WATCHING THE Weather Report

Weather derivatives based on temperatures and precipitation can enhance portfolio diversification and asset allocation.

BY MELANIE CAO, ANLONG LI AND JASON WEI

Weather derivatives date back to the mid-'90s, when deregulation of the energy and utility industries started in the U.S. Faced with growing competition and uncertain demand, energy and utility companies sought effective hedging tools to stabilize earnings. As monopolies gave way to competitive wholesale markets, hedging price alone was no longer adequate since volumetric risk also came into play.

In the deregulated environment, energy merchants quickly realized that weather conditions were the main source of revenue uncertainties. Weather affects both short-term demand and long-term supply of electricity and natural gas. A particular weather pattern (e.g., a strong global warming trend) can also affect the long-term supply as energy producers re-adjust their production levels.

This close association created a natural impetus for the development of weather derivatives. The first publicized deal was signed in 1997 between Koch Energy and Enron on a temperature index for Milwaukee, Wisconsin, for the winter of 1997-1998.

In September 1999, the Chicago Mercantile Exchange (CME) began listing futures and options on temperature indices. At the time of writing, the CME lists contracts on temperature indices for 15 U.S. cities and five European cities. The cities were chosen based on population, the variability in their seasonal temperatures and the activities seen in OTC markets. The active OTC markets and the appointment of a market maker[1] helped to boost the trading volume on the CME. The total number of contracts traded was 4,165 in 2002 and 14,234 in 2003.

Weather conditions affect not only the energy and utility sectors, but also many other sectors such as agriculture, retail, entertainment, and tourism. In fact, nearly 20% of the U.S. economy is directly affected by weather (see Challis [1999] and Hanley [1999]). Weather derivatives therefore play an important role in integrated risk management and diversification.

Product Descriptions

Although deals have been struck on such underlying variables as rainfall, snowfall and humidity, the vast

FIGURE 1: NUMBER OF CONTRACTS BY TYPE

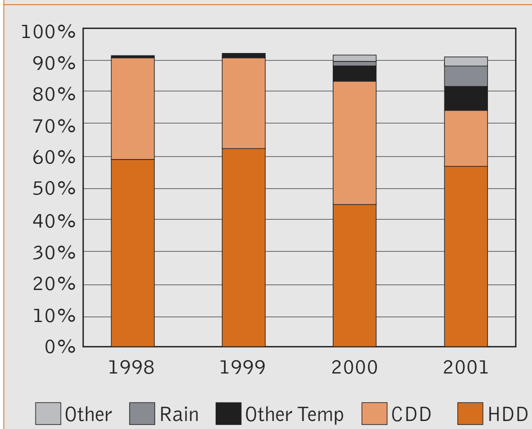

Source: WRMA/PricewaterhouseCoopers

Melanie Cao is assistant professor at York University's Schulich School of Business in Toronto. Anlong Li is director, research at XL Weather & Energy Inc. in Stamford, Conn. Jason Wei is associate professor at the University of Toronto's Rotman School of Management.

majority of contracts are on temperature. As Figure 1 shows, the two most popular contract variables are the heating degree day (HDD) and cooling degree day (CDD). A degree day measures how much a day's average temperature deviates from 65°F (or 18°C), a level of indoor temperature deemed to be comfortable by the utility industry. Precisely speaking, daily HDD = max [0, 65°F - daily average temperature] and daily CDD = max [0, daily average temperature -65°F], where "daily average temperature" is the average of the maximum and minimum temperatures of the day.

Most contracts are written on the accumulation of HDDs or CDDs over a calendar month or a season so that one contract can hedge against revenue fluctuations over a period. At the CME, for instance, contracts on monthly HDDs or CDDs are listed for all 15 U.S. cities, while contracts on seasonal HDDs and CDDs are listed for 10 of the 15 cities.

In some cases, "energy degree day" (EDD) is used in lieu of HDD or CDD. EDDs are the absolute deviation from a benchmark temperature which can be different from 65°F. When the benchmark temperature is 65°F, EDDs are simply the sum of HDDs and CDDs. EDDs can help manage the temperature risk not just for a season but for the whole year. Table 1 contains an example of calculating the daily and cumulative HDDs, CDDs and EDDs.

Broadly speaking, there are three types of temperature derivatives: futures/forwards, swaps, and options. Besides the underlying variable HDD or CDD, a contract must specify such basic elements as the accumulation period, the index station which records temperatures used to construct the underlying variable, and the tick size, i.e., the dollar amount attached to each HDD or CDD. The following examples illustrate forward and option contracts based on HDDs and CDDs.

In the HDD forward example, Power Supply Ltd. agreed to short the cumulative HDDs for the month of February using the Philadelphia International Airport as the index station. The settlement level was 850 HDDs. The tick size was set at $4,000 per HDD, which measures the sensitivity of Power Supply's revenue with respect to changes in the HDD level. With a realization of 650 HDDs, ABC Bank paid Power Supply Ltd. $800,000 to settle the contract.

The CDD put option works in a similar fashion. A cap or maximum payoff is typically specified for an option contract. For instance, the payoff function for a CDD put option with a cap would be specified as min[cap, tick x max (0, strike - CDD)]. In our example, had a cap of $350,000 been specified, the settlement payoff would have been $350,000.

As any other derivative securities, weather derivatives serve the ultimate purpose of risk transfer. Power and utility companies are interested in smoothing their earnings by engaging in price and volumetric hedges. Also, insurance companies, power/energy brokers, and brokerage firms are in a position to act as counterparties thanks to their ability to effectively pool the weather risk and eventually lay it off in the organized market such as the CME. Of course, they earn a fee or a mark-up in the process. Practical examples of weather derivatives being used as a hedge against inclement weather might include a natural gas or electric utility entering into a forward contract on a series of HDDs or CDDs (either with or without a put option to eliminate downside risk), or a restaurant, bar, resort, or other recreational establishment buying a weather hedge against adverse conditions that might have a negative effect on business.

Examples	HDD Forward	CDD Put Option
Current time	December 1, 2001	
Location	Phil. Int'l Airport, Philadelphia	Hartsfield Airport, Atlanta
Long Position	ABC Bank	Air Conditioning Ltd.
Short Position	Power Supply Ltd.	XYZ Bank
Accumulation Period	February 2002	July 2002
Tick Size	$4,000 per HDD	$10,000 per CDD
Settlement Level	850 HDDs	
Strike Level		550 CDDs
Actual Level	650 HDDs	510 CDDs
Payoffs at Maturity	(650 - 850)(4,000)	(550 - 510)(10,000)
(Long Position)	=-$800,000	=$400,000

Temperature Modeling and Derivatives Valuation

The valuation of temperature derivatives has some unique features. To start with, the underlying is a meteorological variable rather than a traded asset—the conventional risk-neutral, arbitrage-free valuation methodology does not apply. In addition, being a meteoro-

TABLE 1: HDDs, CDDs AND EDDs FOR JUNE 6 - JUNE 17, 2002, AT LAGUARDIA INTERNATIONAL AIRPORT, NY

Day	Daily Max.	Daily Min.	Daily Average	Daily HDDs	Cumulative HDDs	Daily CDDs	Cumulative CDDs	Daily EDDs	Cumulative EDDs
June 6	77	59	68.0	0.0	0.0	3.0	3.0	3.0	3.0
June 7	69	57	63.0	2.0	2.0	0.0	3.0	2.0	5.0
June 8	69	55	62.0	3.0	5.0	0.0	3.0	3.0	8.0
June 9	85	57	71.0	0.0	5.0	6.0	9.0	6.0	14.0
June 10	78	65	71.5	0.0	5.0	6.5	15.5	6.5	20.5
June 11	92	64	78.0	0.0	5.0	13.0	28.5	13.0	33.5
June 12	90	60	75.0	0.0	5.0	10.0	38.5	10.0	43.5
June 13	66	57	61.5	3.5	8.5	0.0	38.5	3.5	47.0
June 14	61	56	58.5	6.5	15.0	0.0	38.5	6.5	53.5
June 15	63	55	59.0	6.0	21.0	0.0	38.5	6.0	59.5
June 16	79	61	70.0	0.0	21.0	5.0	43.5	5.0	64.5
June 17	78	63	70.5	0.0	21.0	5.5	49.0	5.5	70.0

logical variable, temperature follows a predictable trend, especially over a longer horizon. The unique nature of the temperature variable brings about two important issues: accurate modeling of the underlying and the assessment of the market price of risk.

The academic literature only begins to make progress in valuing this new class of derivative securities (see Cao and Wei [2003]). In the industry, regardless of the valuation methodologies, a mark-up is usually attached to the model price as a cushion for errors. From the modeling perspective, the existing valuation methods can be loosely classified into three categories: 1) insurance or actuarial valuation, 2) historical burn analysis, and 3) valuation based on dynamic models.

Insurance or Actuarial Method. This method is widely used by insurance companies, and its backbone is statistical analysis based on historical data. A probabilistic assessment is attached to the insured event and a fair premium is calculated accordingly. In the case of weather derivatives, this method is less applicable for most contracts since the underlying variables (e.g., temperature) tend to follow a recurrent, predictable pattern. Nonetheless, if the contract is written on rare weather events such as extreme heat or coldness, then this method will be very useful. In fact, one may even argue that this is the *only* appropriate method in this case, since using a diffusion process to model the temperature will be misguided if the main interest is in extreme events.

Historical Burn Analysis. This method is perhaps the simplest to implement, and as a result, is most prone to large pricing errors. In a nutshell, this method evaluates the contract against historical data and takes the average of realized payoffs as the fair value estimate (see Dischel [1999] for further discussions). The key assumption is that the past always reflects the future on average. This is a strong requirement in most cases. It is not true that a longer time series will always enhance valuation accuracy. Although more data will cover more temperature variations, the future temperature behavior, which drives the derivative security's value, may be quite different from history. This is especially important when the derivative security's maturity is short. Ultimately, it boils down to a trade-off between statistical power and representativeness. The commonly accepted sample length in the industry appears to be between 20 and 30 years.

Similar to the insurance or actuarial method, historical burn analysis is incapable of accounting for the market price of risk associated with the temperature variable. These methods are only useful from the perspective of a single dealer. We need a dynamic and forward looking model to establish a unique market price which incorporates a risk premium.

Dynamic Valuation Models. In contrast to previous methods, a dynamic model directly simulates the future behavior of temperature as a continuous or discrete stochastic process. The continuous process usually takes the following mean-reversion form:

$$dY(t)=\beta[\theta(t) - Y(t)]dt + \sigma(t)dz(t)$$

where $Y(t)$ is the current temperature, $\theta(t)$ is the deterministic long-run level of the temperature, β is the speed

TABLE 2: MEAN, STANDARD DEVIATION, CORRELATION MATRIX

	Correlations						Mean	Std
	N.America	Europe	Pacific	Bond	Commodity	Temp.		
N. America	1.0000						0.0861	0.1612
Europe	0.3815	1.0000					0.0480	0.1508
Pacific	0.1043	0.3456	1.0000				-0.0223	0.1847
Bond	0.0014	-0.0121	-0.0523	1.00000			0.0126	0.0437
Commodity	-0.0103	-0.0277	0.0220	-0.0755	1.0000		-0.0027	0.1742
Temp.	0.0192	-0.0118	-0.0202	-0.0078	-0.0076	1.000	-0.0002	0.0958

Note: There are six index series: equity indices for North America, Europe and Pacific, JPMorgan Bond Index based on U.S. government bonds, Goldman Sachs Commodity Index, and the Temperature Index. The Temperature Index is for New York City, constructed as 1,000 plus the deviation of daily EDDs from the seasonal average EDDs. The sample period is from January 1, 1991, to December 31, 2002, with daily frequency. The mean and standard deviation (Std) for each index are annualized from their daily counterparts.

at which the instantaneous temperature reverts to the long-run level $\theta(t)$, $\sigma(t)$ is the volatility which is season-dependent, and $z(t)$ is a Wiener process which models the temperature's random innovations (see Dischel [1998] and Brody, Syroke and Zervos [2002] for similar specifications). The above process needs to be discretized in order to estimate β and the parameters imbedded in $\theta(t)$ and $\sigma(t)$. The functional forms for $\theta(t)$ and $\sigma(t)$ can be specified based on careful statistical analyses. Once the process is estimated, one can then value any contingent claim by taking expectation of the discounted future payoff.

The above continuous setup usually does not admit closed-form valuation formulas. Additionally, a risk-neutral valuation is *imposed* without any theoretical justification, and the market price of risk is rendered irrelevant. Moreover, the process cannot reflect the persistent serial correlations typically present in daily temperatures.

With the above in mind, researchers (see Cao and Wei [2000], Campbell and Diebold [2003], and Cao and Wei [2003]) have proposed discrete processes. In addition, Cao and Wei [2003] weave the temperature uncertainty and the economy's aggregate output into an equilibrium framework. Motivated by the significant influence of weather on the overall economy, Cao and Wei [2003] propose a serially correlated bivariate-process for the temperature and the aggregate output, and address the market price of weather risk therein.

The temperature process proposed by Cao and Wei [2003] possesses the following features, all of which are based on their careful study of the temperature behavior for U.S. cities. The daily temperature has two components, the first being the seasonal pattern plus a global warming trend and the second being a random innovation; the innovation is serially correlated; and the standard deviation of the innovation is higher in the winter and lower in the summer, captured by a sine wave function. The aggregate output follows a mean-reverting process that is correlated with the current and past temperature innovations. The last building block is the representative agent's preference, which Cao and Wei [2003] specify as constant relative risk aversion (CRRA). Given the temperature risk embedded in the aggregate output, the risk aversion determines the risk premium via equilibrium valuations.

The market price of risk associated with the temperature variable is found to be significant in most cases. Risk premium can represent a significant portion of the derivative's value. Using the risk-free rate to discount the expected payoff will lead to a sizeable error. It is also found that the market price of risk affects option values much more than forward prices. This result is mainly due to the non-linearity in option payoffs. Intuitively, the market price of risk tends to be integrated out in the linear payoffs of forward contracts.

Asset Allocation and Weather Derivatives

Purely from a diversification perspective, weather derivatives may indeed hold some potential. To gain some insight, we now examine the efficient frontiers consisting of the following asset classes: equity, fixed income, commodities and temperature instruments.

Three indices are included in the equity class, namely the regional indices for North America, Europe and Pacific. Fixed income is represented by the JPMorgan Bond Index, while commodities are represented by the Goldman Sachs Commodity Index which covers industrial metals, precious metals, live stocks, agriculture and energy. For the weather investment class, we construct a hypothetical index using New York City's historical temperatures. The index is the 30-day lead daily EDD residual which is constructed by subtracting the historical average EDDs from that day's realized EDDs. For instance, if the historical average EDDs for January 1 is 27 and the realized EDDs for January 1, 2001, and January 1, 2002, are 24.5 and 30.6 respectively, then the two residuals will be -2.5 and 3.6. To avoid a negative index level (for the purpose of calculating returns) and to scale the variance of the return series, we add 1,000 to the EDD residual to arrive at the temperature index. The 30-day lead is used to rule out predictability.[2] Finally, all the time series are for the period of January 1, 1991, to December 31, 2002, with daily frequency.

Table 2 presents the annualized means, standard deviations, and the correlation matrix for the six indices. The North American market fared the best in terms of annualized average returns (8.61%), followed by the European market (4.80%). The Pacific market saw an average loss of 2.23%, and also exhibited the largest standard deviation. The bond index produced a modest gain (1.26%), while both the commodity and the temperature indices experienced a minor loss. As for correlations, the equity indices exhibit relatively larger correlations among themselves. The highest correlation, 0.3815, is between the North American market and the European market.[3] By and large, the correlation of the temperature index with other indices is close to zero, suggesting potential for diversification.

To see how much each asset class can contribute to overall diversification in addition to equities, we constructed efficient frontiers using the above data. We start with the three equity indices and add one index at a time. Adding the fixed-income asset class lowers the minimum variance substantially, and improves the risk-return trade-off for the higher range of standard deviations. As a matter of fact, adding the commodity or temperature asset class to the three equity indices leads to similar results. Now, once the bond index is included, introducing the commodity index does not lead to noticeable improvements in risk-return trade-off. However, including the temperature index over and above all other indices does lead to a notable improvement, especially in the higher range of the standard deviation. For instance, at a portfolio standard deviation of 0.3, the expected return improves by about 40 basis points after adding the temperature index; the improvement is about 60 basis points for a standard deviation of 0.4. We can therefore conclude that weather derivatives, as an alternative class of financial instruments, do hold potential in asset allocation and portfolio management.

Summary

This article offers an introduction to the emerging weather derivatives market. The structure and usage of various weather derivative products are surveyed, and the modeling and pricing issues are discussed. The literature on weather derivatives valuation is still in its infancy, and much more research needs to be done to accurately model weather variables. The article also demonstrates the role of weather derivatives in portfolio management. As an alternative class of financial instruments, weather derivatives can improve the risk-return trade-off in asset allocation decisions. ∎

Acknowledgments

The authors gratefully acknowledge the financial support from the Social Sciences and Humanities Research Council of Canada. The views expressed here are solely the authors' and are not those of the XL Weather & Energy Inc.

References

Brody, D.C., J. Syroke and M. Zervos, "Dynamic Pricing of Weather Derivatives," *Quantitative Finance*, 2 (2002), 189-198.

Campbell, S. and F. X. Diebold, "Weather Forecasting for Weather Derivatives," NBER Working Paper, No. W10141 (2003).

Cao, M. and J. Wei, "Pricing the Weather," *Risk*, May (2000), 67-70.

Cao, M. and J. Wei, "Weather Derivatives Valuation and Market Price of Weather Risk," *Journal of Futures Markets*, 2003, forthcoming.

Challis, S., "Bright Forecast for Profits," *Reactions*, June (1999).

Dischel, B., "Black-Scholes Won't Do," *Risk*, October (1998), 8-9.

Dischel, B., "Shaping History," *WeatherRisk* (Special Report of Risk, 1999), 13-15.

Hanley, M., "Hedging the Force of Nature," *Risk Professional*, 5 July/August (1999), 21-25.

Locke, Jane, "A Cure for Power Price Spikes?" *WeatherRisk* (Special Report of Risk, October 1998), 4-5.

Notes

1. In May 2002, Wolverine Trading, L.P. was named as the lead market maker in CME's weather futures contracts. Wolverine posts continuous bids and offers each day to ensure the liquidity of the market.
2. It is well known that the current technology cannot produce reliable and accurate daily temperature forecasts beyond 10 days into the future.
3. The correlation is higher if the indices are measured in local currencies. Here, the Europe and Pacific indices are measured in U.S. dollars since we take the perspective of a U.S. investor.