Pricing Weather Derivatives

Lixin Zeng

Lixin Zeng is a vice president in the risk analysis and technology services division at E.W. Blanch Company, in Minneapolis, Minnesota.

The impact of weather on business activities is enormous and varies both geographically and seasonally. For example, the 1982-1983 and 1997-1998 El Niño conditions were associated with warm winters in the eastern and midwestern U.S., resulting in significant energy cost savings for consumers and businesses. In addition, these conditions suppressed hurricane activities in the Atlantic and led to minimal economic losses due to hurricanes. However, the same weather pattern was also associated with extreme floods in California, resulting in both economic loss and loss of life. In general, almost all businesses, and in particular, the energy, property and casualty (re-)insurance, and agricultural industries, are either adversely or favorably affected by the weather.

Faced with these challenges and opportunities, a new financial instrument — the weather derivative — has emerged in recent years (e.g., Jovin [1998]). Weather derivatives are structured as swap, call, and put contracts[1] based on weather indexes. Commonly referenced weather indexes include, but are not limited to, heating degree day (HDD), cooling degree day (CDD), precipitation, and snowfall. For example, a major energy user can hedge the risk associated with lower-than-average winter temperatures by buying a winter HDD call. As another example, a snowmobile retailer can hedge against lower-than-expected revenue by purchasing a snowfall put. The flexibility of defining weather indexes allows innovative structures to be developed using these instruments to manage a wide variety of weather-related risks. Sellers of weather derivatives usually include major energy companies which use the instruments to hedge their own risks and to make trading profits. Insurance and reinsurance companies are also important providers of capacity, as they look for alternative ways to deploy their capital. Although data limitations (due to the relatively short history of the weather derivative market) do not allow a thorough analysis of the correlation between weather derivatives and longer established financial instruments, it is widely perceived that the correlation is negligible. Thus, weather derivatives appeal to a wide array of investors as an uncorrelated asset class.

These opportunities to trade on the weather also pose challenges for financial and risk management professionals (e.g., Kaminski [1998]; Stix [1998]), which include pricing, analysis, and portfolio management. First, the no-arbitrage option pricing model is not a practical pricing tool for weather derivatives because the underlying weather indexes are not traded instruments. Second, there also exist difficulties in implementing actuarial techniques for pricing, because the underlying weather indexes are non-stationary. Rather, they are characterized by long-term variations and trends with scales greater than the length of the historical record. In addition, historical weather indexes exhibit a high degree of autocorrelation. This further reduces the number of independent observations. These characteristics make it impossible for traditional actuarial techniques to provide reliable statistical inferences based on the data. Third, reliably measuring and accounting for volatility in the absence of pricing data is problematic. In fact, there is no universal method for estimating volatility associated with weather derivatives, because the economic exposure to weather risk varies greatly among businesses and individuals, but is not yet commonly traded.

The goal of this article is to develop a new pricing scheme that accommodates and reflects these unique characteristics of weather derivatives. The background and generic formulation of weather derivatives are first reviewed. Next, the new pricing scheme is introduced and discussed. The article closes with some future research ideas.

FORMULATION OF WEATHER DERIVATIVES

In the weather derivatives market, as in other derivative markets, there are three commonly traded contracts: calls, puts, and swaps. The buyer of a call pays the seller a premium at the beginning of the contract. In return, if the underlying weather index (denoted W) for the contract period is greater than the pre-negotiated strike price (denoted S), the seller will pay the buyer an amount equal to $P = k(W - S)$, where k is the pre-negotiated tick value of the contract. In addition to this commonly used linear payout scheme, there also exists a so-called binary payout scheme, which stipulates that a fixed amount P_0 will be paid if W is greater than S. The value of a put is equivalent to the value of a call except that the seller pays the buyers when W is less than S. In contrast to a put or a call, a swap requires no up-front premium and, at the conclusion of the contract, the seller makes a payment in the amount of $P = k(W - S)$ to the buyer. In the case of a negative P, the payment is actually made by the buyer to the seller. S is known as the "exercise index" or "reference index" for the swap.

A generic weather derivative contract can be formulated by specifying the following seven parameters:

- Contract type (swap, call, or put);
- Contract period (e.g., from November 1, 1999 to March 31, 2000);
- An official weather station from which the meteorological record is obtained;
- Definition of the underlying weather index (W);
- Strike for put/call or exercise index for swap (both denoted S);
- Tick (k) for a linear payout scheme or the fixed payment (P_0) for a binary payment scheme (defined in the previous paragraph); and
- Premium (for call and put).

The payout (P) of the contract (the amount that the seller must pay to the buyer) for a linear payment scheme is determined by:

$$\begin{aligned} P_{swap} &= k \bullet (W - S) \\ P_{call} &= k \bullet \max(W - S,\ 0) \\ P_{put} &= k \bullet \max(S - W,\ 0) \end{aligned} \qquad (1\text{-A})$$

where the function max(x, y) returns the greater of values x or y. For a binary payment scheme:

$$\begin{aligned} &P_{swap} = P_0 \text{ if } W - S > 0;\ P_{swap} = -P_0 \text{ if } W - S \le 0 \\ &P_{call} = P_0 \text{ if } W - S > 0;\ P_{call} = 0 \text{ if } W - S \le 0 \\ &P_{put} = P_0 \text{ if } W - S < 0;\ P_{put} = 0 \text{ if } W - S \ge 0 \end{aligned} \qquad (1\text{-B})$$

The structure of a weather derivative with the linear payout scheme is almost identical to that of a derivative that references other financial indexes, e.g., equities or commodities. The only exception is that the parameter k is introduced to link the monetary value of the contract to the value of the weather index since the latter is not traded.

Since HDD (heating degree day) and CDD (cooling degree day) are the most frequently used weather indices in the market, their definitions are provided below:

$$\begin{aligned} HDD &= \sum_{i=1}^{N} \max(0,\ 65^{o}F - T_i) \\ CDD &= \sum_{i=1}^{N} \max(0,\ T_i - 65^{o}F) \end{aligned} \qquad (2)$$

where N is the number of days over the contract period and T_i is the arithmetic average of the observed daily maximum and minimum temperatures on the i-th day of the contract. HDD and CDD measure the average temperature variation, during a given period of days, above or below the threshold level of 65 degrees Fahrenheit, set by market convention.

EXHIBIT 1
Cumulative Probability Distribution Function of a Sample Call Contract — the Solid Curve

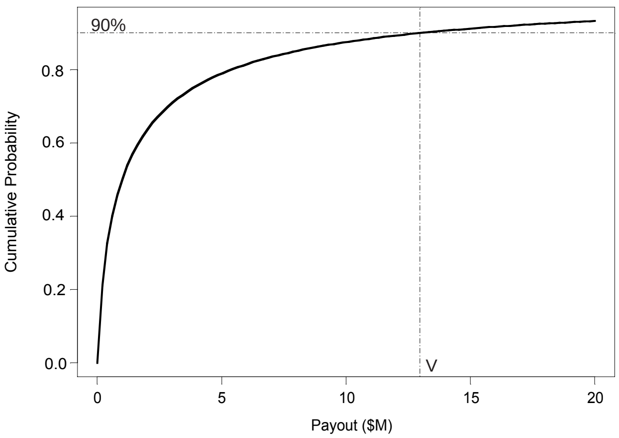

The horizontal dashed line marks the 90% probability of non-exceedance and the vertical dashed lines marks V. It shows that the probability that the payout is less than or equal to V is 90% (i.e. the probability that the payout is greater than V is 10%).

PRICING WEATHER DERIVATIVES

Applicability of Existing Approaches

The actual premium of a call/put (or the exercise index of a swap) is ultimately determined by the supply and demand in the market. The goal of a pricing scheme is to determine their respective fair values.

Traditionally, financial derivatives such as options on equities are priced using *no-arbitrage models* such as the Black-Scholes pricing model (Black and Scholes [1973]), which requires the underlying equity index to be traded. However, since the underlying weather indexes are not traded, a no-arbitrage argument/model cannot be directly applied to price the weather derivatives.

An alternative approach employs *the actuarial method*. The goal is to derive the cumulative probability distribution function (F_p) of the contract payout. The sum of the expected value of the payout (denoted μ) and the overhead expense (e) is the expected cost of the contract, denoted E[C]:

$$E[C] = \mu + e \qquad (3)$$

This is equivalent to the breakeven value such that the long-term profits[2] are non-negative. Thus, the seller of a call or put must demand a premium greater than the expected cost to remain profitable in the long term. In the case of a swap, since no premium is paid at the beginning, an exercise index S must be chosen such that the expected net payout is zero. The chosen index is known as the "fair exercise index."

In addition, the volatility of the payout can be used to determine the level of expected profit that a seller demands. In a rational market, the higher the volatility, the higher the expected profit demanded (and vice versa). Nevertheless, both the measurement of volatility and its effects on the seller's feasible profit vary greatly among businesses and individuals. The market price is primarily determined by aggregating their respective risk tolerances and expectations. Thus, it is not feasible to calculate a unique equilibrium or expected value i.e., implied volatility associated with the contracts. In this study, the following two alternative measures of volatility are calculated:

- The standard deviation of the profit (denoted σ): This is a frequently used measure of the volatility of investment returns. In this particular context, σ closely approximates the standard deviation of the payout because the premium and the overhead cost are usually assumed to be fixed or constant.
- Value at risk with a 10% probability (denoted V): This is defined such that the probability that the payout is greater than or equal to V is 10%. The percentage can be different depending on individual different risk tolerances among market participants. Value at risk is also known as the probable maximum loss or PML.

As noted by D'Arcy [1999], no single measure can provide a complete description of the risk. The only way to fully understand the risk is study the area underlying the cumulative distribution F_p (Exhibit 1). In fact, as shown in the figure, V simply represents a point on this curve. Nevertheless, as summary statistics, σ and V provide useful inferences related to the risk process.

For the sake of consistency, a July CDD call for Phoenix, Arizona, is used throughout the text. It has a strike of 900 degree days and a tick of $5000 per degree day. This example is referred to hereafter as the Phoenix CDD call.

A traditional actuarial technique is a purely statistical approach in which historical records of the observed physical process (e.g., temperature or precipitation) are the main source for the calculation of F_p as well as the associated μ, σ, and V. To analyze the Phoenix CDD call, historical records of July CDD are collected and used to calculate the hypothetical payout of the contract had it been in place in the corresponding years (Exhibit 2) using Equation 1-A or 1-B.

The μ, σ, and V of the contract payout can be estimated either based exclusively on the historical data or by using long-run estimates that are generated via a Monte Carlo simulation approach. In the latter, a parametric probability distribution function (e.g., a normal distribution) is fitted to the weather index data. A large volume of sample weather index values is simulated based on the parameters of the fitted distribution. Based on these sample weather indexes, synthetic payout values are calculated using Equation 1-A or 1-B. These are then used to estimate μ, σ, and V. The results are summarized in Exhibits 3A and 3B.

Predictably, the expected values of μ and σ are the same based on the historical estimates and the Monte Carlo estimates, given the probability distribution is fitted with respect to μ and σ. However, V is significantly different. This is primarily because the small size of the sample historical data does not allow a reliable estimation of the values at the tail of the distribution. This is a common difficulty encountered when pricing with only a limited amount of historical data available. Solely based on the data, it is impossible to determine whether this is because the true distribution is not asymptotically normal or because the historical estimates are simply not statistically significant.

Another unique problem with weather derivatives arises due to long-term variations frequently observed in weather indices. For example, the estimates of μ, σ, and V are all sensitive to the number of years in the historical data included in the calculation primarily because of the upward CDD trend (Exhibit 2). Based on the historical data alone, it is impossible to determine whether such variations are persistent trends or oscillations with scales greater than the length of the historical record. Consequently, there does not exist an objective specification to determine how time variation of the weather processes (e.g., temperature) should be modeled statistically.

Finally, a relatively minor limitation with the pure statistical approach is due to the autocorrelation present in most weather indexes. For example, the historical CDD data for Phoenix CDD call demonstrates significant autocorrelation for lags up to eight years (Exhibit 4). This effectively reduces the number of independent observations (thus the degrees of freedom) available to estimate the values of μ, σ, and V.

EXHIBIT 2

Historical Records of Weather Index (dashed line) and Payout (solid line) of the Phoenix CDD Call

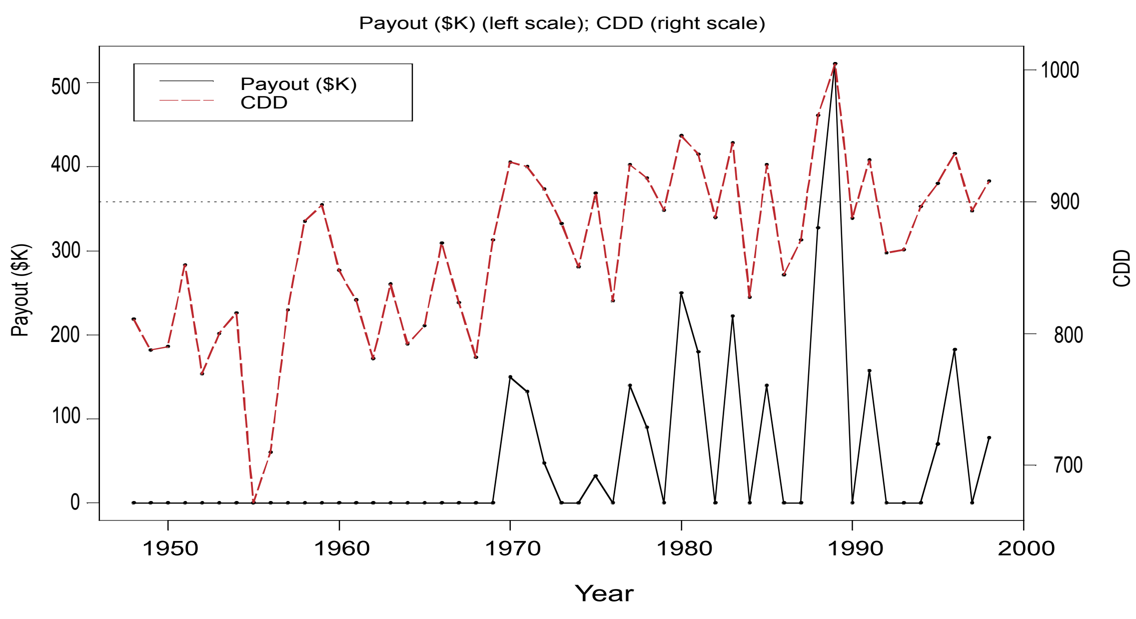

The tick is $5,000 per degree day. The dotted horizontal line marks the strike (900 degree days).

EXHIBIT 3A

Statistics of the Payout of the Phoenix CDD Call Directly Based on Historical Data

Years Included in Calculation	μ	σ	V
All 51 in the historical data	53	104	180
Latest 30	91	123	225
Latest 20	106	143	258
Latest 10	101	163	216

EXHIBIT 3B

Statistics of the Payout of the Phoenix CDD Call Based on Monte Carlo Approach

Years Included in Calculation	μ	σ	V
All 51 in the historical data	59	131	231
Latest 30	90	126	275
Latest 20	107	139	314
Latest 10	114	141	324

Prediction-Based Pricing Approach

The climate system is complex and nonlinear, and its precise prediction is currently infeasible. For example, no climate model can predict a precise CDD value for Phoenix, Arizona, during July 2000. Instead, seasonal predictions of the probabilities that the temperature or the precipitation will be above, near and below the climate norm (denoted p_A, p_N, and p_B, respectively) during three-month periods are provided. The climate norm is based on the historical data from 1961 to 1990. The forecasts extend to the next twelve months. For example, for June-July-August (JJA) 2000 in Phoenix, Arizona, the latest NCEP release (November 18, 1999) predicts p_A, p_N, and p_B to be approximately 0.41, 0.33, and 0.26, respectively.

This study proposes a scheme for objectively including these forecasts in the pricing of a weather derivative. Two assumptions are made in order to establish that p_A, p_N, and p_B are acceptable approximations of the probabilities that the July CDD for Phoenix, Arizona will be above, near, and below the climate norm, respectively. First, the JJA forecast is used as a proxy to that for July since there is no prediction specifically for this month. This is not perfect but acceptable because the rank-order correlation (i.e., the linear correlation of the order statistics) between the mean July and JJA temperatures is 0.82. A rank order correlation close to one indicates that a high July temperature tends to be associated with a high JJA temperature. Secondly, the predicted anomaly probabilities for the temperature (p_A, p_N, and p_B) are assumed to approximate the probabilities that the CDD will be above, near, and below the climate norm, respectively. This is justified by the fact that the rank order correlation between the July monthly mean temperature and CDD data is one.[3]

The proposed method first fits a normal distribution (Exhibit 5) to the historical July CDD data between 1961 and 1990. This is done by assuming that the July CDD follows a normal distribution with the mean and the standard deviation equal to the sample mean and standard deviation, respectively, of the historical data. The probability of non-exceedance is evenly divided into the highest, middle, and lowest thirds, separated by the two vertical dashed lines in Exhibit 5. The fitted distribution is then sampled such that the numbers of samples corresponding to the highest, middle and lowest thirds are proportional to p_A, p_N, and p_B, respectively. For the Phoenix CDD call, 4100, 3300, and 2600 samples are drawn from the highest, middle, and lowest thirds, respectively. This differs from a traditional Monte Carlo method, which samples the fitted distribution evenly across the probability distribution. This difference allows the method to incorporate the probabilistic cli-

EXHIBIT 4

Autocorrelation Function (ACF) of CDD Data Shown in Exhibit 2

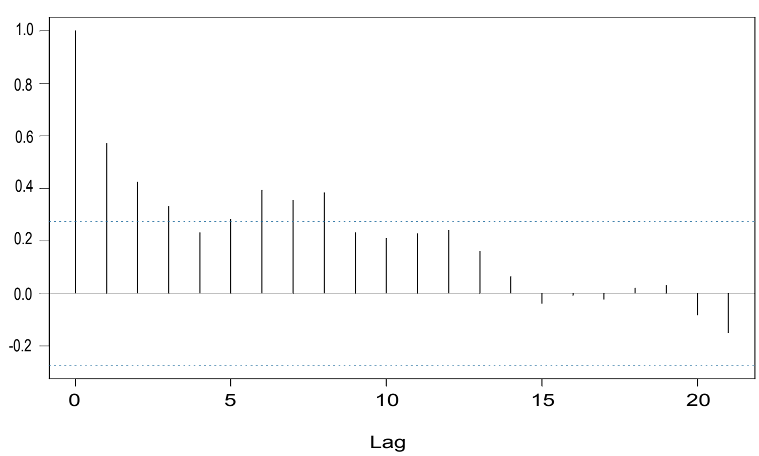

The lag is in years. The dotted horizontal line marks the 95% confidence level.

Exhibit 5

Sample Distribution of Historical Data Between 1961-1990 (solid curve) and Fitted Normal Distribution (dashed curve)

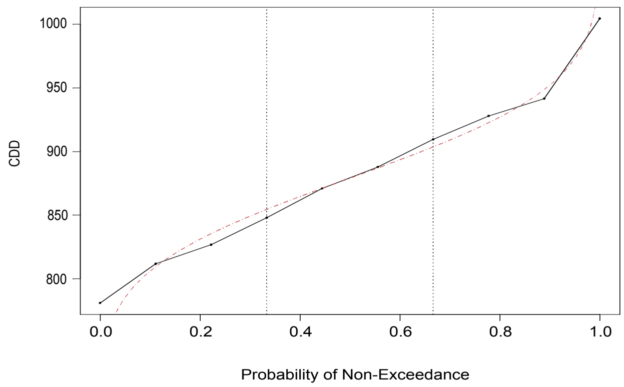

The two vertical dashed lines mark the 33% and 67% percentiles.

mate prediction into the sample CDD values. This new scheme will now be referred to as the biased sampling Monte Carlo approach.

The next steps are the same as in the traditional Monte Carlo approach: the synthetic contract payout values are calculated by Equation 1-A or 1-B. The statistics are summarized in Exhibit 6.

Exhibit 6

Statistics of the Payout of July CDD Call Directly Based on Traditional and Biased Sampling Monte Carlo Approaches

Method	μ	σ	V
Traditional	70	131	264
Biased Sampling	85	142	293

The higher μ and V values reflect the predicted warmer summer, which usually leads to higher CDD and higher potential payout of a CDD call contract. Since the NCEP prediction has been operational for only a year, it is currently impossible to quantify the degree of improvement of the biased sampling over a traditional Monte Carlo approach. However, because such a model explicitly includes the NCEP prediction reflecting the physical processes of the climate system, it is believed to be a better tool for evaluating weather derivatives in force during the prediction period. The reliability of the results relies on the accuracy of the NCEP prediction and quality of the historical data. Thus, any future improvement of the NCEP prediction and data quality will enhance the reliability of our method.

THE GLOBAL CLIMATE SYSTEM

The understanding of the global climate system has been improved as a result of technological advancements achieved in both satellite remote sensing and dynamic modeling techniques (e.g., numerical models on super computers, hybrid dynamic-statistical models). These make feasible the prediction of the interannual variation of the global climate, including the El Niño/La Niña cycle (the periodic warming and cooling of the tropical eastern Pacific Ocean surface), that has profound influences on the weather patterns in many parts of the world. Combining the El Niño/La Niña prediction with the teleconnections between the El Niño/La Niña pattern and the local climate impacts (e.g., Shukla [1998]), local climate patterns can be predicted. For the latest summary and review of the prediction methodology and validation of the prediction produced by the National Center for Environment Prediction (NCEP), the reader is referred to Barnston et al. [1999].

CONCLUSION

This study introduces the basic structures of weather derivatives and reviews the applicability of traditional pricing methods to these relatively new instruments. The no-arbitrage model is not a practical choice because the underlying weather indexes are not traded. The ability of the traditional actuarial approach is limited by the particular statistical properties of the underlying weather indexes, namely the existence of a high degree of autocorrelation and long-term variations in the historical data. A biased sampling Monte Carlo approach is proposed to alleviate these problems by taking advantage of the NCEP seasonal climate forecast, based on a hybrid dynamic-statistical model. Future study will be needed to evaluate the track record of this scheme as NCEP forecasts become continuously available in the coming years.

ENDNOTES

[1]Most weather derivatives are currently traded over the counter as individually negotiated contracts. Exchange-based weather futures and options on futures started trading at the Chicago Mercantile Exchange in September 1999.

[2]Defined as the premium minus expected cost. Since the weather derivatives are usually short-term, the investment income from premium is not discussed here for brevity.

[3]Similar investigations show that, in the case of the weather index being HDD, p_A, p_N, and p_B accurately represent the probabilities that the HDD is below, near, and above the HDD climate norm, respectively. In addition, these assumptions are shown to be acceptable for the approximately 250 weather stations commonly used in most weather derivatives.

REFERENCES

Barnston, A.G., A. Leetmaa, V.E. Kousky, R.E. Livezey, E.A. O'Lenic, H.V. den Dool, A.J. Wagner, and D.A. Unger. "NCEP Forecasts of the El Niño of 1997-98 and Its U.S. Impacts." *Bulletin of the American Meteorological Society*, 80 (1999), pp. 1829-1852.

Black, Fischer, and Myron Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81 (1973), pp. 637-659.

D'Arcy, Stephen. "Don't Focus on the Tail: Study the Whole Dog." *Risk Management and Insurance Review*, 2 (1999), pp. 4-14.

Jovin, Ellen. "Advances on the Weather Front (An Overview of the U.S. Weather Derivatives Market, Global Energy Risk)." *Electrical World*, 212 (9) (1998), p. S4 (2).

Kaminski, Vincent. "Pricing Weather Derivatives (Global Energy Risk)." *Electrical World*, 212 (9) (1998), p. S6 (2).

Shukla, J. "Predicability in the Midst of Chaos: A Scientific Basis for Climate Forecasting." *Science*, 282 (1998), pp. 728-731.

Stix, G. "A Calculus of Risk." *Scientific American*, 278 (5) (1998), p. 92 (6).