Pricing Weather Derivatives by Marginal Value

Mark Davis[*]
Imperial College London

8 December 2000

Abstract

Weather derivatives are a classic incomplete market. This paper gives a preliminary exploration of weather derivative pricing using the 'marginal substitution value' or 'shadow price' approach of mathematical economics. Accumulated heating degree days (HDD) and commodity prices are modelled as geometric Brownian motion, leading to explicit expressions for swap rates and option values.

1 Introduction

Many companies are exposed to 'weather risk'. For concreteness, we shall think in terms of an energy company supplying gas to a retail distributor. If a winter month such as January is unusually warm then the company's profits are adversely affected because of the reduced *volume* of gas sold. Note that this is a separate issue from *price risk* which may also be present. The company can partially hedge the volume risk by trading in weather derivatives, which are normally defined as follows (see Geman [4] for extensive background information). Let T_i, the 'temperature on day i' be the average of the maximum and minimum temperatures in degrees Celsius on that day at a specific location (London Heathrow airport in the UK). The daily number of 'heating degrees' is $\mathrm{HDD}_i = \max(18 - T_i, 0)$ and the accumulated 'heating degree days' (HDDs) over a 1-month (31-day) period ending at date t is $X_t = \sum_{i=0}^{30} \mathrm{HDD}_{t-i}$. Over-the-counter contracts are written with X_t as the 'underlying asset'. These may be swaps, the payment at time T being $A(\kappa - X_T)$ where A is the point value and κ a fixed number of accumulated HDDs, or they may be options with exercise value $A\max(X_T - K, 0)$ for a given strike K. The question is what is the value of these contracts, i.e. the level of the fixed side κ such that the swap has zero value, or the premium to be paid at time 0 for the call option.

---

[*] www.ma.ic.ac.uk/∼mdavis

Since there is no liquid market in these contracts, Black-Scholes style pricing is inappropriate. Valuation is generally done on an 'expected discounted value' basis, discounting at the riskless rate but under the physical measure, which throws all the weight back onto the problem of weather prediction. This approach is to some extent justified in an interesting recent paper by Cao and Wei [1] who analyse the problem in an equilibrium representative agent setting based on that of Lucas' classic paper [6].

Our analysis here is based on the idea that agents in the weather derivative market are not 'representative' but face very specific risks related to the effect of weather on their business. From this point of view it makes sense to say that agents will buy or sell weather derivatives if it increases their utility to do so. We can now apply the general pricing formula given in Davis [2], which is an application of a general theory of valuation based on optimal consumption and investment rules. An excellent exposition of this theory is given by Foldes [3]. The objective of this paper is to formulate the problem in the simplest possible setting, putting the emphasis on analytic tractability. Thus we take accumulated HDDs and commodity prices to be log-normal, an assumption that – as we argue below – is quite reasonable based on empirical analysis. In this framework we can get explicit formulas for swap rates and option values. We can then ask qualitative questions, for example, does the swap rate depend on volatility? (Answer: yes it does).

The main results are contained in the next section. Section 3 gives a brief analysis of some weather data, in order to investigate whether a log-normal distribution for accumulated HDDs is at all realistic. Section 4 gives numerical results for an example. Concluding remarks are given in section 5

2   Pricing Formulas

We model the the accumulated HDDs (over, say, a 1-month sliding window ending at time t) by a log-normal process X_t satisfying

$$dX_t = \nu X_t dt + \gamma X_t dw_1(t). \tag{1}$$

Thus at time T,

$$X_T = \exp(m(T) + \gamma w_1(T)) \tag{2}$$

where

$$m(T) = \log X_0 + (\nu - \frac{1}{2}\gamma^2)T. \tag{3}$$

For pricing a weather derivative maturing at time T the main object of concern is simply the 1-dimensional random variable X_T, and our basic assumption is that this is log-normal, as indicated by (2). We suppose that the volume of gas sold per unit time is some function $v(t) = v(X_t)$ and suppose that – at least over some range – we can take $v(\cdot)$ as linear:

$v(t) = \alpha X_t$. The profit is therefore $Y_t = \alpha X_t S_t$, where S_t is the spot price. As is conventional, we suppose the price to be log-normal:

$$dS_t = \mu S_t dt + \sigma S_t dw_2(t). \tag{4}$$

In these equations, w_1, w_2 are standard Brownian motions with correlation $E[dw_1 dw_2] = \rho dt$. From (1) and (4), Y_t satisfies

$$dY_t = \theta Y_t dt + \xi Y_t dw(t) \tag{5}$$

with $Y_0 = \alpha S_0 X_0$, where

$$\theta = \nu + \mu + \rho \sigma \gamma$$

and

$$\xi = \sqrt{\gamma^2 + \sigma^2 + 2\rho\gamma\sigma}.$$

The new Brownian motion is

$$dw = \frac{1}{\xi}(\gamma dw_1 + \sigma dw_2).$$

Suppose the weather derivative has exercise value $B(X_T)$ at time T. In [2] we gave a valuation formula for an investor whose overall objective is to maximize the expected utility $E[U(H_T)]$ of his portfolio value H_T at time T. This value is

$$\hat{p} = \frac{E[U'(H_T^*)B(X_T)]}{V'(\eta)} \tag{6}$$

where H_T^* is an optimal portfolio of tradable assets with initial endowment η and $V(\eta) = E[U(H_T^*)]$. In the present case our producer has no investment decisions: he simply produces up to the level of current demand and sells at market price. Thus $H_T^* = Y_T$, the profit at time T. We will assume utility is logarithmic, $U(y) = \log y$, and then it is easy to see that $V(y) = \log y + $const. Thus $V'(y) = 1/y$ and the pricing formula (6) becomes

$$\hat{p} = E\left[\frac{Y_0}{Y_T}B(X_T)\right] \tag{7}$$

Proposition 1 *The zero-cost swap rate at time 0 is*

$$\hat{\kappa} = e^{(\nu - \gamma^2 - \rho\sigma\gamma)T}X_0. \tag{8}$$

The option value (7) with $B(x) = [x - K]^+$ is given by

$$\hat{p} = \mathrm{BS}(x_0, K, r, q, \gamma, T), \tag{9}$$

the Black-Scholes call-option formula, in which the 'riskless rate' r and 'dividend yield' q are given by

$$\begin{aligned} r &= \mu + \nu - \gamma^2 - \sigma^2 - \rho\sigma\gamma, & (10)\\ q &= \mu - \sigma^2. & (11) \end{aligned}$$

PROOF: Defining $Z_t = Y_0/Y_t$ we find using (5) and the Ito formula that

$$dZ_t = -rZ_t dt - \xi Z_t dw_t, \quad Z_0 = 1,$$

where r is given by (10). Thus

$$
\begin{aligned}
\hat{p} &= E[e^{-rT} \exp(-\xi^2 T/2 - \xi w_T) B(X_T)] \\
&= \hat{E}[e^{-rT} B(X_T)]
\end{aligned}
\tag{12}
$$

where $\hat{E}$ denotes expectation with respect to the measure $\hat{P}$ defined by

$$\frac{d\hat{P}}{dP} = \exp(-\xi^2 T/2 - \xi w_T)$$

Recall that X_t satisfies (1). We find that $E[dw dw_1] = \rho_1 dt$ where $\rho_1 = (\gamma + \rho\sigma)/\xi$, and $d\hat{w} = dw + \xi dt$ is Brownian motion under $\hat{P}$ by the Girsanov theorem. It follows that under $\hat{P}$ there is a Brownian motion $\hat{w}_1$ such that

$$dX_t = (\nu - \rho_1 \gamma \xi) X_t dt + \gamma X_t d\hat{w}_1(t). \tag{13}$$

We note that the 'drift' is $\nu - \rho_1 \gamma \xi = \nu - \gamma^2 - \rho\sigma\gamma = r - q$ with q defined by (11). Thus when $B(X_T) = X_T - \kappa$ we have $\hat{p} = e^{-qT} X_0 - e^{-rT} \kappa$, so the zero-cost swap rate is $\hat{\kappa} = e^{(r-q)T} X_0$; this is (8). In the case of a call option, $B(X_T) = [X_T - K]^+$, and the result (9) follows from (12) and (13). $\quad \diamond$

2.1 Comments

- The swap rate $\hat{\kappa}$ is not equal to the physical measure forward HDD $e^{\nu T} X_0$ but is equal to $e^{\tilde{r}T} X_0$ where $\tilde{r} = \nu - \gamma^2 - \rho\sigma\gamma$ depends on both HDD and price volatility.

- If the price is constant ($\mu = \sigma = 0$) then the 'dividend yield' q of (11) is zero and the option price (9) is just the no-dividend Black-Scholes price with 'riskless rate' $\tilde{r}$. Note that q depends only on the *price* parameters. For general μ, σ the discount rate is $r = \tilde{r} + \mu - \sigma^2$. The effect of price vol on option value is explored in section 4 below.

- The pricing formulas do not involve the demand sensitivity α, so it is unnecessary to estimate this parameter. Since $Y_t = \alpha X_t S_t$, adjusting α is equivalent to changing the units of the price process S_t. The pricing formula is invariant under such changes; it only depends on the drift and volatility parameters of S_t.

- The riskless rate of interest does not come into the picture in view of the absence of any trading involving the riskless asset.
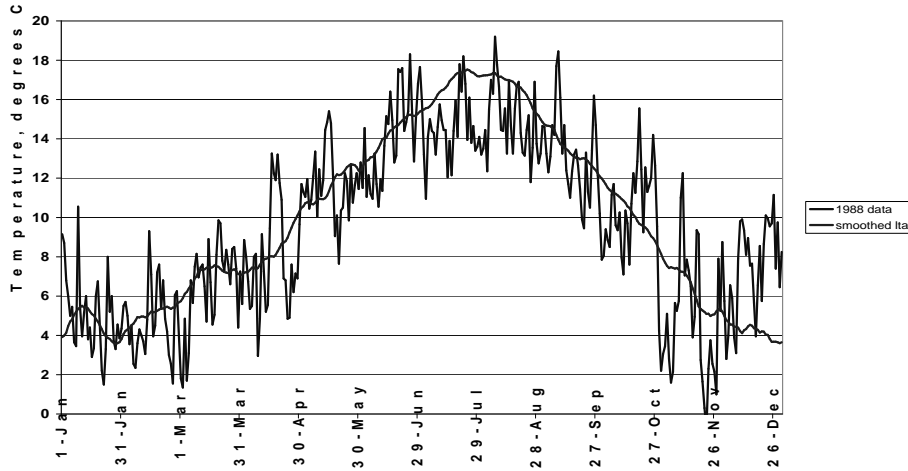
Figure 1: Long-term average temperature and temperatures for 1988

3 HDD Modelling

Weather prediction is a big subject. Nevertheless, some simple things can be said that provide an adequate basis for at least some derivative pricing problems. The objective of this section is to provide just enough evidence to convince the reader that a log-normal model for accumulated HDDs is not at all unreasonable, and to give easily-implemented parameter estimation methods. We do not claim to be providing an exhaustive analysis of the data.

The data set[1] consists of daily temperatures (average of maximum and minimum) at Birmingham, England for the 11 years 1988-1998. We denote this series by $\{T_i, i = 1, \ldots, 4015\}$, while $\{\overline{T}_i, i = 1, \ldots, 4015\}$ denotes the long-term average temperature. For each i, $\overline{T}_i$ is obtained by taking the average of the 11 temperatures on the corresponding date and then smoothing the series by moving-average smoothing. Thus the $\overline{T}_i$ series is periodic. Figure 1 shows the two series for the year 1988.

As many researchers have noted, the deviation $D_i = T_i - \overline{T}_i$ is accurately modelled as a low-order autoregression (AR):

$$D_i = \sum_{k=1}^{n} a_k D_{i-k} + b\epsilon_i, \tag{14}$$

where ϵ_i is a sequence of independent unit-variance gaussian residuals. Here we restrict ourselves to the first-order case $n = 1$. The least-squares estimates $\hat{a}, \hat{b}$ of the parameters a_1, b based on the whole data set are $\hat{a} = 0.70, \hat{b} = 1.99$. These estimates are quite stable when estimated over,

---

[1]kindly provided to Tokyo-Mitsubishi International by John Kings of the School of Geography, Birmingham University
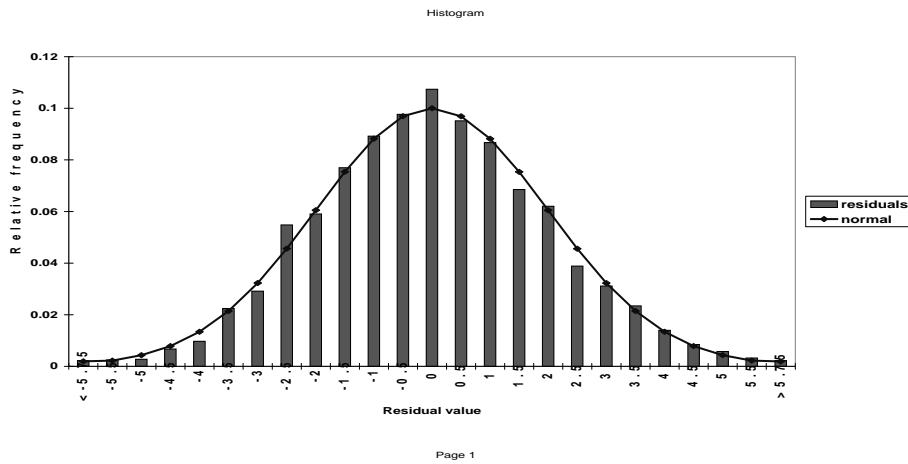
5

Figure 2: Empirical distribution of residuals and best normal fit

say, 3-year windows of data. A more sophisticated analysis would allow for seasonally-dependent variability b, but we have stuck to a constant-parameter model. The residual sequence is then $\hat{\epsilon}_i = (D_i - \hat{a}D_{i-1})/\hat{b}$. The first 10 estimated correlation coefficients – again based on all the data – of the residuals are all in the range ± 0.045, indicating that the residuals are reasonably "white". What is more striking is the residual empirical distribution, shown in figure 2 along with the normal density with the same mean and variance. The fit is astonishingly good. No financial time series behaves like this! We are thus happy to represent the deviation from long-run average temperature as a gaussian first-order AR.
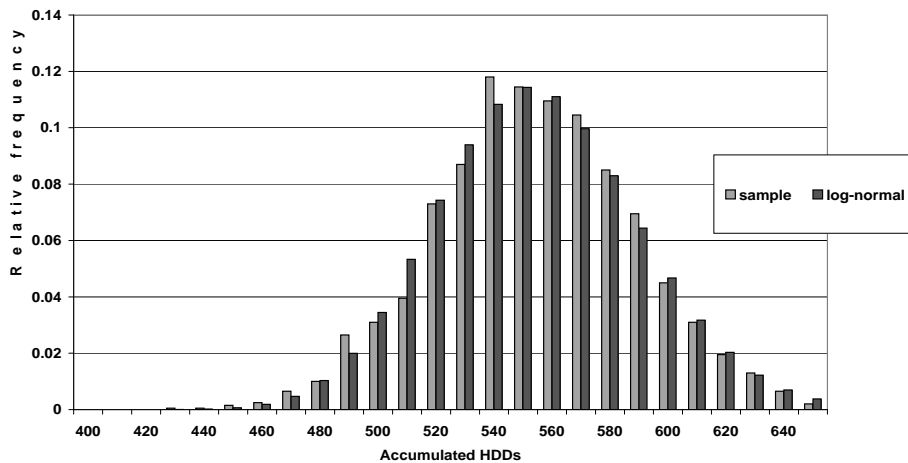

Figure 3: Simulated HDD distribution and best log-normal fit

The AR (14) with $n = 1$ and $|a_1| < 1$ converges to a stationary dis-

6

tribution with mean zero and standard deviation $\Sigma = b/\sqrt{1 - a_1^2}$. The correlation coefficient at lag k is a_1^k. Since $0.7^{15} = 0.0047$ we see that the deviations from long-run average at any time more than two weeks ahead are essentially independent of today's value. Thus if we want to estimate the distribution of accumulated HDDs over a 1-month period starting at any time more than 2 weeks ahead we can simulate D_i from the stationary distribution and take the simulated temperature as $T_i = \overline{T}_i + D_i$. Figure 3 shows the empirical distribution and best log-normal fit for accumulated HDDs over the month of May, using the estimated parameters $\hat{a}, \hat{b}$. The fit is excellent, and similar results are obtained for other months. In fact, this is not surprising: the mean temperature in May is around 11 degrees and with the estimated parameters $\Sigma = 2.77$. Thus the 18 degree barrier is 2.5 standard deviations away from the mean, so that the accumulated HDD is close to being normally distributed. The log-normal distribution with the same mean and variance gives an excellent approximation to standard option values, although of course the tail behaviour is radically different.
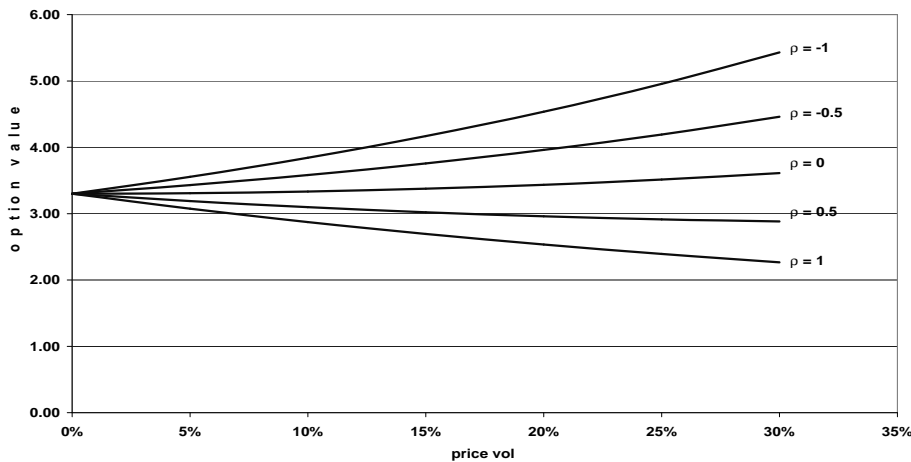
4   Example

Figure 4: Option value as function of price volatility, for different values of correlation

As an example, consider a call option on the accumulated HDDs for May 2001, written on 1 November 2000 with strike $K = 560$. From our simulations, we know that the mean and standard deviation of the accumulated May HDDs are 577 and 35 respectively. (The option is thus 'at the money'.) Referring to the representation (2) we find by calculating the mean and variance that $\gamma = 8.82\%$ and $m(T) = 6.325$. If we take $X_0 = 560$ then this implies $\nu = -0.13\%$. For the price process we take $\mu = 0$, so

there is no drift in the price. However, as can be seen in Figure 7, the value of the option depends significantly on the *price* volatility. Under the measure $\hat{P}$, X_t has drift $r - q = \nu + \gamma^2 - \rho\sigma\gamma$, while the discount factor is $r = (r - q) + \mu - \sigma^2$. If $\rho = 0$ the drift is independent of σ and the option value increases with σ because the discount factor is reduced. For $\rho > 0$ both drift and discount factor are reduced with increasing σ; the net effect is decreasing option value except for very small ρ, as the chart shows. When $\rho < 0$ the effects are in the same direction: less discounting and higher drift lead to increasing value.

5   Concluding remarks

We have offered a simple pricing formula for weather derivatives in terms of their economic value to the purchaser, based on credible models for HDD and taking into account price variability. The mathematics is in the spirit of the Margrabe [7] exchange option formula (see the treatment given by Karatzas in [5], page 24). The results show in particular that pricing by taking the real-world expected value discounted at the riskless rate is incorrect.

   Various foundational issues remain. While taking the HDD X_T at the exercise time T as log-normal seems harmless, modelling the whole *process* $(X_t)_{t \in [0,T]}$ as geometric Brownian motion – while convenient – is far less defensible and does affect the pricing. We can see the problem by noting that moment-matching in equation (2) determines the volatility parameter γ but not the drift ν since various combinations of X_0 and ν give the same value of $m(T)$ in (3). In the example of section 4 we took X_0 to be the forward value, but this is to some extent an arbitrary choice. Taking other values in a reasonable range varies the option price by around $\pm 10\%$. In fact the model (1) is simply an indirect way of specifying the correlation between the final HDD X_T and the price process (S_t). Further work will be directed towards better models that specify this dependence in an unambiguous way.

References

[1] Melanie Cao and Jason Wei, Pricing weather derivatives: an equilibrium approach, working paper, Dept. of Economics, Queen's University, Kingston Ontario, October 1999

[2] M.H.A. Davis, Option pricing in incomplete markets, in *Mathematics of Derivative Securities*, eds. M.A.H. Dempster and S.R. Pliska, CUP 1998

[3] L. Foldes, Valuation and martingale properties of shadow prices: an exposition, J. Economic Dynamics and Control 24 (2000) 1641-1701

[4] H. Geman (ed.) *Insurance and Weather Derivatives: from Exotic Options to Exotic Underlyings*, Risk Books, London 1998

[5] I. Karatzas, *Lectures on the Mathematics of Finance*, American Mathematical Society, Providence, 1997

[6] R.E. Lucas, Asset prices in an exchange economy, Econometrica 46 (1978) 1429-1445

[7] W. Margrabe, The value of an option to exchange one asset for another, J. Finance 33 (1978) 177-186