Weather derivatives

Introduction

Evaluating weather derivatives requires a different approach to that used for evaluating common financial products. One reason is the difficulty of replication since temperature, rainfall or wind is not a traded asset. Consequently delta neutral techniques cannot be used and, in addition, there is a lack of liquidity in some temperature contracts. Therefore a number of market participants have started to use an actuarial approach when dealing with weather derivatives. Extracting and detrending HDD or CDD[1] from the data, and then fitting a distribution to the events makes valuation possible based on the expectation of the loss (example below) plus a given risk premium that reflects the sensitivity to risk. However, in doing so, a number of problems arise. These stem from the fact from that, in most cases, a maximum of 40 years of data is available. Some of these issues include:

- How many years should I take into account?
- Should I detrend the temperature before or after extracting HDD /CDD?
- Were the data correctly detrended and was the forward accurately extracted?
- Is the fitted distribution accurate or appropriate?
- Has there been any change in the distribution in recent years?

Example of "Actuarial" pricing

"Actuarial" pricing methodology is based on extracting the distribution of risk from historical values. The three necessary steps to obtain the fair value of a weather contract are illustrated below using Heathrow as a reference site.

Step 1 - Filtration:

First the data are filtered taking into account any trend and extrapolated to the next year.

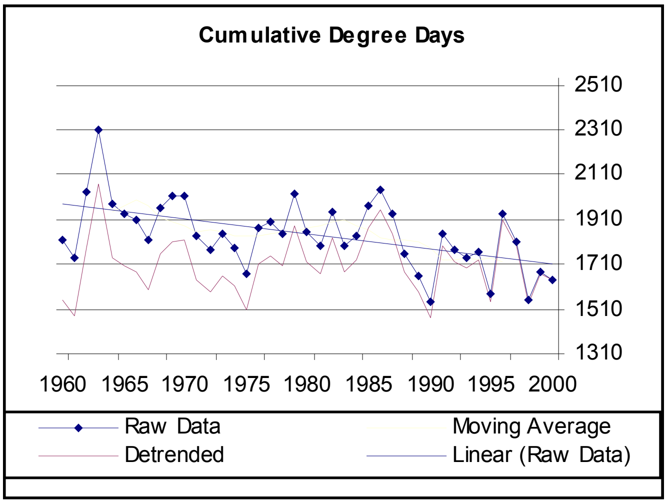

[1] HDD (Heating Degree Days) and CDD (Cooling Degree Days) are indices representing the coldness and the warmth of winter and summer respectively. Other indices like CTD+ CTD- (Critical Temperature Days above or below a reference) or GDD (Growing Degree Days) are also commonly used indices.

Step 2 – Fitting distribution:
Then, a parametric distribution is fitted to the discrete distribution of risk like the normal distribution in the next figure:

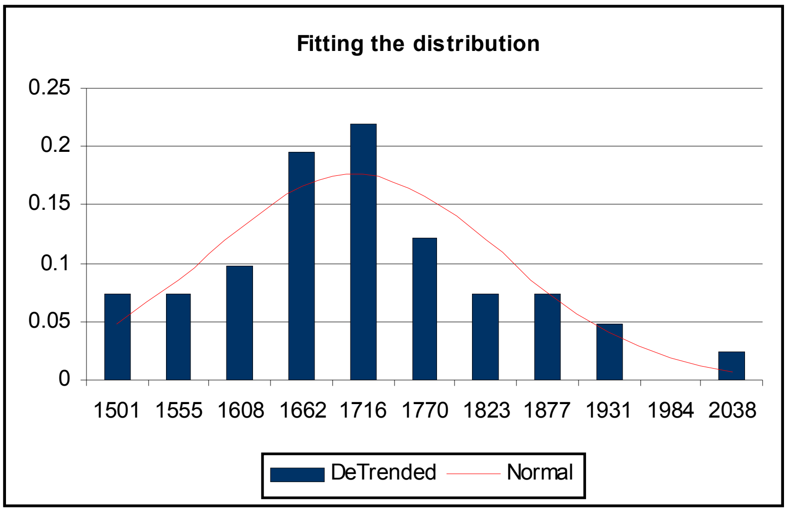

Step 3 – Evaluation:

Finally, the price is evaluated using closed formulae.

Because an HDD is a cumulative index, it is possible to break down the price of a call option[2] into an up & out call and a digital call. This makes them easy to value.

So, assuming a normal distribution, the price of an up & out call is:

$$HDD_{Value} \times \frac{\sigma}{\sqrt{2\pi}} \left(e^{-\frac{\alpha^2}{2}} - e^{-\frac{\beta^2}{2}} \right) + HDD_{Value} \times \gamma \left[N(\beta;0;1) - N(\alpha;0;1) \right]$$

and a digital call is :

$$Cap \left[1 - N\left(\frac{HDD_{Sup} - \mu}{\sigma};0;1 \right) \right]$$

where HDD_{Value} is the value for each recorded degree day below the temperature reference (usually 65°F), HDD_{inf} and HDD_{sup} are the strikes of the call option, Cap = (HDD_{sup} – HDD_{inf}) × HDD_{Value} is the maximum money value of the pay off, $\alpha = \frac{HDD_{\inf} - \mu}{\sigma}$; $\beta = \frac{HDD_{Sup} - \mu}{\sigma}$; $\gamma = (\mu - HDD_{Inf})\sigma$; μ & σ are the mean and the standard deviation of the HDD distribution ; N(X;0;1) the standard normal cumulative distribution function evaluated in X.

Of course, the evaluation strongly depends on the treatment of the data (filtration) and the fitted distribution. All these referred problems and pitfalls are exacerbated in portfolio management due to the requirement to estimate multivariate distributions.

[2] In the weather market all options are capped. Therefore a call option refers to a call spread .

Temperature simulation

Simulating the underlying index (temperature, rainfall, etc) becomes extremely important because of the constraints above. Instead of merely having forty data points available, one can recreate the temperature process extracted from 40 * 365 = 14,600 data points. The process can then be simulated thousands of times to simulate the real distribution of the pay offs of the weather deal.

One can simulate the temperature using the following process:

$$T_i = m_i + S_i + \gamma_i AR(p)$$

Where T_i is the temperature value at time i, S_i the seasonality of the temperature at time i, m_i the trend of the temperature, γ_i a sinusoid function of i and AR(p) an autoregressive process of order p with non identically distributed noise [2].
Fitting it to London data for the period starting the 1st of November and ending the 31st of March (simulation started the 30th October) we obtain the following distribution:

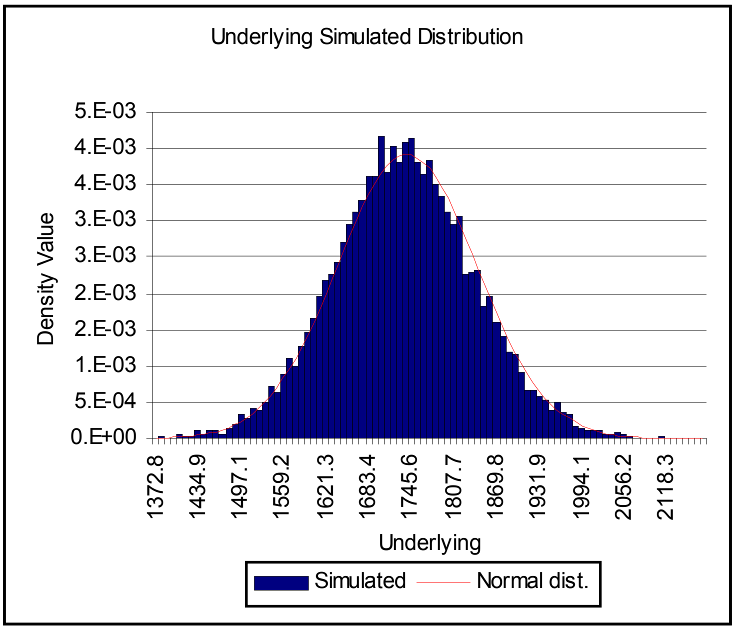

The main information that we extract from this distribution is that the **forward is different** from the one extracted using detrended HDD (1739 instead of 1702) and that the **volatility is lower** than the one extracted (100 instead of 130). There are also some limited differences in the Skewness and Kurtosis.
Some advantages of the simulation are obvious. **When the option has started, there is no need to value any conditional distribution, as one needs to do in "actuarial" analysis; the information is contained in the process itself**. Other benefits of temperature simulation lie in portfolio risk management as explained below.

Portfolio Analysis

E.g. 1) Consider two sites, London and Manchester. These are separated by approximately 180 Miles (roughly 300 km). The correlation one should model for the combined profit and loss balance of a portfolio of temperature contracts based on these cities is not necessarily that based on the correlation of the HDD indices. Supposing, for example, that the portfolio consists of two Critical Temperature (CTD) options, one referring to London and the other one to Manchester. In this example, the critical temperature reference is set to 34C (93.2F). It can be seen that over the last forty years the maximum temperature measured in Manchester

is 33.7C and in London 36.5C. Therefore, on a descriptive point of view, one can conclude that the Pearson's correlation between these two indices is zero. In reality we know this cannot be the case: the underlying processes in both are temperature processes and we know these to be correlated.
In this example, simulation is key since one can extract from the previous equation the residues for both sites and separate out the seasonal effects. Clearly, when it is wintertime in London it is wintertime in Manchester and so the temperature is correlated. But looking at the residues the view is somewhat different. **If we know that the temperature is below its average in London does this give us any information about the temperature value in Manchester relative to its own average?** With CTD structures this becomes an interesting correlation since we may not want them to happen simultaneously in both locations.

E.g. 2) Suppose now that your portfolio contains two options related to the same location, one based on CTD, the other one based on HDD. We can see that the correlation between these two indices may be very weak even though both indices are derived from the same temperature processes. This is shown in the figure below:

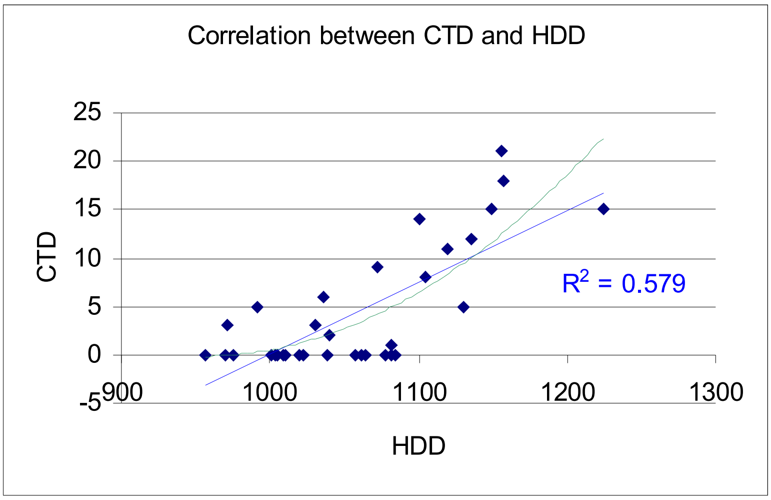

Therefore one needs to use the same temperature values for both locations and not consider them separately. No convexity would be apparent here if using Pearson's correlation.

References

[1] Brockwell P. & Davis R., "Time Series: Theory and Methods", Springer Series in Statistics, Second Edition,1991.

[2] Moreno, "Riding the Temp", FOW, Special Supplement Weather Derivatives, December 2000.

This week's Learning Curve was written by Michael Moreno associate director of ***Speedwell Weather Derivatives*** *Limited and lecturing teacher at the French Management & Actuarial School* ***ISFA****.*