THE PRICING OF DEGREE-DAY WEATHER OPTIONS

by

Calum G. Turvey

WORKING PAPER 02/05

Department of Agricultural Economics and Business
University of Guelph
Guelph, Ontario

November 2, 2001

Calum Turvey is a professor in the Department of Agricultural Economics and Business and the Faculty of Management at the University of Guelph, Canada. This research was partially funded by the Ontario Ministry of Agriculture, Food and Rural Affairs. The author is grateful for significant comments from an anonymous referee and the journal editor.

Abstract

This paper presents a model and framework for pricing degree-day weather derivatives when the weather variable is a non-traded asset. Using daily weather data from 1840-1996 it is shown that a degree-day weather index exhibits stable volatility and satisfies the random walk hypothesis. The paper compares the options prices from the recommended model and compares it to a typical insurance-type model. The results show that the insurance model overprices the option value at-the-money and this may explain why the bid-ask spreads in the weather derivatives market is sometimes very large.

Keywords: weather derivatives, degree-day options, weather risk.

An Options Pricing Model for Degree-day Weather Derivatives

The role of weather as a source of business risk has resulted in an emerging market for weather based insurance and derivative products. In the U.S.A. companies such as World Wide Weather Insurance Inc., American Agrisurance Inc., Natsource Inc. (a New York City brokerage), Enron (a U.S. utility company), and SwissRe (a U.S. reinsurance companies) all offer weather risk products, such as swaps and/or options. In the fall of 1999 the Chicago Mercantile Exchange (CME) began trading in degree-day futures and options. Applications are wide spread among the natural gas, oil, and electricity sectors, and more and more such products are being used for agricultural and other weather sensitive industries such as ski resorts and snow mobile manufacturing. The main attraction of weather derivatives is that it insures volume rather than price. Too cool or too hot, too dry or too wet affects energy demand in utilities, production of crops and processing inventory in agriculture.

The types of contracts used to insure weather events are varied and include both swaps and options. In terms of heat-based options there are two different types. First, there are multiple event contracts. A utility company may want to insure against a specific event such as daily high temperatures being below 5^{o}F for 3 days straight, and the contract might stipulate that up to 4 events would be insured over a 90 day period, or an agribusiness firm may want to insure against multiple events of the daily high temperature exceeding 90^{o}F for 4 days straight in order to compensate for yield and/or quality loss.

Second, are straight forward derivative products based upon such notions as cooling degree-days above 65^{o}F (an indication of electricity demand for air conditioning), heating degree-days below 65^{o}F (an indication of electricity, oil, and gas demand required for heating), and growing degree-days or crop heat units above 50 degrees Fahrenheit (an indication of maximum crop yield potential in agriculture). These contracts are described as cooling degree-day call (put) spreads or heating degree-day call (put) spreads. The options contract (or ticket) has several sections including a general description of the product and the insured event; the specific weather location; the weather units being measured (e.g. degree-days or rainfall); the weather index being used (e.g. cooling, heating, or growing degree-days for the contract term); the contract term (e.g. June 1 to August 31); the index strike (e.g. 400 cooling degree-days); the unit price and currency (e.g. $5,000 per cooling degree-day); the settlement terms which indicate the specific source of weather information, the timing for payment, and any adjustments made due to revisions from the weather authority; and the buyer's premium[1]. Likewise, the CME has specified CDD and HDD futures contracts for each month of the year for 20 cities across the U.S.A. Thinly traded, the futures contracts are based on cumulative degree-days within the month, and settlement is based on readings supplied by Earth Satellite Corporation. The Notional value of the futures contract is equal to $100 times the CDD or HDD index value. CME degree-day options on CDD and HDD futures are listed in units of degrees Fahrenheit.

One of the problems facing the weather derivatives markets is how these derivatives should be priced. In the absence of a tradeable contract in weather and equilibrium price cannot be established using conventional means (Dischel 1998). At one end of the pricing spectrum, Cao and Wei (2000) develop a pricing model based on expected utility maximization with an equilibrium developed from Lucas's (1978) model. Davis (2001) also concludes that a Black-Scholes type framework is not appropriate for pricing weather derivatives as a matter of course, but under the assumptions of Brownian motion, expected utility maximization, a drift rate that includes the natural growth rate of the degree day measure, the natural growth rate in the spot price of a commodity (e.g. fuel price) and the natural growth rate in firm profits, then degree day options can be priced by a Black-Scholes analogue. Considine (undated) provides some simpler formulas based on the historical probability distribution of weather outcomes as well as a gaussian (normality) model that he claims can be sufficient at times. Turvey (2001) presents a number of flexible rainfall and heat related option contracts based upon historical probabilities.

There are empirical issues related to weather derivatives and a large part of this paper is dedicated toward resolving these issues in general, and the pricing of degree-day options in particular. First, until the CME started trading weather futures there was no forward market for weather. Individuals speculate on what a heat index might be 90 days hence, but unlike stock market indexes there is no mechanism for transparent price discovery on which to base such a prediction, and nature is under no obligation to comply with subjective market assessments. Second, rain or heat or any other insurable condition does not have a tangible form that is easily described (in contrast with common stock or a commodity futures contract). Third, for cities in

[1] To insure against excessive heat, World Wide Weather uses the following wording "The Company will insure from July 1, 1999 through August 31, 1999 that the temperature will not be 100 degrees Fahrenheit or above at the National Weather Service Station located in Santa Barbara, California". For growing degree-days "The Company will insure from April 1, 1999 to May 31, 1999 that there will be 1000 or more GDDs at the National Weather Service Station located in Fresno, California. Everyday where the average temperature is X degrees over 50 degrees Fahrenheit, there are x GDDs for that day".

the U.S.A. and elsewhere that are not listed on the CME, there does not exist a forward market weather index that would allow brokers, traders, and insurers to price such derivatives on an ongoing and transparent basis. This can impact liquidity in the OTC and insurance markets and can also have an impact on the appropriate market price of risk with which to price the contract. Fourth, the mechanics of brokering weather contracts depends specifically on the nature of the contract. A common approach is to use historical data and from this use traditional insurance 'burn-rate' methods to determine actuarial probabilities of outcomes. This convention limits trade. For the most part counterparties must agree on a price prior to the opening contract date and are in general restricted by lack of data to efficiently price and trade the contract during the period in which it is active.

For pricing put and call options on cumulative weather outcomes a limiting factor is in the transparency of a forward weather index. A forward weather index such as the HDD and CDD futures at the CME would operate like any other index and would be used to provide a current estimate of what the final weather index settlement would be. In so doing it would provide a mechanism for counterparties to trade on a continuous basis, and would also provide a mechanism for the continual pricing of the options' intrinsic values.

This paper develops an option pricing model based on such an index even if it is not traded. However, unless the index is formally traded, deriving option values from it will require consideration of the natural diffusion rate and the market price of risk as per lemma 4 in Cox, Ingersoll and Ross (1988). This paper discusses the properties of such an index, shows the evolution of the index in a dynamic context, and develops an options pricing model.

The theoretical development of an option pricing model for cumulative degree-day call and put spreads is the focus of this paper. The next section briefly establishes the economic motivation for weather derivatives. This is followed by an explanation of the insurance-type model and then the pricing of options in a dynamic framework is presented. The theoretical model is then applied to the pricing of degree-day derivatives for Toronto, using daily mean temperatures from 1840 to 1996.

The Economics of Weather in the Profit Maximizing Firm

The purpose of this section is to provide some economic intuition behind the role of weather and its effects on firm profits. In a classical economic context the profits of the firm can be defined by

(1) Π (P, Q(W), C(W),t) = P(Q) Q(W) - C(Q(W))

where P is the output price as a function of Quantity, Q; Q represents a quantity of sales or total productivity and it is a function of the weather variable; and the cost function C( ) is an increasing function of Q. The derivative of the profit function with respect to weather gives

(2) $$d\Pi() = \left[Q\frac{\partial P}{\partial Q}\frac{\partial Q}{\partial W} + P\frac{\partial Q}{\partial W} - \frac{\partial C}{\partial Q}\frac{\partial Q}{\partial W} \right] dW$$

In equation (2) $\partial Q/\partial W$ can be viewed as the marginal productivity of weather and being industry specific it can take on values that are negative, positive, or zero.

Equation (2) indicates three not mutually exclusive impacts on firm profitability. The first part on the right hand side of equation (2) is the direct price effect. The direct price effect is a consequence of changes in supply or demand in the market place due to weather impacts. It can be positive or negative. Sustained drought conditions in agriculture will at least affect prices in localized markets, if not national or international markets. We have seen, for example, crop failures in Eastern Europe lead to dramatic increases in local and international wheat prices. In other cases mild winters in the northern hemisphere will lead to excess supply and price reductions in vacation spots at winter holiday destinations. Energy wholesalers can see dramatic increases or decreases in prices depending on weather driven demand impacts.

The second component is the increase or reduction in quantity produced or sold. For the farmer facing localized drought conditions profit losses will result from decreased yields. For the ski resort the number of lift tickets sold will be an increasing function of snowfall.

The third component is the impact on costs. In some industries the cost effect is subtle. Reseeding, fertilizing or herbicide spraying will often result from extreme weather events in agriculture. In the energy sector peak load pricing resulting from excess demand in extreme heat or cold conditions is a significant cost borne by the municipal utility. In many regulated electrical energy markets any price increase to be transferred to consumers is regulated so that cost increases cannot easily be transferred to consumers in the short run.

The common element across examples is that the weather impact is transmitted to firm profitability through the quantity variable. The uncommon element is the precise functional form of Q(W). In energy markets it will be a convex function with quantity affects rising at weather extremes due to excess demand. In agriculture the Q(W) function is likely concave with crop losses occurring at either extremes of weather conditions. Still, other industries will exhibit strictly increasing, or strictly decreasing functions of weather. In fact, it is the heterogeneity of weather impacts across firms and industries that has given rise to a market for weather-risk derivative products including swaps and options.

It is this heterogeneity that has also given rise to a market based on specific weather events rather than firm cashflow. Since by equation (2) different firms are impacted differently it is impractical to even attempt to insure cashflow directly. Instead, firms engaged in trade in the weather derivatives market will attempt to hedge cashflow risks by adjusting the notional value or hedge position in specific-event weather derivatives. To gain some intuition as to why this market evolved, consider the expected value of equation (1) at which some point in time T is given by

(3) $E[\Pi(P, Q(W), C(Q),T) = \int[P(Q)\,Q(W) - C(W)]\,g(W)dW.$

Where g(W) is the probability distribution function associated with the specific weather event. Define W_z as the certainty equivalent value of a weather index, below which profits will fall and above which profits will rise. Then

(4) $E[\Pi(\,)] = \int^{Wz}\Pi\,(W)\,g(W)dW + \int_{Wz}\Pi\,(W)\,g(W)dW.$

The first part of equation (4) gives the expected value of losses below W_z and the second gives the expected value of gains. If the firm wants to insure losses below $\Pi(W_z)$ then it could buy a cashflow-based insurance contract with a value

(5) $V(W_z) = \int^{Wz}(\Pi(W_z) - \Pi(W))\,g(W)dz.$

However, such a contract would be difficult to price in practice because the cashflows will not generally be observable. Furthermore, such a contract implies a specific business to business transaction, and as indicated by the discussion around equation (2), heterogeneous weather impacts suggests that each weather affected firm will need to negotiate a separate contract. This comes with substantial transactions costs.

To avoid these transactions costs capital markets are designing a set of homogenous weather-based derivative products with a payoff structure contingent on specific weather outcomes rather than firm-specific cashflow. These include swaps and options. In this market it is the affected firm that defines its weather risk and the quantity of standardized weather contracts to be bought (or sold). From equation (5), at some point in time T, we can equate the value of losses due to weather events to a number of specific-event weather derivative products, i.e.

(6) $\int^{W_z} (\Pi(W_z) - \Pi(W))\, g(W)dW = \Delta\theta\int^{W_z} (W_z - W)\, g(W)dW$

or

(7) $V(W_z) = \Delta V^*(W_z)$

Equation (6) sets the expected value of cashflow losses equal to the expected payoff from a weather option. The (put) option value is measured by the right hand side in equation (6) and is denoted as $V^*(W_z)$ in equation (7). The expected value of the option at time T is given by Max $E[W_z - W, 0]$. The intrinsic value of the weather index is multiplied by θ which has units converting the value of W to \$/W (for example \$100/cooling degree-day). The value Δ represents the number of contracts required to insure the certainty equivalent value of profits, that is $\Delta = V(W_z) / V^*(W_z)$.

While the economics guiding the advent of the weather derivatives market is rational the problem of pricing weather derivatives remains unresolved. In the next section a common insurance type solution - referred to as the 'burn-rate' model or insurance model is presented and then a general solution to the problem along the lines of modern options pricing is presented.

The Pricing of Weather Options

This section describes a pricing methodology for weather derivatives. First, a ‘burn rate’ method employed by many brokers and insurers is described. Second, based on the assumption that a forward degree-day weather index exists for any location a general proof that such contracts can be priced using a formula similar to Black’s formula for pricing European options on futures contracts is presented; third, a simple approach to creating a forward weather index is provided; fourth, the underlying assumptions of volatility and a random walk in a weather index are empirically evaluated, and fifth the option model and insurance ('burn rate') model are compared.

The 'Burn Rate' Method for Pricing Weather Derivatives

In the absence of a forward weather index the pricing of weather derivatives is relatively straightforward. Using historical data cumulative degree-days (heating days, cooling days or crop heat units) are calculated for the time period in question and the options are priced as

(8) $$V_p = e^{-pT} E\{Max[Z - W_T^*, 0]\}$$

for a put option, and

(9) $$V_c = e^{-pT} E\{Max[W_T^* - Z, 0]\}$$

for a call option where p is the appropriate risk adjusted discounted rate, T is time or duration of contract in years, Z is the strike level in degree-days, and W^*_T is the value of the index at expiration also measured in degree-days. Since V measures the expected value of in-the-money degree-days, the actual price of the option is calculated by multiplying V by a dollar value with units $/degree-day. In equations (8) and (9) it is assumed that the payoff is $1/ degree-day. The probabilities that establish V are assumed to be stationary priors drawn from historical weather patterns and can be defined as either discrete or continuous.

The conventional methodology used in the industry is the 'burn rate' model which uses discrete observations of the n=1, N sampling distribution. That is, for a put option,

(10) $$V_p = \frac{e^{-pT}}{N} \sum_{n=1}^{N} Max[Z - W_n^*, 0]$$

where each in-the-money observation is given equal weight. The burn-rate approach draws from statistical inference over time, which assumes that history will repeat itself with the same likelihood as the past events described by the data used. In the alternative a long enough time series could be used to fit a known continuous probability distribution (e.g. a normal distribution) and the put option price could be obtained from

(11) $$V_p = e^{-pT} {}_L\int^Z (Z - W^*) f(w^*) dw$$

where L is a lower bound to the distribution.

If a continuous probability distribution is used there is an underlying assumption that the data series used is sufficient to fully describe the limiting probability distribution of outcomes with all asymptotic properties intact. Whether the probabilities of specific weather events are described by discrete or continuous distributions it should be noted that the measure of variance represents independent, cross-sectional inter-year risks. The use of this variance may or may not be representative of the current year's (intra-year) risk, and as will be shown later, what is assumed about the underlying stochastic structure is a critical element in distinguishing between the burn-rate model and modern options pricing.

Weather Indices, Futures Hedging, and Options Pricing

The burn-rate models will typically be purchased prior to the insured period, and will be traded infrequently, if at all. The reason that such contracts will not be traded results from the fact that there is no transparent mechanism to update or revise the probabilities during the insured period and hence no opportunity to arbitrage risk. The opportunity to arbitrage requires liquidity and liquidity requires observable volatility in an expected weather index W^*_T. IF W_T is the value of a degree-day weather index at expiration then for any $t<T$ there must exist an expectation about W_T, that is $W^*_T = E[W_T|t]$, conditional on weather information up to and including time t. Observable volatility in W^*_t requires first the existence of a forward weather index, and secondly that it be defined by an inter-temporal stochastic process.

The continuous time stochastic differential equation for the weather index can be described by Brownian motion and the Ito process

$$(12) \quad dW^*_t = \mu W_t^* dt + \sigma W_t^* dZ_t$$

The stochastic process described by (12) describes a random walk and is fundamental to the design of new derivative products for entities that follow a Markov process. As shown by Merton (1993), Black and Scholes (1973), Black (1976) and others, if the underlying assumptions in (12) hold then it can be used to price options. In Equation (12) μ is the mean change in cumulative degree-days and σ is the variance of the daily change in degree-days. The assumptions, which are empirically tested in this paper, are that the diffusion rate μ is constant over time and σ^2 increases linearly in time.

Equilibrium Pricing Formulas for Degree-Day Derivatives

With the introduction of the CME degree-day future contracts there will be, at least for specific locations, a spanning asset for which a classical options pricing formula can be derived. However, there are more jurisdictions without contracts than with, and this implies that not all risks can be spanned and risk-neutral valuation techniques cannot readily be used without modification. Under such circumstance it is necessary to apply a different set of rules to price options on non traded assets. In particular, an options pricing model when the underlying asset cannot be spanned by traded assets requires including the market price of risk. This has lead some practitioners to declare that modern options theory in the form of Black (1976) or Black-Scholes (1973) will not work (Nelken 1999, Dischel 1998) for pricing weather derivatives.

To capture the market price of risk, equation (12) is represented by

(13) $dW^*_t = (\mu - \lambda\sigma) W^*_t \, dt + \sigma W^*_t \, dZ$

where λ represents the market value of risk, and $\lambda\sigma$ the risk premium. Equation (13) is not consistent with Black-Scholes, but is consistent with the risk-neutral solution of Cox and Ross. The market price of risk as an economic entity results from Lemma 4 in Cox, Ingersoll and Ross and ensures that in equilibrium the rate of return on the option equals the risk-free rate in equilibrium. Therefore the term $(\mu-\lambda\sigma)$ is called the risk-neutral growth rate and it can be used to derive equilibrium prices of options on non-traded assets.

To price these options we modify the dynamic programming approach presented in Dixit and Pyndick. The Bellman equation is

(14) $F(W,t) = E\,[F(W,t) + dF(W,t)]e^{-pt}$

where F(w,t) is the value of the option, and p is the appropriate (risk adjusted) discount rate. Using equation (5) suppressing the * in W^*, and applying Ito's lemma

(15) $dF(W,t) = (1/2\, F''_w \sigma^2 W^2 + F'_w (\mu-\lambda\sigma) W + F'_t)\, dt + F'_w \sigma\, Wdz$

Setting $e^{-pt} = (1-pdt)$, substituting dF(W,t) into equation (14), and solving yields the stochastic differential equation.

(16) $\frac{1}{2} F''_w \sigma^2 W^2 + F'_w (\mu-\lambda\sigma) W + F'_t = pF.$

Equation (16) is a common-form partial differential equation. The call option value of F(W, X, t) that solves this equation for a strike price $X=W_z$ is

(17) $C(W,t) = F(W,t) = \theta\, [e^{-pt} N(d_1) X - e^{-(P-(\mu-\lambda\sigma))t} N(d_2) W]$

where t is time remaining until option enquiry, θ is the value per tick, X is the strike price in degree-days, p is the discount rate, N( ) is the value of the standard normal cumulative distribution function evaluated at d_1 or d_2,

$d_1 = [\ln (W/X) + (\mu - \lambda\sigma + .5\sigma^2)t]/\sigma\sqrt{t}$

and

$d_2 = d_1 - \sigma\sqrt{t}$

Since the market price of risk is explicitly included in the solution, the appropriate discount rate 'p' for a risk-neutral valuation is the risk free rate, r. However this still leaves unresolved the problem of determining the market price of risk λ. In a more general framework the diffusion $\mu - \lambda\sigma = r$ is called the risk neutral growth rate (Cox and Ross, 1976) and is a necessary condition for equilibrium pricing. In contrast μ is viewed as the natural growth rate in the value of the underlying. The value $\lambda = (\mu\text{-}r)/\sigma$ is then the market price of risk.

If the market price of risk so defined is applied to freely traded assets then $p = r = \mu\text{-}\lambda\sigma$ can be substituted into equation (17) and the resulting formula is identical to Black-Scholes. A more general argument is required for assets that are not-traded. For this we appeal to the security market line of the capital asset pricing model where

$\mu = r + \beta\, [r_m - r]$

or

(18) $\mu - r = \beta\, [r_m - r].$

Then we can define the market price of risk λ as

(19) $\lambda = \beta\, [r_m - r]/\sigma$

so that in equation (17), d_1 becomes

(20) $d_1 = [\ln (W/X) + (\mu - \beta\, [r_m - r] + .5\sigma^2)t\}/\sigma\sqrt{t}.$

To use equation (20) we need to interpret the Sharpe-Lintner model in the broadest sense. Roll's critique of the CAPM reminds us that the basic theory of pricing assets in equilibrium does not only apply to traded assets but non-traded assets as well. With this, the true market portfolio is unobservable and broad based indices such as the S&P500 used to proxy the true market portfolio return may be biased. Nonetheless, the theory provides for the equilibrium pricing of capital assets so that (20) must hold in (theoretical) equilibrium. In fact Stambaugh's follow-up to Roll indicates that inferences about the CAPM model are consistent with theory even when assets are not traded.

As indicated above, equation (21) is a general solution to pricing all assets in equilibrium. For the particular case of weather derivatives its form becomes simplified. If the underlying is a futures contract such as those traded on the CME, then we can use Dusak's argument that since the number of long positions equal the number of short positions then the outstanding value of the futures market is always zero and therefore excluded from the market portfolio.

Consequently the β coefficient for all futures contracts is zero. Furthermore, we know from Black that to avoid riskless arbitrage the growth rate in the futures contract must equal the risk free rate.

For futures contracts on weather variables (20) becomes

(21) d_1 [ln (W/X) + (r + .5σ^2)t]/$\sigma\sqrt{t}$

and (17) becomes the standard Black-Scholes pricing model with p=r. Using $W(t) = e^{-rt} W(T)$ and substituting this into (17) gives Black's model for pricing options on futures.

When the weather index is not a traded variable we rely on the direct relationship between the non-traded weather index and the market portfolio. Since the impact of weather events in localized regions will not be correlated with the market portfolio, then it to will have a beta of zero. This is consistent with the empirical findings in Cao and Wei (2000). The result and conclusion does not imply that the conditional underlying risks of economic outputs are zero, butthat in equilibrium the source of the risk can be diversified away. However, unlike a futures contract the non-traded weather variable will not grow at the risk-free rate. In fact the spot value at time t will simply equal the expected value at time T, that is $W^*_t = E[W^*_T]$. This implies a natural tendency towards mean reversion so $E[\mu]=0$. By substituting $\beta=0$ and $\mu=0$ into equation (17) and setting p=r to account for risk neutral valuations, the pricing model for call option on a non-traded weather index is given by

(22) $C(W,t) = \theta e^{-rt} [N(d1)X-N(d2)W]$

where

$d1 = [\ln (W/X) + .5\sigma^2 t]/\sigma\sqrt{t}$

and

$d_2 = d_1 - \sigma\sqrt{t}$

As a reminder the parameter θ is the tick value measured in \$/degree and the bracketed term is measured in degrees. The equivalent put option value is

(23) $P(W,t) = \theta e^{-rt} [N(-d_2)W - N(-d_1)X]$

The solution value of the option pricing models rests on three assumptions that are evaluated in the empirical section. Assumption 1 is that the natural dynamics of dW originates from a random walk and hence unanticipated changes in W are not serially correlated. If strong and predictable autocorrelation is present then asymmetric information between buyers and sellers of the option will allow for risk free arbitrage opportunities. In a later section I provide strong evidence that W* evolves as a random walk that is consistent with geometric Brownian motion.

The second assumption is that volatility is non-stochastic. If volatility evolves in an unpredictable way then equation (15) is misspecified and a more complex solution would be required. Hence the third assumption is that $\sigma = E[\sigma(t)]$ in equation (15) which means that volatility is set equal to the mean from a sample of measured volatility. I will show later that volatility is stable within and between years. This assumption is consistent with the assumption of time dependence in Merton (1993) and Wilmott (1998).

The fourth assumption is that $E[\mu]=0$ in equation (15). This assumption simply states that $W^*_O = E[W_T]$ and investors in weather options will use the mean of the historical sampling distribution as an unbiased estimate of the initial condition for dW. This is exactly how the opening prices of the CME exchange traded degree-day future prices are set. A less naïve condition is that $W^*_O = E[W_T|\Omega]$ where the expectation is now based on the conditional mean based on the information set Ω at time t=0. This is relevant when counterparties believe that degree-days will be higher or lower than the historical average. This may or may not come about as a variance preserving shift. However, I will provide evidence that the volatility of the degree-day index is remarkably stable even with large swings in the historical value of the indices and will also show how differing expectations affect option prices.

Defining a Weather Index

In the previous section the existence of a forward weather index was presumed. While possibility rather than existence is sufficient to support the development of an option pricing model, it is obviously a limitation to implementation and practice. The CME futures contracts will satisfy the spanning requirement of a correlated underlying derivative security, but CME contracts do not exist for many regions or cities. Hence the foregoing is a generalized solution that can be used to price options even if a formal futures contract does not exist. In this section a general approach to constructing a weather index using historical data is presented. In the next section the index model will be applied to a case study of degree-day contracts for Toronto.

The challenge for any broker or exchange to accurately price weather options is in the construction of an appropriate weather index which can be observed on a daily basis, and provide representative measures of volatility. To construct such an index it is useful to draw on the unique characteristic that the weather index cannot be influenced by human speculation. In this context the index is observable, objective, and representatively transparent. For example, settlement of the CME contracts is based exclusively on the data collected by Earth Satellite Corporation. Furthermore, a consistent characteristic of weather is that it is seasonal and systematic; summer, for example, always starts of with low temperatures that rise to a peak, and then decreases towards autumn. A naïve hedger planning a hedge in early spring would naturally assume that the summer weather pattern would follow the average pattern as dictated by history. Critical to this is the additional assumption that temperature is mean reverting: In the absence of any contrary information it is not unreasonable to assume that if the average temperature on June 30th is 70 degrees Fahrenheit, then in the current year the best unbiased estimate is that it will also be 70 degrees. The notion of mean reversion is also a natural phenomenon; the tendencies for temperature to fall to within a normal range following a heat wave, or to rise to normal temperatures following a cold snap is clearly the norm rather than the exception.

The absence of predictability and the assumption of mean reversion suggest that the best initial (t=0) unbiased estimate of the forward index is the historical average of the index over the specified contract time horizon. Indeed the opening trade on the CME futures contracts will most likely be the long-run average cumulative degree-day with some adjustment for long-term forecasts or revised expectations. The initial index value is given by equation (18):

$$(24) \qquad W_0^* = E[W_T] = \sum_{t=0}^{T} E[W_t]$$

where W represents the weather index (e.g. cooling degree-days, heating degree-days, growing degree-days or cumulative rainfall). After 1 day the observed weather condition at t=0 is recorded and the index value is appropriately adjusted to include the actual outcome plus the projected outcome;

$$(25) \qquad W^*_1 = W_0 + \sum_{t=1}^{T} E[W_t] \,.$$

Similarly at t=2

$$(26) \qquad W^*_2 = W_0 + W_1 + \sum_{t=2}^{T} E[W_t] \,,$$

and for any time increment k in the sequence

$$(27) \qquad W^*_k = \Sigma^k_{t=0} W_t + \sum_{t=k+1}^{T} E[W_t].$$

As the index evolves with time the instantaneous percentage change in the weather index can be calculated as

(28) $E[\mu] = E[(W^*_k - W^*_{k-1})/ W^*_{k-1}]$

and daily volatility is

(29) $\sigma^2 = E[\mu - E[\mu]]^2$.

Finally, under the assumption of mean reversion the path described by $\sum_{t=0}^{T} E[W_t]$ needs to be estimated. This can be done by using historical data directly but since this has to be recalculated for each day in the contracts life it is computationally intense. In the alternative, $\sum_{t=0}^{T} E[W_t]$ can be estimated from a simple regression equation to get the same result. In this study the estimated equation describing the evolution of temperatures during the summer months was quadratic.

The Pricing of Cooling Degree-Day Options

In this section option premia are calculated for Toronto Ontario using Environment Canada daily mean temperatures from 1840 to 1996. The contracts examine summer cooling degree-day call (put) spreads. With this option the buyer agrees to pay a fixed premium in exchange for payment from the seller if the defined Weather Index settles above (below) the Index Strike for the Contract Term. The payment equals the number of Weather Units the Weather Index falls above (below) the Index Strike times the Unit Price. There may be a payout limit but this is not considered in this study.

First the temperature history from June 1 through August 31 is described from a historical perspective. As history will always be the source of weather patterns it is important to understand how more recent trends compare to past trends.

Second, using a cooling degree-day measure of heat above 65 degrees Fahrenheit, degree-days are calculated for each day and cumulative degree-days are calculated for each year.

Third, a quadratic regression equation is estimated with mean daily degree-days as the dependent variable and time and time squared (within the contract term) as the independent variables.

Fourth, using mean cumulative cooling degree-days as the initial index value, observed daily degree-days, and the regression equation, the forward index value for each day, in each year was calculated.

Fifth, using the daily forward index values, the empirical volatility of the index is calculated from the variance of the daily percentage change in index values. This is done for each year.

Sixth, assuming a discount rate of 6.5%, the historical mean volatility, 92 days to expiration, and a strike price (which is varied), call and put option premiums are calculated. As a point of comparison premiums using the 'burn rate' approach are also calculated.

Toronto's Weather History

This section describes the weather history from June 1 to August 31 for the years 1840-1996 in Toronto. The data used were obtained from Environment Canada and represents one of the longest available weather data series in Canada. Figure 1 plots the data. The plot shows an overall increase in mean daily temperature over this time period, with temperatures increasing at an increasing rate until approximately 1930 and then increasing at a decreasing rate. Since Approximately 1950 there does not appear to be a significant rise in mean daily temperatures.

Figure 2 shows the cumulative cooling degree-days in Toronto between 1840 and 1996. The cooling degree-days increase with the mean temperature as would be expected, but the graph also illustrates the variability and unpredictability of the measure. The graph shows that cooling degree-days increased at an increasing rate throughout most of the 19th century but appear to be quite stable or decreasing in terms of mean value towards the end of the 20th century. Table 1 summarizes the key statistics for the entire 1840-1996 period and the sub period from 1930 to 1996. From 1840 the average cooling degree-days ranged from 107 to 787 with a mean of 379 and a standard deviation of 147. The period since 1930 has cooling degree-days ranging from 186 to 787 with a higher mean of 489 and a standard deviation of 114.

Figure 3 illustrates the mean actual and predicted daily degree-days within the 92-day period from June 1 to August 31. The pattern is parabolic and the statistical fit (using a quadratic equation) of predicted to average was approximately 93% (R-squared)[2]. Figure 4 illustrates the

---

[2] With daily temperatures about 65°F as the dependent variable the equation is $Temp = -.38 + .21T - .002T^2$ where T is day number (e.g. 1-92). Only the intercept is not statistically significantly different from zero.

cumulative degree-day effect throughout the time period. The degree-day value used in options pricing is the total sum recorded on the 92nd day.

Calculating the Cooling Degree-Day Weather Index

This section describes how the CDD weather index was calculated. The index was calculated for each year in order to assess the range of CDDs and to measure volatility. The cooling degree-day weather index was generated from a combination of observed daily data in each year, the seasonal regression equation, and the average cumulative degree-day value across all years. The initial index value at t=0 is assumed equal to the average cumulative degree-day value. This is identical to the sum of the marginal degree-days illustrated in figure 3. The smooth parabola in figure 3 illustrates how the regression equation smoothes the variability in daily degree-day measures and acts as an unbiased predictor of the most likely temperature path based on the assumption that weather patterns are mean reverting. To calculate the index the degree-day above 65f is calculated from the first observation (day 1). Then the sum of the predicted daily degree-days is calculated along the parabola from day 2 through to day 92. Assuming that the day one degree-day measure is small this will provide a day 1 index value very close to the long run average. On day 2, the actual degree-day measure is taken and is added to the day 1 value. The sum of the predicted is then taken from day 3 to day 92 and added to the actual day 1 plus day 2 values (see equation 17). The procedure is repeated for each of the 92 days (see equation 18), and is repeated for each year in the sample.

Figure 5 illustrates the results for three recent years in Toronto; 1986 was an average year with cooling degree-days of 386. The summer started of quite cool and this caused the index to fall below the average until about day 55 where a warming trend caused a slight increase in the value of the index; 1988 was a hot year and the index was above average throughout the season. A short cooling spell from day 31 to about day 40 caused the index to decrease but beyond that cooling degree-days were significantly higher than average. The 1988 index peaked at approximately 750 on day 80, but a cooling trend caused the index to fall to 725 by day 92; In contrast to 1986 and 1988, 1992 was unusually cool with cumulative degree-days of 186 by day 92. The index was average for the first 3 weeks of June, but after that a long cooling trend caused the index to fall to a low of about 180 before ending at 186.

Calculating Volatility

Volatility is measured relative to the percentage change in the value of the index on a daily basis and then converted to an annualized (365 day) basis for convenience. Table 2 and Figure 6 show the estimated average volatility for Toronto cooling degree-days from 1840 to 1996 and from 1930 to 1996.

The results indicate that the weather has actually been less variable since 1930 than in the previous 90 years. From 1840 to 1996 annualized volatility was .2063 or 20% per year, but this decreased to .1739 or 17% per year in the mid to latter part of the 20th century. For the entire period the minimum volatility was found to be 16.62% with a maximum of 29.61%, while the latter part of the century the range was as low as 14.14% but only went as high as 23.5%. Combined with the information in Table 1, weather averages in Toronto saw an increase in mean summer temperatures and degree-days, but this increase did not come with increased variability. In fact, the standard deviation of cumulative degree-days (Table 1) is lower for the 1930-1996 period than the 1840-1996 period. Importantly, these observations signify that when options on weather are being priced it is important to match recent weather trends on index values and volatility. In the next section, which calculates option premia, an approach, which mitigates this problem, is discussed.

Volatility Stability

Use of the options pricing model requires stability in the index's volatility within a given year and across years. The first item is important because if daily volatility is a function of time or is characterized by discernable jumps the proposed pricing model will be misspecified. The second is important because stability in volatility across years means that the sample volatility can be used as an unbiased estimate of volatility.

Volatility stability was measured by calculating the percentage daily change in the weather index in each year (91 days), i.e. In $[W_T/W_{t-1}]$. To determine the stability of volatility rolling 30-day standard deviations of the percentage change were calculated and annualized to a 365 day year. Thus for 91 days used in this study there were 61 volatility estimates for each year. Table 4 shows the results from this evaluation over the 1840-1996 period and two subperiods 1840-1935 and 1936-1996.

The annualized volatilities have been stable across years, with the average 30 day volatility being about 20%. This compares to the average volatility over the whole 91 days of .2063 as shown in Table 2. The results also show that the standard deviations are low relative to the mean. For example a standard deviation of .023 for 1840-1996 indicates that the average 30 day volatilities ranged from .178 to .223 approximately 67% of the time. The within year coefficient of variation (mean/standard deviation) reveals that the means are 6.42, 5.98 and 7.13 times the within-year 30-day standard deviations for each of the periods. These numbers imply that not only is volatility stable across years but they are quite stable within each year as well.

The Variance Ratio Test for Random Walks

The pricing of weather derivatives requires that the weather index W^*_T evolves over time as a random walk described by geometric Brownian Motion. Failure to support the random walk hypothesis would vitiate the model structure. In addition, the model assumes that the volatility of W^*_T is fairly stable. Failure to show stability in volatility would require expanding the model to include a volatility diffusion or jump process.

A general test for a random walk as presented by Lo and MacKinlay (1999) is the variance ratio test. Under the normal definition of a diffusion process the expectations are that the mean diffusion rate is constant and volatility is linear in time. Hence the mean return on an asset with two time steps will be twice that for a single time step and likewise the variance of the two time step will be double that of a single time step. These conditions can be stated as follows[3]

a) $E[W_{t+1} - W_t] = \mu_1$

b) $E[W_{t+k} - W_t] = k\mu_1$

c) $VAR[W_{t+1} - W_t] = \sigma_1^2$

d) $VAR[W_{t+k} - W_t] = k\sigma_1^2 = \sigma^2_k$

[3] If levels data are used the conditions apply to arithmetic Brownian motion, and if In(Wt) is used they apply to geometric Brownian motion. The tests in this paper use the logarithmic conversion, but I do not change the notation of Wt.

where k represents the number of time steps, μ_k represents the mean of a k-step (or k day) diffusion, and σ^2_k represents the variance of a k-step diffusion. This property allows for a simple test of the random walk by testing the null hypothesis Ho: $\frac{\sigma_k^2}{k\sigma_1^2} - 1 = 0$.

Lo and MacKinlay (1999) provide a measure of the asymptotic variance for the variance ratio. If

$$\text{(30)} \quad Z = \frac{\sigma_k^2}{k\sigma_1^2} - 1$$

then

$$\text{(31)} \quad \sigma_z^2 = \frac{2(2k-1)(k-1)}{3(N-1)k}$$

is the asymptotic variance of z when overlapping lags of length k are drawn from N observations. Thus the null hypothesis can be tested against the standardized normal z test with mean and variance σ^2_z. For example, if z<1.96 in equation (30) we would fail to reject the null hypothesis at the 5% level if (31) was used as the asymptotic population variance. In the alternatively one could also use an F-Test for the differences in variances when the means are equal. In this case the numerator is $\hat{\sigma}_k^2 = \sigma_k^2 / k$ so that the ratio $\hat{\sigma}_k^2 / \sigma_1 \sim F(N-k, N-1)$. The null hypothesis would be rejected if the ratio fell outside the two-tailed range of the F-distribution over a specified acceptance region.

Seasonality and the Variance Ratio Test

One of the concerns about pricing weather derivatives is the influence of seasonality on the random outcomes. Autocorrelation brought about by seasonal weather patterns can lead to rejections of hypothesis using the variance ratio test even if autocorrelation is spurious as is found in heat waves and so on.

The impact of systematic seasonal influences cannot be ignored but the effects can be removed. Removing systematic weather patterns leaves a path dependent residual that resembles a random walk. To see this define the daily temperature path above 65°F as a function of time as g(t). In the current study for example, g(t) = a+bt-ct^2 is a quadratic which fits nicely the summer weather patterns in Toronto. The function g(t) is a deterministic function of time. Estimated from daily data across many years non-systematic weather events are removed and captured by the residuals.

Given the definition of g(t), the expected value of cumulative cooling degree-days at time t=0 is

(32) $$W_0^* = \int_1^T g(t)dt$$

After day 1 the actual degree-day is calculated and its deviation from expected is recorded. At the end of day 1

(33) $$W_1^* = g(1) + e_1 + \int_2^T g(t)dt$$

and at the end of day 2

(34) $$W_2^* = [\int_1^2 g(t)dt + e_1 + e_2] + \int_3^T g(t)dt\,.$$

This can be generalized to any data t≤T as

(35) $$W_T^* = \int_1^t g(t)dt + \int_1^t e_t dt + \int_{t+1}^t g(t)dt$$

(36) $$W_t^* = \int_1^T g(t)dt + \int_1^t e_t dt$$

One can see from this process that the random part follows a walk over each time step, and is independent of the systematic influence of time. That is

(37) $$E[W^*_2 - W^*_0] = E[W^*_2 - W^*_1 + W^*_1 - W^*_0]$$
$$= E[e_1 + e_2]$$

or more generally

(38) $$E[W^*_t - W^*_{t-k}] = ke_t$$

as required for a random walk.

However, while this process is consistent with a random walk, the existence of a random walk must still be treated as a hypothesis. Recall that the function g(t) removed non-systematic weather events such as heat waves or other extraordinary items that could cause year specific transitory autocorrelation.

Transitory autocorrelation can be removed by averaging across years, but to do this the definition of the variance ratio must be modified accordingly. From the above results the mean value of the k-lag error in each year was shown to be ke_t.

For a variety of reasons the value of e_t will not necessarily equal $e_{t, n-1}$ or $e_{t, n+1}$ so across N years the expected value of ke_t is

(39) $$k\mu_1 = e[ke_1] = \frac{1}{N}\sum_{n=1}^{N} ke_{1,n}$$

where the subscript n has been added to denote the year of the observation.

The variance of equation (39) is given by

(40) $$\sigma^2_k = E[ke_1 - E[ke_1]]^2$$

$$= E\,[ke_1 - k\mu_1]^2$$

$$= k^2\,[e_1 - \mu_1]^2$$

(41) $$\sigma^2_k = k^2\sigma^2_1$$

This measure of σ^2_k contrasts with the standard measure of $\sigma^2_k = k\sigma^2_1$ in that the step multiplier k is squared. This results from averaging the mean lags from each year. However, since seasonality and spurious autocorrelation have been removed the variance measure is unbiased and asymptotically efficient and the variance ratio is

(42) $$VR_a = \frac{\sigma_k^2}{k^2\sigma_1^2}.$$

Moreover, the Lo and MacKinlay asymptotic estimator can still be used to test the null hypothesis Ho: $\frac{\sigma_k^2}{k^2\sigma_1^2} - 1 = 0$

Long-Run Versus Short-Run Effects

Having resolved problems of seasonality and asynchronous autocorrelation a final question to consider is whether the random walk hypothesis holds across a smaller number of years. Rather than averaging across the entire history of weather records (e.g. 1840-1996) a useful examination would be to examine shorter (e.g. overlapping 30-year) time horizons. The benefits to doing this are to first determine if acceptance or rejection of a random walk is due to long versus short time horizons, and second to examine the persistence or frequency of the random walk over time. Each 30-year sample can be considered an unbiased estimate of the larger population, but the asymptotic population variance is known. Therefore the standard errors can easily be estimated.

The standard error of the sample n<N is

$$(43) \qquad S_n = \left[\frac{2(2k-1)(k-1)}{3(N-1)k} \right]^{1/2} / \sqrt{n}$$

and this can be used in the denominator of a t-statistic with n-1 degrees of freedom.

Variance Ratio Test Results

The variance ratio hypothesis was tested using daily data from 1840-1996 for 92 days in each year. The tests were conducted by first calculating the value W^*_t,n for each day and across years, and then converting this data to logarithms. The results are presented in Table 4 for all years and 5 subperiods for lags of 1-10 days and lags of 35-40 days.

According to theory a random walk would be rejected if the means of the k-lag difference or the variance of the k-lag difference are significantly different than the values of k in the first column. Using N=92 in the calculation of the asymptotic variance in equation (31) and calculating the test statistic Z in equation (30) there are no instances where the random walk is rejected. Using N=92 the F-test fails to reject the null hypothesis that the variance ratio is significantly different than 1.0 in all case.

Failure to reject the null hypothesis on the variance ratio occurs even though there is a visible departure in the computed value of k in Table 3 from the theoretical value of k. The reason for this is that the asymptotic variance increases with k. For example when k=2 the asymptotic variance for N=92 is .011, but for k=40 it is .34. Therefore even though the 48.27 value of σ^2_{40}/σ^2_1 in the 1840-1870 subperiod is 8.27 points above the theoretical value of 40, the normalized variance ratio test $(\sigma^2_{40}/k^2\sigma^2_1 - 1)/\sigma_z$ is equal to 1.20 which falls below the critical value of 1.96 at the 5% level.

To examine whether the results in Table 4 are a consequence of chance or sampling the variance ratios were also calculated for overlapping 30 year periods and the null hypothesis was tested using the t-statistic which accounts for possible sampling error. The standard error is defined in equation (37) which divides the Lo and Mackinlay asymptotic variance measure by the square root of 30 (years). Of 128 overlapping time periods in no case was the null hypothesis rejected at the 5% level for up to 29 lags, and only 1 violation beyond that. Repeating the analysis for 20 and 10-day lags revealed that at 20 lags there were 29 rejections for an acceptance rate of 80.4% and at 10 lags there were 37 of 148 rejections for an acceptance rate of 75%.

The results offer strong support for the random walk hypothesis even when a small number of years are considered. But in this result also resides the caveate that to truly smooth individual year effects at least 30 years should be considered in practice. Notwithstanding this assertion there is sufficient evidence to conclude that indeed the index of cooling degree-days follows a random walk about the seasonal trend. Failure to reject the random walk also implies that volatility jumps are probably not of great concern. This does not imply volatility is a constant value, but it does imply that an average value of volatility across years, $E[\sigma]$ is an unbiased estimate of volatility. Furthermore given the evidence in Table 3 the estimate $E[\sigma]$ will be consistent and efficient. The evidence suggests that the option pricing model proposal in this paper is appropriate for the pricing of degree-day weather options.

Estimates of Cooling Degree-Day Option Premia

This section reports actual option premiums calculated for Toronto, Ontario. The contracts considered are 92-day put and call options with contract terms from June 1 with an expiry on August 31. Each tick in-the-money (θ) was valued at \$5,000 per degree-day. Several empirical considerations are illustrated in the results. First, premium estimates are calculated using the both the inter-year 'burn-rate' method used in the insurance industry (equations 8 and 9) and the intra-year Black's option pricing model (equations 22 and 23). Second, in order to illustrate the importance of 'relevant time horizon', estimates are provided for the 1840-1996 data period and the 1930 to 1996 sub-period. Third, the options pricing model is sensitive to the initial index value, W^*_0, and using a simple average in all cases would not be prudent. For the options pricing model only, a range of initial values of W^*_0 are examined. This type of sensitivity analysis is important because weather agencies such as Environment Canada and the U.S. Weather Service cannot generally predict forward temperatures with reasonable accuracy. However, they can and do provide three or four-month forecasts that state whether conditions are going to be normal, below normal, or above normal. If the prediction is above normal, for example, the buyer of a call may want to increase the initial expectation of W^*_T to match the forecast and reduce the premium.

Tables 5 for 1840-1996 and 6 for 1930-1996 present results for base case at-the-money option pricing calculations as well as a range of strike prices above and below this value. The at-the-money strike is defined as the average cooling degree-days across the years sampled. This is

379.39 for 1840-1996 and 489.50 for 1930-1996. The option premiums differ between the options model and the burn rate model as well as across the two time periods. When the sampling period was represented from 1840 the at-the-money put and call price was $77,073 for the 379-CDD strike option model and approximately $297,030 for the burn rate model (Table 5). The maximum payoff for the put option under either case would have been $1,361,450 for the put option and $2,038,150 for the call option. As the strike price was increased put options would be issued in-the-money and the put option premiums would rise as the call premiums fell. For a strike of 600 CDD the option model put premium was $1,085,126 while the burn-rate model was $1,136,421. The maximum put payoff increased to $2,464,500. The corresponding call option for the option model was $0 and for the burn-rate model it was $33,405. The maximum payoff that would have possibly occurred with this strike over this period was $935,100. A lower than average strike implies that put options are issued out-of-the-money, while call options are issued in-the-money. At a strike of 250 CDD the put options price is negligible, while the call option price is $636,438. Using the burn-rate model the corresponding put and call prices were $63,947 and $710,420 with maximum payoffs being $714,500 for the puts and $2,685,100 for the calls.

A similar pattern was observed for the 1930-1996 period (Table 6). The at-the-money option price (489.5 CDD) for the put and call was $83,835 and using the burn-rate model the put-call price was approximately $220,358. The maximum put and call payoffs would have been $1,516,900 and $1,487,600 respectively. For in-the-money calls with a strike of 250 CDD the call option was $1,178,041 and the corresponding put was $0. The burn-rate put and call prices were $4,767 and $1,202,279 respectively, with maximum payoffs of $319,400 and $2,685,100. For in-the-money puts at 600 CDD the put option price was $544,298 and the call price was only $776. The burn-rate premiums were $624,900 and $72,412 for the put and call respectively.

These results illustrate some important and critical details regarding the pricing of degree-day derivatives and the selection of a time period over which to analyze heat. The difference between options pricing and burn-rate models is striking, especially when priced at-the-money. Using the 1840-1996 period the burn-rate model prices the insurance at 3.85 times the option pricing model whereas the 1930-1996 period the pricing multiple is 2.63. The ratio converges to 1 for policies that are in-the-money and infinite for options out-of-the-money. The results illustrate why different approaches to pricing weather options can result in large bid-ask spreads.

The explanation for these differences lies in how risk is measured and what risks are actually being traded. The burn rate model assumes that history will repeat itself and the variability and probability distribution of the past will be replicated in the future. It rests upon an actuarial structure, which is seemingly predictable, but one, which also carries with it some significant variability. In contrast the options pricing model is not backward looking in the sense of a memorized historical probability distribution. It assumes an infinite of random weather patterns, which can occur in any season. The role of history is vague only in its use to establish seasonal norms and a range of volatility measures, but once these are established history's role is done.

Another key difference is the assumption of a starting point. The options pricing model assumes a numerical starting position from which variability in a weather index is measured, and the price of the option is sensitive to this initial position. For example the further the index strike is below the initial index value the higher will be the value of a call option and lower will be the value of a put option. Because the burn-rate model's principal Gaussian assumption is that history will repeat itself, the burn-rate model does not require an estimate of the initial weather index value perse.

As discussed above the initial assumption regarding the forward weather index is crucial to the accurate pricing of options. Tables 5 and 6 present results assuming that the initial index value is equal to the historical mean. In reality this may not be the case. Weather forecasts may predict higher or lower than normal temperatures and this will have a conditional impact on what the initial index value is. For example mean growing degree-days for the 1930-1996 period was 489.5 with a range from 186 to 787 and a standard deviation of 114 (Table 1). If the long-range weather forecast was for warmer than usual weather, then it would be prudent to increase the initial index value accordingly so that the likelihood that a put option would end in-the-money is lower and the likelihood that the call would end in-the-money is higher. Likewise, if the long-range forecast was for cooler than normal weather then the index would be decreased such that the likelihood of a put ending in-the-money would increase and the likelihood of a call ending in-the-money would decrease.

Table 7 presents option pricing sensitivity results for the 1840-1996 period and Table 8 presents results for the 1930-1996 period. Since the burn-rate model does not rely on initial conditions only the option pricing model is considered. Each column in Tables 7 and 8 represent a percentage of the mean with 1.0 representing the mean, .50 being 50% of the mean and 150% being 50% higher than the mean. As predicted by options theory as the initial condition decreases the put option value increases and the call option value increases holding the strike level constant. For example if the strike level is 400 CDD the values in Table 7 for the 1840-1996 period for a put option is $1,034,468 and call price is $0 if the weather forecast implies that cumulative degree-days will be 50% less than average. If the cooling trend is believed to be less severe, say 75%, then the put value for a 400 CDD strike decreases to $507,929 and the call price is $23. If the weather prediction calls for a 50% increase in cooling degree-days then the likelihood that a 400 CDD put option will expire in-the-money is nil, and the put option is priced at $21. In contrast, the likelihood that the call option expires in-the-money rises and the call option premium increases to $831,708. A similar pattern is illustrated for the 1930-1996 period in Table 8.

Conclusions

This paper addressed the pricing issue of degree-day weather derivatives. The market for weather insurance products has increased dramatically in past years for several reasons. First weather derivatives are directed at hedging production or volume versus price risk. In the natural gas and energy sectors, utilities will often fix prices to the consumer or face regulated prices to consumers. Electrical utilities must of ten pay peak-load prices when energy demand exceeds contracted supplies, and natural gas and oil companies must pay higher spot prices when extreme cold causes excess demand in those markets. Agriculture is also an industry that faces weather related production risk. A crop insurer might have to pay increased indemnities if weather is either too hot or too cool, and might use weather derivatives as a reinsurance product, or a food processor might require a hedge against undeliverable forward contracts resulting from weather conditions.

An important driver of the weather derivatives market is the relationship between economic damage and specific events. Electrical utilities know with 100% certainty that prolonged above normal heat or below normal coolness will create an increased demand for their products, and they can also determine statistically at what level, in cumulative heating or cooling degree-days, this occurs. What is unknown is when the specific event will occur and with this uncertainty routine hedging of weather risks can provide economic stability and increase share values.

This paper examined the pricing methods of degree-day derivatives. It was shown that many of the underlying assumptions used in modern options pricing are relevant to weather conditions. The critical, and justifiable, assumption is that weather risks follow a Martingale, and based on this assumption the stochastic differential equation which drives weather dynamics follows an Ito process. It was shown that applying arbitrage free arguments to this stochastic process results in a pricing formula similar, but not identical to Black's formula for pricing European options on futures. The key difference between the pricing model developed in this paper is that price per degree-day is held constant while the quantity variable (degree-days) varies, whereas the original Black's model holds quantity constant (e.g. 1 bushel) while allowing price to vary. Another difference is that Black's (1976) model is derived from trading in an underlying futures contract which is subject to many supply and demand influences which create volatility and liquidity whereas the weather option relies on a non-traded forward weather index. An approach to defining such an index was discussed at length in this paper.

The approach used in this paper differs markedly from an insurance approach to pricing weather derivatives. The 'burn-rate' approach, prices premiums based upon what would have occurred over a recent time period. It was pointed out that the key difference between the burn-rate model and the options pricing model is in how risk is defined. Under the burn-rate model it is assumed that history will repeat itself with the same likelihood, but not necessarily the same order, as the time horizon selected for pricing. In other words, the approach assumes that the relevant measure of risk is the inter-year variability in weather. The options pricing model developed in this study makes no such assumption and is in fact based on intra-year risks. As with conventional options pricing, volatility and the initial value of the weather index are the key drivers of risk. History is used only to measure volatility and determine a range of index values, but once a measure of volatility is selected and the initial condition determined, history has no further role to play in the pricing process. For example the 1840 to 1996 period had mean cooling degree-days (above 65f) of 379 CDD and an annualized volatility of 20.63% for the period June 1 to August 31. Using the 1930-1996 period the average cooling degree-days was

489 CDD with a volatility of 17.39%. Under no year was volatility found to exceed 29.6%, yet the implied volatility that would equate the options pricing model to the burn rate model was 80% for the 1840-1996 period and 45.8% for the 1933 to 1996 period.

It was shown that there is a significant and often large difference between the burn-rate model and the options pricing model, particularly for products priced at or near-the-money. It was shown that the burn rate model prices options as much as 2 to 3 times higher than the options pricing model. The two approaches converge only for options that are priced in the money or out of the money. It is consistent with the various theories of pricing non-traded assets in equilibrium, and in a risk-neutral economy. Statistical analyses confirmed that the underlying assumptions required for pricing degree-day weather options are empirically valid.

The options pricing model presented in this paper is new. On one hand it is an improvement over the traditional burn-rate approach in that it places much more emphasis on risk and for a derivatives market which is essentially designed to manage the buying and selling of risk there can be efficiency and liquidity gains if the model is implemented in practice. On the other hand the traditional approach is easy to implement and even easier to comprehend. However, if a formal derivatives market for weather insurance is going to emerge it is very likely that the approach developed in this study will provide foundation for pricing weather derivative products.

Table 1: Historical Summary of Toronto Cooling Degree-Days				
	Mean	**Standard Deviation**	**Minimum**	**Maximum**
1840-1996	379.39	146.67	107.10	787.02
1930-1996	489.50	114.69	186.12	787.02

Table 2: Historical Summary of Toronto Cooling Degree-Days' Volatility				
	Mean	**Standard Deviation**	**Minimum**	**Maximum**
1840-1996	.2063	.0012	.1662	.2961
1930-1996	.1739	.0009	.1414	.235

Table 3: Seasonality and Stability in Volatility			
	1840-1996	1840-1936	1936-1996
Mean (365 days)	.201	.207	.193
Standard Deviation	.023	.022	.021
Coefficient of Variation Mean	6.42	5.98	7.13
Coefficient of Variation Standard Deviation	3.19	3.03	3.32

Table 4: Estimated Value of Means and Variance Ratios (VR)*												
	1840-1996		1840-1870		1871-1900		1901-1930		1931-1960		1961-1996	
Lag K	Mean	VR	Mean	VR	Mean	VR	Mean	VR	Mean	VR	Mean	VR
1	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.0
2	1.99	2.01	2.00	2.01	2.01	2.00	2.00	2.01	2.01	2.03*	2.01	2.00
3	2.99	3.02	3.00	3.03	3.03	3.01	3.01	3.02	3.02	3.08*	3.04	3.01
4	3.99	4.04	4.00	4.06	4.06	4.03	4.00	4.04	4.03	4.14	4.07	4.02
5	4.98	5.06	5.00	5.09	5.09	5.07	4.95	5.06	5.03	5.23	5.11	5.03
6	5.97	6.09	6.00	6.14	6.14	6.11	5.87	6.09	6.03	6.32	6.15	6.05
7	6.96	7.13	7.00	7.21	7.18	7.16	6.79	7.12	7.05	7.41	7.18	7.10
8	7.95	8.17	8.01	8.30	8.22	8.21	7.71	8.17	8.08	8.52	8.20	8.19
9	8.95	9.23	9.02	9.40	9.26	9.26	8.65	9.23	9.12	9.65	9.22	9.22
10	9.97	10.30	10.03	10.54	10.32	10.32	9.62	10.29	10.18	10.80	10.23	10.24
35	35.47	31.05	36.07	42.07	38.99	40.99	30.39	38.50	37.27	41.56	37.74	40.85
36	36.47	40.22	37.11	43.28	40.22	42.27	31.31	39.67	38.35	42.75	39.01	42.03
37	37.46	41.58	38.14	44.50	41.44	43.56	32.23	40.82	39.41	43.91	40.28	43.20
38	38.47	42.55	39.16	45.75	42.69	44.86	33.12	41.94	40.47	45.06	41.54	44.38
39	39.47	43.73	40.16	47.02	43.95	46.15	34.01	43.18	41.51	46.23	42.79	45.57
40	40.48	44.90	41.5	48.27	45.20	47.46	34.95	44.37	42.56	47.41	44.02.	46.77

* The variance ratio (VR) is $\sigma^{*}_{k}/k\sigma^{2}_{1}$ and the mean ratio is $\mu k_{k}/\mu_{1}$.

Table 5: European Options and Burn Rate Premiums: 1840-1996, Tick = $5,000						
	Option Value		Burn Rate		Maximum Payoff	
	Put	Call	Put	Call	Put	Call
200	0	882,374	18,215	915,190	464,500	2,935,100
250	0	636,438	63,647	710,420	714,500	2,685,100
300	692	391,264	135,264	533,239	964,500	2,435,100
350	23,156	167,718	229,910	376,885	1,214,500	2,185,100
379.39	77,073	77,073	297,054	1,361,450	2,038,150	2,038,150
400	139,950	38,574	352,121	249,096	1,464,500	1,935,100
450	351,674	4,361	508,943	155,918	1,714,500	1,685,100
489.50	542,100	497	657,13	106,788	1,912,000	1,487,600
500	539,513	263	698,560	95,534	1,964,500	1,435,100
550	839,198	10	908,806	55,781	2,214,500	1,185,100
600	1,085,126	0	1,136,421	33,405	2,464,500	935,100
650	1,331,062	0	1,370,114	17,089	2,714,500	685,100

Table 6: European Options and Burn Rate Premiums: 1930-1996, Tick = $5,000						
	Option		Burn Rate		Maximum Payoff	
	Put	Call	Put	Call	Put	Call
200	0	1,423,978	1,035	1,448,548	69,400	2,935,100
250	0	1,178,041	4,767	1,202,279	319,400	2,685,100
300	0	932,103	8,498	956,010	569,400	2,435,100
350	2.52	686,168	20,119	717,631	819,400	2,185,100
379.39	94	541,697	841,749	39,261	2,319,400	685,100
400	670	440,897	50,519	498,031	1,069,400	1,935,100
450	17,974	212,266	121,639	319,150	1,319,400	1,685,100
489.5	83,835	83,835	220,358	220,370	1,516,900	1,487,600
500	113,047	61,400	249,622	197,134	1,569,400	1,435,100
550	306,996	9,411	417,813	115,325	1,819,400	1,185,100
600	544,298	776	624,900	72,412	2,069,400	935,100
650	789,497	37	841,749	39,261	2,319,400	685,100

Table 7: Sensitivity of Options Prices to Initial Conditions, 1840-1996

	.50		.75		1.0		1.25		1.50	
Strike	**Put**	**Call**	**Put**	**Call**	**Put**	**Call**	**Put**	**Call**	**Put**	**Call**
200	69,975	19,287	10	415,853	0	882,375	0	1,348,906	0	1,815,437
250	296,757	132	6,873	176,779	1	636,438	0	1,102,968	0	1,569,499
300	542,563	0	104,962	28,931	692	391,192	0	857,031	0	1,323,562
350	788,500	0	323,336	1,367	23,156	167,718	100	611,193	0	1,077,624
400	1,034,468	0	567,929	23	139,950	38,574	4,644	369,800	21	831,708
450	1,280,375	0	813,844	0	351,674	4,361	46,048	165,266	1,039	586,788
500	1,526,313	0	1,059,781	0	593,513	263	174,937	48,218	13,747	353,559
550	1,772,250	0	1,305,719	0	839,198	10	381,556	8,900	72,838	166,712
600	2,018,188	0	1,551,656	0	1,085,126	0	619,676	1,081	209,925	57,861
650	2,264,125	0	1,797,594	0	1,331,063	0	864,624	92	412,954	14,593

Table 8: Sensitivity of Options Prices to Initial Conditions, 1930-1996

	.50		.75		1.0		1.25		1.50	
Strike	**Put**	**Call**	**Put**	**Call**	**Put**	**Call**	**Put**	**Call**	**Put**	**Call**
200	335	220,449	0	822,046	0	1,423,978	0	2,025,910	0	2,627,842
250	56,523	30,700	0	576,109	0	1,178,041	0	1,779,973	0	2,381,904
300	272,149	388	502	330,673	0	932,103	0	1,534,035	0	2,135,967
350	517,699	0	28,254	112,488	2.52	686,168	0	1,288,098	0	1,890,030
400	763,636	0	175,874	14,170	670	440,897	0	1,042,160	0	1,644,092
450	1,009,573	0	408,223	582	17,974	212,266	12	796,234	0	1,398,155
500	1,255,510	0	653,588	9	113,047	61,400	837	551,122	0	1,152,217
550	1,501,448	0	899,516	0	306,996	9,411	13,374	317,721	33	906,313
600	1,747,386	0	1,145,454	0	544,298	776	77,173	135,584	1,004	661,346
650	1,993,323	0	1,391,391	0	789,497	37	226,707	35,180	10,968	425,373

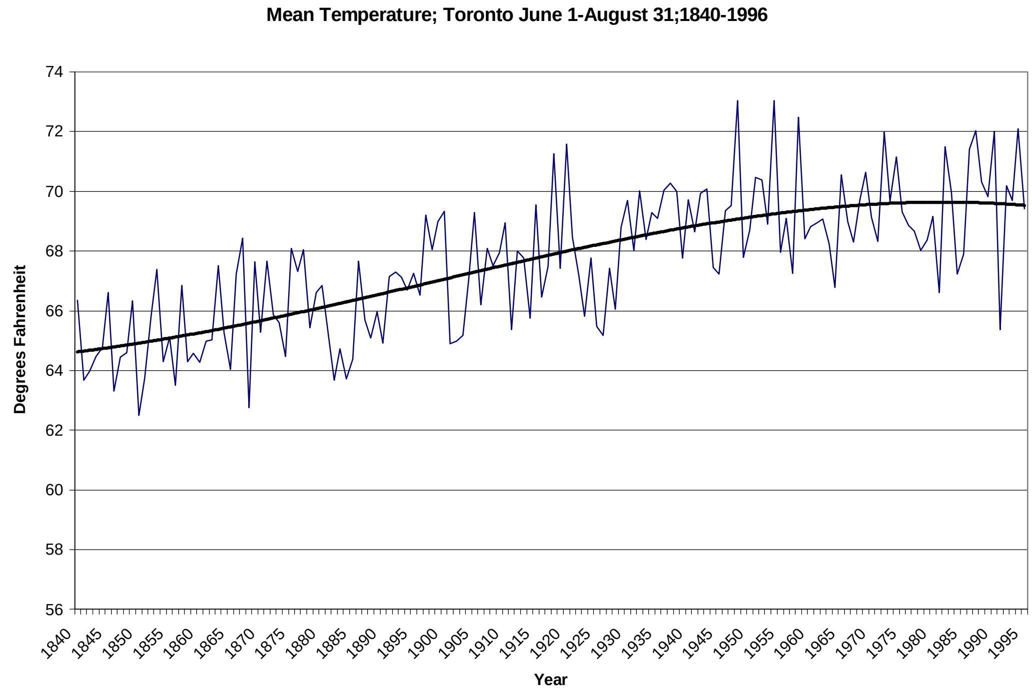

Figure 1:Mean Seasonal Temperature, Toronto, June 1 to August 31

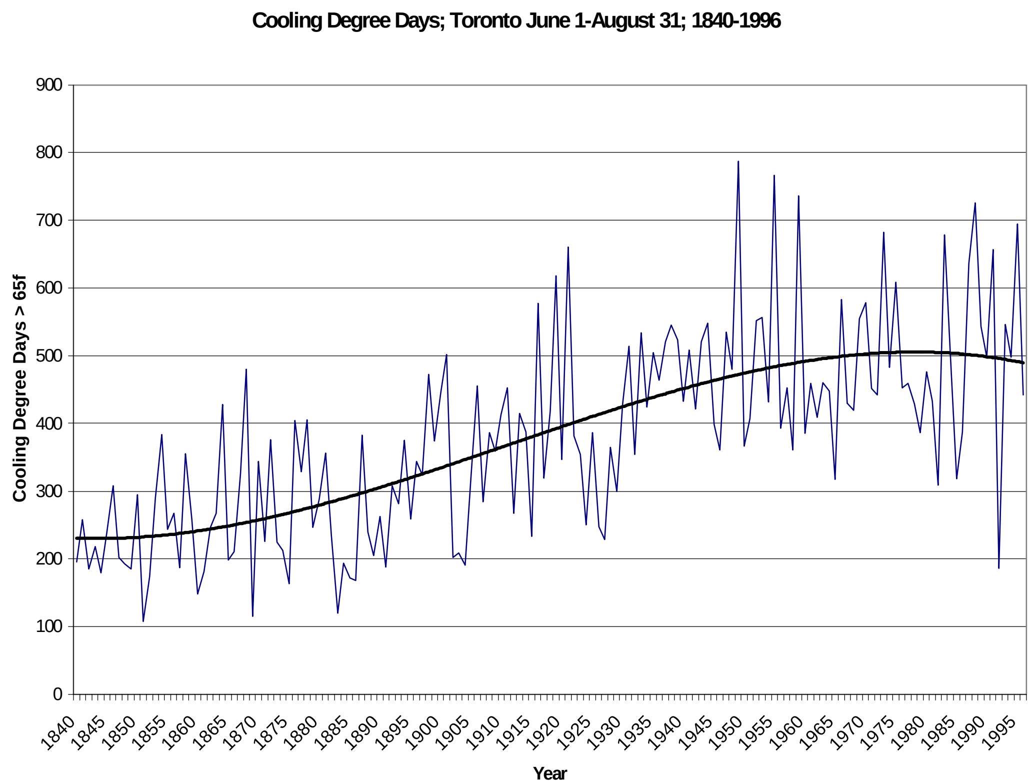

Figure 2:Mean Actual and Predicted Daily Degree-Days, Toronto, June 1 to August 31

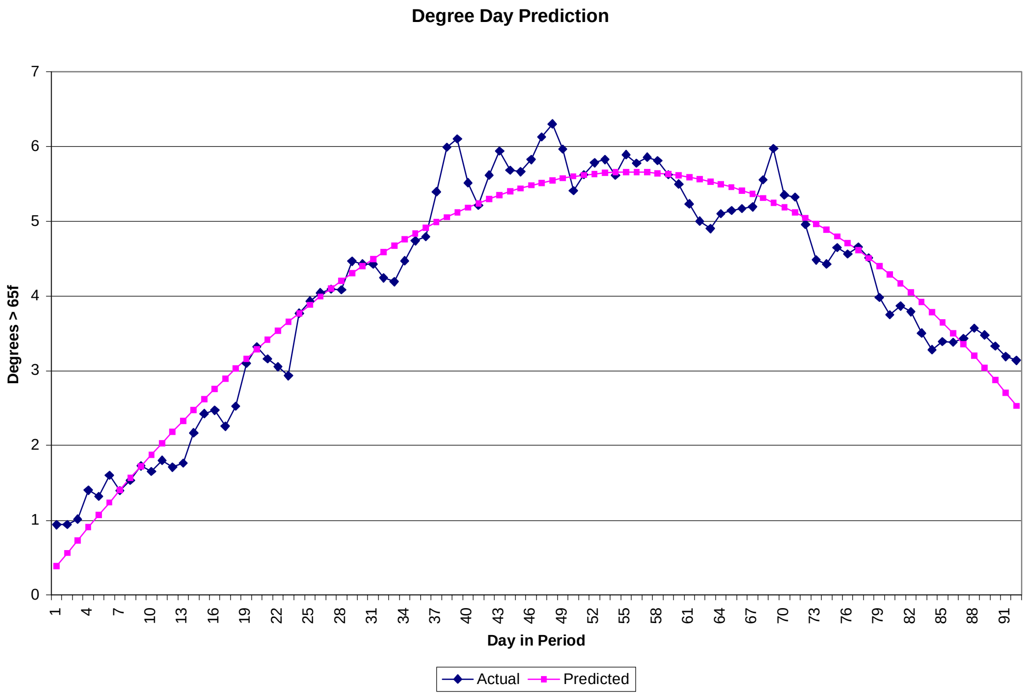

Figure 3:Mean Daily Cooling Degree-days, actual and predicted, Toronto, June 1 to August 31

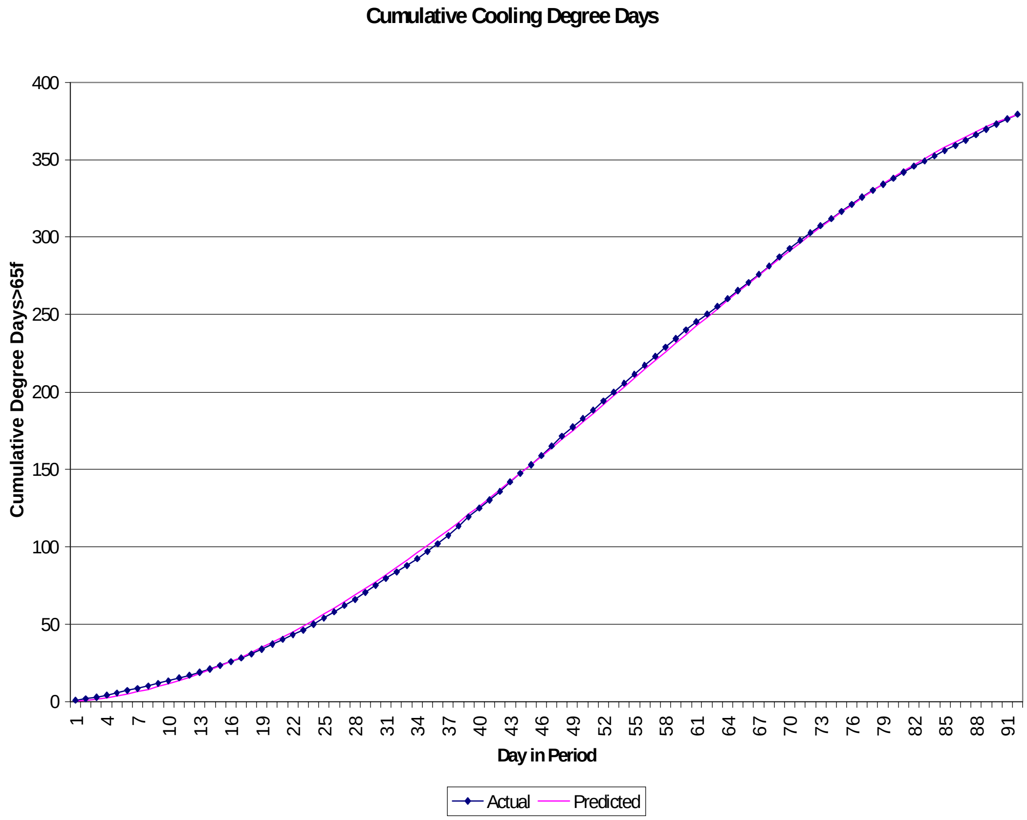

Figure 4:Cumulative Cooling Degree-days, Actual and Predicted, Toronto, June 1 to August 31

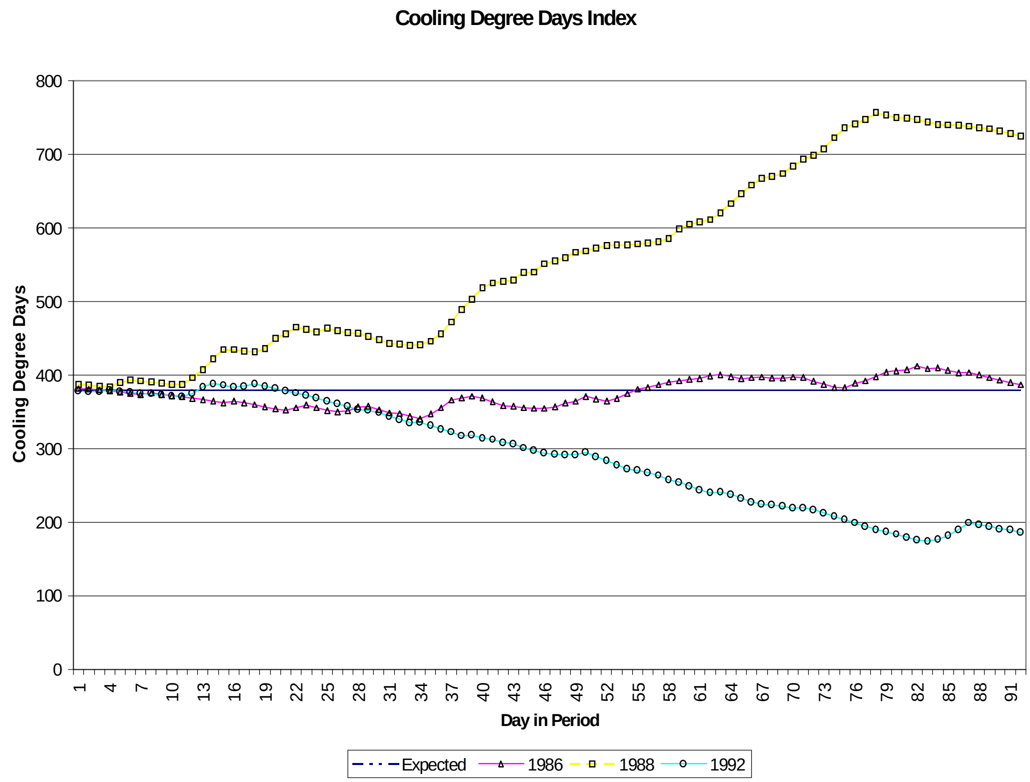

Figure 5: Cooling Degree-Day Weather Indexes for 1986 (average), 1988 (above average) and 1992 (below average), Toronto, June 1 to August 31

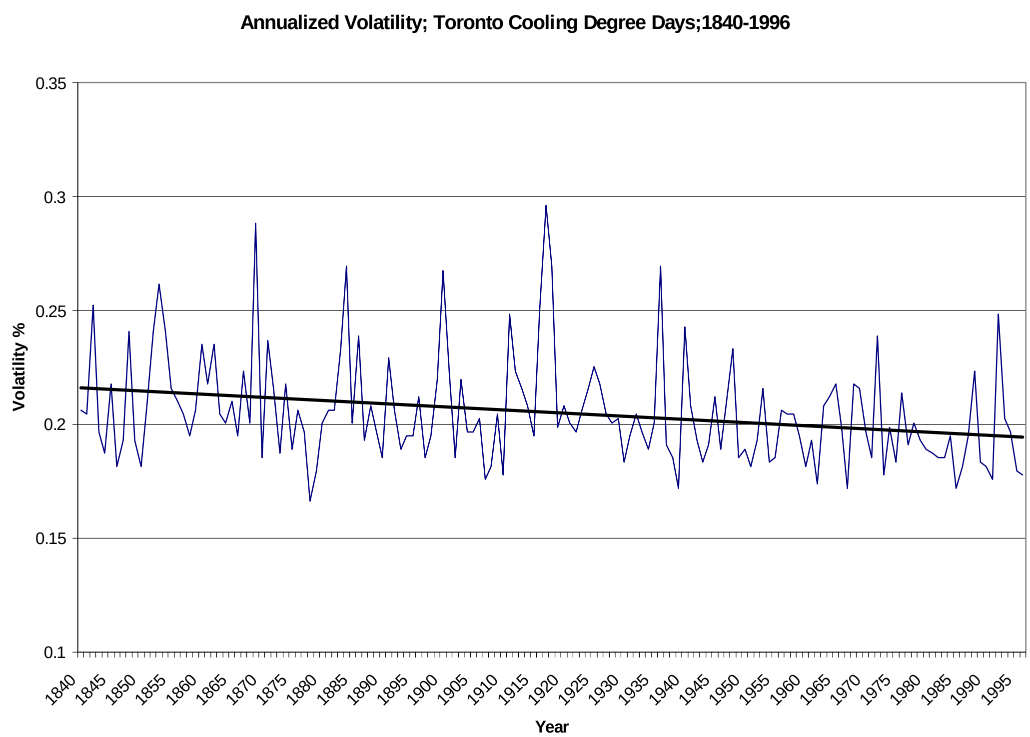

Figure 6:Mean Annualized (365 day) Volatility, Toronto, June 1 to August 31. 1840-1996

References

Black, F. and M. Scholes (1973). "The Pricing of Options and Corporate Liabilities". Journal of Political Economy. 81:637-654.

Black, F. (1976). "The Pricing Commodity Contracts". Journal of Financial Economics. 3:167-179.

Cao, M, and J. Wei (1999) "Pricing Weather Derivatives: An Equilibrium Approach" Working Paper, The Rotman Graduate School of Management, University of Toronto, October

Considine, G. (undated) "Introduction to Weather Derivatives" CME Working paper at http://www.cme.com/allaire/spectra/system/securemediastore/introweather.pdf

Cox, J.C, J.E. Ingersoll, and S.A. Ross (1985) "An Intertemporal General Equilibrium Model of Asset Prices" Econometrica 53:363-384.

Cox, J.C. and S.A. Ross (1976). "The Valuation of Options for Alternative Stochastic Processes." J. Financial Economics. 3:145-166.

Davis, M.(2001) "Pricing Weather Derivatives by Marginal Value" Quantitative Finance 1:305-308.

Dixit, A. and R.S. Pindyck (1994) Investment Under Uncertainty Princeton University Press, Princeton N.J..

Dischel, Bob (1998) "Black-Scholes Won't Do" Weather Risk Special Report, Energy & Power Risk Management, Risk Publications October.

Dusak, K. (1973). "Futures Trading and Investor Returns: A Investigating of Commodity Market Risk Premiums." J. Political Economy. 81:1387-1405.

Hull, J.C. (1993) Options, Futures and Other Derivative Securities, Prentice Hall, New Jersey.

Lintner, J. (1965). "The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets." Review of Econ. and Stat. 47:13-37.

Lo, Andrew W. and A.C. MacKinlay (1999) A Non-Random Walk Down Wall Street. Princeton University Press, Princeton, N.J.

Lucas, R.E. (1978) "Asset Prices in an Exchange Economy" Econometrica 46:1429-1445.

Merton, R.C. (1993). Continuous Time Finance. Basil Blackwell, Cambridge Ma.

Nelken, Izzy (1999) "Weather Derivatives – Pricing and Hedging" Super Computer Consulting Inc., Mundelein, Illinois. www.superc.com

NatScource Inc. (1998/1999). New York, N.Y. Propriety Discussion.

Roll, R. (1977). "A Critique of the Asset Pricing Theories Test." J. Financial Economics. 4:129-176.

Sharpe, W.F. (1964). "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." J. of Finance. 191(3):425-442.

Turvey, C.G. (2001) “Weather Insurance and Specific Event Risks in Agriculture” Rev. Of Agric. Econom. Fall/Winter

World Wide Weather Inc. (1999). www.weatherins.com

Weather Risk Advisory Ltd. (2000) "The European Weather Derivatives Market". erivativesreview.com, October.

Wilmott, P. (1998) Derivatives: The Theory and Practice of Financial Engineering. John Wiley and Sons, U.K.