

Pricing Temperature Derivatives with a Filtered Historical Simulation Approach

Rui Zhou, Johnny Siu-Hang Li, and Jeffrey Pai

September 12, 2017

Abstract

In this paper, we propose pricing temperature derivatives using a filtered historical simulation (FHS) approach that amalgamates model-based treatment of volatility and empirical innovation density. The FHS approach implicitly captures the risk premium of the entire risk-neutral model (except the innovation distribution), thereby providing significantly more flexibility than existing methods that use only one designated parameter to capture the risk premium. Additionally, instead of relying on the fitted innovation distribution, the FHS approach uses empirical innovations to capture excess skewness, excess kurtosis, and other non-standard features in the temperature data, all of which are important for the correct pricing of temperature derivatives. We apply the FHS approach to pricing derivatives written on the temperature of Chicago, and demonstrate that this approach yields better in-sample and out-of-sample pricing performance than the constant market price of risk method and the consumption-based method.

Keywords: Weather derivatives; Filtered historical simulation; Risk-neutral pricing

1 Introduction

Weather plays a significant role in many business activities, and unexpected weather events may cause large financial losses. For example, a mild summer reduces electricity demand and thus lowers the profit margins of utility companies, a winter with less snowfall increases the operating costs of ski resorts, and a drought leads to reduced agricultural production. The adverse financial consequences of climate variation are called weather risk. According to Allianz (2013), over 30% of US GDP (\$5.7trn of \$15.7trn) is directly or indirectly affected by weather and climate, and 70% of US companies are exposed to severe weather risk. The same report also reveals that many businesses are sensitive to even small changes in weather. It is estimated that in 2012, 3.4% of US GDP - or \$534bn - was affected by routine weather variation.

Weather risk can be managed using insurance, catastrophe bonds and weather derivatives. Severe weather risk is traditionally managed by insurance on an indemnity basis. A drawback of this approach is that the claims process is usually time-consuming and costly. Moreover, as Hess et al. (2002) note, private insurance also suffers from information asymmetry, a lack of acceptable forms of collateral, and government interventions, leading to high unit transaction costs, a limited spread of institutions, and reduced access for the poor.

Catastrophe bonds were first launched in the mid-1990s. Through these bonds, weather risk can be transferred to capital market investors, who collectively have a much larger capacity than the insurance industry. Catastrophe bonds address the information asymmetry problem by linking bond payments to an industry-wide loss or a weather index. Additionally, they provide more timely payments than insurance contracts.

Similar to catastrophe bonds, weather derivatives also transfer weather risk to capital market investors. Their payoffs depend on meteorological events, which are easily measurable. In contrast to catastrophe bonds, weather derivatives are designed to hedge the adverse financial consequences of routine weather variation instead of those of extreme weather-related events. The first weather derivative transaction - a weather option embedded in a power contract - was executed in 1997 by Aquila (Considine, 2000). Since then, the market for weather derivatives has grown rapidly due to energy market deregulation in the US. Deregulation allows for more competition among energy providers, thereby calling for improved weather risk management. To satisfy the needs of various market participants, weather derivatives have been expanded to cover many different types of weather variables such as temperature, rainfall, snow, wind speed, and humidity. Their terms range from one week to several years, whereas their transaction sizes range from a few tens of thousands of dollars for small risks to several hundred million dollars or more for greater exposures. Although most weather derivative transactions are still over-the-counter, standardized weather derivatives are now listed on the Chicago Mercantile Exchange (CME), the Intercontinental Exchange (ICE), and the London International Financial Futures and Options Exchange (LIFFE).

In this paper, we focus on the pricing of temperature derivatives, the most frequently traded weather derivatives in the market. The procedure for pricing temperature derivatives typically involves two key steps.

The first step is to construct a model for temperature dynamics. Because many temperature derivatives have times-to-maturity of several months, long-term temperature models are more appropriate for pricing purposes than are meteorological weather models (which focus on short-term point forecasts). Long-term temperature models have been considered by several researchers, including Campbell and Diebold (2005), Alaton et al. (2010) and Schiller et al. (2012), who use time-series models with seasonality effects; Cabellero et al. (2002) and Pai and Ravishanker (2009), who also incorporate long memory effects; Brody et al. (2002), Benth and Såltyt -Benth (2007), and Zapranis and Alexandridis (2008, 2009), who employ the Ornstein–Uhlenbeck process and its extensions; and Jewson and Cabellero (2003) and Ritter et al. (2010), who incorporate information from climatological models into probabilistic forecasts.

The second step is to use the assumed model to derive estimates of weather derivative prices either analytically or numerically. This step is particularly challenging, partially because the underlying variable (weather) is not a tradable asset. Additionally, sophisticated temperature dynamics imply market incompleteness, and therefore, there exist multiple equivalent martingale (risk-neutral) probability measures. Of course, the well-known Black-Scholes formula is not applicable to weather derivative pricing, given its significant departures from the Black-Scholes assumptions. Various pricing methods have been developed to overcome the pricing challenge. For example, Alaton et al. (2010), Benth and Såltyt -Benth (2005), Benth et al. (2007), and H rdle and Cabrera (2012) employ risk-neutral pricing methods adapted from those used in financial markets. Davis (2001), Cao and Wei (2004), Richards et al. (2004), Platen and West (2005), Lee and Oren (2010), Chang et al. (2009), and Hamisultane (2010) consider economic methods based on the utility of a representative agent. Additionally, Xu et al. (2008) and Brrockett et al. (2006) apply an indifference pricing approach, in which the price of a temperature derivative is determined in such a way that

it is indifferent to both hedgers of weather risk and the issuers of the derivative. Finally, Brix et al. (2002) and Jewson and Brix (2005) use actuarial methods in which the price of a weather derivative is estimated by discounting its expected payoff under the real-world probability measure with a risk-adjusted discount rate. We refer readers to Hamisultane (2008) and Geman (2005) for comparisons of various approaches to pricing weather derivatives.

This paper focuses on the risk-neutral route. Most existing risk-neutral weather derivative pricing methods (Alaton et al., 2010, Benth and Såltytè-Benth, 2005 and Benth et al., 2007) assume that, up to the risk premium, the historical (real-world) and pricing (risk-neutral) volatility dynamics are governed by the same parameters. Chernov and Ghysels (2002) and Christoffersen and Jacobs (2004) empirically show that this restriction leads to poor pricing performance in the context of equity options. Additionally, Barone-Adesi et al. (2008) argue that in an incomplete market, investors do not necessarily price options using this assumption, so a more flexible approach is necessary to achieve better pricing performance. In addition, we note that in several of the articles cited above (Cao and Wei, 2004, Richards et al., 2004, and Alaton et al. 2010), only one parameter is used to capture the risk premium. However, this approach does not seem to be adequate because the market price of risk calibrated from weather derivative transaction prices is time-varying (Härdle and Cabrera, 2012). Although Benth and Såltytè-Benth (2005) and Benth et al. (2007) represent the market price of risk by a continuous function of time, their studies do not clearly indicate how to optimally specify the function.

In this paper, we attempt to develop an improved approach to the risk-neutral pricing of weather derivatives. On the modelling front, we choose to extend Campbell and Diebold's (2005) model because of its simplicity and rich intuition. The original model contains trend, seasonal, and cyclical components for the conditional mean dynamics of temperature, and allows for conditional variance dynamics with both seasonal and cyclical components. Nevertheless, it does not capture the asymmetric response of (conditional) volatility to positive and negative shocks in temperature, a phenomenon that seems to be particularly pronounced in summer. To overcome this limitation, we consider replacing the generalized autoregressive conditional heteroscedasticity (GARCH) process for the cyclical volatility component with an exponential GARCH (E-GARCH) process (Nelson, 1991). The empirical results indicate that the modified version not only captures the asymmetry but also provides significantly better goodness-of-fit. Another drawback of the original model is that it ignores the possible non-normality of residuals, a phenomenon that is observed by Benth and Såltytè-Benth (2005) and Zapranis and Alexandridis (2009) and is found to result in incorrect pricing by Ahčan (2011). To preserve possible departures from normality (e.g., excessive kurtosis), we propose using the empirical distribution (instead of the fitted normal distribution) of the residuals when estimating temperature derivative prices.

In terms of bridging the real-world and risk-neutral temperature dynamics, we consider the filtered historical simulation (FHS) approach to introduce more flexibility to the treatment of risk premiums. The FHS approach is semi-parametric, amalgamating model-based treatment of volatility (e.g., GARCH) and empirical innovation density. It has been used by Barone-Adesi et al. (2002), Kuuster (2006) and Nozari et al. (2010) to calculate portfolio risk measures and by Barone-Adesi et al. (2008) to price equity options. Although equity option pricing under GARCH may be accomplished by other methods, such as those proposed by Duan (1995), Heston and Nandi (2000), and Christoffersen et al. (2008), these methods cannot be transferred to the context of temperature derivatives because they require a tradable underlying but the underlying of a temperature derivative is not tradable.

We draw on the work of Barone-Adesi et al. (2008). In more detail, we assume that real-world and risk-

neutral temperature models have the same structure and innovation distribution. We first fit the chosen model structure to historical temperatures to obtain the (empirical) distribution of innovations and then (re-)calibrate the parameters in the chosen model structure using observed market prices to obtain the risk-neutral model for pricing purposes. The FHS approach has a number of advantages over the existing risk-neutral methods for pricing temperature derivatives. First, instead of capturing the risk premium by a single designated parameter, the FHS approach implicitly captures the risk premium in the entire risk-neutral model (except the distribution of the innovations), thereby providing significantly more flexibility. Second, the FHS approach uses empirical innovations to capture excess skewness, excess kurtosis, and other non-standard features of the temperature data, all of which are important for correct temperature derivative pricing (Ahčan, 2011).

We propose two versions of the FHS method: a full version for use with a longer calibration period and a reduced version for use with a shorter one. We compare the proposed approach with the constant market price of risk (CMPR) method (Alaton et al., 2010) and the consumption-based pricing method (Cao and Wei, 2004) and demonstrate that the FHS method yields significantly better in-sample and out-of-sample pricing accuracy than these two methods.

The remainder of this paper is structured as follows. Section 2 reviews the basic concepts of temperature derivatives, while Section 3 presents the daily temperature model. Section 4 details the FHS pricing method, and Section 5 compares our pricing method with the CMPR method and the consumption-based pricing method. Concluding remarks are given in Section 6.

2 Temperature Derivative Basics

The payoff of a temperature derivative is linked to an underlying temperature index that is based on temperatures measured at specific weather stations. The most common indexes in the temperature derivative market are Heating Degree Day (HDD) and Cooling Degree Day (CDD)¹. A degree day is a measure of how much a day's average temperature deviates from 65 degrees Fahrenheit (i.e., 18 degrees Celsius). A threshold of 65 degrees Fahrenheit is used because furnaces (air conditioners) are typically turned on when the temperature is lower (higher) than this threshold.

Let T_i be the average temperature in day i . The HDD and CDD values for day i can be expressed as

$$\text{HDD}_i = \max(0, 65 - T_i)$$

and

$$\text{CDD}_i = \max(0, T_i - 65),$$

respectively. A high HDD value indicates a high demand for heating, whereas a high CDD suggests a high demand for cooling.

The payoff of a temperature derivative is often based on the cumulated HDD/CDD over a period of time. The cumulative HDD and CDD over the sample period of $[t_0, t_1]$ are given by

$$\text{cHDD}_{t_0, t_1} = \sum_{i=t_0}^{t_1} \text{HDD}_i$$

¹The CME also trades contracts written on Cumulative Average Temperature (CAT) for some European cities. In this paper, we focus on contracts written on HDD and CDD.

Index	Call payoff	Put payoff	Future payoff
HDD	$\max(\text{cHDD}_{t_0,t_1} - K, 0)$	$\max(K - \text{cHDD}_{t_0,t_1}, 0)$	$\text{cHDD}_{t_0,t_1} - F$
CDD	$\max(\text{cCDD}_{t_0,t_1} - K, 0)$	$\max(K - \text{cCDD}_{t_0,t_1}, 0)$	$\text{cCDD}_{t_0,t_1} - F$

Table 1: Payoffs of options and futures written on cumulative CDD and HDD indexes

and

$$\text{cCDD}_{t_0,t_1} = \sum_{i=t_0}^{t_1} \text{CDD}_i,$$

respectively.

There are three types of cumulative CDD/HDD indices traded on the CME: weekly, monthly, and seasonal strip. Monthly cumulative CDD contracts can be written for April, May, June, July, August, September and October, while monthly cumulative HDD contracts can be written for October, November, December, January, February, March and April. Seasonal cumulative CDD contracts can be written for two to seven consecutive calendar months from April to October, whereas seasonal cumulative HDD contracts can be written for two to seven consecutive calendar months from October to April.

Two forms of weather derivatives are traded on the CME: European options and futures. For both, the contract unit is \$20 times the underlying Degree Days (CDD/HDD) Index. The minimum price fluctuation is 1 index point or \$20 per contract. In practice, futures are settled daily at the market price, but for simplicity, we assume that futures are settled only on expiration.

Let K be the option strike and F be the future price. Table 1 summarizes the payoff functions of the options and futures written on the cumulative CDD/HDD over a sample period of $[t_0, t_1]$.²

3 The Temperature Model

3.1 Establishing the Model Structure

We obtained from the National Centres for Environmental Information daily average temperature data collected at Chicago O’Hare airport from 2002 to 2008.³ As with typical temperature data series, the data we consider exhibit seasonality and volatility clustering, which can be captured by Campbell and Diebold’s (2005) model, in which the conditional mean dynamics are modelled by a seasonal component, a polynomial deterministic trend and autoregressive lags, and the conditional variance dynamics are modelled by a Fourier series seasonal component and a GARCH cyclical component.

In addition to seasonality and volatility clustering, we note that volatility has an asymmetric response to positive and negative shocks in temperature, particularly during summer times. This phenomenon can be observed from Figure 1, in which the solid line shows the actual daily average temperatures in summer 2016 (from 07/01/2006 to 08/31/2006), while the dashed line represents the corresponding fitted daily average temperatures based on a Fourier series with one sine and cosine term. We observe that the volatility of temperature is higher when the actual temperature is above the fitted value (i.e., positive temperature shocks) and vice versa.

²We assume that the final settlement occurs on the last day (t_1) of the HDD/CDD accumulation period.

³<https://www.ncdc.noaa.gov/cdo-web/>




Figure 1: The actual daily average temperatures from 07/01/2006 to 08/31/2006 (solid line) and the corresponding fitted daily average temperatures based on a Fourier series with one sine and cosine term (dashed line)

To capture this property, we replace the GARCH process in Campbell and Diebold's (2005) model with an EGARCH process that allows the sign and magnitude of shocks to have asymmetric effects on volatility. The modified model can be expressed as follows:

$$\begin{aligned}
T_i &= Y_i + \text{Trend}_i + \text{Seas}_i, \\
Y_i &= \sum_{l=1}^L \rho_l Y_{i-l} + \sigma_i \epsilon_i, \\
\text{Trend}_i &= \sum_{m=0}^M \beta_m t^m, \\
\text{Seas}_i &= \sum_{p=1}^P \left(\sigma_{c,p} \cos \left(2\pi p \frac{d(i)}{365} \right) + \sigma_{s,p} \sin \left(2\pi p \frac{d(i)}{365} \right) \right), \\
\log \sigma_i^2 - \text{sq}_i &= c + \sum_{r=1}^R \alpha_r \left(|\epsilon_{i-r}| - \sqrt{\frac{2}{\pi}} \right) + \sum_{r=1}^R \xi_r \epsilon_{i-r} + \sum_{s=1}^S \eta_s (\log \sigma_{i-s}^2 - \text{sq}_{i-s}), \\
\text{sq}_i &= \sum_{q=1}^Q \left(\gamma_{c,q} \cos \left(2\pi q \frac{d(i)}{365} \right) + \gamma_{s,q} \sin \left(2\pi q \frac{d(i)}{365} \right) \right), \\
\epsilon_i &\stackrel{\text{i.i.d.}}{\sim} N(0, 1),
\end{aligned}$$

where ρ_l , σ_i , β_m , $\sigma_{c,p}$, $\sigma_{s,p}$, α_r , ξ_r , η_s , $\gamma_{c,q}$ and $\gamma_{s,q}$ are parameters to be estimated and $d(i)$ is a repeating step function that cycles through $1, \dots, 365$.⁴ The trend component in the model is a simple polynomial function of time i . The seasonal components for both the conditional mean and conditional variance are modelled by the Fourier series. The cyclical component for the conditional mean follows a L -lag

⁴We ignore February 29 in leap years.

autoregressive process, while that for the conditional variance follows EGARCH(S, R). The Fourier series captures the volatility seasonality, whereas the EGARCH component captures the remaining non-seasonal volatility persistence.

We do not observe any apparent time trend in the six years of temperature data, so we set $M = 0$. On the basis of standard time-series analysis methods, we set $L = 3, P = 1, Q = 2, R = 1$, and $S = 1$.

3.2 Estimating the Model

We fit the chosen model to the Chicago daily average temperature data from January 1, 2002 to December 31, 2007 using the pseudo maximum-likelihood approach (Bollerslev and Wooldridge, 1992) with the nominal assumption of normal innovations. This technique provides consistent parameter estimates even if the true innovation density is non-normal. Table 2 shows the parameter estimates and their standard errors. It is noteworthy that parameter ξ_1 , which measures the asymmetry response of the conditional volatility to positive and negative shocks, is highly significant (with a p -value of less than 10^{-6}).

ρ_1	ρ_2	ρ_3	β_0	$\sigma_{c,1}$	$\sigma_{s,1}$	c
0.8682 (0.0161)	-0.2399 (0.0198)	0.0876 (0.0145)	50.8679 (0.3144)	-23.0271 (0.4021)	-8.9596 (0.4226)	1.8050 (0.2364)
α_1	ξ_1	η_1	$\gamma_{c,1}$	$\gamma_{s,1}$	$\gamma_{c,2}$	$\gamma_{s,2}$
0.0556 (0.0293)	0.1919 (0.0232)	0.4769 (0.0686)	0.4592 (0.0327)	0.3232 (0.0323)	-0.0514 (0.0323)	-0.2106 (0.0313)

Table 2: Estimates of the parameters in the chosen daily average temperature model (standard errors are shown in parentheses)

To further examine the benefit of using the EGARCH process, we compare our chosen model with two alternatives. Same as the original model of Campbell and Diebold (2005), the first alternative uses the simplest version of the GARCH process (Bollerslev, 1986) for the cyclical component of the conditional variance:

$$\sigma_i^2 - \text{sq}_i = c + \sum_{r=1}^R \alpha_r \epsilon_{i-r}^2 + \sum_{s=1}^S \eta_s (\sigma_{i-s}^2 - \text{sq}_{i-s}).$$

The second alternative uses the asymmetric GARCH process (Glosten et al., 1993) for the cyclical component of the conditional variance:

$$\sigma_i^2 - \text{sq}_i = c + \sum_{r=1}^R \xi_r \epsilon_{i-r}^2 \mathbb{I}_{\{\epsilon_{i-r} < 0\}} + \sum_{r=1}^R \alpha_r \epsilon_{i-r}^2 + \sum_{s=1}^S \eta_s (\sigma_{i-s}^2 - \text{sq}_{i-s}),$$

where $\mathbb{I}_{\{\epsilon_{i-r} < 0\}}$ is an indicator function that equals one if $\epsilon_{i-r} < 0$ and zero otherwise. The asymmetric GARCH captures the asymmetric impact of negative and positive shocks on conditional volatility but does not avoid the possibility of having a negative conditional variance because it is applied to σ_i^2 instead of $\ln \sigma_i^2$.

Table 3 shows the Bayesian Information Criteria (BIC)⁵ for the chosen model and the two alternatives. The BIC values support the use of an EGARCH process for the cyclical component of the conditional variance.

⁵The BIC is calculated as $k \ln(n) - 2 \ln(\hat{L})$, where n is the number of sample points, k is the number of parameters and $\hat{L}$ is the likelihood of the fitted model. A model with a lower BIC value is more preferred.

	EGARCH	GARCH	Asymmetric GARCH
BIC	19,570	19,648	19,595

Table 3: The BIC values of the chosen model (which uses EGARCH for the cyclical component of the conditional volatility) and of the two alternative models (which, respectively, use GARCH and asymmetric GARCH for the cyclical component of the conditional volatility)

We now examine the goodness-of-fit of the chosen model. Figure 2 shows the standardized residuals ϵ_i . No obvious pattern is observed, indicating that the model adequately captures the pattern in the historical temperatures. Figure 3 presents the standardized residuals' sample autocorrelation and partial autocorrelation functions, both of which suggest that the serial correlation in the historical temperatures is adequately captured by the model. Table 4 displays the p -values (for lags 1 to 10) of the Ljung-Box Q-test applied to the standardized residuals. The null hypothesis of no serial correlation is not rejected at the 5% level of significance, further confirming the model's adequacy in capturing serial correlation.

Lag	1	2	3	4	5	6	7	8	9	10
p -value	0.3951	0.6797	0.8560	0.4202	0.2269	0.3002	0.3911	0.2040	0.2054	0.2132

Table 4: The p -values (lags 1 to 10) of the Ljung-Box Q-test applied to the standardized residuals




Figure 2: The standardized residuals of the estimated model

The distribution of ϵ_i is required when simulating future temperature paths. Therefore, although the estimation method used provides consistent estimates even if the true distribution of ϵ_i is not normal, it is still necessary to examine the normality of ϵ_i . The sample skewness and excess kurtosis of the standardized residuals are -0.0649 and 0.2899 , respectively. The excess kurtosis appears to be significant, which means that the distribution of ϵ_i is heavier-tailed than normal distribution. The Q-Q plot of ϵ_i (Figure 4) also points to the same conclusion.

We further examine the normality of the standardized residuals using the Anderson-Darling test, the Jarque-Bera test and the Lilliefors test. The p -values of these tests are tabulated in Table 5. All three tests reject the null hypothesis that the standardized residuals are normally distributed at the 5% significance level. It is clear that assuming a normal distribution for ϵ_i is inappropriate when simulating sample paths of future temperatures for pricing purposes. To preserve the departures from normality, we use the empirical distribution (instead of the fitted normal distribution) of ϵ_i when estimating temperature derivative prices. Further details are provided in Section 4.2, where the pricing procedure is discussed.




Figure 3: The sample autocorrelation and partial autocorrelation functions of the standardized residuals

Test	Anderson-Darling	Jarque-Bera	Lilliefors
p -value	0.0006	0.0119	0.0030

Table 5: The p -values of the normality tests for the standardized residuals




Figure 4: The Q-Q plot of the standardized residuals

4 Filtered Historical Simulation Pricing

4.1 The Idea

The temperature model described in Section 3 is specified under the real-world measure. It cannot be directly used for risk-neutral pricing because it does not consider the market price of risk. To perform risk-neutral pricing, the model must be adjusted to form a model defined under a risk-neutral probability measure. This step is challenging, partly because temperature is not (directly) tradable and partly because the sophisticated dynamics of temperature imply market incompleteness and hence multiple risk-neutral probability measures. As mentioned in Section 1, existing risk-neutral methods for pricing temperature derivatives (Benth and Såltytè-Benth, 2005, 2007; Alaton et al., 2010; Benth et al. 2007;) are overly restrictive and may thus lead to poor pricing performance.

We attempt to improve the existing risk-neutral weather derivative pricing methods by adapting the FHS approach that was considered previously by Barone-Adesi et al. (2008) in the context of equity options. In more detail, we assume in our proposed method that the real-world temperature dynamics follow the process defined in Section 3, with innovations following the empirical (instead of the fitted normal) distribution to retain the departures from the normality assumption, which are found to be significant in Section 3. We further assume that the model for the pricing (risk-neutral) dynamics is the same as that for the real-world dynamics, except that the parameters in the risk-neutral model are instead calibrated to the market prices of existing temperature derivatives. From a technical viewpoint, one difference between our approach and the original work of Barone-Adesi et al. (2008) is that we do not impose the constraint that the expected return on the underlying asset in the risk-neutral probability measure is equal to the risk-free interest rate because temperature is not (directly) tradable.

4.2 The Pricing Formulas

Having defined the risk-neutral dynamics, the price of a temperature derivative can be estimated by discounting at the risk-free interest rate the expected payoff from the derivative under the risk-neutral probability measure. Consider an HDD contract whose payoff is linked to the cumulative heating degree days over the period of $[t_0, t_1]$. The option and future prices at time t for $t < t_1$ can be expressed as follows:

$$\begin{aligned} C(t, \text{HDD}, t_0, t_1) &= e^{-\frac{r(t_1-t)}{365}} \mathbb{E}_t^Q [\max(\text{cHDD}_{t_0, t_1} - K, 0)], \\ P(t, \text{HDD}, t_0, t_1) &= e^{-\frac{r(t_1-t)}{365}} \mathbb{E}_t^Q [\max(K - \text{cHDD}_{t_0, t_1}, 0)], \\ F(t, \text{HDD}, t_0, t_1) &= \mathbb{E}_t^Q [\text{cHDD}_{t_0, t_1}], \end{aligned}$$

where C , P and F denote the call price, put price and futures price, respectively; r is the annualized continuously compounded risk-free interest rate; and $\mathbb{E}_t^Q[\cdot]$ represents the risk-neutral expectation given information up to and including time t . The pricing formulas for CDD contracts are similar.

The day- t values of C , P and F are calculated numerically with the following procedure.

1. Simulate a temperature path from day t to t_1 under the risk-neutral probability measure by executing the following loop for $s = t, t + 1, \dots, t_1$:
 - (a) Draw a value from the empirical distribution of the standardized residuals to obtain a realization of ϵ_s ;

- (b) Using the realization of ϵ_s and the risk-neutral parameters, obtain realizations of the conditional variance (σ_s^2) and the temperature (T_s) in day s .
2. Repeat Step (1) many times to obtain a collection of simulated temperature paths under the risk-neutral probability measure. For each of the paths, calculate the payoff from the derivative being priced.
3. Determine the day- t price by averaging the simulated payoffs and discounting the average payoff at the assumed risk-free interest rate.

We allow the valuation date t to be later than the beginning date t_0 of the sample period. If $t \leq t_0$, then the cHDD/cCDD values in the expectation are based entirely on simulated temperatures under the risk-neutral probability measure. If $t > t_0$, then the cHDD/cCDD values in the expectation are based in part on the historical temperature data (from t_0 to t) and in part on the simulated temperatures (from $t + 1$ to t_1) under the risk-neutral probability measure.

4.3 Calibrating the Pricing Process

The FHS approach assumes that the real-world and risk-neutral temperature models have the same structure but different parameters. The following steps are taken to calibrate the risk-neutral temperature model.

1. Set the initial pricing parameters (i.e., parameters in the risk-neutral model) to the parameters estimated using historical temperature data.
2. Given the current estimates of the pricing parameters, determine the theoretical price of the i -th temperature derivative transaction for $i = 1, 2, \dots, N$, where N is the total number of observed transactions.
3. Calculate mean absolute percentage pricing error, defined as $\sum_{i=1}^N e_i / N$, where e_i is the absolute percentage difference between the simulated theoretical price and the observed price of the i -th transaction.
4. Repeat the previous steps until the mean absolute percentage error is minimized.⁶

The criterion of minimizing the mean absolute percentage pricing error can be viewed as a method to select a unique risk-neutral probability measure among the infinitely many; it does not entail a direct specification of the state-price density while allowing the calibrated risk-neutral model to give a good fit to the observed market prices.

⁶We use the 'fmincon' function with the interior-point algorithm in MATLAB to solve the optimization problem. Because of the nature of our market price data, we choose to use the mean absolute percentage error as the optimization criterion instead of following Barone-Adesi et al. (2008), who use the mean squared pricing error. In our dataset, the observed prices span a wide range of values (with the minimum of 20 and the maximum of 1308), and a higher observed price is often associated with a higher absolute difference between the observed and theoretical prices. As such, if the mean squared error is used, then transactions with high observed prices will be given much heavier weights than other transactions and the calibrated risk-neutral model may be distorted.

4.4 Market Price Data

To calibrate the pricing process, we use the data concerning the Chicago temperature futures traded in the CME from February 2008 to October 2008. We discard derivatives cleared within 14 days of expiration because long-term temperature models such as ours generally do not produce sufficiently accurate temperature forecasts (Campbell and Diebold, 2005). A total of 76 future transactions are used in the actual calibration.

4.5 Pricing results

We now present the pricing results and evaluate the performance of the FHS pricing method. The evaluation is based on the in-sample and the out-of-sample pricing errors. To compute these quantities, we divide the sample of market prices into two portions. The first portion (which contains market prices of the derivatives cleared before a certain cut-off date) is used to calibrate the pricing process. The in-sample pricing error is the minimized mean absolute percentage difference between the fitted and actual prices (before the cut-off date). The second portion (which contains the market prices of the derivatives cleared after the cut-off date) is used to examine the accuracy of the prices estimated by the calibrated pricing model. The out-of-sample pricing error is the mean absolute percentage difference between the estimated prices and the actual prices (after the cut-off date). The following four cases are considered.

	First portion	Second portion
Case 1	02/01-04/30	05/01-10/31
Case 2	02/01-07/31	08/01-10/31
Case 3	02/01-09/15	09/16-10/31
Case 4	02/01-10/31	–

For the last case, all market price data are used to calibrated the pricing process, and therefore, the out-of-sample pricing error is unavailable. All the calculations are performed using an assumed continuously compounded risk-free interest rate of 2% per annum.

Table 6 tabulates the in-sample pricing error, the out-of-sample pricing error and the calibrated parameters in the pricing process for each of the four cases we consider. The fit of the pricing process (measured by the in-sample pricing error) is particularly good for Cases 1 and 2 because the numbers of transactions for these two cases (27 and 39, respectively) are **small** relative to the number of calibrated parameters (14 in total) in the pricing model. However, the out-of-sample pricing errors for Cases 1 and 2 are high, indicating that the pricing process may have overfitted the observed market prices. For Case 3, both the in-sample and out-of-sample pricing errors are reasonable. We do not observe the overfitting problem for Case 3 because the pricing process for this case is calibrated to a much larger (56) number of observed market prices.

To address the overfitting problem, we propose a reduced FHS method in which only a subset of the 14 parameters is calibrated to observed market prices. We choose to calibrate the EGARCH parameters in the conditional volatility component because they are closely related to the risk premium of the temperature derivatives. We also choose to calibrate the trend parameter in the conditional mean to capture the possible change in the conditional mean component as the probability measure changes from real-world to

	Case 1	Case 2	Case 3	Case 4
Calibration period	02/01-04/30	02/01-07/31	02/01-09/15	02/01-10/31
β_0	50.0691	51.0553	50.8759	51.0370
$\sigma_{c,1}$	-22.1113	-22.3529	-22.7689	-22.6883
$\sigma_{s,1}$	-8.0815	-8.6675	-8.6335	-8.5699
ρ_1	1.0454	1.2208	1.2033	0.1910
ρ_2	-0.2443	-0.3322	-0.4848	-0.2725
ρ_3	-0.4814	-0.4011	-0.1591	0.1813
$\gamma_{c,1}$	2.8398	2.9238	3.2423	0.0851
$\gamma_{s,1}$	-1.0463	-0.3041	3.0862	1.8499
$\gamma_{c,2}$	1.2082	0.3523	-0.0645	-0.6008
$\gamma_{s,2}$	-0.4785	-0.8826	-1.3765	0.6011
c	1.4723	1.2030	0.6867	3.8051
η_1	0.3320	-0.0753	-0.0536	-0.3027
ξ_1	-1.0722	-0.3685	0.0446	-0.1643
α_1	0.1827	0.4143	0.7927	0.2305
In-sample pricing error	3.09%	7.29%	11.10%	9.79%
out-of-sample pricing error	42.95%	23.13%	6.35%	-

Table 6: The in-sample pricing error, out-of-sample pricing error and calibrated pricing process parameters for each of the four calibration periods under consideration. The values are obtained using the full FHS method with all 14 parameters in the pricing process calibrated to observed market prices

risk-neutral. All the other pricing parameters are assumed to be the same as their real-world counterparts. There are 5 parameters to be calibrated in the reduced FHS method.

Table 7 presents the in-sample pricing errors, the out-of-sample pricing errors and the calibrated parameters in the pricing process when the reduced FHS method is used. For Cases 1 and 2, the fit of the pricing process in reduced FHS method is not as good as that in the full FHS method; however, the reduced FHS pricing method results in much lower out-of-sample pricing errors, indicating that the overfitting problem is mitigated. For Cases 3 and 4, the reduced FHS method results in marginally larger in-sample and out-of-sample pricing errors than the full version.

5 Comparison with Other Pricing Methods

In this section, we compare the FHS method with two alternative pricing methods: Alaton et al.'s (2010) constant market price of risk (CMPR) method and Cao and Wei's (2004) consumption-based pricing method.

5.1 The CMPR Method

In the CMPR method, the pricing process contains an additional constant market price of risk parameter compared to the assumed real-world process. Alaton et al. (2010) connects the real-world and risk-neutral processes by specifying the state price density and derives an analytical formula for pricing HDD options.

	Case 1	Case 2	Case 3	Case 4
Calibration period	02/01-04/30	02/01-07/31	02/01-09/15	02/01-10/31
β_0	51.3955	50.7242	50.8014	50.9968
c	2.9012	2.2633	2.9086	0.4318
η_1	0.0940	0.1655	-0.1524	-0.0317
ξ_1	-0.3343	0.9178	0.2227	-1.3140
α_1	0.4612	0.5865	0.6633	2.7952
In-sample pricing error	9.82%	9.74%	12.63%	10.79%
Out-of-sample pricing error	14.68%	13.33%	6.66%	-

Table 7: The in-sample pricing error, out-of-sample pricing error and calibrated pricing process parameters for each of the four calibration periods under consideration. The values are obtained using the reduced FHS method with 5 of the 14 parameters in the pricing process calibrated to observed market prices

Following Alaton et al.(2010), we add a constant market price of risk to our assumed historical temperature model and obtain the following pricing process:

$$\begin{aligned}
T_i &= Y_i + \text{Trend}_i + \text{Seas}_i, \\
Y_i &= \sum_{l=1}^L \rho_l Y_{i-l} + \sigma_i(\epsilon_i - \lambda),
\end{aligned}$$

where λ is the market price of risk parameter. The pricing process is fully defined by the parameters of the assumed historical temperature model and the market price of risk parameter.

An analytic pricing formula is not available because of the complexity of our assumed temperature model. We use simulations instead to determine the prices of temperature derivatives. To ensure comparability with the FHS approach, we simulate pricing paths using the empirical distribution of the innovations. Similar to the calibration procedure described in Section 4.3, the value of λ is chosen such that the absolute percentage difference between the fitted and observed market prices within the calibration window is minimized.

For the four cases under consideration, Table 8 shows the in-sample and out-of-sample pricing errors when the CMPR method is used. The in-sample and out-of-sample errors produced by the CMPR method are significantly larger than those produced by the reduced FHS method. The poor performance of CMPR method likely occurs because a single, constant parameter is not sufficiently flexible to capture the characteristics of the market price of risk. Table 8 also shows the calibrated values of λ . A positive λ means that the short position holder of a CDD future contract pays a risk premium to the long position holder, and vice versa.

	Case 1	Case 2	Case 3	Case 4
Calibration period	02/01-04/30	02/01-07/31	02/01-09/15	02/01-10/31
λ	-0.0061	0.0147	0.0355	0.0066
In-sample pricing error	10.35%	12.01%	18.76%	16.06%
Out-of-sample pricing error	19.42%	19.56%	10.96%	-

Table 8: The in-sample pricing error, out-of-sample pricing error and calibrated value of λ when the CMPR method is applied to each of the four calibration periods under consideration

5.2 The Consumption-Based Method

5.2.1 The framework

Cao and Wei (2004) price weather derivatives using an extension of the Lucas's (1978) pure exchange economy. It is assumed that there are two fundamental sources of uncertainty in this economy - the aggregate dividend and temperature - and that the representative agent has a risk preference characterized by a constant relative risk aversion, such that the utility of his/her day i consumption c_i is given by

$$U(c_i, i) = \frac{c_i^{\gamma+1}}{\gamma+1},$$

where γ is the risk aversion parameter. The equilibrium conditions for the pure exchange economy imply that c_i is equal to the day- i aggregate dividend δ_i .

Cao and Wei (2004) use the following process to model δ_i :

$$\ln \delta_i = \alpha + \mu \ln \delta_{i-1} + v_i,$$

where α and μ are parameters that are estimated to historical data. The error term v_i of the aggregate dividend process is assumed to satisfy

$$v_i = \sigma z_i + \sigma \frac{\varphi}{\sqrt{1-\varphi^2}} \epsilon_i + \sigma \zeta_1 \epsilon_{i-1} + \dots + \zeta_m \epsilon_{i-m},$$

where $z_i \stackrel{\text{i.i.d.}}{\sim} N(0, 1)$, ϵ_i is the corresponding error term from the daily temperature model, and φ , σ and ζ_i for $i = 1, \dots, m$ are parameters that are estimated to historical data. It is further assumed that ϵ_i and z_i are independent. Because Cao and Wei (2004) have shown that the lagged effect of temperature on the aggregate dividend is not significant, we set $\zeta_i = 0$ for $i = 1, 2, \dots, m$, which implies that the correlation between v_i and ϵ_i is φ .

Under the consumption-based asset-pricing model, the day- t price of a contingent claim with a payoff q_{t_1} on a future date t_1 can be expressed as

$$p_t = \mathbb{E}_t \left[e^{-\frac{\rho(t_1-t)}{365}} \frac{U'(c_{t_1}, t_1)}{U'(c_t, t)} q_{t_1} \right],$$

where ρ is the annualized rate of time preference, $\mathbb{E}_t[\cdot]$ represents the real-world expectation given information up to and including time t , and $U'(c_i, i)$ is the partial derivative of $U(c_i, i)$ with respect to c_i . In the expectation,

$$e^{-\frac{\rho(t_1-t)}{365}} \frac{U'(c_{t_1}, t_1)}{U'(c_t, t)}$$

represents the intertemporal marginal rate of substitution of consumption, which is also referred as the stochastic discount factor or the pricing kernel.

Since $U'(c_i, i) = c_i^\gamma$ and $c_i = \delta_i$, we have

$$p_t = \mathbb{E}_t \left[e^{-\frac{\rho(t_1-t)}{365}} \frac{\delta_{t_1}^\gamma}{\delta_t^\gamma} q_{t_1} \right],$$

which implies that the day- t price of a zero-coupon bond that pays \$1 on day- t_1 is

$$e^{-\frac{\rho(t_1-t)}{365}} \mathbb{E}_t \left[\frac{\delta_{t_1}^\gamma}{\delta_t^\gamma} \right].$$

As before, we use r to denote the assumed (constant) continuously compounded annualized risk-free interest rate. It follows that ρ and r satisfy the following relation:

$$e^{-\frac{\rho(t_1-t)}{365}} \mathbb{E}_t \left[\frac{\delta_{t_1}^\gamma}{\delta_t^\gamma} \right] = e^{-\frac{r(t_1-t)}{365}}. \quad (1)$$

Given r , γ and the aggregate dividend process, ρ can be determined by solving equation (1).

Finally, using the consumption-based method, the theoretical time- t prices of options and futures written on the cumulative HDD over $[t_0, t_1]$ can be computed using the following formulas:

$$\begin{aligned} C(t, \text{HDD}, t_0, t_1) &= \mathbb{E}_t \left[e^{-\frac{\rho(t_1-t)}{365}} \frac{\delta_{t_1}^\gamma}{\delta_t^\gamma} \max(\text{cHDD}_{t_0, t_1} - K, 0) \right], \\ P(t, \text{HDD}, t_0, t_1) &= \mathbb{E}_t \left[e^{-\frac{\rho(t_1-t)}{365}} \frac{\delta_{t_1}^\gamma}{\delta_t^\gamma} \max(K - \text{cHDD}_{t_0, t_1}, 0) \right], \\ F(t, \text{HDD}, t_0, t_1) &= \frac{\mathbb{E}_t [\delta_{t_1}^\gamma \text{cHDD}_{t_0, t_1}]}{\mathbb{E}_t [\delta_{t_1}^\gamma]}. \end{aligned}$$

The pricing formulas for CDD contracts are similar.

5.2.2 Parameter calibration and pricing

Cao and Wei (2004) suggest using personal consumption expenditures as a proxy for aggregate dividend. The consumption-based pricing approach requires the same data frequency for the aggregate dividend and temperature, but only monthly, quarterly and annual consumption expenditure data are available. To work around this problem, Cao and Wei (2004) use monthly data for both quantities, whereas Hamisultane (2008) linearly interpolates monthly consumption expenditures to obtain daily consumption expenditures. Hamisultane's (2008) method does not preserve the correlation between temperature and consumption expenditure.

We address the data frequency problem with a new method based on the assumption that the correlation between daily temperature and consumption is the same as that between monthly temperature and consumption. In more detail, we first estimate a monthly average temperature model (detailed in Appendix A), which is then connected to the following model for monthly annualized personal consumption expenditures:

$$\begin{cases} \ln \delta_i^{(m)} &= \alpha^{(m)} + \mu^{(m)} \ln \delta_{i-1}^{(m)} + v_i^{(m)} \\ v_i^{(m)} &= \sigma^{(m)} z_i^{(m)} + \sigma^{(m)} \frac{\varphi}{\sqrt{1-\varphi^2}} \epsilon_i^{(m)} \end{cases}, \quad (2)$$

where $\delta_i^{(m)}$ represents the annualized personal consumption expenditure for month i , $\alpha^{(m)}$, $\mu^{(m)}$, $\sigma^{(m)}$ and φ are model parameters, $z_i^{(m)} \stackrel{\text{i.i.d.}}{\sim} N(0, 1)$, and $v_i^{(m)}$ and $\epsilon_i^{(m)}$ are the error terms for the monthly personal consumption expenditure and the monthly temperature models, respectively. As before, φ can be interpreted to mean the contemporaneous correlation between the dividend process and the temperature process. It is assumed to be the same regardless of the data frequency. Estimates of the parameters in (2) are obtained by regressing $\ln \delta_i^{(m)}$ on $\ln \delta_{i-1}^{(m)}$ and $\epsilon_i^{(m)}$.

Table 9 shows the estimates of the parameters in (2). The estimate of μ is close to 1, an outcome that can be attributed to the fact that the log monthly annualized personal expenditures exhibit a steady trend rather than being mean-reverting (see Figure 5). The null hypothesis that $\mu^{(m)} = 1$ cannot be rejected at the 5% significance level. Therefore, we set $\mu = 1$ and perform a linear regression of $\ln \delta_i^{(m)} - \ln \delta_{i-1}^{(m)}$ on

$\epsilon_i^{(m)}$ to obtain revised parameter estimates. The revised parameter estimates (shown in Table 10) yield 0.2707 as an estimate of the contemporaneous correlation φ .

Finally, the parameters in the daily consumption expenditure model are determined by setting $\alpha = \alpha^{(m)}/30$ and $\sigma = \sigma^{(m)}/\sqrt{30}$. Table 11 displays the parameter estimates for the daily consumption expenditure model.

Parameter	Estimate	Standard error
$\alpha^{(m)}$	0.0396	0.0356
$\mu^{(m)}$	0.9961	0.0039
$\sigma^{(m)} \frac{\varphi}{\sqrt{1-\varphi^2}}$	0.0010	0.0004
$\sigma^{(m)}$	0.0032	—

Table 9: Parameter estimates for the monthly personal consumption expenditure model (when no constraint is applied $\mu^{(m)}$)

Parameter	Estimate	Standard error
$\alpha^{(m)}$	0.0045	0.0004
$\sigma^{(m)} \frac{\varphi}{\sqrt{1-\varphi^2}}$	0.0009	0.0004
$\sigma^{(m)}$	0.0032	—

Table 10: Parameter estimates for the monthly personal consumption expenditure model (when $\mu^{(m)}$ set to 1)

Parameter	α	φ	σ
Estimate	1.4824×10^{-4}	0.2707	5.6602×10^{-4}

Table 11: Parameter estimates for the daily personal consumption expenditure model (with $\mu = 1$)

When $\mu = 1$, the solution to equation (1) is

$$\rho = r + 365\gamma\alpha + \frac{365\gamma^2\sigma^2}{2(1-\varphi^2)}. \quad (3)$$

A proof of this result is provided in Appendix B. Equation (3) specifies the one-to-one correspondence between ρ and γ given parameters in the daily consumption expenditure model.

The value of γ (and hence ρ) is calibrated to the observed prices of temperature futures using the following procedure:

1. Assume an arbitrary initial value for γ .
2. Determine the theoretical price of the i -th temperature derivative transaction for $i = 1, 2, \dots, N$, where N is the total number of observed transactions.
 - (a) Suppose that the transaction occurred on day t . Simulate 2000 temperature paths for the period of $[t, t_1]$ from the daily temperature model defined under the real-world probability measure. To ensure comparability with the FHS method, the simulations are performed using the empirical distribution (instead of the fitted normal distribution) of the innovations.




Figure 5: The monthly annualized personal consumption expenditures from January 2002 to January 2008

- (b) Simulate 2000 paths of daily consumption expenditures for the period of $[t, t_1]$.
 - (c) Calculate the theoretical price of the temperature derivative using the given value of γ and the simulated paths of daily temperature and daily consumption expenditures.
3. Calculate the mean absolute percentage pricing error (in-sample) defined as the average absolute percentage difference between the simulated theoretical price and the observed price across the N observed transactions.
 4. Repeat Steps 1 to 3 with different values of γ until the mean absolute percentage pricing error (in-sample) is minimized.

For all of the four cases we consider, Table 12 shows the in-sample and out-of-sample pricing errors produced by the consumption-based method. Compared to the reduced FHS method, the consumption-based method yields larger pricing errors. The inferior performance of the consumption-based method can be attributed to the fact that its pricing kernel $e^{-\frac{\rho(t_1-t)}{365}} \delta_{t_1}^\gamma / \delta_t^\gamma$ is fully defined by the aggregate dividend (consumption expenditure) process, the risk aversion parameter and the rate of time preference. Because the aggregate dividend process is driven by historical data and the rate of time preference must satisfy equation (1), there is only one degree of freedom (the risk aversion parameter) to capture the market price of risk. Therefore, similar to the CMPR method, the consumption-based method suffers from an inflexibility problem. Furthermore, the consumption-based pricing method is affected by the choice of the utility function, the specification of temporal non-separability, and the measurement error associated with the consumption data. Empirical work has shown that the consumption-based pricing method does not always perform satisfactorily in asset pricing (Campbell and Cochrane, 2000; Mehra, 2012).

	Case 1	Case 2	Case 3	Case 4
Calibration period	02/01-04/30	02/01-07/31	02/01-09/15	02/01-10/31
γ	-3.7689	-69.7247	-68.0605	-27.9411
In-sample pricing error	10.35%	11.75%	18.27%	15.90%
Out-of-sample pricing error	19.19%	19.67%	11.98%	-

Table 12: The in-sample pricing error, out-of-sample pricing error and calibrated value of γ when the consumption-based method is applied to the four calibration periods under consideration

6 Concluding Remarks

The FHS approach studied in this paper possesses two important properties: (1) the real-world and risk-neutral temperature processes are identical only up to the model structure and the innovation distribution and (2) the payoffs of temperature derivatives are simulated using the empirical distribution rather than the fitted normal distribution of the innovations. The former property gives the risk-neutral process more flexibility to capture the underlying risk premium, whereas the second property capture excess skewness, excess kurtosis, and other non-standard features in the temperature data, all of which are important for correct temperature derivative pricing.

We have applied the FHS pricing method to derivatives written on the temperature of Chicago and found that the full FHS method is subject to an overfitting problem when the calibration period is short because the number of free parameters in this case is too large relative to the number of observed prices. To mitigate this problem, we proposed a reduced FHS method in which 5 of 14 model parameters are calibrated to market prices. We also compared the proposed approach with the CMPR and consumption-based pricing methods and found that the reduced FHS method consistently yields lower in-sample and out-of-sample pricing errors than the two well-cited alternative methods.

A limitation of the proposed approach is that time-series models - including the one used in this paper - do not produce sufficiently accurate short-term temperature forecasts. As such, the pricing errors produced by the proposed approach (and the two considered alternatives) are large when the derivatives being priced are short-dated. Several researchers demonstrate that meteorological forecasts can improve the accuracy of weather derivative pricing (Zeng, 2000a, b; Jewson and Caballero, 2003; Dorfleitner and Wimmer, 2010). Therefore, a direction for further research is to improve the FHS method by incorporating meteorological forecasts. This extension is challenging because it involves additional parameters which may complicate the calibration procedure.

References

- Ahćcan A. (2011). Statistical analysis of model risk concerning temperature residuals and its impact on pricing weather derivatives. *Insurance: Mathematics and Economics*, 50, 131-138.
- Alaton, P. , Djehiche, B and Stillberger, D. (2010). On modelling and pricing weather derivatives. *Applied Mathematical Finance*, **9**, 1-20.
- Allianz (2013). The Weather Business - How companies can protect against increasing weather volatility. Retrieved from [http://www.agcs.allianz.com/assets/PDFs/Reports/2484%20Allianz%](http://www.agcs.allianz.com/assets/PDFs/Reports/2484%20Allianz%20Weather%20Business.pdf)

- Barone-Adesi, G., Engle, R.F. and Mancini, L. (2008). A GARCH option pricing model with filtered historical simulation. *Review of Financial Studies*, **21**, 1223-1258.
- Barone-Adesi, G., Giannopoulos, K., and Vosper, L. (2002). Backtesting derivative portfolios with filtered historical simulation (FHS). *European Financial Management*, **8**, 31-58.
- Benth, F.E. and Såltytè-Benth, J. (2005). Stochastic modelling of temperature variations with a view towards weather derivatives. *Applied Mathematical Finance*, **12**, 53-85.
- Benth F.E. and Såltytè-Benth, J. (2007). The volatility of temperature and pricing of weather derivatives. *Quantitative Finance*, **7**, 553-561.
- Benth F.E., Såltytè-Benth, J., and Koekebakker, S. (2007). Putting a price on temperature. *Scandinavian Journal of Statistics*, **34**, 746-767.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroscedasticity. *Journal of Econometrics*, **31**, 307-327.
- Bollerslev, T. and Wooldridge, J.M. (1992). Quasi maximum likelihood estimation and inference in dynamic models with time varying covariances. *Econometric Reviews*, **11**, 143-172.
- Brix, A., Jewson, S. and Ziehmman, C. (2002). Weather derivative modelling and valuation: A statistical perspective. *R.S. Dischel (Ed.), Climate risk and the weather market*, Risk Books, London.
- Brockett, P.L., Wang, M., Yang, C., Zou, H. (2006). Portfolio effects and valuation of weather derivatives. *The Financial Review*, **41**, 55-76.
- Brody, D.C., Syroka, J., and Zervos, M. (2002). Dynamical pricing of weather derivatives. *Quantitative Finance*, **2**, 189-198.
- 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**, 127-140.
- Campbell, J.Y. and Cochrane, J.H. (2000). Explaining the poor performance of consumption-based asset pricing models. *Journal of Finance*, **55**, 2863-2878.
- Campbell S. D. and Diebold, F.X. (2005). Weather Forecasting for Weather Derivatives. *Journal of the American Statistical Association*, **100**, 6-16.
- Cao, M. and Wei, J. (2004). Weather derivatives valuation and market price of weather risk. *Journal of Futures Markets*, **24**, 1065-1089.
- Chang, C.-C, Lin, J.-B., and Shen, W.-M. (2009). Pricing weather derivatives using a predicting power time series process. *Asia-Pacific Journal of Financial Studies*, **38**, 863-890.
- Chernov M. and Ghysels E.. (2000). A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. *Journal of Financial Economics*, **56**, 407-458.
- Christoffersen P. and Jacobs K.. (2004). Which GARCH Model for option valuation. *Management Science*, **50**, 1204-1221.

- Christoffersen, P., Jacobs, K., Ornathanalai, C., and Wang, Y. (2008). Option valuation with long-run and short-run volatility components. *Journal of Financial Economics*, **90**, 272-297.
- Considine, G. (2000). Introduction to weather derivatives. Retrieved from https://www.cmegroup.com/trading/weather/files/WEA_intro_to_weather_der.pdf.
- Davis, M. (2001). Pricing weather derivatives by marginal value. *Quantitative Finance*, **1**, 305-308.
- Dorflleitner, G. and Wimmer, M. (2010). The pricing of temperature futures at the Chicago Mercantile Exchange. *Journal of Banking and Finance*, **34**, 1360-1370.
- Duan, J.C. (1995). The GARCH option pricing model. *Mathematical Finance*, **5**, 13-32.
- Geman, H. and Leonardi, M. (2005). Alternative approaches to weather derivatives pricing. *Managerial Finance*, **31**, 45-71.
- Glosten, L.R., R. Jagannathan, and D. Runkle. (1993). Relationship between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, **48**, 1779-1801.
- Hamisultane, H. (2008). Which method for pricing weather derivatives? Retrieved from <https://hal.archives-ouvertes.fr/halshs-00355856/document>.
- Hamisultane, H. (2010). Utility-based pricing of weather derivatives, *The European Journal of Finance*, **16**, 503-525.
- Hansen, L.P. and Singleton K.J. (1983). Stochastic consumption, risk Aversion, and the temporal behavior of asset returns. *Journal of Political Economy*, **91**, 249-265.
- Härdle, W.K. and Cabrera, B.L. (2012). The implied market price of weather risk, *Applied Mathematical Finance*, **19**, 59-95.
- Hess, U., Richter, K. and Stoppa A. (2002). Weather risk management for agriculture and agribusiness in developing countries. *R.S. Dischel (Ed.), Climate risk and the weather market*, Risk Books, London.
- Heston S. and Nandi S. (2000). A closed-form GARCH option valuation model, *Review of Financial Studies*, **13**, 585-625.
- Jewson, S. and Caballero, R. (2003). The use of weather forecasts in the pricing of weather derivatives. *Meteorological Applications*, **10**, 377-389.
- Jewson, S. and Brix, A. (2005). Weather derivative valuation. Cambridge University Press, Cambridge.
- Kuester, K., Mittnik, S., and Paolella, M.S. (2006). Value-at-Risk prediction: A comparison of alternative strategies. *Journal of Financial Econometrics*, **4**, 53-89.
- Lee, Y. and Oren, S.S.. (2010). A multi-period equilibrium pricing model of weather derivatives. *Energy Systems*, **1**, 3-30.
- Lucas, R.E. (1978). Asset prices in an exchange economy. *Econometrica*, **46**, 1429-1445.
- Mehra, R. (2012). Consumption-based asset pricing models. *Annual Review of Financial Economics*, **4**, 385-409.

- National Centres for Environmental Information. (2016). Retrieved from <https://www.ncdc.noaa.gov/cdo-web/>.
- Nelson, D.B. (1991). Conditional heteroskedasticity in asset returns: a new approach. *Econometrica*, 59, 347-370.
- Nozari, M., Raei, S.M., Jahangiri, P., and Bahramgiri, M. (2010). A comparison of heavy-tailed VaR estimates and Filtered Historical Simulation: Evidence from emerging markets. *International Review of Business Research Papers*, 6, 347-359.
- Pai, J. S. and Ravishanker, N. (2009). A multivariate preconditioned conjugate gradient approach for maximum likelihood estimation in vector long memory processes. *Statistics and Probability Letters*, 79, 1282-1289.
- Platen, E. and West, J. (2005). A fair pricing approach to weather derivatives. *Asia-Pacific Financial Markets*, 11, 23-53.
- Richards, T.J., Manfredo, M.R., and Sanders, D.R. (2004). Pricing weather derivatives. *American Journal of Agricultural Economics*, 86, 1005-1017.
- Ritter, M., Muckhoff, O., and Odening, M. (2010). Meteorological forecasts and the pricing of weather derivatives. Working paper. Retrieved from <http://sfb649.wiwi.hu-berlin.de/papers/pdf/SFB649DP2010-043.pdf>
- Schiller, F., Seidler, G., and Wimmer, M. (2012). Temperature models for pricing weather derivatives. *Quantitative Finance*, 12, 489-500.
- Wheatley, S. (1988). Some tests of international equity integration. *Journal of Financial Economics*, 21, 177-212.
- Xu, W., Odening, M., and Musshoff, O. (2008). Indifference pricing of weather derivatives. *American Journal of Agricultural Economics*, 90, 979-993.
- Zapranis, A. and Alexandridis A. (2008). Modelling the temperature time-dependent speed of mean reversion in the context of weather derivatives pricing. *Applied Mathematical Finance*, 15, 355-386.
- Zapranis, A. and Alexandridis A. (2009). Weather derivatives pricing: Modeling the seasonal residual variance of an Ornstein–Uhlenbeck temperature process with neural networks. *Neurocomputing*, 73, 37-48.
- Zeng, L. (2000a). Pricing weather derivatives. *Journal of Risk Finance*, 1, 72-78.
- Zeng, L. (2000b). Weather derivatives and weather insurance: concept, application and analysis. *Bulletin of the American Meteorological Society*, 81, 2075-2082.

Appendix A: The Monthly Temperature Model

The daily and monthly temperature models differ only in their seasonality terms. For the latter, the seasonality terms of the conditional mean and conditional volatility are

$$\text{Seas}_i^{(m)} = \sum_{p=1}^P \left(\sigma_{c,p} \cos \left(2\pi p \frac{d(i)}{12} \right) + \sigma_{s,p} \sin \left(2\pi p \frac{d(i)}{12} \right) \right)$$

β_0	$\sigma_{c,1}$	$\sigma_{s,1}$	$\gamma_{c,1}$	$\gamma_{s,1}$	c
50.6033 (0.2924)	-19.3814 (0.4005)	-14.4907 (0.4025)	0.4461 (0.1721)	0.4822 (0.1715)	2.3999 (0.1179)

Table 13: Estimates of the parameters in the chosen monthly average temperature model (standard errors are shown in parentheses)

and

$$\text{sq}_i^{(m)} = \sum_{q=1}^Q \left(\gamma_{c,q} \cos \left(2\pi q \frac{d(i)}{12} \right) + \gamma_{s,q} \sin \left(2\pi q \frac{d(i)}{12} \right) \right),$$

respectively, where $d(i)$ cycles through $1, 2, \dots, 12$. The EGARCH effect is not significant when considering monthly temperature data. We set $L = 0$, $M = 0$, $P = 1$, $Q = 1$, $R = 0$ and $S = 0$. The parameter estimates for the monthly temperature model are shown in Table 13.

Appendix B: A Proof of Equation (3)

Proof. When $\mu = 1$, the day- t_1 aggregate dividend (consumption expenditure) can be written as

$$\ln \frac{\delta_{t_1}}{\delta_t} = \alpha(t_1 - t) + \sum_{i=t+1}^{t_1} v_i.$$

Since $v_i \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma^2/(1 - \varphi^2))$, we have

$$\frac{\delta_{t_1}}{\delta_t} \sim \text{log-normal} \left(\alpha(t_1 - t), \frac{(t_1 - t)\sigma^2}{1 - \varphi^2} \right)$$

given δ_t , and hence

$$\mathbb{E}_t \left[\frac{\delta_{t_1}}{\delta_t} \right] = e^{\gamma \alpha(t_1 - t) + \frac{\gamma^2(t_1 - t)\sigma^2}{2(1 - \varphi^2)}}.$$

Using this result, equation (1) can be rewritten as

$$e^{-\frac{\rho(t_1 - t)}{365}} e^{\gamma \alpha(t_1 - t) + \frac{\gamma^2(t_1 - t)\sigma^2}{2(1 - \varphi^2)}} = e^{-\frac{r(t_1 - t)}{365}}.$$

It follows that

$$-\frac{\rho(t_1 - t)}{365} + \gamma \alpha(t_1 - t) + \frac{\gamma^2(t_1 - t)\sigma^2}{2(1 - \varphi^2)} = -\frac{r(t_1 - t)}{365},$$

or equivalently,

$$\rho = r + 365\gamma\alpha + \frac{365\gamma^2\sigma^2}{2(1 - \varphi^2)}.$$

□