Equilibrium Valuation of Weather Derivatives*

Melanie Cao
Department of Economics
Queen's University
Kingston, Ontario, Canada
Canada, K7L 3N6
caom@qed.econ.queensu.ca

Jason Wei
Faculty of Management
University of Toronto
Toronto, Ontario, Canada
M5S 3E6
wei@mgmt.utoronto.ca

First Draft: October, 1998
This Draft: May, 2000

Abstract

This paper proposes and implements an equilibrium valuation framework for weather derivatives. We generalize the Lucas model of 1978 to include the weather as a fundamental variable in the economy. The model is specialized to temperature derivatives. Temperature behavior for the period of 1979-1989 is closely studied for five major cities in the U.S., and a model is proposed for the daily temperature variable which incorporates all the key properties of temperature behavior including seasonal cycles and uneven variations throughout the year. The temperature variable affects the aggregate output both contemporaneously and in a lagged fashion. The temperature system is estimated using the 20-year data and numerical analyses are performed for forward and option contracts on heating degree days (HDD's) and cooling degree days (CDD's). The key advantages of our model include the use of weather forecasts as inputs and the ability to handle contracts of any maturity, for any season. Numerical analyses show that the market price of risk associated with the temperature variable is insignificant in most cases, especially when the aggregate dividend process exhibits mean reversion. The market price of risk becomes important when the risk aversion is high or when the aggregate dividend process is close to a random walk. Finally, we show that the so-called historical simulation method can lead to significant pricing errors due to its erroneous implicit assumptions.

* The previous version of this paper was titled "Pricing Weather Derivatives: an Equilibrium Approach". We are grateful to the Social Sciences and Humanities Research Council of Canada for financial support. We, respectively, also acknowledge Queen's Edith Whyte research grant and the University of Toronto Connaught Fund. We thank workshop participants at the Fields Institute, Michigan State University, Ontario Power Generations Company, Queen's University, and York University, and especially Ted Neave, Raymond Ross and Sam Tian for their comments.

Equilibrium Valuation of Weather Derivatives

Abstract

This paper proposes and implements an equilibrium valuation framework for weather derivatives. We generalize the Lucas model of 1978 to include the weather as a fundamental variable in the economy. The model is specialized to temperature derivatives. Temperature behavior for the period of 1979-1989 is closely studied for five major cities in the U.S., and a model is proposed for the daily temperature variable which incorporates all the key properties of temperature behavior including seasonal cycles and uneven variations throughout the year. The temperature variable affects the aggregate output both contemporaneously and in a lagged fashion. The temperature system is estimated using the 20-year data and numerical analyses are performed for forward and option contracts on heating degree days (HDD's) and cooling degree days (CDD's). The key advantages of our model include the use of weather forecasts as inputs and the ability to handle contracts of any maturity, for any season. Numerical analyses show that the market price of risk associated with the temperature variable is insignificant in most cases, especially when the aggregate dividend process exhibits mean reversion. The market price of risk becomes important when the risk aversion is high or when the aggregate dividend process is close to a random walk. Finally, we show that the so-called historical simulation method can lead to significant pricing errors due to its erroneous implicit assumptions.

1. Introduction

It is estimated that about \$ 1 trillion of the \$ 7 trillion U.S. economy is weather sensitive (Challis 1999 and Hanley 1999). Weather conditions directly affect agricultural outputs and the demand for energy products, and indirectly affect retail businesses. For instance, the inventory of winter coats at a department store depends on weather forecast for the coming winter and the eventual sales depend on the actual weather condition (Agins and Kranhold, 1999). Likewise, earnings in the power industry depend on the retail prices and the sales quantities of electricity. Although weather conditions may not directly determine the retail price of electricity, it is certainly one of the most important factors affecting the demand. Until 1997, earnings stabilization for utility firms was primarily achieved through price hedging mechanisms while volumetric risks were largely left unhedged. However, increasing competition due to ongoing deregulations has made it necessary for companies to hedge the volumetric risk caused by unexpected weather conditions. Such needs have created a new family of over-the-counter weather derivatives. Meantime, the Chicago Mercantile Exchange has introduced futures contracts on the temperature of many U.S. cities. Recently, the London International Financial Futures and Options Exchange (Liffe) and internet companies Wire and Intelligent Financial Systems have joined forces to create a web-based service aimed at establishing trading of European weather derivatives on the internet (*Risk*, March 2000).

The underlying variables for weather contracts include temperature, rainfall, snowfall and humidity.[1] However, the most commonly contracted weather variable in the financial industry is temperature. Specifically, most contracts are written on the so-called heating degree day (HDD)

[1]For a complete survery, see Hanley (1999).

and cooling degree day (CDD) defined on daily average temperatures. The daily average temperature, in turn, is the arithmetic average of the maximum and minimum temperature recorded on a midnight-to-midnight basis. The precise expressions of HDD and CDD are defined below:

Daily HDD = max(65 degree Fahrenheit - daily average temperature, 0);

Daily CDD = max(daily average temperature - 65 degree Fahrenheit, 0).

To simplify the language in the text, we refer to the daily average temperature as daily temperature. Intuitively, HDD measures the coldness of the day compared to a benchmark of 65 Fahrenheit for the winter season while CDD measures the extent to which a summer day is hot. For a typical Northern or Midwest city, an HDD season includes winter months from November to March and the CDD season (or summer season) from May to September. April and October are commonly referred to as the shoulder months.

Because of the high correlation between the electricity consumption and HDD/CDD, most contracts are written on the accumulation of HDD or CDD over a certain period (e.g., a calendar month or a season) so that one contract can be used to hedge a particular period. The popular transactions in the OTC market include HDD / CDD swaps and options for large cities like Atlanta, Chicago, Dallas, New York and Philadelphia. [2] The swap contracts or forward contracts are similar to the exchange-traded futures contracts. There are four basic elements in these contracts: (i) the underlying variable: HDD or CDD; (ii) the accumulation period: a season or a calendar month; (iii) a specific weather station reporting daily temperatures for a particular city; and (iv) the tick size: the dollar amount attached to each HDD or CDD. Table 1 presents the typical transactions of an HDD swap (or forward) for New York and a CDD option for Chicago. In the New York HDD swap, the tick size is set at $5,000 per HDD. XYZ Co. agrees to pay ABC Co. a fixed rate of

[2] See Smithson and Choe (1999) for a brief survey of the market.

1,000 HDD and in return for a floating rate which is the actual accumulated HDD during January, 1999. The realized HDD for January, 1999 is 956. Then the payoff for XYZ Co. at maturity is $\$5000 \times (956 - 1000) = -\$220,000$. Similarly, the Chicago CDD option has a tick size of \$5,000 and a strike level of 190 CDD. The actual accumulated CDD in June, 1999 is 196, which is higher than the strike level. Thus, xyz Co. would exercise the call option at maturity and receive a payoff of $\$5,000 \times (196 - 190) = \$30,000$. These contracts can be used by a power generation company to hedge against revenue losses due to abnormal temperatures.

Table 1: Examples of HDD- and CDD-based Swap and Option

	HDD Swap (or Forward)	CDD Call Option
Location	La Guardia Airport, New York	O'Hare Airport, Chicago
Buyer	YYY Co. (paying fixed rate)	xyz Co. (paying call premium)
Seller	AAA Co. (paying floating rate)	abc Co.
Accumulation Period	January 1 - 31, 1999	June 1 - 30, 1999
Tick Size	\$5,000 per HDD	\$5,000 per CDD
Fixed Rate	1,000 HDD	
Strike Level		190 CDD
Floating Rate	the actual HDD for January, 1999 = 956 HDD	
Settlement Price		the actual CDD for June, 1999 = 196 CDD
Payoffs at Maturity for the Buyer	$(956 - 1000) \times 5000 = -\$220,000$	$(196 - 190) \times 5000 = \$30,000$

Despite the rapid growth of weather derivatives, the bid / ask spread is still very large and there is not yet an effective pricing method. In addition, some key questions are yet to be answered. For example, insofar as weather variables are not tradeable, is the market price of risk a significant factor in valuations?

The objective of this paper is four-fold. First, we will propose a general equilibrium model to price weather derivatives, and we specialize it to temperature derivatives. Second, we will present a realistic model for the dynamics of the daily average temperature, which is very different from that of a security price. For example, temperatures are seasonal and cyclical, can be predicted with reasonable accuracy for the very near future, and will vary within a well-defined range in the long run. Third, we would like to establish whether the market price of risk associated with the temperature variable significantly affects the valuation of weather derivatives. Fourth, we will develop a pricing framework for derivatives based on the accumulated HDD or CDD. Here, challenges arise because the underlying is the accumulation of the daily HDD/CDD which are non-linear in daily temperatures.

The key contributions of this paper lie in the accomplishments of the aforementioned objectives. First, we extend Lucas' (1978) equilibrium asset-pricing model where fundamental uncertainties in the economy are generated by aggregate dividend and a state variable representing the weather condition. When specialized to temperature derivatives, it is shown that the equilibrium derivative prices depend on the agent's risk preference and the correlations, both contemporaneous and lagged, between the temperature variable and the aggregate dividend. Second, we propose an auto-regressive, mean-reverting dynamic system for the daily temperature. This system is capable of capturing the seasonality and the global warming trend and can incorporate weather forecast. We use Maximum Likelihood method to estimate the temperature dynamics for Atlanta, Chicago, Dallas, New York and Philadelphia. Lastly, based on the estimated parameters, we simulate derivative prices and analyze in an equilibrium framework the market price of risk associated with the temperature variable. We attempt to answer an important question: is it valid to ignore the market price of risk when pricing weather derivatives, as the industry commonly does?

The paper is organized as follows. Section 2 contains the general setup of the economy. Section 3 specializes the setup to a constant relative risk aversion preference and an autoregressive temperature variable. Section 4 sets out the formulas and key results for HDD and CDD contracts. Estimations for the temperature process for five U.S. cities are presented in Section 5. Numerical analyses with an emphasis on market price of risk are presented in Section 6. A brief summary and concluding remarks are given in Section 7. Proofs and exhibits are collected in appendices.

2. The Setup of the Economy

In a discrete-setting, consider an extension of the Lucas (1978) pure exchange economy where the fundamental uncertainties in the economy are driven by two state variables: the aggregate dividend (δ) and the weather condition (Y). Aggregate dividends can be viewed as aggregate outputs or dividends on the market portfolio; the weather variable could be temperature, rainfall, snowfall, or humidity. The dynamics governing the aggregate dividend and the weather variable are exogenous processes on a given probability space (- , $\mathcal{F}$, $\mathcal{P}$). There is a representative investor whose information structure is given by the filtration $\mathcal{F}_t \equiv \sigma(\delta_\tau, Y_\tau;\ \tau \in (0, 1, 2, ..., t))$. The agent has an infinite lifetime horizon. In the financial market, the representative agent can trade a single risky stock, pure discount bonds and a finite number of other contingent claims at any time. The risky stock can be viewed as the market portfolio. Therefore, its dividend stream $\{\delta_t\}$ is understood as aggregate dividends in the economy. The total supply is normalized to one share and the contingent claims are written on the risky stock, the pure discount bond or the weather variable. The net supply of all contingent claims and the riskless bond is zero.

The agent's preference is described by a smooth time-additive expected utility function

$$V(c) = E_0 \left(\sum_{t=0}^{\infty} U(c_t, t) \right), \tag{2.1}$$

where $U : \mathcal{R}_+ \times (0, \infty) \to \mathcal{R}$ is smooth on $(0, \infty) \times (0, \infty)$ and, for each $t \in (0, 1, 2, ..., \infty)$, $U(\cdot, t) : \mathcal{R}_+ \to \mathcal{R}$ is increasing, strictly concave, and has a continuous derivative $U_c(\cdot, t)$ on $(0, \infty)$. Initially, the agent is endowed with one share of the risky stock. Denote his portfolio holdings at time t as $\theta_t = (\theta_t^s,\, \theta_t^B,\, \theta_t^{x\prime})$, where θ_t^s, θ_t^B and $\theta_t^{x\prime}$ represent the number of shares invested in the risky stock, the discount bond and other contingent claims, respectively. Also denote the security prices at time t by a vector X_t and the corresponding vector of dividends by D_t. The agent's consumption over time is financed by a trading strategy $\{\theta_t, t \geq 0\}$. His decision problem is to choose an optimal trading strategy so as to maximize his expected lifetime utility. The first order conditions yield the standard Euler equation:

$$X_t = E_t \left(\sum_{\tau=t+1}^{\infty} \frac{U_c(c_\tau, \tau)}{U_c(c_t, t)} D_\tau \right). \tag{2.2}$$

Thus, the price of any security equals the sum of expected dividends, discounted at the stochastic marginal rate of substitution.

In equilibrium, both the financial market and the goods market clear so that aggregate consumption equals the dividends generated from the risky stock. Therefore, the time t price of a contingent claim with a payoff q_T at a future time T, denoted by $C_t(t, T)$, is

$$C_t(t, T) = \frac{1}{U_c(\delta_t, t)} E_t \left(U_c(\delta_T, T) q_T \right), \quad \forall \quad t \in (0, T). \tag{2.3}$$

In particular, at time t the equilibrium price of a riskless bond paying 1 unit of consumption goods at T and 0 at all other times, is

$$B(t,T) = \frac{1}{U_c(\delta_t,t)} E_t\left(U_c(\delta_T,T)\right), \quad \forall \quad t \in (0,T). \tag{2.4}$$

Contingent claims based on a weather variable can be valued via (2.3) once the agent's preference, the dividend process, and the weather variable process are specified. In the following, we will specialize the above setup to temperature related derivatives since they are by far the most commonly traded products.

3. Valuation Framework for Temperature Derivatives

3.1. Dynamics of the Temperature Variable

3.1.1. Description of Temperature Behavior

To ensure accurate modeling of the temperature variable, we first examine its behavior. Historical daily temperature data, covering the period from 1979 to 1998 (inclusive), for Atlanta, Chicago, Dallas, New York and Philadelphia are obtained from the National Climate Data Center (NCDC), a subsidiary of the National Oceanic Atmospheric Administration (NOAA). Exhibits 1 and 2 summarize sample statistics and Exhibit 3 depicts the warming trend in Atlanta, which is typical of all cities studied.[3] The following remarks are in order.

Remark 1. *The sample means of the two Southern cities (Atlanta and Dallas) are higher than those of the three Northern counterparts. The highest and lowest sample means are 66 and 50 of Dallas and Chicago, respectively.*

[3] To simplify the analysis, we have omitted the observations for February 29 from the sample. Therefore, each year consists of 365 days and the sample size for 20 years is 7300.

Remark 2. *Northern cities generally have larger standard deviations. Chicago has the highest sample standard deviation (20 degrees), indicating large temperature swings. Atlanta has the lowest sample standard deviation (15 degrees).*

Remark 3. *Correlations among the five cities are very high and are above 0.84. New York and Philadelphia, the two nearest cities, present the highest correlation, 0.9853.*

Remark 4. *Daily temperatures exhibit strong auto-correlations.*

Remark 5. *Standard deviations of monthly CDD's for the two Southern cities are higher than those of their Northern counterparts. The reserve is true for HDD standard deviations (Exhibit 2).*

To facilitate further discussions, let us index the years in the sample period by yr, thus $yr = 1$ for 1979, $yr = 2$ for 1980, ..., $yr = 20$ for 1998. Also, index January 1 as $t = 1$, January 2 as $t = 2$, and so on for 365 days in a year. Denote $Y_{yr,t}$ as the temperature on date t in year yr. Below, we define the mean $(\overline{Y}_t)$ and the standard deviation (ψ_t) for date t as

$$\overline{Y}_t = \frac{1}{20} \sum_{yr=1}^{20} Y_{yr,t} \qquad \& \qquad \psi_t = \sqrt{\frac{1}{20} \sum_{yr=1}^{20} \left(Y_{yr,t} - \overline{Y}_t\right)^2}; \qquad \forall \quad t = 1, 2, ..., 365.$$

We plot in Exhibit 4 the daily standard deviations for Atlanta and Chicago, which show a clear seasonal pattern: the temperature variation in the HDD season is larger than that in the CDD season. This is common for all cities in consideration.

3.1.2. Modelling Daily Temperature Behavior

In light of the properties identified in the previous section, a model for the daily temperature must possess the following features. First, it must capture the seasonal cyclical patterns; second, the daily variations in temperature must be around some average "normal" temperature, to be elaborated on later; third, it should allow forecasts to play a key role in projecting temperature paths in the future; fourth, it should incorporate the autoregressive property in temperature changes (i.e., a warmer day is most likely to be followed by another warmer day, and vice versa); fifth, the extent of variation must be bigger in the winter and smaller in the summer; sixth, a projected temperature path into the future should not wander outside of the normal range of the temperature for each projected point in time (for instance, it is unlikely for a summer day in Chicago to see a temperature of -10 Fahrenheit).; and seventh, the model must reflect the global warming trend.

Although a diffusion process, especially a mean-reverting process, is capable of accommodating most of the required features, we decide against it for one key reason: a one-factor diffusion process can not incorporate autocorrelations in temperature innovations, especially for lags beyond one. In addition, with a Markov diffusion, there is a non-zero probability that a particular path of temperatures does not resemble a temperature evolution at all (violating the sixth requirement). For these reasons, we resort to a discrete, autoregressive model. To this end, define the de-meaned and de-trended residual of the daily temperature as $U_{yr,t}$,

$$U_{yr,t} = Y_{yr,t} - \widehat{\widehat{Y}}_{yr,t}, \quad \forall \quad yr = 1, 2, ..., 20 \;\&\; t = 1, 2, ..., 365. \tag{3.1}$$

Assumption 1. *The daily temperature residual, $U_{yr,t}$, follows a k-lag autocorrelation system:* [4]

$$\begin{aligned} & U_{yr,t} = \textstyle\sum_{i=1}^{k} \rho_i U_{yr,t-i} + \sigma_{yr,t} * \xi_{yr,t} \\ & \sigma_{yr,t} = \sigma_0 - \sigma_1 \mid \sin(\pi t/365 + \phi) \mid, \\ & \xi_{yr,t} \sim i.i.d.\ N(0,1), \\ & \forall \quad yr = 1, 2, ..., 20; \quad \& \quad t = 1, 2, ..., 365. \end{aligned} \tag{3.2}$$

where $\xi_{yr,t}$ models the randomness in the temperature changes.

[4] We confess a slight abuse of notation here. Notice that at the beginning of year yr, we must use the data from the end of the previous year, $(yr - 1)$ to calculate the auto-regressive terms. It is understood that the index yr will automatically take appropriate values when required.

In the above, the volatility specification using the sine wave reflects the fifth requirement and the feature in Exhibit 4. The parameter ϕ captures the proper starting point of the sine wave. The autocorrelation setup reflects the fourth feature. The other features (including the global warming trend) are captured by the specification of $\widehat{\overline{Y}}_t$,which we delineate next.

The variable $\widehat{\overline{Y}}_t$ serves the purpose of de-meaning and de-trending. Ideally, it should reside in the middle of the band that contains temperature fluctuations. It is therefore tempting to use the historical daily average, $\overline{Y}_t$ as a proxy for $\widehat{\overline{Y}}_t$. But this will be a poor choice, because for a particular year, all the realized temperatures could be below or above the historical averages, representing a very cold or very warm year. This point is illustrated in Panel A of Exhibit 5 for New York. It is seen that although the historical average $\overline{Y}_t$ was more or less in the center of realized temperatures before the winter months of 1980, it was well above the realized temperatures for the last part of 1980 and the early part of 1981. This suggests that we could use $\overline{Y}_t$ as a starting point and make some season specific adjustments so that the eventual anchor points would be roughly in the middle of the variation band.

We propose to adjust the historical average $\overline{Y}_t$ in the following steps. 1) For each month of the year, we calculate the average of the daily averages $\overline{Y}_t$, and there will be twelve such monthly averages; 2) for each particular year, we calculate the realized, average temperature of each month; 3) for each month, we find the difference between the actual monthly average from Step 2 and the average from Step 1; and 4) for each day of the month, we adjust the historical average $\overline{Y}_t$ by the quantity calculated in Step 3, and this adjusted average is $\widehat{\overline{Y}}_{yr,t}$ and will be referred to as the adjusted mean temperature. [5] To illustrate the adjustment mechanism, suppose our concern

[5] The period of one month is chosen as a trade-off. Too long a period will not solve the non-centering problem and too short a period will unnecessarily exaggerate the short term fluctuations and diminish the meaning of "average"

is March 1986, and suppose the mean of the average daily temperature for the month of March, calculated from Step 1, happens to be 50 F. We now calculate the average of the 31 realized daily temperatures for March 1986 and suppose it is 45F, which indicates a colder than normal March. This is the average from Step 2. We then follow Step 3 to find the difference between the two averages: 45 F - 50 F = -5 F. Finally, following Step 4, we adjust each of the historical average temperature $\overline{Y}_t$ for March by -5 F. Suppose the historical daily average temperatures for March 1, 2, 3, 4, ..., 31 are 48 F, 53 F, 50 F, 55 F, ...,60 F, then the adjusted mean averages for March 1, 2, 3, 4, ..., 31 of 1986 will be 43 F, 48 F, 45 F, 50 F, ..., 55 F. The $\widehat{\overline{Y}}_{yr,t}$ will assume these values for March 1986 in actual estimations. We postpone estimation details to Section 5. [6]

3.2. Agent's Preference and Aggregate Dividend Behavior

For analytical tractability, the agent is assumed to have a constant relative risk aversion:

Assumption 2. *The representative agent's period utility is described by*

$$U(c_t, t) = e^{-\rho t} \frac{c_t^{\gamma+1}}{\gamma+1}, \tag{3.3}$$

with the rate of time preference, $\rho > 0$ and the risk parameter $\gamma \in (-\infty, 0]$. [7]

For the aggregate dividend process, we appeal to Marsh and Merton (1987), whose estimation results suggest mean-reversion in the rate of aggregate dividend changes. Specifically,

or "mean". Needless to say, one could get more sophisticated in making the adjustments. For instance, rather than following the calendar months, one could always center the day in question in a 30 day (e.g.) period and make the above adjustments on a rolling basis. But as shown in Panel B of Exhibit 5, the simple adjsutment already works well.

[6] In the valuation context which by necessity is forward looking, $\widehat{\overline{Y}}_{yr,t}$ can naturally be considered as daily temperature forecasts. If the forecasts were of 100% accuracy, then weather derivatives won't exist since perfect planning is achievable. In reality, it is precisely the uncertainty in the forecasts that drives the value of weather derivatives. The random term in the temperature dynamic captures this uncertainty. Indeed, this is one of the key advantages of our model: it allows forecasts and their uncertainties to be built into the valuation of derivatives.

[7] Hereinafter, for the consumption and dividend variables, we use the time subscript "t" in the usual way (i.e. natural progression through time) without refering to a specific year.

Assumption 3. *The aggregate dividend, δ_t, evolves according to the following Markov process:*

$$\ln \delta_t = \alpha + \mu \ln \delta_{t-1} + \nu_t, \quad \forall \quad \mu \leqslant 1 \tag{3.4}$$

where $1 - \mu$ measures mean reversion, and ν_t is the error term and takes the following form

$$\nu_t = \sigma \epsilon_t + \sigma[\frac{\varphi}{\sqrt{1-\varphi^2}} \xi_t + \eta_1 \xi_{t-1} + \eta_2 \xi_{t-2} + \eta_3 \xi_{t-3} + ... + \eta_m \xi_{t-m}], \qquad 0 \leqslant m \leqslant +\infty \tag{3.5}$$

In the above, ϵ_t is an i.i.d. standard normal variable which captures the randomness due to all factors other than the temperature; ξ_t and its lagged terms are innovations of the temperature variable as defined in (3.2). By construction, the contemporaneous correlation between the dividend process and the temperature process is φ. The lagged terms capture the lagged effects of the temperature on the aggregate dividend or output of the economy. By necessity and assumption, $\sum_{j=1}^{m} \eta_j^2$ ($\forall$ m) is bounded. When t represents a future time and when all the lags are beyond the present time, the conditional variance of ν_t is $\sigma^2 \left[1 + \frac{\varphi^2}{1-\varphi^2} + \sum_{j=1}^{m} \eta_j^2\right]$, which breaks down the overall variation in the aggregate dividend into three parts: the part due to all factors other than the temperature, σ^2; the part due to the contemporaneous impact of the temperature, $\sigma^2 \frac{\varphi^2}{1-\varphi^2}$; and the part due to lagged impact of the temperature, $\sigma^2 \sum_{j=1}^{m} \eta_j^2$. The correlation between the dividend innovation at time t and the temperature innovation at time $t - j$ $\forall$ j is $\eta_j \ / \ \left[1 + \frac{\varphi^2}{1-\varphi^2} + \sum_{j=1}^{m} \eta_j^2\right]$. When $\varphi = 0$ and $\eta_j = 0, \forall$ j, the dividend process is totally independent of the temperature innovation.

The specifications for the representative agent's preference in (3.3), the dividend process in (3.4) and (3.5), and temperature dynamics in (3.1) and (3.2), together with the general pricing equation in (2.3), will enable us to derive a value for any claim contingent upon the temperature variable. In the next section, we will apply the framework to HDD / CDD derivatives. However, we will first state below a general result for the interest rates.

3.3. Term Structure of Interest Rates

Proposition 3.1. *Under the CRRA utility in (3.3) and the dividend process in (3.4) and (3.5), the price of a pure discount bond at time t with maturity T is*

$$B(t,T) = e^{-\rho(T-t)} \exp\left(\gamma\Upsilon(t,T) + \frac{1}{2}\gamma^2\Sigma(t,T)\right) \tag{3.6}$$

where

for $T - t \geq m + 1$

$\Upsilon(t,T) = \alpha \sum_{i=t+1}^{T} \mu^{T-i} + (\mu^{T-t} - 1)\ln\delta_t + \mu^{T-t-m}\sigma\left(\sum_{i=1}^{m} \mu^{i-1}(\sum_{j=1}^{i} \eta_{m-j+1}\widehat{\xi}_{t-i+j})\right)$

$\Sigma(t,T) = \sigma^2\left(\sum_{i=t+1}^{T} \mu^{2(T-i)} + (\sum_{i=0}^{m} \eta_i\mu^{m-i})^2(\sum_{i=t+1}^{T-m} \mu^{2(T-m-i)}) + \sum_{i=0}^{m-1}(\sum_{j=0}^{i} \eta_j\mu^{i-j})^2\right)$

and

for $T - t \leq m$

$\Upsilon(t,T) = \alpha \sum_{i=t+1}^{T} \mu^{T-i} + (\mu^{T-t} - 1)\ln\delta_t + \sigma\left(\sum_{i=1}^{\tau} \mu^{\tau-i}(\sum_{j=i}^{m} \eta_j\widehat{\xi}_{t+i-j})\right)$

$\Sigma(t,T) = \sigma^2\left(\sum_{i=t+1}^{T} \mu^{2(T-i)} + \sum_{i=0}^{\tau-1}(\sum_{j=0}^{i} \eta_j\mu^{i-j})^2\right)$

and $\widehat{\xi}_{t-l}$ $(0 \leqslant l \leqslant m)$ are the realized error terms for the temperature process.

The yield to maturity defined via $e^{-R(t,T)(T-t)} = B(t,T)$ is

$$R(t,T) = -\frac{\ln B(t,T)}{T-t} = \rho - \frac{\gamma\Upsilon(t,T) + \frac{1}{2}\gamma^2\Sigma(t,T)}{T-t}. \tag{3.7}$$

Proof: see Appendix A.

For the special case $\mu = 1$ where changes in the dividend growth rate follow a random walk, the terms $\Upsilon(t,T)$ and $\Sigma(t,T)$ reduce to

for $T - t \geq m + 1$

$\Upsilon(t,T) = \alpha(T-t) + \sigma\left(\sum_{i=1}^{m}(\sum_{j=1}^{i} \eta_{m-j+1}\widehat{\xi}_{t-i+j})\right)$

$\Sigma(t,T) = \sigma^2\left(T - t + (\sum_{i=0}^{m} \eta_i)^2(T-t-m) + \sum_{i=0}^{m-1}(\sum_{j=0}^{i} \eta_j)^2\right)$

and

for $T - t \leq m$

$$\Upsilon(t,T) = \alpha(T-t) + \sigma\left(\sum_{i=1}^{\tau}(\sum_{j=i}^{m}\eta_j\widehat{\xi}_{t+i-j})\right)$$

$$\Sigma(t,T) = \sigma^2\left(T - t + \sum_{i=0}^{\tau-1}(\sum_{j=0}^{i}\eta_j)^2\right)$$

The parameters in the bond price formula must be collectively restricted so that the yield in (3.7) is positive.

4. Valuing HDD / CDD Derivatives

4.1. Valuation of HDD / CDD Forward Contracts and Options

Consider an HDD forward contract with a tick size of \$1 and delivery price, K. The accumulation period starts at T_1 and ends at $T_2 > T_1$. Denote $HDD(T_1,T_2) = \sum_{\tau=T_1}^{T_2}\max(65 - Y_\tau, 0)$. [8] Then, by (2.3) and (3.3), the value of the HDD forward contract at time t, $f_{HDD}(t,T_1,T_2,K)$, can be expressed as

$$\begin{aligned} f_{HDD}(t,T_1,T_2,K) &= E_t\left(\frac{U_c(\delta_{T_2},T_2)}{U_c(\delta_t,t)}[HDD(T_1,T_2) - K]\right) \\ &= e^{-\rho(T_2-t)}E_t\left(\frac{\delta_{T_2}^{\gamma}}{\delta_t^{\gamma}}[HDD(T_1,T_2) - K]\right). \end{aligned} \tag{4.1}$$

By definition, the forward price at time t, $F_{HDD}(t,T_1,T_2)$, is the value of K which makes $f = 0$. That is,

$$F_{HDD}(t,T_1,T_2) = \frac{E_t\left(\delta_{T_2}^{\gamma}HDD(T_1,T_2)\right)}{E_t(\delta_{T_2}^{\gamma})} = \frac{e^{-\rho(T_2-t)}\delta_t^{-\gamma}E_t\left(\delta_{T_2}^{\gamma}HDD(T_1,T_2)\right)}{B(t,T_2)}. \tag{4.2}$$

Similar expressions can be derived for CDD contracts:

$$\begin{aligned} f_{CDD}(t,T_1,T_2,K) &= e^{-\rho(T_2-t)}E_t\left(\frac{\delta_{T_2}^{\gamma}}{\delta_t^{\gamma}}[CDD(T_1,T_2) - K]\right), \\ F_{CDD}(t,T_1,T_2) &= \frac{e^{-\rho(T_2-t)}\delta_t^{-\gamma}E_t\left(\delta_{T_2}^{\gamma}CDD(T_1,T_2)\right)}{B(t,T_2)}. \end{aligned} \tag{4.3}$$

[8] For brevity, we will drop the index "yr" for the temperature variable, Y. For example, Y_τ is understood as the daily temperature of day τ for a particular year. For our simulations to be presented later, $yr = 1999$.

Now consider a European option written on $HDD(T_1, T_2)$ with maturity T_2 and a strike price X. Denote the call and put prices at time t as $C_{HDD}(t, T_1, T_2, X)$ and $P_{HDD}(t, T_1, T_2, X)$, respectively. Again, by (2.3) and (3.3), the call and put values can be expressed as

$$C_{HDD}(t, T_1, T_2, X) = e^{-\rho(T_2-t)}\delta_t^{-\gamma} E_t \left(\delta_{T_2}^{\gamma} \max(HDD(T_1, T_2) - X, 0) \right), \tag{4.4}$$

$$P_{HDD}(t, T_1, T_2, X) = e^{-\rho(T_2-t)}\delta_t^{-\gamma} E_t \left(\delta_{T_2}^{\gamma} \max(X - HDD(T_1, T_2), 0) \right). \tag{4.5}$$

Similarly, call and put options written on $CDD(T_1, T_2)$ can be priced as

$$C_{CDD}(t, T_1, T_2, X) = e^{-\rho(T_2-t)}\delta_t^{-\gamma} E_t \left(\delta_{T_2}^{\gamma} \max(CDD(T_1, T_2) - X, 0) \right), \tag{4.6}$$

$$P_{CDD}(t, T_1, T_2, X) = e^{-\rho(T_2-t)}\delta_t^{-\gamma} E_t \left(\delta_{T_2}^{\gamma} \max(X - CDD(T_1, T_2), 0) \right). \tag{4.7}$$

4.2. Market Price of Risk

Without further restrictions to the dividend and temperature processes, it is virtually impossible to obtain closed-form solutions to the pricing formulas derived above. However, it is possible to make some general statements about the relevance of market price of risk.

Proposition 4.1. *The risk premium in the value of a derivative security is zero if the dividend process and the temperature process are completely independent, i.e., $\varphi = 0$ and $\eta_j = 0, \forall\ j$. In this case, any contingent claim can be valued by discounting its payoff at the risk-free rate.*

Proof: combining (2.3), (2.4), and (3.3) by requiring $\varphi = 0$ and $\eta_j = 0, \forall\ j$ in (2.3) and (2.4) leads to the result.

In certain special cases, it is possible to make some specific statements about the market price of risk for forward prices. Let's take the forward price for HDD in (4.2) as an example. The forward price can be re-written as

$$F_{HDD}(t,T_1,T_2) = \frac{E_t\left(\delta_{T_2}^{\gamma} HDD(T_1,T_2)\right)}{E_t(\delta_{T_2}^{\gamma})} = E_t\left(HDD(T_1,T_2)\right) + \frac{Cov\left(\delta_{T_2}^{\gamma},\ HDD(T_1,T_2)\right)}{E_t(\delta_{T_2}^{\gamma})}, \quad (4.8)$$

where $cov(\cdot,\cdot)$ stands for covariance. The first term represents the expected future spot value of HDD, and the second term represents forward premium. Similar results can be obtained for the CDD forward price. Clearly, (4.8) is consistent with Proposition 4.1. Realizing that HDD (CDD) is negatively (positively) related to the temperature, we can easily summarize via (4.8) the relationship between the forward prices and the future expected spot prices in the following table.

Table 2: Comparative Statics

$\varphi < 0,\ \eta_j \leqslant 0, \forall\ j$	$\varphi > 0, \eta_j \geqslant 0, \forall\ j$
$F_{HDD} < E_t\left(HDD(T_1,T_2)\right)$	$F_{HDD} > E_t\left(HDD(T_1,T_2)\right)$
$F_{CDD} > E_t\left(CDD(T_1,T_2)\right)$	$F_{CDD} < E_t\left(CDD(T_1,T_2)\right)$

Further more, when the lagged correlations between the aggregate dividend and the temperature, and the autocorrelations in the tempertuare process are all zero (i.e., $\eta_j = 0, \forall\ j$ and $\rho_i = 0$ $\forall\ i$), it is possible to obtain closed form formulas for the forward prices.

Proposition 4.2. *When $\eta_j = 0, \forall\ j$ and $\rho_i = 0\ \forall\ i$, the equilibrium forward prices at time $t < T_1$ before the accumulation period are*

$$F_{HDD}(t,T_1,T_2) = \sum_{\tau=T_1}^{T_2}\left([65-\mu'_Y(\tau)]\cdot N\left(\frac{65-\mu'_Y(\tau)}{\sigma_{yr,\tau}}\right) + \frac{\sigma_{yr,\tau}}{\sqrt{2\pi}}\exp\left[-\frac{(65-\mu'_Y(\tau))^2}{2\sigma_{yr,\tau}^2}\right]\right); \quad (4.9)$$

$$F_{CDD}(t,T_1,T_2) = \sum_{\tau=T_1}^{T_2}\left([\mu'_Y(\tau)-65]\cdot N\left(\frac{\mu'_Y(\tau)-65}{\sigma_{yr,\tau}}\right) + \frac{\sigma_{yr,\tau}}{\sqrt{2\pi}}\exp\left[-\frac{(65-\mu'_Y(\tau))^2}{2\sigma_{yr,\tau}^2}\right]\right); \quad (4.10)$$

$$\textit{with} \quad \mu'_Y(\tau) = \widehat{\overline{Y}}_\tau + \gamma\varphi\mu^{\tau-t}\sigma/\sqrt{1-\varphi^2}\sigma_{yr,\tau}\ .$$

The equilibrium forward prices at time $t \in (T_1, T_2)$ during the accumulation period are

$$F_{HDD}(t, T_1, T_2) = HDD(T_1, t) + \sum_{\tau=t+1}^{T_2} \left([65 - \mu'_Y(\tau)] \cdot N(\frac{65 - \mu'_Y(\tau)}{\sigma_{yr,\tau}}) + \frac{\sigma_{yr,\tau}}{\sqrt{2\pi}} \exp\left[-\frac{(65 - \mu'_Y(\tau))^2}{2\sigma^2_{yr,\tau}} \right] \right);$$

$$F_{CDD}(t, T_1, T_2) = CDD(T_1, t) + \sum_{\tau=t+1}^{T_2} \left([\mu'_Y(\tau) - 65] \cdot N(\frac{\mu'_Y(\tau) - 65}{\sigma_{yr,\tau}}) + \frac{\sigma_{yr,\tau}}{\sqrt{2\pi}} \exp\left[-\frac{(65 - \mu'_Y(\tau))^2}{2\sigma^2_{yr,\tau}} \right] \right).$$

The equilibrium values of the forward contracts at time t are

$$\begin{aligned} f_{HDD}(t, T_1, T_2, K) &= B(t, T_2)(F_{HDD}(t, T_1, T_2) - K) \quad \text{and} \\ f_{CDD}(t, T_1, T_2, K) &= B(t, T_2)(F_{CDD}(t, T_1, T_2) - K). \end{aligned}$$

Proof: (see Appendix B).

Remark 6. *The equilibrium value of an HDD/CDD forward contract in Proposition 4.2 is the present value of the difference between the forward price and the delivery price discounted at the riskfree rate, which is consistent with results for stock or currency forward contracts when interest rates are non-stochastic.*

It is seen from (4.9) and (4.10) that when $\varphi = 0$, the parameters pertaining to the dividend process completely drops out, and it can be verified that the forward prices in this case are simply the expected future spot prices:

$$F_{HDD}(t, T_1, T_2) = E_t(HDD(T_1, T_2)) \quad \text{and} \quad F_{CDD}(t, T_1 T_2) = E_t(CDD(T_1, T_2)).$$

No risk premium is required. This corroborates the general prediction of Proposition 4.1.

5. Empirical Estimation

The setup in Section 3 calls for joint estimation of the aggregate dividend and the daily temperature processes. The aggregate dividend can be approximated by such macroeconomic variables as GNP or aggregate consumption. Unfortunately, the frequency of such data is usually low (at most monthly), making the joint estimation impossible since the temperature dynamic is daily. To get around this difficulty, we will independently estimate the temperature process, and then perform simulations for different scenarios of risk aversion and dividend process parameters.

Let Θ be the vector containing all the parameters (ρ_1, ρ_2, ρ_3, ..., ρ_k, σ_0, σ_1, ϕ) pertaining to the temperature process, then the log-likelihood function is

$$l(\Theta; Y) = -\frac{1}{2} \sum_{yr=1}^{20} \sum_{t=1}^{T} \left(\frac{[Y_{yr,t} - E_{yr,t-1}(Y_{yr,t})]^2}{\sigma_{yr,t}^2} + \ln\left(2\pi\sigma_{yr,t}^2\right) \right)$$

with

$$\begin{aligned} E_{yr,t-1}(Y_{yr,t}) &= \widehat{\widehat{Y}}_{yr,t} + \sum_{i=1}^{k} \rho_i U_{yr,t-i}; \qquad \sigma_{yr,t} = \sigma_0 - \sigma_1 \mid \sin(\pi t/365 + \phi) \mid, \\ \forall \quad yr &= 1, 2, ..., 20; \quad \& \quad t = 1, 2, ..., 365 \end{aligned}$$

Given the large sample size, the estimation variances can be computed based on the asymptotic distribution of $\widehat{\Theta}$:

$$\sqrt{T}(\widehat{\Theta} - \Theta) \quad \tilde{} \quad N[0, \quad \Im(\Theta)^{-1}] \quad with \quad \Im(\Theta) \equiv \lim_{T\to\infty} -E\left(\frac{1}{T}\frac{\partial^2 l(\theta; Y)}{\partial\theta\partial\theta'}\right).$$

In order to determine k, the number of lags, we estimate the system sequentially for $k = 1, 2, 3, ...$, and perform maximum likelihood ratio tests (i.e. χ^2 tests) along the way. We stop when the maximum likelihood value ceases to improve. It turns out that three lags describe the data the best. For brevity, we only report the estimation and testing results for $k = 3$, 4 in Exhibit 6.

Several observations are in order. First, for $k = 3$, almost all parameters are estimated with very low standard errors, implying the proper specification of the estimation system. This is by no means a fluke since our specification is guided by the properties of the daily temperature. Second, standard errors of the parameter σ_1 is very small, implying the appropriateness of using the sine wave to fit the overall volatility structure. (When the system is estimated by specifying a constant volatility throughout the year (i.e. $\sigma_1 = 0$), the likelihood value is much smaller.) Third, the first order auto-regressive behavior tends to be stronger for Southern cities, and ρ_1 has the highest value for Atlanta. Roughly, a stronger auto-correlation means less dramatic changes in temperature, and vice versa. As shown in Exhibit 1, Atlanta does have the lowest overall standard deviation in the sample period. Fourth, the parameter ϕ in the sine function is negative for all cities, indicating that the coldest days in the winter tend to come after January 1. For example, $\phi = -0.2014$ for Chicago, which means that, on average, the coldest day is on January 23 ($0.2014 * 365/\pi \simeq 23$). [9]

6. Numerical Analysis

6.1. Simulation Design

To begin with, we will require some general guidance on setting the parameter values for simulation purposes. First, what value do we assume for μ, the mean reversion parameter for the dividend process? Shiller (1983) estimated μ to be 0.807. Marsh and Merton (1987), in a series of estimations, estimated μ to be as high as 0.945. We will set $\mu = 0.9$, roughly as a middle point. Sensitivity analysis will be performed later for other values of μ.

[9] It should be pointed out that, strictly speaking, using the so called mean adjusted average $\widehat{\overline{Y}}_\tau$ as inputs in the estimation will lead to an overfitting problem. However, since our primary objective is to estimate the volatility and autoregressive behaviour of the residual, the potential overfitting should not cause a fundamental problem. In valuations which we will turn to next, $\widehat{\overline{Y}}_\tau$ are treated as forecasts.

Second, how many lagged error terms to keep in (3.5)? We will simulate two cases, one with only the contemporaneous correlation and the other with 90 lagged error terms. We will adopt a simple geometric decay function for the coefficients, $\eta_j \ \forall\ j$, to be discussed later.

Third, what value to assume for volatility, σ? Given a structure of η_j and a contemporaneous correlation φ which we will use as comparative static variable, we will set σ so that the overall volatility of the dividend process, $\sigma\sqrt{1+\frac{\varphi^2}{1-\varphi^2}+\sum_{j=1}^{90}\eta_j^2}$, is 20%, a magnitude similar to that of a stock market index.

Fourth, what about the rate of time preference, ρ? Since it is typically close to the real riskfree interest rate, we will set it at 0.03.

Fifth, the risk aversion parameter, γ will also be used as a comparative static variable. Given the choice of volatility parameters and a value for γ, we will set the average dividend growth rate α and the initial dividend δ_t according to (3.7) such that the riskfree interest rate or yield is maintained at 6%. Keeping the same level of interest rate across different scenarios will ensure meaningful comparisons and analyses. Note that as long as the yield is fixed, the choices of the combination of α and δ_t are not limited. We arbitrarily set α at 7%, and solve for δ_t. The results are invariant to the combinations, due to the fact that only the ratio of the future dividend over today's dividend matters in our valuations, and the fact that both α and δ_t only contribute to the mean of the ratio. It is apparent from (3.6) and (3.7) that fixing the yield is equivalent to fixing the mean of the ratio once other parameter values are given.

Finally, we must decide on the inputs for the weather forecasts $\widehat{\widehat{Y}}_{yr,t}$. Since our primary objective is to study the price behavior (as opposed to accurate valuation), we will simply use the historical daily average temperature $\overline{Y}_t$.[10]

With the above in place, the simulation procedure can be summarized as: 1) generating bivariate paths for the dividend process in (3.4) and (3.5), and the daily temperature process in (3.1) and (3.2) (using the parameters in Exhibit 6 and $k = 3$); 2) tracking realized HDD/CDD values of each path; 3) calculating the discounted payoff of the derivative security in question according to the formulas in Section 4.1; and 4) repeating steps 1 through 3 a large number of times (i.e., 10,000) and averaging the payoffs to obtain the desired derivative security value.

To reduce simulation errors, we employ the antithetic variable technique. In addition, we also employ a procedure equivalent to the control variate technique. Notice that the fundamental variable in our framework is the daily temperature, and the underlying variable for most weather derivatives is HDD or CDD, which is a nonlinear function of daily temperatures. While our model will produce almost "unbiased" temperature forecasts in that the average temperature for a future point will be almost equal to the input forecast, it can not guarantee an unbiased forecast for the HDD / CDD. To further illustrate this point, we compare some forward prices in Exhibit 7. The first column reports the average seasonal CDD and HDD for each city averaged across the twenty years. The third column contains the CDD and HDD calculated from the average daily temperatures. (The difference between the two columns is due to the different sequence of averaging.) The fifth column is the counterpart of Column 3 except that the inputs are the adjusted daily average temperatures of the last season. Columns 2 and 4 report the simulated forward CDD's and HDD's using the corresponding forecast inputs when the correlations, both

[10] We also tried the adjusted average temperature of 1998, $\widehat{\overline{Y}}_{1998,t}$ $(t = 1, 2, ...)$ as forecasts. All conclusions remain qualitatively the same.

contemporaneous and lagged, between the dividend process and the temperature are set to zero. [11] Let us compare Column 2 with Column 3, and Column 4 with Column 5. Notice that the difference in some cases is very small, almost entirely attributable to simulation errors; but in some other cases, the difference is too big to be explained by simulation errors. The reason for the sizable differences lies in the aforementioned fact that the simulated quantities are nonlinear functions of the underlying variable. To ensure correct pricing, we perform a two-stage simulation for each value estimate. In the first stage, we simulate forward prices by setting the correlation parameter and the moving average coefficients to zero (i.e., $\varphi = 0$, $\eta_j = 0, \forall\, j$), and record the difference between the model price and the implied forward price from the forecasts (e.g., in Exhibit 7, when average temperatures are used as forecasts, this is the difference between Column 2 and Column 3). Then, in the second stage, we simulate the derivative security's price by setting the correlation parameter back to its actual value. Here, before calculating the derivative's payoff for each path, we first adjustment the realized CDD or HDD by the difference found in the first stage. In a nutshell, the above procedure amounts to ensuring unbiased paths of the CDD and HDD which are underlying variables for weather derivatives.

6.2. Temperature Derivative Prices and Market Price of Risk

For all subsequent simulations, we will examine four risk-aversion cases: $\gamma = -0.5$, -1.0, -2.0, and -4.0, in the order of increasing risk aversion. $\gamma = -1.0$ corresponds to log utility. Most empirical studies indicate that γ ranges between zero and -2.0, although some recent studies suggest an even higher risk aversion to accommodate the so-called equity premium puzzle.[12] We include $\gamma = -4.0$ as

[11] To make the point clear, we set $\mu = 1.0$ in the dividend process, which as a random walk produces a bigger variation.

[12] See Mehra and Prescott (1985) for a brief survey of emirical studies on estimating the risk aversion parameter.

a "high" risk aversion case. For each risk aversion scenario, we examine two correlation levels, each of which can be positive or negative. Since qualitatively, HDD contracts are almost mirror image of CDD contracts, for brevity, we will only report results for CDD contracts. In the next section, we will first examine the simple case where the dividend process and the temperature process are only contemporaneously correlated. We then introduce the lagged impact of the temperature on the output. Throughout the discussions we will attempt to assess the importance of market price of risk. (Recall that, $\varphi = 0.0$ and $\eta_j = 0, \forall\, j$ amount to a zero market price of risk for the temperature variable, irrespective of the risk-aversion level.)

6.2.1. Contemporaneous Correlations Only

Here, the two correlation levels are 0.15 and 0.25. Since the lagged the correlations are zero, a contemporaneous correlation of 0.15 or 0.25 (positive or negative) means that about 2.3% or 6.3% of the total variance in output is due to the temperature variations. Exhibit 8 reports CDD forward prices. Several observations are in order. First, since the forecast forward prices are simple expectations corresponding to $\varphi = 0$ and $\eta_j = 0, \forall\, j$, comparing the forecast forward prices with those under different values of φ confirms the predictions in Table 2. The percentage differences are an indirect measure of the magnitude of the market price of risk, or risk premium of the temperature variable. It is interesting to observe that the impact of the market price of risk is insignificant for all cases. The largest percentage price difference, -0.06%, is for Chicago when the correlation is 0.25 and the risk aversion parameter is -4.0. Second, For a fixed correlation, a higher risk aversion leads to a bigger risk premium, which makes intuitive sense. Third, other things being equal, a stronger correlation leads to a bigger risk premium, which again makes intuitive sense. (This and the previous observation will become more manifest in later cases.) Fourth, depending on the sign of the correlation, a higher risk-aversion may lead to a higher or lower forward price. Specifically, with a negative correlation φ, a higher risk-aversion leads to a higher forward price, and vice versa.[13] The intuitive reason lies in the valuation equations in Section 4.1. Notice that the future payoff is "discounted" back at a rate which is a function of certain fixed parameters and the stochastic dividend ratio raised to the power of γ. Since $\gamma < 0$ and since the average dividend growth rate is positive (i.e. $\alpha > 0$), a higher dividend ratio leads to a smaller discount factor (i.e. a lower present value), and vice versa. Now, with a negative correlation, a higher CDD ending value is most likely accompanied by a lower dividend ratio, and a lower dividend ratio will lead to a higher present value. The more negative is γ, the more manifest the above effect, and hence the pattern. Of course, this also explains why, under a particular risk-aversion parameter, the relationship between the forward price and the correlation is negative.

Exhibit 9 contains CDD call and put option prices. We first explain the so-called historical simulation prices. Some authors (e.g. Hunter, 1999) have discussed the use of historical simulation in pricing weather derivatives. The idea is similar to the historical simulation used in some of the VaR calculations. Specifically, a derivative contract's payoff is calculated using realized, historical underlying variable values, and the average payoff over a sample period (say 10 years) is taken as the estimate of the derivative's value. In our case, it boils down to evaluating the payoff for each of the twenty years and then averaging the payoffs to arrive at a value. (The discounting is done using the 6% constant rate.) The strike price is set equal to the historical average CDD for each city. It is seen that the call and put option values are equal for each city. This is expected since the exercise price is set at the (realized) forward price level.

[13] In some cases the percentage difference is too small to retain the negative sign with two percetange decimal places. But the absolute prices clearly show that the zero percentages have the right sign.

For all other option values, the strike prices are in Column 3, which are the forecast forward prices from the first column of Exhibit 8. The option prices in Column 4 are "risk neutral" option values, calculated by setting the correlation to zero. As such they do not contain risk premiums. The rest of the setup is the same as Exhibit 8.

The first observation is the significant difference between the historical simulation prices and other simulated prices. To understand the reason, realize that in regular simulations, the overall level of the temperature one year into the future is more or less contained by the forecasts, and variations are around this overall level. But with historical simulations, we are implicitly assuming that next year's temperatures can have extremely large variations, even in terms of the general level; the very cool and the very hot summers in the past 20 years command the same probability in realization. In reality however, meteorologists are at least able to forecast the general level of temperatures (i.e. cool, normal or hot) for the next summer with some accuracy. Historical simulations tend to exaggerate the impact of extreme realizations when the sample size is not large, and in our case, we have only 20 observations for each season. This points out a serious drawback of historical simulations used in pricing weather derivatives.

Another observation is that all the patterns associated with the forward prices when varying parameter values apply to call options. [14] The opposite patterns apply to put options, which is intuitive since the value of a put option is inversely related to the level of CDD. In fact, all the explanations for the patterns in forward prices also apply to options. However, the percentage differences are very different in the two exhibits. While they are small in magnitude in both exhibits, those for the options are many-fold larger than the forward percentage differences. The

[14] Notice that the percentages are calculated before the prices are rounded off. This is why we sometimes observe non-zero percentages, although the prices apear to be the same.

marked difference in risk premium impacts between the two types of instruments is mainly due to the nature of the payoffs: linear for forward contracts and non-linear for options.

The largest percentage difference in price, 1.28%, is observed for the combination of $\gamma = -4.0$ and $\varphi = -0.25$ for the city of Dallas. It means roughly that the CDD call option contains a risk premium of 1.28% purely due to the market price of risk associated with the temperature variable. By and large, it appears that the market price of risk is not a significant factor if the dividend and temperature processes are only contemporaneously correlated.

6.2.2. Contemporaneous and Lagged Correlations

To fully assess the significance of risk premium, we now introduce lagged correlations. As stated earlier, we examine 90 lags. We now need to specify the moving average coefficients, η_j. For simplicity, we assume a simple geometric decay. Specifically, once the contemporaneous correlation, φ is set, we will calculate η_1 as $q\varphi$, η_2 as $q^2\varphi$, η_3 as $q^3\varphi$, and so on with $0 < q < 1$. The decay factor q is chosen such that $|\eta_{90}| = |q^{90}\varphi| = 0.0001$, a level arbitrarily chosen to signify the eventual diminution of the lagged effect. With such a structure and recalling that the total variance of the dividend growth is $\sigma^2 \left[1 + \frac{\varphi^2}{1-\varphi^2} + \sum_{j=1}^{90} \eta_j^2\right]$, we could easily calculate the portion of the variance attributable to the temperature variations: $\frac{\frac{\varphi^2}{1-\varphi^2} + \sum_{j=1}^{90} \eta_j^2}{1 + \frac{\varphi^2}{1-\varphi^2} + \sum_{j=1}^{90} \eta_j^2}$. Since 1/7th of the GNP is believed to be weather sensitive, a rough guidance for the variance proportion is that it be smaller than 1/7 (i.e., 0.143). With this restriction, setting the number of lags is equivalent to setting the contemporaneous correlation. Given $m = 90$, $|\varphi| = 0.25$ would lead to a proportion of 0.284, which is too high. In light of this, for subsequent simulations with lagged correlations, we will examine two levels of $|\varphi|$: 0.075 and 0.15, which correspond to a variance proportion of 0.04 and 0.13 respectively. It becomes apparent that there is no need to examine longer or shorter lags as long as the variance proportion attributable to the temperature is kept constant (e.g. at 0.13). Lastly, without loss of generality, we set to zero the realized, lagged innovations of the temperature process, $\widehat{\xi}_{t-m} = 0$ $(1 \leqslant m \leqslant 90)$.

Exhibits 10 and 11 are counterparts of Exhibits 8 and 9, with the lagged correlations. To begin with, it is seen that the qualitative relations between forward / option prices and the model parameters remain the same when lagged correlations are introduced. But there is a marked increase in the impact of the temperature variable's risk premium. For forward prices, comparing Panel A of Exhibit 8 with Panel B of Exhibit 10, we see that the percentage differences (between the "risk-neutral" prices and the prices containing risk premiums) increase by many folds. Comparing Panel A's in both exhibits, it is seen that, with lagged correlations, even a lower general level of correlation (0.075) would lead to a bigger risk premium. Nonetheless, the risk premium in forward prices is still negligible in that all percentage differences are within one per cent.

Very similar results are obtained for option prices, except that most of the percentage differences are no longer negligible. For both correlation levels, as long as $\gamma \leqslant -1$, the percentage differences are generally bigger than one per cent in absolute terms. With $|\varphi| = 0.15$ and $\gamma \leqslant -2$, all option values contain a risk premium of more than 3%. In those cases, the market price of risk for the temperature variable can not be ignored. The common industry practice of discounting temperature derivative payoffs at the riskfree rate is valid in an approximate sense only when the correlations, both contemporaneous and lagged, are very low and / or the risk aversion of investors is low.

6.2.3. Other Scenarios

Increasing Lagged Correlations. So far, we have been assuming diminishing lagged correlations, i.e., the impact of contemporaneous temperature movements on the aggregate output is the strongest, and it diminishes over time. It may be intuitively argued (but need to be empirically verified) that the impact takes a reverse order in the form of increasing lagged correlations. For example, if it is extremely hot for the time being, then the impact of this abnormal temperature may only show up later on. (The impact of temperature on crops is a clear example.) To assess the importance of risk premium in this scenario, we repeat the calculations for Exhibits 10 and 11 by reversing the order of the moving average coefficients, $\eta_j \ \forall \ j$. For brevity, we only report the results for option prices in Exhibit 12. Comparing Exhibit 12 with Exhibit 11, we observe that the percentage price differences decrease slightly for Atlanta and Dallas, and increase slightly for Chicago, New York and Philadelphia. Very similar results hold for forward prices. The reason is complex and lies in the standard deviation pattern in Exhibit 4 and the mean reversion feature of the dividend process. To facilitate understanding, we explain the logic in steps. 1) The size of the risk premium depends on the covariance between the error terms of the dividend and temperature proccesses, as shown in (4.8). 2) The terminal dividend level, δ_{T_2} is affected by all error terms between t and T_2, but only those between T_1 and T_2 affect the covariance. 3) With mean reversion in the dividend process, looking back from T_2, more distant error terms exhibit less impacts (e.g., the factor in front of the error term ν_{T_1} in the final expression of $\ln\delta_{T_2}$ is $\mu^{(T_2-T_1)}$). 4) It is apparent from Exhibit 4 that, for the last $m = 90$ day period up to T_2 (which is the end of September), reversing the order of $\eta_j \ \forall \ j$ will lower the covariance, since the higher η_j's are now associated with lower variances for the temperature variable. 5) For all other days within the contract period $T_2 - T_1$, the order reversion of $\eta_j \ \forall \ j$ will increase the covariance. To see this, let's look at $\mu^{(T_2-m)}\nu_{T_2-m}$ and ν_{T_2}, in which form they enter the experssion of $\ln\delta_{T_2}$ as error terms. Before reversing the order of $\eta_j \ \forall \ j$, the temperature error term ξ_{T_2-m} receives almost a zero weight in ν_{T_2} (by design), but receive a non-zero weight of $\eta_1\mu^{(T_2-m)}$ in ν_{T_2-m}. When we reverse the order, the weight in ν_{T_2-m} is almost zero, but the weight in ν_{T_2} is η_1. Since $|\eta_1| > |\eta_1\mu^{(T_2-m)}|$, the net effect is to increase the covariance. 6) Extending the logic in 5) to other lagged terms within each ν_τ ($T_1 \leqslant \tau \leqslant T_2 - m$), we see that the order reversion will increase the covriance for all days other than those between $T_2 - m$ and T_2. 7) The overall result is a trade-off between the negative effect in 4) and positive effect in 6). The difference between the "colder" and "warmer" cities is mainly due to the standard deviation profiles. As shown in Exhibit 6, the sine wave has a bigger amplitude for Chicago, New York and Philadelphia than for Atlanta and Dallas. It becomes clear that when $\mu = 1$, only the negative effect in 4) prevails. In sum, unless the dividend process follows a random walk, the relationship between the size of risk premium and the structure of the lagged correlations depends on the specific profile of the termperature volatilities.

Sign of the Lagged Correlations. One may also question the appropriateness of assuming the same sign for all coefficients. In the absence of empirical observations, we can only hypothesize. Insofar as "normal" is desirable as far as temperature is concerned, abnormal variations can only have negative impacts on the economy. Nonetheless, some cases may conceivably exist where abnormal variations have different directional impacts, depending on how far the lag is. What is certain is that as long as the signs of η_j are not uniform, the impact of market price of risk will decline, simply due to a lower covariance between the dividend process and the temperature process. We did repeat the calculations in Exhibits 10 and 11 by assuming alternating signs in η_j $\forall$ j, and found that the percentages went down substantial in magnitude. For example, for options, all percentage numbers are smaller than 1% in absolute terms. For brevity, we omit the exhibits.

Mean Reversion of the Aggregate Dividend. So far, we have assumed $\mu = 0.9$, which corresponds to a mean reversion rate of 0.1. It is useful to know how sensitive our results are to the level of mean reversion in the dividend process. To this end, we repeat the calculations in Exhibits 10 and 11 by assuming four other levels of μ: 0.80, 0.85, 0.95, 1.00. Note that $\mu = 1.0$ corresponds to a random walk. Since results are very similar for all five cities, we only report those for Chicago (which has the smallest percentage differences for option values in Exhibit 11) in Exhibit 13. It is seen that a higher value of μ leads to a bigger risk premium in forward and option values. This makes intuitive sense since a higher μ means bigger variations in the aggregate dividends. What is striking is the nonlinearity of the impact. To illustrate, for options, when μ increases by 0.1 from 0.8 to 0.9, the percentages roughly double for all cases; when μ increases by another 0.1 from 0.9 to 1.0, the percentages increase in magnitude by as much as 20 folds. With a random walk, the risk premium is more than 10% for all option values. An obvious conclusion is that, in determining the significance of the market price of risk for the non-traded temperature variable, we must carefully determine the degree of mean reversion in the aggregate dividend process.

Risk Aversion. We have arbitrarily defined the case of $\gamma = -4.0$ as "high risk aversion". Indeed, Mehra and Prescott (1985) decide to set γ to -10 as an arbitrary maximum in their study of the equity premium. They find that even this extreme value of risk aversion could not explain the relative magnitude of the equity return and the riskfree rate. Their observation does not necessarily suggest a higher empirical value of γ. In fact, to reproduce the seemingly low riskfree rate, they need a risk aversion level much lower than what $\gamma = -2.0$ would suggest. On the other hand, such a low risk aversion will be completely inconsistent with the higher level of observed equity returns. Hence the "puzzle". They conjecture that market friction may be the answer. In this sense, we may still take comfort in the empirical range of 0 and -2.0 for the pamameter γ. Nonethelsss, we do not claim in anyway that we put a closure to this issue. What we can say though is that risk aversion is indeed one of the important factors affacting the risk premium in our study.

7. Summary and Conclusion

In this paper, we propose and implement an equilibrium valuation framework for weather derivatives. We specialize the framework to temperature contracts. The framework is the generalized Lucas's model of 1978. The underlying economic variable is the aggregate dividend and the underlying variable for weather derivatives is the daily temperature, and the two are correlated both contemporaneously and in a lagged fashion. We study the temperature behavior for the period of 1979 to 1998 for five major cities in the U.S., and derive key properties of the temperature dynamics. Our model not only allows easy estimation, but also incorporates key features of the daily temperature dynamics such as seasonal cycles and uneven variations throughout the year. The temperature system is estimated using the maximum likelihood method, and temperature contracts are priced accordingly.

Our framework has many advantages. It allows the use of weather forecasts in modelling the future temperature behavior. In addition, since our starting point is the daily temperature, the framework is capable of handling temperature contracts of any maturity, for any season, and it requires only a one-time estimation. In contrast, if one starts by modelling the cooling degree days (CDD's) or heating degree days (HDD's) directly, then by nature of the temperature behavior, the CDD's or HDD's will necessarily be season and maturity specific, which implies that each contract will require a separate estimation procedure. This will not only create potential inconsistency in pricing, but also render the whole idea impractical if many different contracts are dealt with or if the valuation is to be done on an on-going basis. Last but not least, our equilibrium framework allows us to answer a very important question: Can one use the riskfree rate to discount payoffs to obtain weather derivative values without incurring too big an error? In other words, is the market price of risk for the non-tradable temperature variable a significant factor in valuation?

Several conclusions can be drawn from the study. First and foremost, the market price of risk associated with the temperature variable appears to be negligible in most cases. Its impact is stronger only if the risk aversion is high or when the mean reversion in the aggregate dividend process is weak. Risk neutral valuation, or using the riskfree rate to discount derivatives payoffs, is strictly valid only when the correlations, both contemporaneous and lagged, are zero. However, with modest correlations and risk aversion, and a mean reversion greater than 0.1 for the aggregate dividend process, the risk premium is not significant. In other words, the industry practice of assuming a zero market price of risk seems to be warranted.

Second, the market price of risk affects option values much more than forward prices, mainly due to the payoff specification. Thanks to forward contracts' linearity in payoffs, much of the impact is "integrated" out. For options, the truncation in payoffs leaves room for the market price of risk to exert its impact.

The third conclusion has to do with a common practice in the industry, which is to use the historical simulation approach to estimate weather derivative values. We show that in most cases, this is not valid. Weather contracts typically cover a period to come and do not extend very far into the future. However, historical simulations implicitly assume that the next season's temperature can resemble any of the past seasons in the sample, including extreme seasons (very cold or very warm). As a result, in most cases, the historical simulation method tends to over estimate option prices.

As for future research directions, one obvious avenue is to adapt the framework to other weather variables such as snowfall and rainfall. This is going to be more challenging in that the weather variable such as rainfall is no longer a continuous variable. Moreover, the cumulation of such a variable in a season is far more important than the realized level within, say, a day. Nonetheless, derivative contracts on such variables will have direct appeal to users such as farmers and ski resort operators. We hope that our work represents the first step toward developing a more comprehensive strand of literature.

References

[1] Agins, T. and K. Kranhold, February 18, 1999, "Coat Peddlers Turn to Forecasters to Beat Winter Heat", *Wall Street Journal*, page B1 and B9.

[2] Black, F. and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637-655.

[3] Challis, Simon, 1999, "Bright Forecast for Profits", *Reactions*, June 1999.

[4] Fama, E. and K. R. French, 1988a, "Permanent and Temporary Components of Stock Market Prices," *Journal of Political Economy*, 96, No. 1-3, 246-273.

[5] Fama, E. and K. R. French, 1988b, "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3-26.

[6] Hanley, M., 5 July/August, 1999, "Hedging the Force of Nature", *Risk Professional*, issue1, 21-25.

[7] Hull, J., 1997, Options, Futures and other Derivative Securities, Prentice Hall, Englewood Cliffs, New Jersey.

[8] Hunter, Robert, 1999, "Managing Mother Nature", *Derivatives Strategy*, February 1999.

[9] Lucas, R. E., 1978, "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429-1445.

[10] Marsh, T. A. and R. Merton, 1987, "Dividend Behavior for the Aggregate Stock Market," *Journal of Business*, 60, 1-40.

[11] Mehra, R and E. Prescott, 1985, "The Quity Premium, a Puzzle", *Journal of Monetary Economics,* 15, 145-161.

[12] Shiller, R.J., 1983, "Do Stock Prices Move too Much to be Justified by Subsequent Changes in Dividends? Reply", *American Economic Review,* 73, No 1, 236-237.

[13] Smithson, C. and C. Choe, 1999, "Whatever the Weather", *Risk*, 12 (9).

Appendices

A. Proof of Proposition 3.1

A.1. Equilibrium Discount Bond Price

Based on the assumption on aggregate dividend, we have

$$\ln \frac{\delta_T}{\delta_t} = \alpha \sum_{i=t+1}^{T} \mu^{T-i} + (\mu^{T-t} - 1) \ln \delta_t + \sum_{i=t+1}^{T} \mu^{T-i} \nu_i.$$

Without loss of generality, suppose there are m lagged error terms in ν_i. Define $\eta_0 = \frac{\varphi}{\sqrt{1-\varphi^2}}$, $\Upsilon(t,T) = E_t(\ln \frac{\delta_T}{\delta_t})$, and $\Sigma(t,T) = Var_t(\ln \frac{\delta_T}{\delta_t})$. Then, the marginal distribution of $\ln \frac{\delta_T}{\delta_t}$ conditional on $\ln \delta_t$ is

$$f(\ln \frac{\delta_T}{\delta_t} \mid \ln \delta_t) = \frac{1}{\sqrt{2\pi\Sigma}} \exp\left(-\frac{(\ln \delta_T/\delta_t - \Upsilon)^2}{\Sigma} \right)$$

with

for $T - t \geq m + 1$

$$\Upsilon(t,T) = \alpha \textstyle\sum_{i=t+1}^{T} \mu^{T-i} + (\mu^{T-t} - 1) \ln \delta_t + \mu^{T-t-m} \sigma \left(\sum_{i=1}^{m} \mu^{i-1} (\sum_{j=1}^{i} \eta_{m-j+1} \widehat{\xi}_{t-i+j}) \right)$$

$$\Sigma(t,T) = \sigma^2 \left(\textstyle\sum_{i=t+1}^{T} \mu^{2(T-i)} + (\sum_{i=0}^{m} \eta_i \mu^{m-i})^2 (\sum_{i=t+1}^{T-m} \mu^{2(T-m-i)}) + \sum_{i=0}^{m-1} (\sum_{j=0}^{i} \eta_j \mu^{i-j})^2 \right)$$

and

for $T - t \leq m$

$$\Upsilon(t,T) = \alpha \textstyle\sum_{i=t+1}^{T} \mu^{T-i} + (\mu^{T-t} - 1) \ln \delta_t + \sigma \left(\sum_{i=1}^{\tau} \mu^{\tau-i} (\sum_{j=i}^{m} \eta_j \widehat{\xi}_{t+i-j}) \right)$$

$$\Sigma(t,T) = \sigma^2 \left(\textstyle\sum_{i=t+1}^{T} \mu^{2(T-i)} + \sum_{i=0}^{\tau-1} (\sum_{j=0}^{i} \eta_j \mu^{i-j})^2 \right)$$

The price of a pure discount bond at time t with maturity T is

$$\begin{aligned} B(t,T) &= E_t \left(\tfrac{U_c(c_T,T)}{U_c(c_t,t)} \cdot 1 \right) = e^{-\rho(T-t)} E_t \left((\tfrac{\delta_T}{\delta_t})^\gamma \right) \\ &= e^{-\rho(T-t)} \exp \left(\gamma E_t(\ln \tfrac{\delta_T}{\delta_t}) + \tfrac{1}{2}\gamma^2 Var_t(\ln \tfrac{\delta_T}{\delta_t}) \right) \\ &= e^{-\rho(T-t)} \exp \left(\gamma \Upsilon + \tfrac{1}{2}\gamma^2 \Sigma \right). \end{aligned}$$

The yield-to-maturity, $R(t,T)$, is

$$R(t,T) = -\frac{\ln B(t,T)}{T-t} = \rho - \frac{\gamma\Upsilon + \frac{1}{2}\gamma^2\Sigma}{T-t}.$$

B. Proof of Proposition 4.2: Equilibrium forward Prices

When there is no autocorrelation in daily temperatures, i.e., $\rho_i = 0$ (for $i = 1, 2, ..., k$), and no lagged impact of the temperature on the aggregate output, i.e., $\eta_j = 0, \forall\, j$, the joint distribution of $(\ln \delta_T,\, Y_T)$ conditional on $(\ln \delta_t,\, Y_t)$ is a bi-variate normal distribution with correlation coefficient φ :

$$f(\ln \delta_T, Y_{yr,T} | \ln \delta_t, Y_{yr,t}) = N[\ln \delta_T, Y_{yr,T}; \mu_\delta(t,T), \Sigma_\delta(t,T), \mu_Y(t,T), \Sigma_Y(t,T), \varphi],$$

where

$$\mu_Y(T) = E_t(Y_{yr,T}) = \widehat{\widehat{Y}}_{yr,T} \qquad and \qquad \Sigma_Y(T) = var_t(Y_{yr,T}) = \sigma^2_{yr,T}.$$

B.1. Forward Prices before the Accumulation Period

Define $\zeta \equiv \epsilon + \frac{\varphi}{\sqrt{1-\varphi^2}}\xi$. For an HDD forward contract with the accumulation period from T_1 to maturity T_2, its forward price at time $t < T_1$ (before the accumulation period) is

$$F_{HDD}(t, T_1, T_2) = \frac{1}{B(t,T_2)} E_t\left(\frac{U_c(c_{T_2}, T_2)}{U_c(c_t, t)} \cdot HDD(T_1, T_2)\right)$$

$$= \frac{e^{-\rho(T_2-t)}\delta_t^{-\gamma}}{B(t,T_2)} E_t\left(\delta_{T_2}^{\gamma} \sum_{\tau=T_1}^{T_2} \max(65 - Y_{yr,\tau}, 0)\right)$$

$$= \frac{e^{-\rho(T_2-t)}\delta_t^{-\gamma}}{B(t,T_2)} \delta_t^{\gamma\mu^{T_2-t}} E_t\left(e^{\sum_{i=t+1}^{T_2} \gamma\mu^{(T_2-i)}[\alpha+\sigma\zeta_i]} \sum_{\tau=T_1}^{T_2} \max(65 - Y_{yr,\tau}, 0)\right).$$

We can rewrite the above as

$$= \frac{e^{-\rho(T_2-t)}\delta_t^{\gamma(\mu^{T_2-t}-1)}}{B(t,T_2)} E_t\left(\sum_{\tau=T_1}^{T_2} e^{\sum_{\substack{i=t+1\\ i\neq\tau}}^{T_2} \gamma\mu^{(T_2-i)}[\alpha+\sigma\zeta_i]} e^{\gamma\mu^{T_2-\tau}(\alpha+\sigma\zeta_\tau)} \max(65 - Y_{yr,\tau}, 0)\right)$$

$$= \frac{e^{-\rho(T_2-t)}\delta_t^{\gamma(\mu^{T_2-t}-1)}}{B(t,T_2)} \sum_{\tau=T_1}^{T_2} E_t\left[e^{\sum_{\substack{i=t+1\\ i\neq\tau}}^{T_2} \gamma\mu^{(T_2-i)}[\alpha+\sigma\zeta_i]}\right] E_t\left[e^{\gamma\mu^{T_2-\tau}(\alpha+\sigma\zeta_\tau)} \max(65 - Y_{yr,\tau}, 0)\right]$$

$$= \frac{e^{-\rho(T_2-t)}\delta_t^{\gamma(\mu^{T_2-t}-1)}}{B(t,T_2)} \sum_{\tau=T_1}^{T_2} \left[e^{\sum_{\substack{i=t+1\\ i\neq\tau}}^{T_2} \gamma\mu^{(T_2-i)}[\alpha+\frac{1}{2}\gamma\sigma^2/(1-\varphi^2)\mu^{(T_2-i)}]}\right] E_t\left[e^{\gamma\mu^{T_2-\tau}(\alpha+\sigma\zeta_\tau)} \max(65 - Y_{yr,\tau}, 0)\right]$$

Tedious calculations show that

$$E_t\left[e^{\gamma\mu^{T_2-\tau}(\alpha+\sigma\zeta_\tau)}\max(65-Y_{yr,\tau},0)\right]$$

$$= e^{\gamma\mu^{T_2-\tau}(\alpha+\frac{1}{2}\gamma\sigma^2/(1-\varphi^2)\mu^{(T_2-\tau)})}\left([65-\mu'_Y(\tau)]\cdot N(\tfrac{65-\mu'_Y(\tau)}{\sigma_{yr,\tau}})+\tfrac{\sigma_{yr,\tau}}{\sqrt{2\pi}}\exp\left[-\tfrac{(65-\mu'_Y(\tau))^2}{2\sigma^2_{yr,\tau}}\right]\right)$$

$$\text{with} \qquad \mu'_Y(\tau)=\widehat{\widehat{Y}}_{yr,\tau}+\gamma\varphi\mu^{\tau-t}\sigma/\sqrt{1-\varphi^2}\sigma_{yr,\tau}\,.$$

Therefore,

$$F_{HDD}(t,T_1,T_2)=$$

$$\frac{e^{-\rho(T_2-t)}\delta_t^{\gamma(\mu^{T_2-t}-1)}}{B(t,T_2)}\sum_{\tau=T_1}^{T_2}\left(\begin{array}{c}\left[e^{\sum_{i=t+1,\, i\neq\tau}^{T_2}\gamma\mu^{(T_2-i)}[\alpha+\frac{1}{2}\gamma\sigma^2/(1-\varphi^2)\mu^{(T_2-i)}]}\right]e^{\gamma\mu^{T_2-\tau}(\alpha+\frac{1}{2}\gamma\sigma^2/(1-\varphi^2)\mu^{(T_2-\tau)})}\\ \times\left[[65-\mu'_Y(\tau)]\cdot N(\frac{65-\mu'_Y(\tau)}{\sigma_{yr,\tau}})+\frac{\sigma_{yr,\tau}}{\sqrt{2\pi}}\exp\left(-\frac{(65-\mu'_Y(\tau))^2}{2\sigma^2_{yr,\tau}}\right)\right]\end{array}\right)$$

$$=\sum_{\tau=T_1}^{T_2}\left(\left[[65-\mu'_Y(\tau)]\cdot N(\tfrac{65-\mu'_Y(\tau)}{\sigma_{yr,\tau}})+\tfrac{\sigma_{yr,\tau}}{\sqrt{2\pi}}\exp\left(-\tfrac{(65-\mu'_Y(\tau))^2}{2\sigma^2_{yr,\tau}}\right)\right]\right).$$

Similarly, the corresponding forward prices on CDD is

$$F_{CDD}(t,T_1,T_2)=\tfrac{1}{B(t,T_2)}E_t\left(\tfrac{U_c(c_{T_2},T_2)}{U_c(c_t,t)}\cdot CDD(T_1,T_2)\right)$$

$$=\sum_{\tau=T_1}^{T_2}\left([\mu'_Y(\tau)-65]\cdot N(\tfrac{\mu'_Y(\tau)-65}{\sigma_{yr,\tau}})+\tfrac{\sigma_{yr,\tau}}{\sqrt{2\pi}}\exp\left[-\tfrac{(65-\mu'_Y(\tau))^2}{2\sigma^2_{yr,\tau}}\right]\right).$$

B.2. Forward Prices during the Accumulation Period

For HDD forward price at time $t\in(T_1,T_2)$ during the accumulation period, we have

$$F_{HDD}(t,T_1,T_2)=\tfrac{1}{B(t,T_2)}E_t\left(\tfrac{U_c(c_{T_2},T_2)}{U_c(c_t,t)}\cdot HDD(T_1,T_2)\right)$$

$$=\tfrac{1}{B(t,T_2)}E_t\left(\tfrac{U_c(c_{T_2},T_2)}{U_c(c_t,t)}\cdot HDD(T_1,t)\right)+\tfrac{1}{B(t,T_2)}E_t\left(\tfrac{U_c(c_{T_2},T_2)}{U_c(c_t,t)}\cdot HDD(t+1,T_2)\right)$$

$$=HDD(T_1,t)\tfrac{1}{B(t,T_2)}E_t\left(\tfrac{U_c(c_{T_2},T_2)}{U_c(c_t,t)}\right)+\tfrac{e^{-\rho(T_2-t)}\delta_t^{-\gamma}}{B(t,T_2)}E_t\left(\delta_{T_2}^{\gamma}HDD(t+1,T_2)\right)$$

$$=HDD(T_1,t)+\sum_{\tau=t+1}^{T_2}\left([65-\mu'_Y(\tau)]\cdot N(\tfrac{65-\mu'_Y(\tau)}{\sigma_{yr,\tau}})+\tfrac{\sigma_{yr,\tau}}{\sqrt{2\pi}}\exp\left[-\tfrac{(65-\mu'_Y(\tau))^2}{2\sigma^2_{yr,\tau}}\right]\right).$$

Also, the CDD forward price at time $t\in(T_1,T_2)$ is

$$F_{CDD}(t,T_1,T_2)$$

$$=CDD(T_1,t)+\sum_{\tau=t+1}^{T_2}\left([\mu'_Y(\tau)-65]\cdot N(\tfrac{\mu'_Y(\tau)-65}{\sigma_{yr,\tau}})+\tfrac{\sigma_{yr,\tau}}{\sqrt{2\pi}}\exp\left[-\tfrac{(65-\mu'_Y(\tau))^2}{2\sigma^2_{yr,\tau}}\right]\right).$$

Exhibit 1: Summary Statistics

	Atlanta	Chicago	Dallas	New York	Philadelphia
Mean	63	50	66	56	56
Median	64	50	67	56	56
Mode	79	70	86	72	75
Standard Deviation	15	20	16	17	18
Minimum	5	-17	9	3	1
Maximum	92	93	97	93	92
Sample Size	7,300	7,300	7,300	7,300	7,300
			Correlation		
Atlanta	1.0000				
Chicago	0.8847	1.0000			
Dallas	0.8777	0.9038	1.0000		
New York	0.8966	0.8964	0.8443	1.0000	
Philadelphia	0.9125	0.8970	0.8455	0.9853	1.0000
			Auto Correlation		
k-lags					
1	0.9402	0.9421	0.9354	0.9448	0.9462
2	0.8690	0.8809	0.8680	0.8896	0.8926
3	0.8281	0.8494	0.8318	0.8654	0.8678
4	0.8069	0.8304	0.8132	0.8533	0.8550
5	0.7952	0.8181	0.8005	0.8470	0.8486
6	0.7867	0.8091	0.7918	0.8431	0.8437
7	0.7804	0.8022	0.7855	0.8394	0.8380
8	0.7764	0.7973	0.7813	0.8346	0.8330
9	0.7728	0.7925	0.7773	0.8297	0.8283
10	0.7687	0.7894	0.7731	0.8246	0.8228
11	0.7665	0.7870	0.7718	0.8197	0.8175
12	0.7652	0.7857	0.7720	0.8164	0.8142
13	0.7614	0.7835	0.7683	0.8124	0.8098
14	0.7562	0.7793	0.7608	0.8099	0.8054
15	0.7534	0.7759	0.7558	0.8070	0.8017

Exhibit 2: Summary Statistics of Monthly HDD and CDD (1979 - 1998)

		Jan.	Feb.	Mar.	Apr.	May	June	July	Aug.	Sept.	Oct.	Nov.	Dec.
							Atlanta						
HDD	Average	679	493	328	143	21	1	0	1	11	117	336	586
	Std. Dev.	125	90	85	63	18	4	0	2	10	43	97	126
	Maximum	882	657	465	261	63	18	0	11	32	188	514	797
	Minimum	462	297	156	36	2	0	0	0	0	22	131	346
CDD	Average	0	2	14	59	198	385	502	452	275	67	8	2
	Std. Dev.	1	3	10	39	66	68	72	63	59	38	13	3
	Maximum	3	8	36	141	322	494	639	589	440	178	49	13
	Minimum	0	0	2	4	62	221	372	349	192	12	0	0
							Chicago						
HDD	Average	1308	1065	857	516	230	51	6	10	112	407	766	1132
	Std. Dev.	189	165	101	80	81	31	6	12	46	79	106	184
	Maximum	1627	1359	1095	643	364	118	19	37	189	583	958	1562
	Minimum	956	733	733	393	87	6	0	0	34	288	598	891
CDD	Average	0	0	2	9	48	161	283	240	93	9	0	0
	Std. Dev.	0	0	3	15	40	62	73	92	36	10	0	0
	Maximum	0	0	13	53	167	254	398	445	158	44	1	0
	Minimum	0	0	0	0	4	38	152	106	19	0	0	0
							Dallas						
HDD	Average	627	442	269	91	10	0	0	0	8	69	305	557
	Std. Dev.	124	101	63	40	9	0	0	0	10	32	66	128
	Maximum	911	635	384	186	29	2	0	0	38	145	414	933
	Minimum	401	299	164	26	0	0	0	0	0	14	190	389
CDD	Average	1	4	28	87	270	491	642	622	399	145	22	4
	Std. Dev.	2	9	20	35	78	74	77	73	64	35	16	5
	Maximum	6	37	77	158	464	668	844	737	563	220	52	13
	Minimum	0	0	4	24	171	382	551	480	302	78	3	0
							New York						
HDD	Average	988	838	702	374	120	13	0	1	35	234	513	824
	Std. Dev.	145	125	80	46	40	12	1	3	16	73	69	139
	Maximum	1241	1173	913	446	187	49	3	9	74	376	651	1198
	Minimum	735	671	613	286	51	1	0	0	11	101	389	644
CDD	Average	0	0	1	4	63	224	392	350	153	22	1	0
	Std. Dev.	0	0	4	8	42	55	56	56	37	23	2	0
	Maximum	0	0	18	34	184	325	490	444	222	95	7	2
	Minimum	0	0	0	0	4	112	271	239	100	0	0	0
							Philadelphia						
HDD	Average	1002	835	676	349	106	10	0	2	39	261	538	852
	Std. Dev.	154	124	95	60	46	11	1	5	19	75	86	136
	Maximum	1243	1170	911	440	191	42	4	21	92	408	704	1219
	Minimum	738	644	542	193	38	0	0	0	14	138	392	701
CDD	Average	0	0	2	9	79	245	409	352	155	22	1	0
	Std. Dev.	0	0	5	11	47	66	71	65	42	21	2	0
	Maximum	0	0	20	34	230	401	540	470	244	83	6	0
	Minimum	0	0	0	0	9	133	283	268	90	0	0	0

Exhibit 3: Global Warming Trend

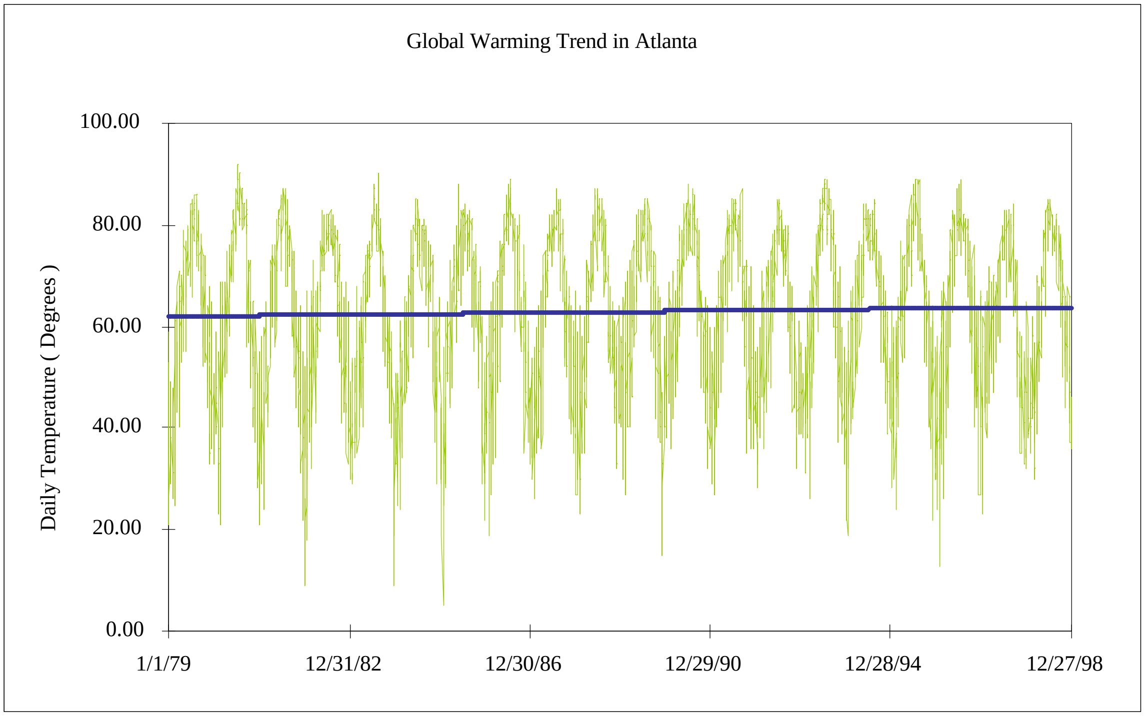

Exhibit 4: Standard Deviation of Date t's Temperature (ψ_t)

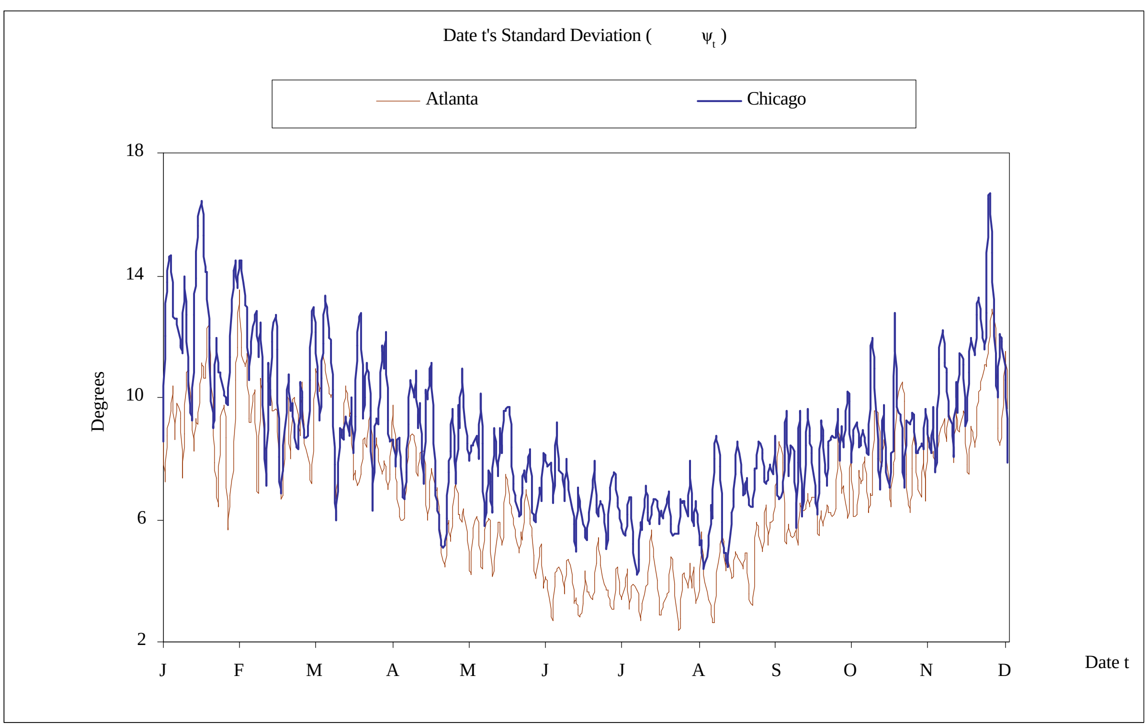

Exhibit 5: Realized Temperatures in New York
For a Colder-Than-Normal Winter (November 1980 - March 1981)

Panel A: Date t's Mean Temperature ($\overline{Y}_t$) vs Realized Temperatures

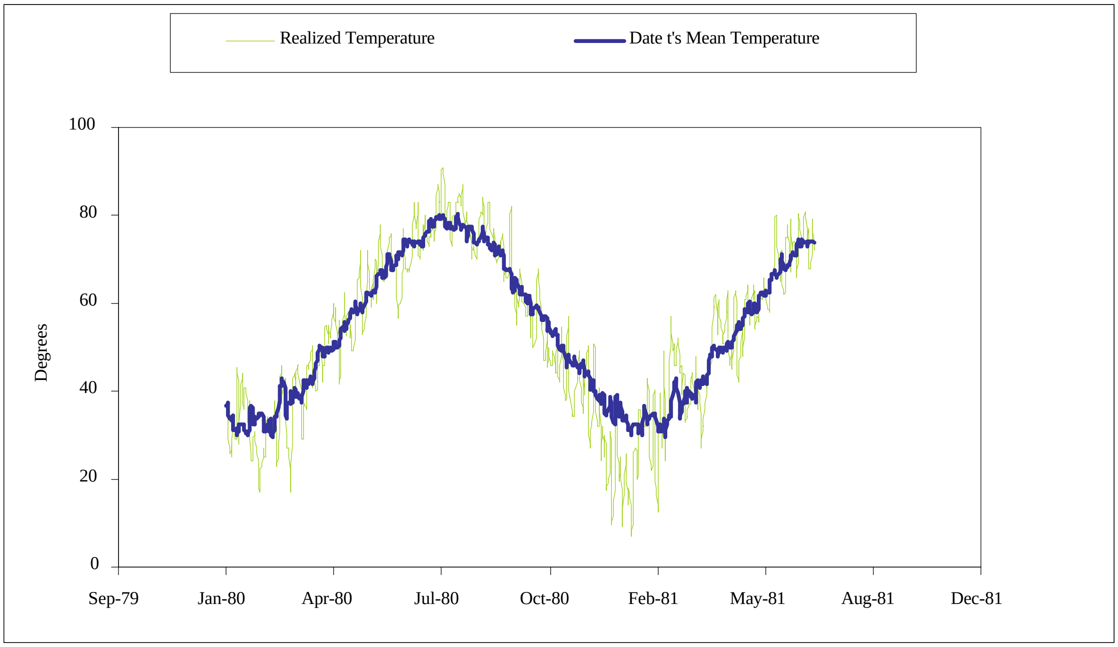

Panel B: Date t's Adjusted Mean Temperature ($\widehat{\overline{Y}}_t$) vs Realized Temperatures

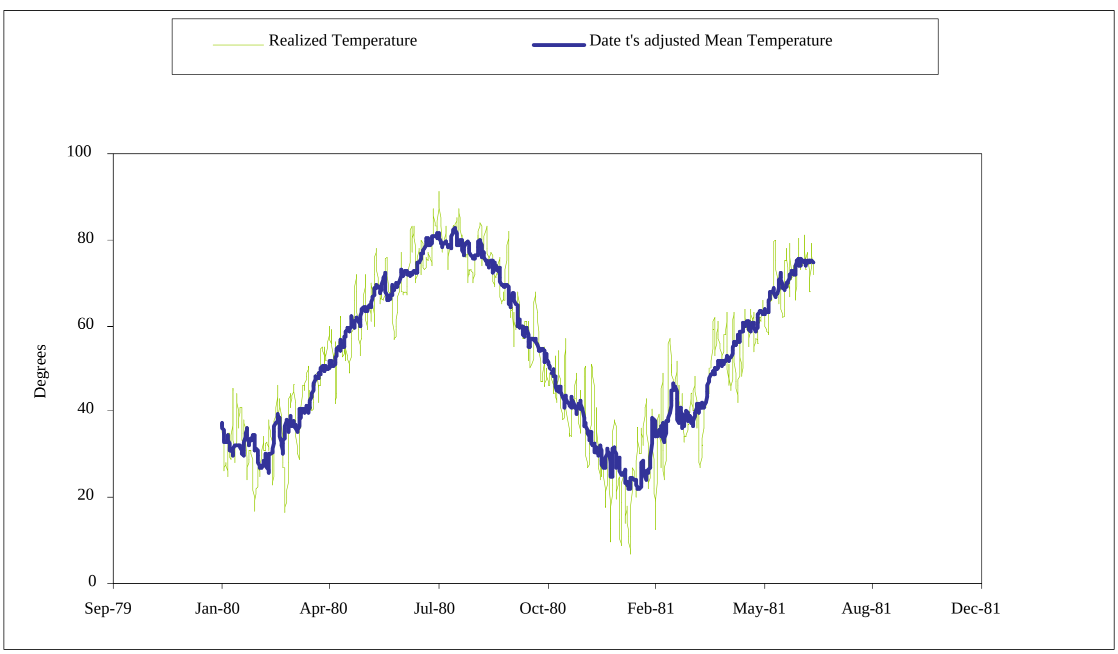

Exhibit 6: Maximum Likelihood Estimation Results

ρ_1	ρ_2	ρ_3	ρ_4	σ_0	σ_1	ϕ	Log - Likelihood	LR = $2\ln(L_1/L_2)$
					Atlanta			
0.8833	-0.3035	0.0322		7.5980	5.0912	-0.1881	$\ln(L_1) =$ -20,626	250 *
(0.01170)	(0.01520)	(0.01169)		(0.12086)	(0.14603)	(0.01067)		
0.9482	-0.3158	0.0760	0.0083	7.6565	5.0781	-0.1940	$\ln(L_2) =$ -20,751	
(0.01166)	(0.01567)	(0.01166)	(0.01288)	(0.12208)	(0.14776)	(0.01088)		
					Chicago			
0.7989	-0.2570	0.0428		7.8289	3.1294	-0.2014	$\ln(L_1) =$ -23,130	268 *
(0.01170)	(0.01467)	(0.01170)		(0.13922)	(0.18181)	(0.02316)		
0.8619	-0.2697	0.1029	-0.0188	7.9315	3.1211	-0.1998	$\ln(L_2) =$ -23,264	
(0.01166)	(0.01509)	(0.01165)	(0.01905)	(0.14057)	(0.18135)	(0.02369)		
					Dallas			
0.8158	-0.2436	0.0201		8.9378	6.3349	-0.1418	$\ln(L_1) =$ -21,381	258 *
(0.01170)	(0.01483)	(0.01170)		(0.14060)	(0.16800)	(0.00953)		
0.8711	-0.2482	0.0567	0.0162	9.0257	6.3488	-0.1497	$\ln(L_2) =$ -21,510	
(0.01168)	(0.01522)	(0.01167)	(0.01267)	(0.14175)	(0.16936)	(0.00970)		
					New York			
0.7558	-0.2631	0.0463		6.5372	2.7035	-0.2432	$\ln(L_1) =$ -21,719	298 *
(0.01169)	(0.01433)	(0.01169)		(0.11241)	(0.14520)	(0.02238)		
0.8117	-0.2612	0.1007	-0.0071	6.7129	2.8152	-0.2420	$\ln(L_2) =$ -21,868	
(0.01166)	(0.01472)	(0.01166)	(0.01374)	(0.11510)	(0.14831)	(0.02201)		
					Philadelphia			
0.7726	-0.2595	0.0473		6.9034	3.1654	-0.2015	$\ln(L_1) =$ -21,792	306 *
(0.01169)	(0.01446)	(0.01169)		(0.11957)	(0.15360)	(0.01932)		
0.8290	-0.2559	0.0973	0.0015	7.0545	3.2360	-0.2024	$\ln(L_2) =$ -21,945	
(0.01166)	(0.01486)	(0.01166)	(0.01459)	(0.12207)	(0.15675)	(0.01932)		

Note: 1. The Estimated Systems:

$$U_{yr,t} = \rho_1 U_{yr,t-1} + \rho_2 U_{yr,t-2} + \rho_3 U_{yr,t-3} + \sigma_{yr,t} * \xi_{yr,t} \quad (1)$$

$$U_{yr,t} = \rho_1 U_{yr,t-1} + \rho_2 U_{yr,t-2} + \rho_3 U_{yr,t-3} + \rho_4 U_{yr,t-4} + \sigma_{yr,t} * \xi_{yr,t} \quad (2)$$

with $U_{yr,t} = Y_{yr,t} - \widehat{\widehat{Y}}_{yr,t}$, $\quad \sigma_{yr,t} = \sigma_0 - \sigma_1 \mid \sin(\pi t/365 + \phi) \mid$,

$$\xi_{yr,t} \sim i.i.d.\ N(0,1), \quad \forall \quad yr = 1, 2, ..., 20 \quad \& \quad t = 1, 2, ..., 365.$$

2. The numbers in the parentheses are standard errors.
3. The null hypothesis (H_0) is $\rho_4 = 0$. The likelihood ratio (LR) test is computed as $LR = 2\ln L_1 - 2\ln L_2$ which is asymptotically distributed as $\chi^2(1)$ under H_0.
4. The 1 percent critical level for χ^2 with 1 degree of freedom is 6.6 and * indicates that the test statistic is significant.

Exhibit 7: Comparison of Forward Prices

	Sample Average	Theoretical Model Price	Price Based on Average Temperature	Theoretical Model Price	Price Based on Adjusted Avg. Temperature of 1998
	(1)	(2)	(3)	(4)	(5)
			CDD Season (May - September)		
Atlanta	1812.00	1797.22	1777.95	1902.17	1893.41
Chicago	823.60	799.26	674.80	1002.25	858.24
Dallas	2424.55	2414.13	2405.65	3154.69	3153.45
New York	1181.80	1169.98	1101.80	1296.11	1226.85
Philadelphia	1239.75	1220.75	1149.35	1351.34	1286.41
			HDD Season (November - March)		
Atlanta	2419.47	2423.70	2396.95	2729.68	2715.65
Chicago	5114.37	5127.72	5126.15	4508.18	4506.00
Dallas	2179.21	2202.51	2141.05	2246.26	2192.98
New York	3859.63	3864.00	3862.35	3419.25	3417.25
Philadelphia	3901.00	3901.62	3899.75	3406.81	3404.08

Note: 1. Theoretical model prices are calculated based on $\gamma = -0.5$, $\phi = 0.0$.
2. When forecasts are the adjusted average temperatures, we use November, and December of 1997 and the first three months of 1998 for HDD calculation.
3. The parameter μ in the dividend process is set to 1.0.