi(B) = 1 - \phi_1 B - \dots - \phi_p B^p$ and $\theta(B) = 1 - \theta_1 B - \dots - \theta_q B^q$ are polynomials in B of degrees p and q respectively, p and q are known integers and B is the backward shift operator. The series is stationary and invertible if $|d| < 1/2$, and the roots of $\phi(z) = 0$ and $\theta(z) = 0$ lie outside the unit circle. The unobserved series $\{a_t\}$ is the fractionally differenced noise process represented by

$$(1 - B)^d a_t = \epsilon_t \quad \text{where} \quad (1 - B)^d = \sum_{j=0}^{\infty} \binom{d}{j} (-B)^j, \quad (3.2)$$

d is the real fractional degree of differencing and ϵ_t are independent and identically distributed as normal random variables with mean 0 and variance σ_ϵ^2 . When $d=0$, (1) reduces to the short memory ARMA process. See also Beran (1994) and Palma (2007) for details on modeling and prediction.

The exact likelihood function based on n observations $Z_n = (z_1, \dots, z_n)^T$ from a Gaussian ARFIMA process $\{z_t\}$ with mean μ is

$$f(Z_n; \Psi) = (2\pi\sigma^2)^{-n/2} |\Omega_n|^{-1/2} \exp \left[\frac{-(Z_n - \mu \mathbf{1}_n)^T \Omega_n^{-1} (Z_n - \mu \mathbf{1}_n)}{2\sigma^2} \right], \quad (3.3)$$

where $\Psi = (\Phi^T, \Theta^T, d, \mu, \sigma^2)^T$ of vector dimension $(p + q + 3)$ for ARFIMA processes or $\Psi = (\Phi^T, \Theta^T, \mu, \sigma^2)^T$ of vector dimension $(p + q + 2)$ for ARMA processes, $\Phi = (\phi_1, \dots, \phi_p)^T$, $\Theta = (\theta_1, \dots, \theta_q)^T$, $\mathbf{1}_n$ is the vector of 1's and $\sigma^2 \Omega_n$ is the covariance matrix of Z_n with elements γ_k^z , the autocovariances of $\{z_t\}$ of lag k .

Computation of Ω_n for ARFIMA(p, d, q) models may be based on the convolution of ACFs of ARFIMA($0, d, 0$) and ARMA(p, q) models (Hosking, 1981)

$$\gamma^z(k) = \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} \psi_i \gamma^a(k + i - j) \psi_j^T,$$

where ψ_i denotes the i^{th} MA weight of the ARMA(p, q) part of the process and $\gamma^a(k)$ is the lag k covariance of the ARFIMA($0, d, 0$) part. The infinite sum may be truncated when terms become smaller than some chosen small constant. This will give accurate results if the roots of $\phi(z) = 0$ and $\theta(z) = 0$ are not too close to the boundaries of the stationarity and invertibility region. The double summation involved could lead to high computation time in some cases. The relation between the ARFIMA(p, d, q) process and the corresponding ARFIMA($0, d, 0$) process may be exploited to provide computationally feasible closed form expressions for the likelihood. Specifically, we avoid the computation of autocorrelations or partial regression coefficients from the ARFIMA(p, d, q) process by a transformation which enables the bulk of the computations to be from the simpler ARFIMA($0, d, 0$) process which admits closed form expressions for the required quantities. Let $y_t = z_t - \mu$, and $u_t = D(B)y_t$; then $\phi(B)u_t = \theta(B)\epsilon_t$, and $(1 - B)^d a_t = \epsilon_t$ is an ARFIMA($0, d, 0$) process. Also, $\phi(B)y_t = \theta(B)a_t$; and

$$a_t = (z_t - \mu) - \sum_{i=1}^p \phi_i (z_{t-i} - \mu) + \sum_{j=1}^q \theta_j a_{t-j}.$$

Let $A_n = (a_1, \dots, a_n)^T$, $A_q = (a_{-q}, \dots, a_0)^T$ and $Z_p = (z_{-p}, \dots, z_0)^T$.

Chen *et al.* (2006) employed the preconditioned conjugate gradient algorithm for fast computation of nearly exact MLEs of ARFIMA(p, d, q) model parameters. Using

Table 1: Parameter estimates of ARFIMA(1,d,1) model for the three cities.

City	ϕ	d	θ	μ	σ	AIC	SBC
Boston	0.208 (.038)	0.194 (.021)	-0.299 (.026)	51.501 (.697)	6.501 (.055)	17194	17255
Dallas	0.358 (.034)	0.190 (.024)	-0.326 (.018)	67.351 (.671)	5.430 (.044)	16107	16138
Houston	0.329 (.032)	0.202 (.022)	-0.281 (.020)	70.241 (.569)	4.830 (.040)	15213	15244

the univariate version of the method and code developed in Pai and Ravishanker (2009), we employ the PCG algorithm to evaluate the quadratic form in the likelihood of an ARFIMA(0, d , 0) process, and then “adjust” to the ARFIMA(p , d , q) likelihood by integrating a joint distribution over the marginal distribution of the history $L = (Z_{-p}, A_{-q})$. We also use the approximation in Böttcher and Silbermann (1999) to evaluate the determinant of the ARFIMA(0, d , 0) covariance matrix. Table 1 presents MLEs for ARFIMA(1, d , 1) model parameters fit to three cities, viz., Boston, Dallas and Houston, as well as model selection criteria such as the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). Based on 10 years of daily data up to Dec 2004, we fit a few low order models ($p = 0, 1, 2; q = 0, 1, 2$) to temperature from each city, and based on the minimum Bayesian Information Criterion (BIC), we chose the (1, d , 1) model in each city.

4 Application to Pricing Temperature Derivatives

The price of the temperature call option, put option, and the temperature futures on the index (CDD or HDD) at time t can be expressed as follows:

$$\begin{aligned}
 C(t, T_t, I_t) &= \delta e^{-r(t_m-t)} (E^Q[\max(I_{t_m} - K, 0) | F_t]), \\
 P(t, T_t, I_t) &= \delta e^{-r(t_m-t)} (E^Q[\max(K - I_{t_m}, 0) | F_t]), \\
 F(t, T_t, I_t) &= \delta (E^Q[I_{t_m} | F_t]),
 \end{aligned}$$

where T_t is the daily average temperature at time t , I_t is the corresponding temperature index at time t , δ is the tick size equal to \$20, r is the annualized continuously compounded risk-free interest rate, t_m is the maturity date of the corresponding contract, K is the strike price, $E^Q[\cdot]$ represents the risk-neutral expectation given information up to and including time t .

Since the underlying variable (temperature) is not tradable, the market for weather derivatives is an incomplete market. Following Alaton *et al.* (2002), we include a constant market price of risk λ to obtain the following pricing process:

$$z_t - \mu = \sum_{i=1}^{\infty} \pi_i (z_{t-i} - \mu) + \sigma(\epsilon_t - \lambda).$$

where π_i is the π -weights of the ARFIMA model, ϵ_t is i.i.d. Gaussian with mean zero and variance 1, and λ is the market price of risk parameter. As shown in Alaton *et al.* (2002), after changing of the risk-neutral measure to the real-world measure using the above pricing process, the prices are estimated by discounting its payoff under the real-world probability measure adjusted by a constant risk premium.

Therefore, we consider an actuarial approach using a simple additive risk loading principle. The price of a contract is the expected outcome plus a risk premium which is the risk parameter (factor) κ times the risk measured by the standard deviation of the future outcome. The actuarial price of the weather call option, put option, and the weather futures on the index (CDD or HDD) at time t are respectively denoted as follows:

$$C^A(t, T_t, I_t) = \delta e^{-r(t_m-t)} (E[\max(I_{t_m} - K, 0) | F_t] + \kappa \sigma_c), \quad (4.1)$$

$$P^A(t, T_t, I_t) = \delta e^{-r(t_m-t)} (E[\max(K - I_{t_m}, 0) | F_t] + \kappa \sigma_p), \quad (4.2)$$

$$F^A(t, T_t, I_t) = \delta (E[I_{t_m} | F_t] + \kappa \sigma_I), \quad (4.3)$$

where T_t is the daily average temperature at time t , K is the strike price, δ is the tick size equal to \$20, t_m is the maturity date of the corresponding contract, $E(\cdot | F_t)$ is the conditional expectation under the true probability, κ is the risk parameter, and σ_c , σ_p , σ_I are the volatilities of the corresponding payoffs and index. The risk premium (for example, $\kappa \sigma_I$ for futures) is assumed to be constant under the same contract (even under different times to maturity).

The following steps describe the use of "best" ARFIMA model for pricing the temperature derivatives:

- Use the estimates from the best model to do stochastic simulation of future weather outcomes.
- Estimate the option price using the simulated forecasts.

The parameter estimates for the ARFIMA(1, d , 1) models are based on 10 years of daily pre-processed temperatures up to Dec 2004. We then price contracts for two

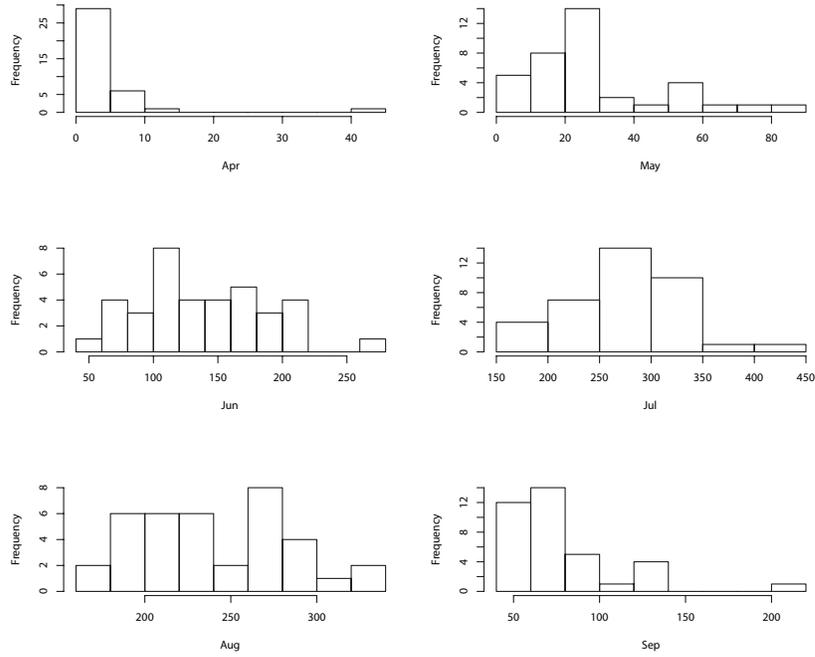


Figure 3: Histogram of Empirical CDD of Boston Futures Contracts.

cycles each for cHDD and cCDD between 2005 and 2007. We check our estimates against actual financial data on weather derivative prices, using observations provided by CME. Details of the stochastic simulation for the contracts are described in sections 4.1 and 4.2.

4.1 Simulating Realizations of cHDD or cCDD

We illustrate the approach on Nov-Mar cHDD contracts. For each city and each contract, we carry out the following steps to simulate $M = 1000$ realizations of cHDD or cCDD.

- (a) Simulate M realizations of the pre-processed data, z_t^{v*} as follows. Sample from the conditional distribution of z_t^v given data up to time $t - 1$ and MLEs for Ψ obtained via the PCG algorithm. Under the ARFIMA(p, d, q) model fit to the time

series z_t^v , the conditional distribution is given by

$$f_t(z_t^v | Z_{t-1}^v, \Psi) \sim N(\mu_t, v_t^2), \text{ where}$$

$$\mu_t = \sum_{j=1}^{t-1} \phi_j^{t-1} z_{t-j}^v$$

$$v_t^2 = \gamma_0^y / \sigma^2 \prod_{j=1}^{t-1} (1 - (\phi_j^j)^2),$$

by substituting the MLE $\widehat{\Psi}$ for Ψ , and where ϕ_j^{t-1} correspond to partial regression coefficients from the ARFIMA(p, d, q) process.

- (b) Use the fitted model and inverse filtering to compute M daily average temperatures:

- (i) $z_t^{v*} = y_t^* + \widehat{\mu}$;
- (ii) $z_t^{a*} = (z_t^{v*} - \bar{z}_a) \sigma_k + \bar{z}_a$;
- (iii) $z_t^* = z_t^{a*} \bar{s}_k$.

For $h = 1, \dots, M$, convert the simulated future daily average temperatures $z_{t,h}^*$ into $\text{HDD}_{t,h} = \max(0, 65 - z_{t,h}^*)$ or $\text{CDD}_{t,h} = \max(0, z_{t,h}^* - 65)$

- (c) For $h = 1, \dots, M$, compute cHDD_h and cCDD_h for seasonal or monthly contracts as follows:
- (i) for HDD seasonal contract of duration November 1 until March 31 (151 days), $\text{cHDD}_h = \sum_{t=1}^{151} \text{HDD}_{t,h}$;
 - (ii) for CDD seasonal contract of duration May 1 until September 30 (153 days), $\text{cCDD}_h = \sum_{t=1}^{153} \text{CDD}_{t,h}$;
 - (iii) for HDD monthly contracts for October, December, January and March (31 days); November and April (30 days); February (28 days), $\text{cHDD}_h = \sum_t \text{HDD}_{t,h}$ where the summation is taken over the number of corresponding days in that month;
 - (iv) for CDD monthly contracts for April, June, and September (30 days); May, July, August, and October (31 days), $\text{cCDD}_h = \sum_t \text{CDD}_{t,h}$ where the summation is taken over the number of corresponding days in that month.
- (d) Obtain the empirical CDF of these M cHDD_h or cCDD_h values, for $h = 1, \dots, M$. At this stage, we are able to estimate σ_I in equation (4.3).

Table 2: Estimated cCDD of Boston May Contract using Historical Burn Analysis.

Month	Apr	May	Jun	Jul	Aug	Sep
Average	3.4	28.8	139.1	275.1	239.8	77.7
Standard Deviation	7.3	20.9	49.3	58.1	44.8	33.1

4.2 Empirical Results

We illustrate our pricing approach using the data from Boston. We first calculate monthly CDD from (April to September) from 1970 to 2006. The historical averages and standard deviations of CDD months are shown in Table 2. The histograms of these empirical distributions relevant to the traditional pricing approach, Burn Analysis, are shown in Figure 3. While the method is very quick and simple, it does not produce reliable estimates, primarily because it does not incorporate statistical modeling of the corresponding weather index. For example, the estimated price based on Burn Analysis (based on data from 1970-2006) was 28.8 but the cCDD was observed to be 66 in 2007. May 2007 was a warm month in Boston, and our statistical model could pick up the higher than usual value of CDD right away, after only a few days into the month.

Suppose that we are at the beginning of a contract month, say May 1 2006, and we predict the CDD for the remaining days in May 2006 using the approach discussed in section 4.1. Therefore, we have one set of $M = 1000$ simulated cCDD values for each day of May. Figure 4 shows box plots for the last 15 days of May 2006. These provide $M = 1000$ simulated future distribution values for each day of each contract month based on the fitted ARFIMA(1, d , 1) model. As expected, the distributions are heavily skewed to the right, and the interquartile ranges get tighter as the time approaches the settlement date. On the last day of May, the cCDD for the May contract actually follows a mixed distribution with a probability mass of $\Pr(T_{31} \leq 65)$ at the value of $\sum_{t=1}^{30} \text{CDD}_t$. That is, the cCDD will remain the same if the temperature on May 31 is below 65. This can also be seen from the Quantile-Quantile plot of the simulated cCDD on May 15 and May 31, 2006, shown in Figure 5. The two variables seem to follow the same distribution except for the probability mass mentioned above.

For simplicity, we set $\delta = 1$. Let t denote a day on which an actual price is available. For $i = 1, \dots, M$, let $\hat{\mu}_t = \hat{E}(I_{t,m} | F_t)$ denote the average of the simulated cCDD (or cHDD) values for day t for the selected city, say Boston. Let F_t^o denote the observed price from the CME database and let κ_0 denote the constant risk premium for all t (κ_0 here is $\kappa \sigma_I$ in (4.3)). In our empirical study, we allow for a non-negative risk premium and

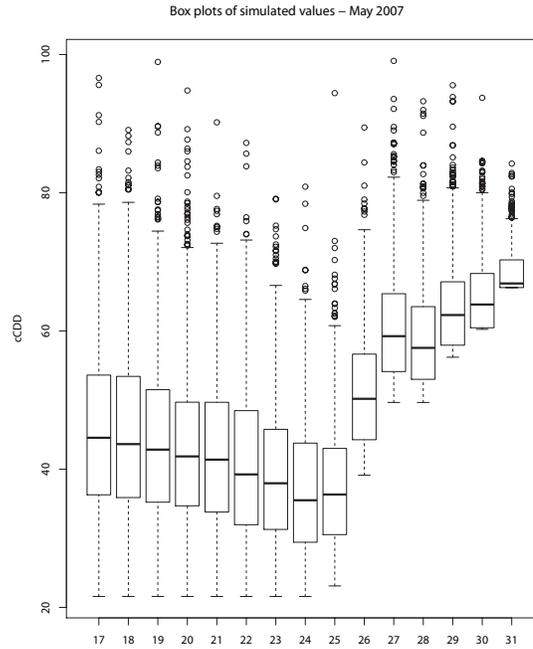


Figure 4: Box Plots of Predicted cCDD's for May 16-31, 2006, Boston Futures Contract.

we estimate the constant risk premium κ_0 by the method of least squares using all the observed price data in the given contract:

$$\hat{\kappa}_0 = \frac{\sum_t (F_t^o - \hat{\mu}_t)_+}{\#obs}, \quad (4.4)$$

where $(x)_+ = x$, if $x \geq 0$, and 0 otherwise. When pricing a monthly contract, we estimate the risk premium from the current contract and use it to price the next contract. The estimated risk premiums for 2006 CDD contracts are respectively 2.0, 11.6, 0.6, and 0.1 for May, June, August, and September (there were no transactions in the entire month of July).

Within the same monthly contract, we estimate the risk premium using the most recent observed transaction to price derivatives at the current time. In Figure 6, we use the end of the day observed price on the previous day to estimate the risk premium, and use it to price the value for the current day. We illustrate the pricing method by using the May 2007 contract (cCDD) for Boston. The settle price is 66 (again, we

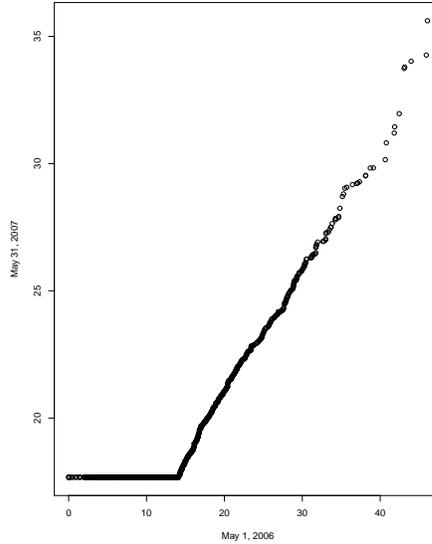


Figure 5: Quantile-Quantile plot of simulated cCDD's of May 1 and May 31, 2006, Boston Futures Contract.

assume $\delta = 1$). The circles indicate the observed price and the dotted line indicates the prediction from our model. The predictions are close to the observed values except for May 29, 2007. The reason for the discrepancy on this date is that there are no transactions between May 25 and May 29, and the observed temperature suddenly increased during those days.

5 Conclusions

In this paper we describe an approach for pricing weather derivatives based on model fitting and prediction of daily average temperatures using an autoregressive fractionally integrated moving average ARFIMA(p,d,q) model. We pre-process the temperature data to capture the seasonality and volatility, and we employ the preconditioned conjugate gradient (PCG) algorithm for fast computation. We show that the proposed pricing model produces dynamic estimates following the observed values closely to the settlement date using 2007 May Boston CDD futures.

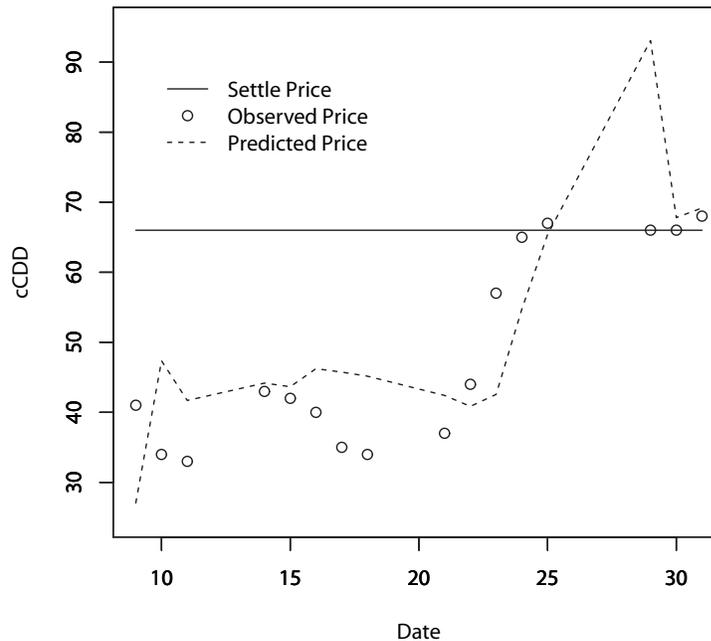


Figure 6: Observed vs Predicted Values for May 2007 Boston CDD Futures Contract.

There may be a few different ways to improve our predictions, and better pricing. For instance, our pricing approach based on the statistical modeling does not incorporate information about temperature forecasts from a meteorological system (Chincarini, 2011). Meteorological forecasts are usually available up to two weeks and information is updated frequently. In practice, traders in the weather derivatives market could combine available information from statistical and meteorological sources in order to improve pricing valuation in trading weather derivatives. The area of effective weather derivatives pricing is certainly poised to attract novel research and implementation.

Acknowledgments.

The authors are grateful to an anonymous referee and the editorial board for their excellent suggestions that improved the discussion and presentation of our results.

References

- Alaton, P., Djehiche, B. and Stillberger, D. (2002), On modeling and pricing weather derivatives. *Applied Mathematical Finance*, **9**(1), 1-20.
- Below, F. (2008), *The seven wonders of the corn yield world, 2008 Illinois Crop Protection Technology Conference*.
- Beran, J. (1994), *Statistics for Long-memory Processes*. Routledge.
- Berlage, K. (2013), The Weather Business - How companies can protect against increasing weather volatility. *Allianz Global Corporate and Specialty Reports*.
- Black, F. and Scholes, M. (1973), The pricing of options and corporate liabilities. *Journal of Political Economy*, **81**, 637-655.
- Böttcher, A., Silbermann, B. (1999), *Introduction to Large Truncated Toeplitz Matrices*. Springer Science and Business Media.
- Brix, A., Jewson, S. and Ziehmann, C. (2002), Weather derivative modeling and valuation: a statistical perspective, climate risk and the weather market. *Climate Risk and the Weather Market, Risk Books, London*, 127-150.
- Caballero, R., Jewson, S. and Brix, A. (2002), *Long memory in surface air temperature: detection, modeling and application to weather derivative valuation*. *Climate Research*, **21**(2), 127-140.
- Campbell, S. D. and Diebold, F. X. (2005), Weather forecasting for weather derivatives. *Journal of the American Statistical Association*, **100**, 6-16.
- Caporin, M. and Pres, J. (2009), *Memory time-varying models for weather derivative pricing*.
- Cao, M. and Wei, J. (2004), Weather derivatives valuation and market price of weather risk. *Journal of Futures Markets*, **24**(11), 1065-1089.
- Chen, W.W., Hurvich, C.M. and Lu, Y. (2006), On the correlation matrix of the discrete Fourier transform and the fast solution of large Toeplitz systems for long-memory time series. *Journal of the American Statistical Association*, **101**, 812-822.
- Chincarini, L. (2011), No chills or burns from temperature surprises: an empirical analysis of the weather derivatives market. *Journal of Futures Markets*, **31**(1), 1-33.

- Chokshi, A. (2012), Emergence of weather derivatives - feasibility in Indian context. *ZENITH International Journal of Business Economics and Management Research*, **2**, 139-152.
- Davis, M. (2001), Pricing weather derivatives by marginal value. *Quantitative Finance*, **1**, 305-308.
- Dischel, B. (1998), *Options pricing: Black-Scholes won't do*, *Risk*, 8-9.
- Geman, H. (1999), The Bermudan Triangle: weather, electricity and insurance derivatives. *Insurance and weather derivatives*, 197-203.
- Hamisultane, B. (2008), Which method for pricing weather derivatives?.
- Hosking, J. R. M. (1981), Fractional differencing. *Biometrika*, **68**, 165-176.
- Jewson, S. (2004), Introduction to weather derivative pricing. *The Journal of Alternative Investments*, **7**(2), 57-64.
- Jewson, S., Brix, A. and Ziehmman, C. (2005), *Weather Derivative Valuation: The Meteorological, Statistical, Financial and Mathematical Foundations*. Cambridge University Press.
- Jewson, S. and Caballero, R. (2002), *The use of weather forecasts in the pricing of weather derivatives*. *Meteorological Applications*, **10**(4), 377-389.
- Leggio, K. B. and Lien, D. (2002), Hedging gas bills with weather derivatives. *Journal of Economics and Finance*, **26**(1), 88-100.
- Pai, J. and Ravishanker, N. (2009), A multivariate preconditioned conjugate gradient approach for maximum likelihood estimation in vector long memory processes. *Statistics and Probability Letters*, **79**(9), 1282-1289.
- Palma, W. S. (2007), *Long-Memory Time Series: Theory and Methods*. New York: John Wiley and Sons.
- Platen, E. and West, J. (2005), A fair pricing approach to weather derivatives. *Asia-Pacific Financial Markets*, **11**(1), 23-53.
- Schiller, F., Seidler, G. and Wimmer, M. (2012), Temperature models for pricing weather derivatives. *Quantitative Finance*, **12**(3), 489-500.
- Zhou, R., Li, J.S.-H. and Pai, J. (2016), Hedging crop yield with exchange-traded weather derivatives. *Agricultural Finance Review*, **76**, 172-186.

