What price the weather?

Outlook fine for temperature derivatives

Weather derivatives are becoming an increasingly important risk management tool for a number of industries. HARVEY STERN presents an empirical approach to the pricing of weather derivatives and concludes that it is appropriate to price weather derivatives by using a combination of data.

DR HARVEY STERN ASIA is a meteorologist with the Bureau of Meteorology. He holds a PhD from the University of Melbourne, for a dissertation entitled "Statistically based weather forecast guidance". The author thanks his colleagues at the Bureau of Meteorology for their suggestions during the preparation of this paper, in particular Mark Williams, Glen Prank and Patrizia Pace.

The earliest published work on weather derivatives is that by the current author (Stern 1992), who employed option pricing theory to establish a measure of the economic consequences of changes in the global mean temperature. These consequences were examined across scales from the macro- to the micro-economy, being replicated by a combination of Global Mean Temperature futures contracts and an associated set of option contracts.

The energy and power industry has so far taken best advantage of the opportunities presented by weather derivatives. Clemmons *et al* (1999) report the first weather derivative contract as a "temperature-related power swap . . . transacted in August 1996".

By the end of 1997, some 150 weather derivative deals had been completed; by the end of 1998, this figure had increased to 500. An Internet weather derivatives exchange was launched in January 2000. The first Australian deal was transacted in April 2000. Of the 3,500 transactions completed by the middle of the year 2000, 98% were based on temperature, of which most were *degree day* contracts constructed for the energy and power industry.

WEATHER DERIVATIVES DEFINED

Jain and Foster (2000) observe that the world is heavily affected by the weather. "But while there is little anyone can do to control the climate, businesses can now mitigate the exposure they face from adverse weather conditions by using weather derivatives."

These businesses include energy and power, agriculture and agrochemicals, viticulture, brewing, clothing, construction, theme parks, retail food and drink, tourism, sporting, outdoor entertainment, and water authorities and irrigation.

Clewlow *et al (*2000) describe weather derivatives as being similar "to conventional financial derivatives, the basic difference coming from the underlying variables that determine the payoffs", such as temperature, precipitation, wind, "heating degree days" and "cooling degree days".

Weather derivative contracts are typically defined by:

- Location;
- Type of asset (eg, heating degree days);
- Strike (value of underlying asset at which one party is obliged to compensate the other);
- Expiry (the time when one party is obliged to compensate the other); and
- Notional ($ per unit of underlying asset).

Dischell (1998) notes that "traditional weather insurance . . . requires a demonstration of loss [whereas] weather derivatives . . . require no demonstration of loss and provide protection from the uncertainty in normal weather".

This protection is achieved because a weather derivative contract, when applied as a hedge, sets limits on how far revenues can fall and expenses can increase. On the other side of such a contract there may be:

- a speculator, to whom the risk has been transferred in return for a reward; or,
- another hedger, who wishes to protect against loss associated with the opposite scenario (eg, high temperatures) to the scenario that the counterparty is concerned about (eg, absence of high temperatures).

So, like all derivatives, weather derivatives may be used to transfer risk from those who are involuntarily exposed to unwanted risk to those who have a traditional familiarity with risk.

A COOLING DEGREE DAY CALL OPTION

The number of "cooling degree days" during a season (which might be regarded as a measure of the requirement for cooling) is the accumulated number of degrees that the daily mean temperature is above a particular base figure, usually 18°C. If the average temperature on a particular day is below 18°C, there is no contribution for that day. "Heating degree days" (which might be regarded as a measure of the requirement for heating) are defined in the reverse way.

Suppose that an electricity retailer holds a "cooling degree day call option". That option may be viewed as placing a maximum limit on the number of degree days during a season (the "strike") before the electricity company is entitled to purchase a "parcel" of cooling degree days at a pre-determined price, regardless of the price of the electricity. Alternatively, if, at the expiration of the contract, the actual number of cooling degree days is greater than the strike, the seller of the option pays the buyer a certain amount.

A typical weather derivative contract may be defined thus:
Location: Not specified in this case
Type of asset: cooling degree days
Strike: 600 cooling degree days
Expiry: Not specified in this case
Notional: $100 per cooling degree day above 600

If, at the expiration of this contract, the accumulated number of cooling degree days is greater than the strike (600), then the seller of the option pays the buyer the notional ($100) for each cooling degree day above the strike. This is illustrated in the pay-off diagram in Figure 1.

For example, if the accumulated cooling degree days at expiry is 1,400, the pay-off is $(1,400-600) x 100 = $80,000.

PRICING APPROACHES

There are three approaches to the pricing of weather derivatives. These are:

- historical simulation, which involves computing the historical pay-off of a derivative;
- direct modelling of the underlying variable's distribution – this involves modelling the underlying as a normally or log-normally distributed variable;
- indirect Monte Carlo modelling of the underlying variable's distribution – this involves simulating a sequence of data, allowing for the incorporation of accurate seasonal patterns, mean reversion, jumps and changing volatility.

Historical simulation

Historical simulation originated in the insurance industry. The approach asks: "What would be the pay-out, on average, had the company sold the option every year for the last *n* years?" For example, given the past 30 years of temperature data, we can calculate 30 samples of the pay-off for a cooling degree day option for the month of January in Melbourne. The approach is useful because it allows development of an indicative pricing methodology.

Direct modelling of the underlying variable's distribution

To establish the appropriate distribution, we look at the underlying variable's mean and standard deviation over the relevant period. In some cases the distribution may be approximately log-normal; in other cases it appears more normal. A model frequently applied to financial market variables is one developed by Black and Scholes (1973).

The key assumption of the Black and Scholes model is that the variable underlying the option is log-normally distributed. The model therefore suggests that the underlying variable can increase without limit. Weather variables, such as temperature, tend to remain within relatively narrow bands – a "mean-reverting" type of behaviour. As a result, it is considered that the Black and

FIGURE 1 **Pay-off chart for a degree day call**

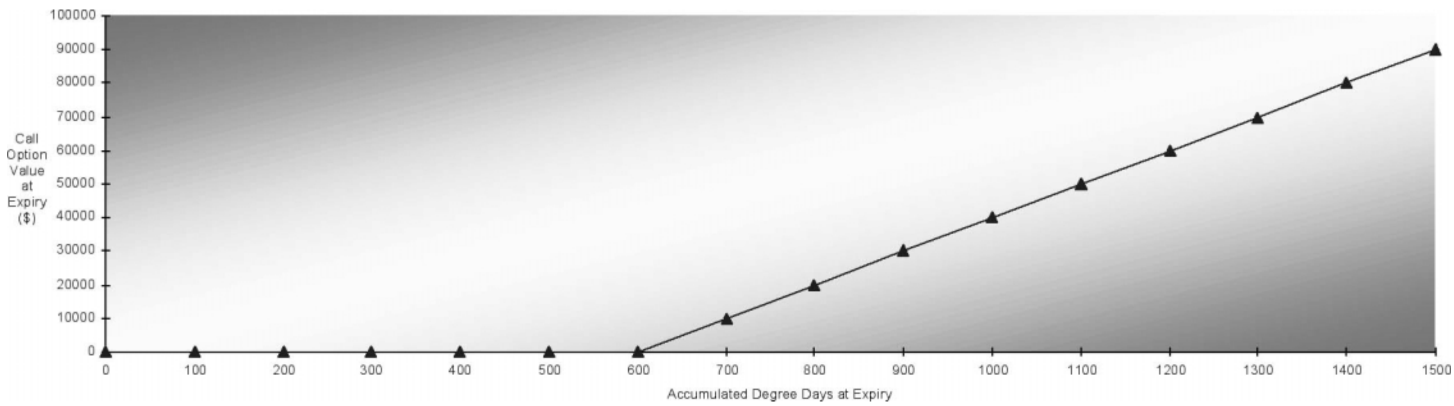

FIGURE 2 **Pay-off chart for a 38° call**

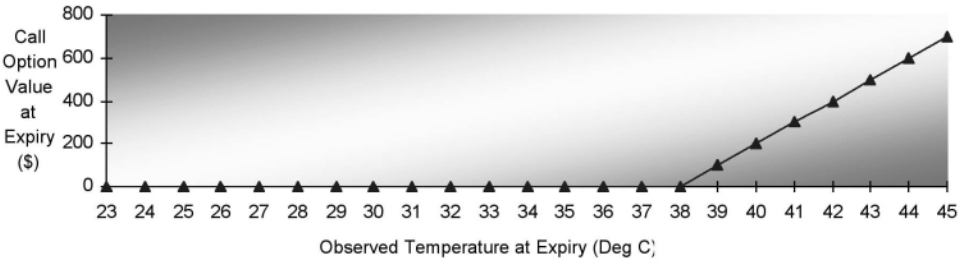

FIGURE 3 **Historical outcomes for a 38°C call**

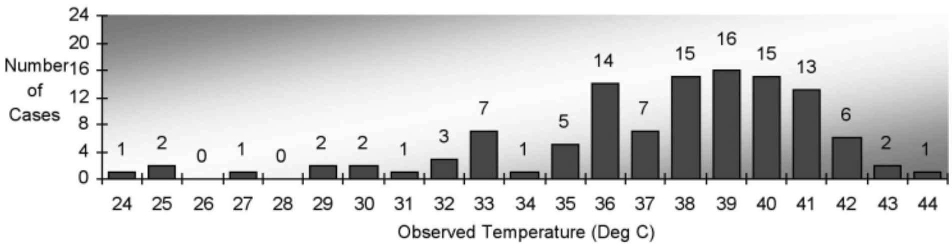

Scholes model has deficiencies when applied to weather variables.

Indirect Monte Carlo modelling of the underlying variable's distribution

Monte Carlo simulation involves simulating a sequence of data. It provides a general and flexible way to price many different weather derivative structures, allowing the use of models that incorporate:

- seasonal patterns;
- forecasts;
- mean-reversion behaviour;
- extreme events;
- jumps; and
- changing volatility.

In a simple example of Monte Carlo modelling (after Dischell 1999), where:

a, *b* and *c* are constants

T(n+1) = projected temperature for day n+1

T(n) = previously projected temperature for day n

M(n+1) = mean temperature for projected day n+1

Ch(n,n+1) = random change from day n to day n+1,

the following equation is employed to described the evolution of a temperature sequence:

$T(n+1) = aT(n) + bM(n+1) + cCh(n,n+1)$

A 38°C CALL

A second illustrative example is that of a *38°C call option*. This example is applied to cases where a temperature of at least 38°C has been forecast. The weather derivative contract is defined thus:

- Location: Melbourne
- Type of asset: Temperature (°C)
- Strike: 38°C
- Expiry: Tomorrow
- Notional: $100 per degree above 38°C

If, at the expiration of a call option contract (that is, tomorrow), the actual maximum temperature is greater than the strike (that is, 38°C), the seller of the option pays the buyer $100 for each 1°C it is above 38°C. This is illustrated in the pay-off diagram in Figure 2.

We now determine the price of our call option contract by employing historical simulation of the outcomes. Between 1960 and 2000, there were 114 forecasts of at least 38°C. The distribution of historical outcomes is shown in Figure 3. The contribution from the historical outcomes to the price of the 38°C call option contract are:

1 x 44°C yields $(44-38) x 1 x 100 =	$600
2 x 43°C yields $(43-38) x 2 x 100 =	$1,000
6 x 42°C yields $(42-38) x 4 x 100 =	$1,600
13 x 41°C yields $(41-38) x 13 x 100 =	$3,900
15 x 40°C yields $(40-38) x 15 x 100 =	$3,000
16 x 39°C yields $(39-38) x 16 x 100 =	$1,600

The other 61 cases (38°C or below) yield nothing. The result is a total contribution of $11,700, and an average contribution over the 114 cases of $103. So, $103 is the price of our call option.

A COOLING DEGREE DAY PUT OPTION

Suppose an electricity retailer holds a cooling degree day put option. That option places a minimum limit on the number of degree days during a season (the "strike") before the electricity company is entitled to sell a "parcel" of cooling degree days at a predetermined price, regardless of the price of the electricity. Alternatively, if, at the expiration of the contract, the actual number of cooling degree days is less than the strike price, the seller of the option pays the buyer a certain amount. The features of the weather derivative contract are:

Location: Not specified in this case
Type of asset: Cooling degree days
Strike: 600 cooling degree days
Expiry: Not specified in this case
Notional: $100 per cooling degree day below 600

This is illustrated in the pay-off diagram in Figure 4. In the case illustrated, suppose the accumulated cooling degree days at the expiry of the option total 300. The pay-off is $(600-300) x 100 = $30,000.

EVALUATING AN ECHUCA MONTHLY RAINFALL DECILE 4 PUT

The fourth illustrative example is that of a *Monthly Rainfall Decile 4 Put Option*. This example applies to cases where the preceding month's Southern Oscillation Index (SOI)[1] is decile 1, 2 or 3. A rainfall total is defined as a decile 1 rainfall total if it ranks among the lowest 10% of all rainfall totals. Similarly, a decile 2 rainfall total ranks among the second lowest 10% of all rainfall totals, and so on, to a decile 10 rainfall total, which is among the highest 10% of all rainfall totals.

The features of the example weather derivative contract are:

Location: Echuca (some 200 km north of Melbourne)

Type of asset: October rainfall (decile)
Strike: Decile 4
Expiry: October
Notional: $100 per decile below decile 4

If, at the expiration of a put option contract (that is, October), the actual rainfall is less than the strike (that is, decile 4), the seller of the option pays the buyer $100 for each decile that it is below decile 4. The pay-off is illustrated in Figure 5.

The price of the put option contract is determined by employing historical simulation of the outcomes. Between 1876 and 1999 there were 119 Octobers with rainfall records, of which 44 were preceded by months with an SOI of decile 1, 2 or 3. The distribution of historical outcomes is shown in Figure 6.

The figure shows that the contributions from the historical outcomes to the price of the decile 4 put option contract are:

9 x decile 1 yields $(4-1) x 9 x 100 =	$2,700
6 x decile 2 yields $(4-2) x 6 x 100 =	$1,200
4 x decile 3 yields $(4-3) x 4 x 100 =	$400

The other 25 cases (decile 4 or above) contribute nothing. The total contribution is $4,300, and an average contribution over the 44 cases is $98. So $98 is the price of the put option.

FIGURE 4 Pay-off chart for a degree day put

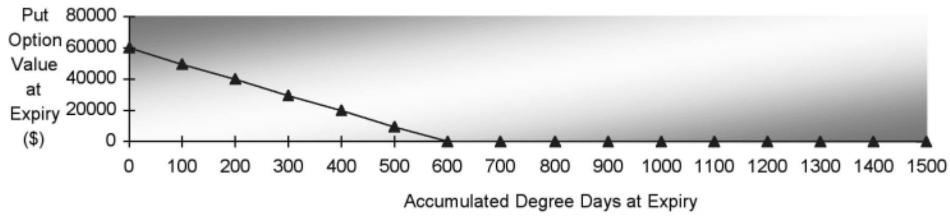

FIGURE 5 Pay-off chart for a decile 4 put

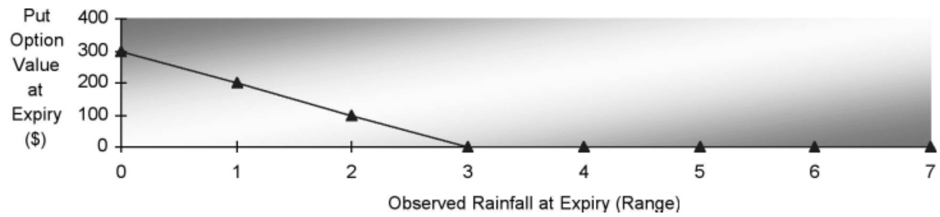

FIGURE 6 Historical outcomes for a decile 4 put

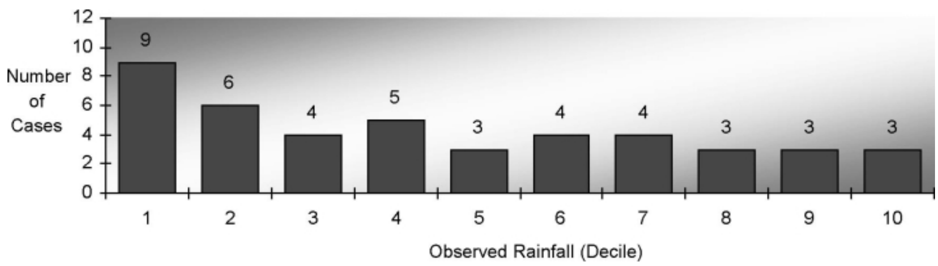

PRICING OTHER DERIVATIVES

One should therefore not be surprised if the modelling and pricing of weather derivatives eventually leads to improved techniques for the pricing of other derivatives. Indeed, the following analysis reports on progress towards that end.

Revisiting Black and Scholes

As already noted, the Black and Scholes 1973 model has deficiencies when applied to weather variables. For this reason, it has been preferred to employ historical simulation and/or Monte Carlo techniques when modelling weather variables (and this paper's focus has been largely on historical simulation).

Notwithstanding the general use of Black and Scholes in the modelling of financial market derivatives, some aspects of the Monte Carlo approach to the modelling of weather derivatives (eg, mean reversion and jumps) may be appropriate to financial derivatives. To test this proposition, the following experiment was designed:

- Australian stock-price data (for the 25 leading stocks) were extracted for an 18-month period (March 1999 to August 2000), noting sequences of at least five consecutive falls or rises (there were 249 such sequences during that period).
- Dates when the sequences reversed were also noted, and it was assumed that on the reversal date an investor bought in the wake of a fall sequence, and short-sold in the wake of a rise sequence.
- It was assumed that the investor closed positions each time an opposite sequence was reversed.

The outcomes of the experiment support the proposition that mean reversion and jumps need to be incorporated into the modelling of stock price derivatives.

Mean reversion

To illustrate the operation of mean reversion, note that:

- the average (mean) return on all 249 sequences was +4.51% with a standard deviation of 12.15%;
- this mean is different from zero at the 0.1% level of significance, the above-zero return reflecting the operation of mean reversion (if it was not operating, the average return would be zero).

These outcomes would be particularly important in the pricing of *American style options*. In the pricing of American style options, because of the flexibility about when positions may be opened and closed, the operation of mean reversion could dramatically reduce (or increase) the value of the option over that which would be suggested by Black and Scholes. This would be also important in the pricing of *European style options*, although to a lesser degree, because of the lack of flexibility about when positions may be closed.

FIGURE 7 **The ratio: (frequency of returns from experiment)/ (frequency of returns if distribution was normal)**

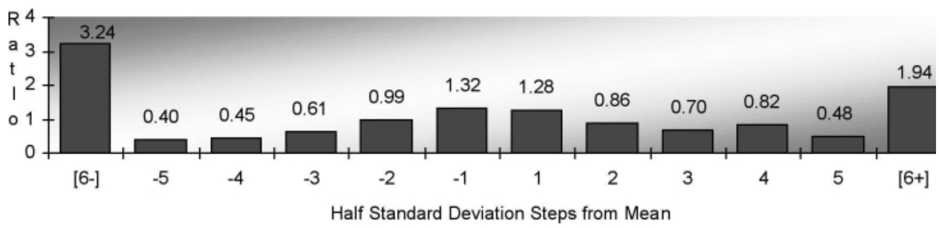

FIGURE 8 **The ratio (without extreme cases)**

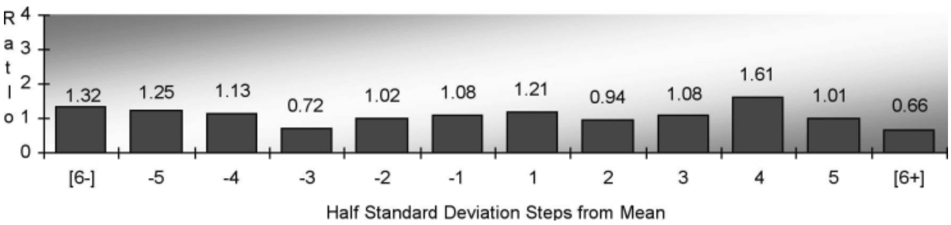

Jumps

To illustrate the operation of jumps, see Figure 7, which presents the ratio

$$\frac{\textit{(frequency of returns from experiment)}}{\textit{(frequency of returns if distribution normal)}}$$

that ratio being presented in half standard deviation steps from the mean. Observe that:

- there was a much higher frequency of "extreme" returns than would be expected had the returns been normally distributed;
- specifically, the high frequency of strongly negative returns (3.24 times what would have been expected) reflects cases where "news" has initiated a sudden re-rating of a stock, so that a trend continues well beyond those trends associated with the typical short-term variations in stock price;
- conversely, the high frequency of strongly positive returns (1.94 times what would be expected) reflects cases where stock prices have "over-shot" on "news", requiring a corrective trend that continues well beyond those trends associated with the typical short-term variations in stock price.

The much higher frequency of "extreme" returns reflects the operation of jumps – if jumps were not present, and the distribution was normal, each of the ratios would have been 1.0. These outcomes would be important in the pricing of options, particularly American style options, where the operation of jumps could dramatically increase the value of the option over that suggested by Black and Scholes. The frequency of returns is unexpectedly high for cases close to the average return, and unexpectedly low for cases moderately distant from the average return. This occurs because the standard deviation has been inflated by the extreme cases. Without the extreme cases, most of the ratios should be closer to 1.0. Figure 8 shows that indeed they would be.

> ***There was a much higher frequency of "extreme" returns than would be expected had the returns been normally distributed.***

A new approach to the pricing of options

To effectively take into account the processes of mean reversion and jumps when pricing options, it is proposed that:

- the frequency distribution of a range of stock-price evolutions be developed using historical analogues to the recent stock-price sequence;
- the price of American style options be determined on the basis of the "best" of a range of closing strategies – current pricing practice results in only two closing strategies being considered, closing at option expiry or closing at dividend payment time.

NOTE

1 The SOI is a representation of the atmospheric pressure difference across the Pacific Ocean and is regarded as an indicator of future rainfall over eastern Australia.

REFERENCES

Black, F., and M. Scholes, 1973, "The pricing of options and corporate liabilities", *Journal of Political Economy*, pp. 637-54;

Clemmons, L., V. Kaminski, and J.H. Hrgovcic, 1999, "Weather Derivatives. Hedging mother nature", *Geman*, pp. 179-82.

Clewlow, L., C. Strickland and M. Booth, 2000, "Weather Derivatives", Short Business Programs, University of Technology, Sydney.

Dischell, R., 1998, "Weather sensitivity, weather derivatives and a pricing model", http://www.adtrading.com (Issue 32).

Dischell, R., 1999, "The Dischell D1 stochastic temperature model for valuing weather futures and options", http://www.adtrading.com (Issue 37).

Geman, H., 1999, *Insurance and weather derivatives. From exotic options to exotic underlyings*, Risk Books.

Jain, G., and D. Foster, 2000, "Beyond Energy. Weather Risk Special Report", *Energy and power risk management*, Vol. 5, No. 5. pp 16-17.

Stern, H., 1992, "The likelihood of climate change: a methodology to assess the risk and the appropriate defence", 5th International Meeting on Statistical Climatology, Toronto, Canada, American Meteorological Society, 22-26 June.