Pricing the weather

Weather derivatives are an eminently sensible risk management tool, yet the market remains small and illiquid. One reason for this is that traditional options pricing techniques fail to capture the unique characteristics of weather. Here, Melanie Cao and Jason Wei propose a new approach that attempts to rectify this problem

About \$1 trillion of the \$7 trillion US economy is weather-sensitive (Challis, 1999, and Hanley, 1999). For example, weather conditions directly affect agricultural output and seasonal demand in the energy sector, and indirectly affect retail businesses. Although the market for weather derivatives has been growing steadily, the bid/ask spread for a typical contract is still very large and there is not yet a widely accepted pricing method used by weather derivatives participants. In this article, we propose a theoretically intuitive and empirically appealing model to price temperature derivatives.[1]

We focus on heating degree-day (HDD) and cooling degree-day (CDD) products. Our starting point is modelling the statistical behaviour of the daily average temperature, which is very different from that of a security price. For example, temperatures are seasonal and cyclical and can be predicted with reasonable accuracy for the next day or two, and vary within a well-defined range in the long run, none of which is a feature of security prices.

The temperature variable

We first examine 20-year historical daily temperatures for Atlanta, Chicago, Dallas, New York and Philadelphia, covering the period 1979–1998. Other than the usual cyclical/seasonal features mentioned above, the data also reveal less conspicuous characteristics. For example, daily temperatures exhibit strong short-term autocorrelations that are significant for the first three lags. Moreover, the temperature tends to vary more in the winter than in the summer. To facilitate further discussions, let us index the years in the sample period by yr, thus $yr = 1$ for 1979, $yr = 2$ for 1980, ..., $yr = 20$ for 1998. Also, we index January 1 as $t = 1$, January 2 as $t = 2$, and so on for 365 days in a year. Denote $Y_{yr,t}$ as the temperature on date t in year yr. Below, we define for date t, the mean temperature $\bar{Y}_t$ and the standard deviation of temperature ψ_t as[2]:

$$\bar{Y}_t = \frac{1}{20}\sum_{yr=1}^{20} Y_{yr,t} \quad \text{and}$$

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

Figure 1 plots the standard deviation for each day of the year for Atlanta and Chicago. The graph shows an overwhelming seasonal pattern in the temperature variations. This seasonal phenomenon is common for all cities in consideration.

With the above in mind, we can postulate the necessary features that a model for the daily temperature ought to possess. First, at the very least, it should capture the seasonal cyclical patterns; second, the daily variations in temperature must be around some average "normal" temperature, to be elaborated on later; third, since temperatures are forecastable (at least for the short run), the model should allow forecasts to play a key role in projecting temperature paths in the future; fourth, it should incorporate the autoregressive property in temperature changes (ie, a warmer day is most likely to be followed by another warmer day, and vice versa); fifth, in light of figure 1, the extent of variation must be bigger in the winter and smaller in the summer; and sixth, a projected temperature path into the future should never wander outside of the acceptable range of the temperature for each projected point in time (for instance, it is conceivably impossible for a summer day in New York City to see a temperature of –15 degrees Fahrenheit).

Although a diffusion process, eg, a mean-reverting process, is capable of accommodating most of the required features, we decide against it for two reasons. First, a one-factor diffusion process cannot incorporate autocorrelation in the temperature innovations; second, due to the continuous (and typically Markovian) nature of a diffusion process, it is possible to have a temperature path that does not resemble a real one. In light of the above, we resort to an autoregressive, discrete model.

Define $U_{yr,t}$ as the daily temperature whose mean and trend have been removed, ie:

$$U_{yr,t} = Y_{yr,t} - \hat{\hat{Y}}_{yr,t} \quad \forall \quad yr = 1,2,\ldots,20 \text{ and } t = 1,2,\ldots,365 \qquad (1)$$

where $\hat{\hat{Y}}_{yr,t}$ is discussed below. The daily temperature residual is assumed to be described by a k-lag autocorrelation system[3]:

$$\begin{aligned} U_{yr,t} &= \sum_{i=1}^{k} \rho_i U_{yr,t-i} + \sigma_{yr,t} \times \xi_{yr,t} \\ \sigma_{yr,t} &= \sigma - \sigma_1 \left|\sin(\pi t/365 + \phi)\right| \\ \xi_{yr,t} &\sim \text{iid } N(0,1) \\ &\forall \quad yr = 1,2,\ldots,20 \quad \text{and} \quad t = 1,2,\ldots,365 \end{aligned} \qquad (2)$$

where ρ_i ($i = 1, 2, \ldots, k$) is the autocorrelation coefficient for the i[th] lag,

1. Standard deviation of date t's temperature (ψ_t)

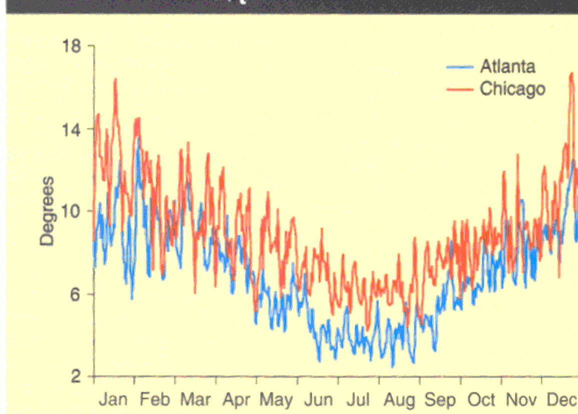

[1] *The literature on weather derivatives pricing is very scanty in both the practitioners world and the academic world. Dischel (1998) used a discussion process to model temperatures. Dischel (1999) and Hunter (1999) briefly examined the so-called "burn rate" method, or the historical simulation method*

[2] *To simplify the analysis, we have deleted the observations for February 29 from the sample. Therefore, each year consists of 365 days and the sample size for 20 years is 7,300*

A. Maximum likelihood estimation results

ρ_1	ρ_2	ρ_3	σ	σ_1	ϕ	Log-l'hood
Atlanta						
0.8833	−0.3035	0.0322	7.5960	5.0912	−0.1881	−20,626
(0.01170)	(0.01520)	(0.01169)	(0.12086)	(0.14603)	(0.01067)	
Chicago						
0.7989	−0.2570	0.0428	7.8289	3.1294	−0.2014	−23,130
(0.01170)	(0.01467)	(0.01170)	(0.13922)	(0.18181)	(0.02316)	
Dallas						
0.8158	−0.2436	0.0201	8.9378	6.3349	−0.1418	−21,381
(0.01170)	(0.01483)	(0.01170)	(0.14060)	(0.16800)	(0.00953)	
New York						
0.7558	−0.263	0.0463	6.5372	2.7035	−0.2432	−21,719
(0.01169)	(0.01433)	(0.01169)	(0.11241)	(0.14520)	(0.02238)	
Philadelphia						
0.7726	−0.2595	0.0473	6.9034	3.1654	−0.2015	−21,792
(0.01169)	(0.01446)	(0.01169)	(0.11957)	(0.15360)	(0.01932)	

Note: Define $U_{yr,t} = Y_{yr,t} - \hat{\hat{Y}}_{yr,t}$, then the estimation system is:

$$U_{yr,t} = \rho_1 U_{yr,t-1} + \rho_2 U_{yr,t-2} + \rho_3 U_{yr,t-3} + \sigma_{yr,t} \times \xi_{yr,t}$$

$$\sigma_{yr,t} = \sigma - \sigma_1 \left|\sin(\pi t/365 + \phi)\right|$$

$$\xi_{yr,t} \sim \text{iid}N(0,1),\ \forall\ yr = 1, 2, \ldots, 20 \text{ and } t = 1, 2, \ldots, 365$$

The numbers in brackets are standard errors

ξ_{yct} is the source of randomness of the daily average temperature, which assumes a standard normal distribution, N(0, 1), and σ_{yct} is the day-specific volatility modelled as a sine wave to reflect the fifth requirement and the feature in figure 1. The parameter ϕ captures the proper starting point of the sine wave. The autocorrelation set-up reflects the fourth requirement. The remaining features are captured by the specification of $\hat{\hat{Y}}_{yct}$, which we will delineate next.

In the context of valuations (which is necessarily forward looking), $\hat{\hat{Y}}_{yct}$ is simply interpreted and treated as the daily temperature forecast, and will be plugged into the system as an input. The model in (1) and (2) will then generate possible realisations of temperatures around the forecasts by incorporating the necessary features postulated above. However, to generate possible temperature realisations from (1) and (2), we must first estimate the model parameters ρ_i, σ, σ_1 and ϕ. In estimations, $\hat{\hat{Y}}_{yct}$ is also an input and must be properly chosen to reflect the required features discussed before. The first and the second requirements would imply that the previously defined average daily temperature, $\bar{Y}_t$, might serve the purpose. However, a moment of reflection would rule it out. Given that we could have an entire season below or above the average seasonable temperature (for instance, the temperature of almost the entire winter of 1980 in New York city was lower than average), using $\bar{Y}_t$ to estimate the system would, for one thing, severely distort the meaning of the standard deviation. We must choose $\hat{\hat{Y}}_{yct}$ such that it is roughly the middle point of variation for any period in question. To this end, we modify the simple average daily temperature, $\bar{Y}_t$ in the following way: 1) for each month of the year, we calculate the average of the daily averages $\bar{Y}_t$, and there will be 12 such monthly averages; 2) for each particular year, we calculate the realised, average temperature of each month; 3) for each month, we find the difference between the actual monthly average from step two and the average from step one; and 4) we adjust the simple mean for each day of the month by the quantity calculated in step three, and this adjusted mean is $\hat{\hat{Y}}_{yct}$ and will be referred to as the adjusted mean temperature.[4] The overall idea is to shift the simple daily means $\bar{Y}_t$ to reflect the range of realised temperatures, which can be quite different from historical averages. To illustrate the adjustment mechanism, suppose our concern is April 1985, and suppose the mean of the average daily temperature for the month of April, calculated from step one, happens to be 60 degrees Fahrenheit. We now calculate the average of the 30 realised daily temperatures for April 1985 and suppose it is 55 degrees Fahrenheit, which indicates a colder than normal April. This is the average from step two. We then follow step three to find the difference between the two averages: 55 − 60 = −5. Finally, following step four, we adjust each of the historical average temperature $\bar{Y}_t$ for April by −5 degrees Fahrenheit. Suppose the historical daily average temperatures for April 1, 2, 3, 4, ... , 30 are 58, 63, 60, 65, ... ,70 degrees Fahrenheit, then the adjusted mean averages for April 1, 2, 3, 4, ... , 30 of 1985 will be 53, 58, 55, 60, ... , 65 degrees Fahrenheit. The $\hat{\hat{Y}}_{yct}$ will assume these values for April 1985 in actual estimations.

Estimation of the temperature system

For a given number of lags, k, the system proposed in (1) and (2) involves the following parameters, $\rho_1, \rho_2, \rho_3, \ldots, \rho_k, \sigma, \sigma_1$ and ϕ. The residual specification allows us to perform a maximum likelihood estimation of the parameters. To determine the cut-off number of lags, we carry out sequential maximum likelihood estimations in the following fashion. We first assume $k = 1$ and proceed to estimate ρ_1, σ, σ_1 and ϕ, and record the resulting maximum likelihood value. We then assume $k = 2$, estimate ρ_1, ρ_2, σ, σ_1 and ϕ, and record the maximum likelihood value, and so on. We halt the estimation when the maximum likelihood value ceases to improve. The optimal number of lags turns out to be three. The estimation results are presented in table A. It can be seen that almost all parameters are estimated with very low standard errors, implying the proper specification of the estimation system. This is by no means a fluke since we have extensively explored and eliminated many other systems. In addition, standard errors of the parameter σ_1 are very small compared with the estimated value of σ_1, implying that the sine wave fitting to the overall volatility structure is indeed appropriate and useful. Finally, the first-order autoregressive behaviour tends to be stronger with southern cities, and ρ_1 has the highest value for Atlanta. Roughly, a stronger autocorrelation means less dramatic changes in temperature, and vice versa. Atlanta indeed has the lowest overall standard deviation in the sample period.

Valuation of HDD/CDD derivatives

Daily HDD and CDD are formally defined as:

$$\text{Daily HDD} = \max\left(65°\text{F} - \text{daily average temperature}, 0\right) \quad \text{and}$$

$$\text{Daily CDD} = \max\left(\text{daily average temperature} - 65°\text{F}, 0\right)$$

Derivatives contracts are typically written on cumulative HDDs or CDDs over a period of, say, one month or a whole season. For a typical northern or midwestern US city (eg, New York and Chicago), an HDD season goes from November to March inclusive, and a CDD season goes from May to September inclusive. April and October are called shoulder months and are excluded in HDD/CDD calculations. Of course, the definition of a HDD or CDD season can vary across cities.

We have shown elsewhere (Cao & Wei, 1999) that, under certain conditions, the market price of temperature risk is insignificant, and risk-neutral valuation can be applied.[5] Consider an HDD forward contract with an accumulation period starting at T_1 and ending at maturity $T_2 > T_1$. Denote $HDD(T_1, T_2) = \Sigma_{\tau=T_1}^{T_2} \max(65 - Y_\tau, 0)$. A CDD contract can be defined similarly. The HDD/CDD forward price will simply be the expected value of the cumulative daily HDD or CDD over the contract period. With the autocorrelation terms, it is difficult to obtain closed-form formulas. However, if we set the autocorrelations to zero, the forward prices can be expressed as:

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

[4] *The period of one month is chosen as a trade-off. Too long a period will not solve the non-centring problem and too short a period will unnecessarily exaggerate the short-term fluctuations and diminish the meaning of "average" or "mean". Needless to say, one could be more sophisticated in making the adjustments. For instance, rather than following the calendar months, one could always centre the day in question in a, say, 30-day period and make the above adjustments on a rolling over basis*

B. Option prices for a CDD season

		Historical simulation		Regular simulation			
				Forecast A		Forecast B	
		Strike price	Option value	Strike price	Option value	Strike price	Option value
Atlanta	Call	1,812.00	87.62	1,777.95	33.16	1,893.41	33.95
	Put		87.62		33.44		34.24
Chicago	Call	823.60	84.88	674.80	39.25	858.24	43.61
	Put		84.88		39.37		43.76
Dallas	Call	2,424.55	86.02	2,405.65	34.35	3,153.45	35.01
	Put		86.02		34.72		35.50
New York	Call	1,181.80	57.72	1,101.80	34.31	1,226.85	35.23
	Put		57.72		34.49		35.43
Philadelphia	Call	1,239.75	80.29	1,149.35	35.70	1,286.41	37.17
	Put		80.29		35.89		37.38

Note: under "historical simulation", we assume that the future will mimic history exactly according to the sample data. The strike prices are the historical average CDDs. Under "simulation forecast A", the simulation is based on (1), (2) and table A, and the forecast is the historical average temperature. The strike price is the seasonal CDD of the averages. Under "simulation forecast B", the simulation is based on (1), (2) and table A, and the forecast is the adjusted average temperature. The strike price is the seasonal CDD of the average

$$F_{HDD}(t, T_1, T_2) = \sum_{\tau=T_1}^{T_2}\left(\left[65-\hat{\hat{Y}}_{yr,\tau}\right]\times N\left(\frac{65-\hat{\hat{Y}}_{yr,\tau}}{\sigma_{yr,\tau}}\right)+\frac{\sigma_{yr,\tau}}{\sqrt{2\pi}}\exp\left[-\frac{\left(65-\hat{\hat{Y}}_{yr,\tau}\right)^2}{2\sigma_{yr,\tau}^2}\right]\right) \quad (3)$$

$$F_{CDD}(t, T_1, T_2) = \sum_{\tau=T_1}^{T_2}\left(\left[\hat{\hat{Y}}_{yr,\tau}-65\right]\times N\left(\frac{\hat{\hat{Y}}_{yr,\tau}-65}{\sigma_{yr,\tau}}\right)+\frac{\sigma_{yr,\tau}}{\sqrt{2\pi}}\exp\left[-\frac{\left(65-\hat{\hat{Y}}_{yr,\tau}\right)^2}{2\sigma_{yr,\tau}^2}\right]\right) \quad (4)$$

where $N(\cdot)$ is the cumulative density function of a standard normal variable.

HDD and CDD options can be handled in a similar way. Consider a European-style option written on $HDD(T_1, T_2)$ with maturity T_2 and a strike price X. Denote the call and put prices at time t as $C_{HDD}(t, T_1, T_2, X)$ and $P_{HDD}(t, T_1, T_2, X)$, respectively. With a risk-free rate r, the call and put values can be expressed as:

$$C_{HDD}(t, T_1, T_2, X) = e^{-r(T_2-t)}E_t\left(\max\left(HDD(T_1, T_2) - X, 0\right)\right) \quad (5)$$

$$P_{HDD}(t, T_1, T_2, X) = e^{-r(T_2-t)}E_t\left(\max\left(X - HDD(T_1, T_2), 0\right)\right) \quad (6)$$

The values of call and put options written on $CDD(T_1, T_2)$ can be expressed in a similar fashion. Closed-form solutions to the above expressions are extremely difficult to obtain due to the double application of the maximum operator. We must resort to simulations.[6]

Broadly speaking, the simulation procedure consists of: 1) generating a path for the daily temperature process in (1) and (2) using the parameters in table A and a set of daily forecasts for the simulation period; 2) tracking realised HDD/CDD values of each path; 3) calculating the payout of the derivative in question; and 4) repeating steps one through three many times and averaging the discounted payouts to obtain the desired derivative value.

A key issue is the choice of inputs for the adjusted mean temperature, $\hat{\hat{Y}}_{yt}$. As mentioned earlier, in the estimation context, $\hat{\hat{Y}}_{yt}$ serves as the "anchoring point". In the valuation context, which, necessarily, is forward-looking, $\hat{\hat{Y}}_{yt}$ can naturally be considered as daily temperature forecasts. The random term in the temperature dynamic will capture the uncertainty in the forecasts. Indeed, this is one of the key advantages of our model, since it allows forecasts as inputs and is capable of accommodating deviations from forecasts commensurate with history, ie, data.

The simulation errors can be reduced in many ways. For example, the antithetic variable technique is readily applicable. In our framework, we propose an additional error reduction measure, which is similar to the control variate technique. Notice that the fundamental variable in our framework is the daily temperature, and the underlying variable for most weather derivatives is HDD/CDD, which are essentially non-linear functions of daily temperatures. While our model will produce almost "unbiased" temperature forecasts in that the average temperature for a future point will be almost equal to the input forecast, it cannot guarantee an unbiased forecast for the HDD/CDD, or forward prices. To ensure correct pricing, a two-stage simulation can be performed. In the first stage, we simulate forward prices and record the difference between the simulated price and the implied forward price from the forecasts. Then, in the second stage, we simulate the derivative's prices whereby for each particular path, we adjust the realised CDD or HDD by the difference obtained in the first stage, and use the adjusted CDD or HDD to calculate the payout. In a nutshell, the above procedure amounts to ensuring unbiased paths of the CDD and HDD, which are underlying variables for weather derivatives. It should be pointed out that when the temperature dynamic, especially the volatility structure, is estimated perfectly, this procedure will not be necessary.

We implement the above procedures for HDD/CDD options by taking January 1, 1999 as the valuation date and by assuming an annual risk-free interest rate of 6%. We use two sets of forecasts, one being the historical daily average temperature and the other being the adjusted mean temperature of 1998. The results are in tables B and C, which also report the so-called historical simulation prices. The historical simulation here uses only the past realised seasonal HDD/CDD, in a way similar to historical simulation in VAR calculations. Essentially, it estimates the average value of the contract if it were written every year in the past years. In the insurance industry, this method is referred to as "burn rate" method.

We can see that the historical simulation approach generates very high option values, compared with regular simulations. For most cities, the CDD option values based on historical simulations are more than twice the regular simulation values, prompting us to argue that the historical simulation,

[5] The key factor that determines the significance of the market price of risk is the correlation between the aggregate output of the economy and the temperature variable. When this correlation is not very high (say, lower than 0.2), then the market price of risk is generally insignificant. Another factor is investors' risk aversion. The significance of the market price of risk tends to increase as investors become more risk-averse

[6] For most OTC HDD or CDD contracts, the payout of, say, a call option, is capped at a certain level. This in no way poses extra difficulty when simulations are used. For simplicity, we ignore payout caps in the numerical illustrations

C. Option prices for an HDD season

		Historical simulation		Regular simulation: Forecast A		Regular simulation: Forecast B	
		Strike price	Option value	Strike price	Option value	Strike price	Option value
Atlanta	Call	2,419.47	108.78	2,396.95	68.44	2,715.65	69.32
	Put		108.78		68.74		69.65
Chicago	Call	5,114.37	144.22	5,126.15	74.96	4,506.00	74.88
	Put		144.23		75.59		75.43
Dallas	Call	2,179.21	60.40	2,141.05	73.17	2,192.98	73.29
	Put		60.40		73.44		73.57
New York	Call	3,859.63	105.53	3,862.35	55.97	3,417.25	55.93
	Put		105.53		56.44		56.35
Philadelphia	Call	3,901.00	109.67	3,899.75	61.14	3,404.08	61.03
	Put		109.67		61.62		61.45

Note: under "historical simulation", we assume that the future will mimic history exactly according to the sample data. The strike prices are the historical average HDDs. Under "simulation forecast A", the simulation is based on (1), (2) and table A, and the forecast is the historical average temperature. The strike price is the seasonal HDD of the averages. Under "simulation forecast B", the simulation is based on (1), (2) and table A, and the forecast is the adjusted average temperature. The strike price is the seasonal HDD of the average. When forecasts are adjusted average temperature, we use November and December 1997 and the first three months of 1998

or burn rate method, should be avoided for weather derivatives valuations. The main reason that this method is likely to fail is its implicit assumption that over the contract period (say, next HDD season), the entire season could be as cold as the coldest season or as warm as the warmest season in the past. But, in reality, in a forecasting context, we would not attach equal probabilities to the two extremes. After all, a typical long run forecast would indicate whether the next season is above, below or around normal. A forecast that attaches equal probabilities to all likelihoods is hardly a forecast!

Discussions and conclusions

We propose and implement a theoretically intuitive and empirically appealing framework for weather derivatives valuations. The proposed temperature system not only allows easy estimation, but also incorporates key features of the daily temperature behaviour such as seasonal cycles and uneven variations throughout the year.

Our valuation framework has many advantages. It allows the use of weather forecasts in modelling the future temperature behaviour. In addition, since our starting point is the daily temperature, the framework is capable of handling temperature contracts of any maturity, for any season, and it requires only a one-time estimation. In contrast, if one starts by modelling CDD or HDD directly, then by nature of the temperature behaviour, the CDD or HDD will necessarily be season and maturity specific, which implies that each contract will require a separate estimation procedure. This will not only create potential inconsistency in pricing, but also render the whole idea impractical if many different contracts are dealt with or if the valuation is to be ongoing.

We show that the historical simulation approach or burn rate method is largely invalid in estimating weather derivative values. Weather contracts typically cover a period due soon and do not extend very far into the future. However, historical simulations implicitly assume that the next season's temperature can resemble any of the past seasons in the sample, including extreme seasons (eg, very cold or very warm). As a result, in most cases, the historical simulation method tends to overestimate option prices.

We would like to issue two caveats before closing. First, our model takes as inputs the daily temperature forecasts. Irrespective of the massive computing power endowed by today's technology, even the most sophisticated forecasting systems cannot produce long-run forecasts to daily precisions. For example, as pointed out by a referee, the US National Weather Services, which sets the standards for long-range weather forecasting, only provides, among other forecasts, three-month forecasts for periods up to one year. Those forecasts only address the seasonal level as opposed to daily level. In the absence of authoritative long-run forecasts, the choices of forecast inputs becomes an art rather than science. But this is precisely what drives the market. When perfect forecasts are available, there will be no weather derivatives markets! A corollary of the above observation is that, unlike with financial derivatives pricing, even if all participants agree on the model structure, large bid-ask spreads for weather derivatives can still exist simply due to different assessments of the unknown weather.

Second, we do not claim perfection and universality of our model. We consider our efforts as the first step in structuring ideas into quantitative methods. By almost any measurement, this field is still in its infancy. We hope that our initial efforts will stimulate more research into the topic, so that a near-perfect model will eventually emerge. ■

Melanie Cao is senior director in the research department at the **Chicago Mercantile Exchange and assistant professor** in the department of economics at Queen's University, Ontario. Jason Wei is associate professor in the division of management at the University of Toronto at Scarborough. We are grateful to the Social Sciences and Humanities Research Council of Canada for financial support. We respectively also acknowledge Queen's Edith Whyte research grant and the University of Toronto Connaught Fund. The analysis and conclusions of this paper are those of the authors and do not indicate concurrence by the Chicago Mercantile Exchange
Comments on this article can be posted on the technical discussion forum on the Risk Web site at http://www.riskpublications.com/risk

REFERENCES

Cao M and J Wei, 1999
Pricing weather derivatives: an equilibrium approach
Working paper, Queen's University and University of Toronto

Challis S, 1999
Bright forecast for profits
Reactions, June

Dischel B, 1998
Black-Scholes won't do
Risk Weather Risk Special Report, pages 8–9

Dischel B, 1999
Shaping history
Risk Weather Risk Special Report, pages 13–15

Hanley M, 1999
Hedging the force of nature
Risk Professional 1, July–August, pages 21–25

Hunter R, 1999
Managing mother nature
Derivatives Strategy, February