A primer on weather derivatives

Pauline Barrieu* and Olivier Scaillet**

28 January 2008

* London School of Economics

** HEC Geneve and Swiss Finance Institute

Abstract: We give a short introduction to weather derivatives. We discuss what they are made for. We describe the markets on which they are exchanged, and how they are used to promote agricultural risk transfer in developing countries via the World Bank program. We also treat some specific issues such as basis risk, pricing and design.

JEL codes: G12, G13, Q54.
Keywords: weather derivatives, World Bank program.

1 Introduction and definition

1.1 What is weather risk?

Weather is not only an environmental issue but also a key economic factor, as recognized by the former US Commerce Secretary, William Daley, in 1998, when he stated that at least $1 trillion of the world economy is weather sensitive. Moreover, it is estimated that 20% of the world economy is directly affected by weather. In a business survey, 160 out of 200 utility companies cited weather as a key determinant of their operating revenues, and half of them as a key determinant under performance below normal. More importantly, the risk exposure is not homogeneous across the globe and some countries, usually those heavily dependent on agriculture, are more sensitive than others. It also includes a large range of phenomenons such as modifications in temperature, wind, rainfall or snowfall.

Weather risk has some specificities compared to other sources of economic risk: in particular, it is a local geographical risk, which cannot be controlled. The impact of weather is also very predictable: the same causes will always lead to the same effects. Moreover, weather risk is often referred to as a volumetric risk, its potential impacts being mainly on the volume and not (at least directly) on the price. This explains why hedging of weather risk via the trading of commodities futures is difficult and imperfect. For example oil futures price does not depend solely on demand (cold winter) and can be high even if demand is low in case of a war for instance. Volumetric risk is imprecisely compensated by the price variation in the futures position.

Usually, when subject to some risk, it is possible to hedge against it by contracting some insurance policies. But, this is not really a possibility for weather risk for two main reasons: first it is more a high frequency - low severity risk but also the same weather event can generate economic losses for some agents and some gains for others. As an example, we can think of a touristic place. If it rains one day during the summer, this is a bad day for providers of outside activities but this will not lead to tremendous losses. On the contrary, if the whole season is cold and wet, this could lead to the bankruptcy for some local businesses. Some rain in this case could be interesting for coffee shops and indoor businesses. Therefore, weather risk is a risk that is part of everyday life, having limited economic consequences on an everyday basis but with some huge potential consequences in its accumulation or repetition. Insurance is not a well-suited solution, as it could be for dramatic weather events such as hail, storm, or drought. To deal with this risk, some financial contracts depending on weather conditions (temperature, rainfall, snowfall...) were created and introduced on the financial market 10 years ago. They are called *weather derivatives*.

Note that weather derivatives differ from Cat (catastrophe) derivative which target rare events of a catastrophic type (high severity with low probability). A natural catastrophe is usually defined as an event yielding a loss of at least five million dollars and affecting a significant number of insured people. Cat derivatives, and in particular cat-bonds, are closely related to insurance contracts and, as such, suffer from the same drawbacks like exposition to moral hazard and adverse selection. In some cases they are considered as true insurance products by tax regulators.

1.2 Weather derivatives

Weather derivatives are financial instruments whose value and/or cashflows depend on the occurence of some meteorological events, which are easily measurable, independently authentifiable, and sufficiently transparent to act as triggering underlyings for financial contracts. Typically location is clearly identified and measurement is provided by independent and reliable sources. The underlying meteorological events can be considered as noncatastrophic. Use of weather derivatives generalizes standard financial risk management practice based on hedging foreign exchange and interest rate risks via financial derivatives.

According to the Weather Risk Management Association (WRMA), the financial market related to weather has two main facets: the management of the financial consequences of adverse weather for those with natural exposure to weather, and the commercial trading of weather risk, both in its own right and in conjunction with a variety of commodities.

The first weather transactions took place in 1997 between Enron and Koch Industries. These transactions were the result of a long thinking process by Koch, Willis and Enron aimed at finding a means of transferring the risk of adverse weather. These deals followed the deregulation of the energy market in the US (transition from an oligopolistic position to a status of mere participant to a competitive market) and was based upon some temperature indices to compensate the energy producer in case of a mild winter. The first European deals took place in 1998 between Enron and Scottisch Hydro Electric on a similar basis.

WRMA conducts every year a survey of the weather market. The results, even if they reflect only a part of the whole market, are interesting and can give an idea of the market trend. The last available survey (2006) shows that even if weather risk can affect the whole economy, the energy sector is still the main part of the market, with companies such as British Gas, Hess Energy or Aquila Energy. Financial institutions and hedge funds, such as ABN Amro or Merrill Lynch, insurance and reinsurance companies, such as AXA Re or Swiss Re are also among the large players of this market, mainly because of their important activity of the energy market (gas and oil) and of the cross-hedging opportunities weather derivatives offer as we will see later. Beside large players, there also exist lots of small couterparties such as municipalities, ski resorts, golf clubs, beverage producers and leasure parks. Success in those markets often requires writing a large number of contracts (economies of scale) and exploiting a large set of sharp skills at the trader and analyst levels. The size of the market was estimated by WRMA to be over \$45 billions in 2006, compared with \$9.7 billions in 2005. About 75% of the transactions are based upon temperatures, and 10% on rainfalls.

The most common underlying is related to the notion of Degree Day which is expressed as the difference between a reference level temperature ($65°F$ or $18°C$) and the average daily temperature T. The average is computed between the maximum and minimum recorded temperature over a particular day. A Heating Degree Day (HDD) is defined as follows:

$$HDD = (65°F - T)^+ ,$$

where $(\cdot)^+ = \max(\cdot, 0)$. The bigger the HDD is, the colder the temperature is, and as a consequence the larger the demand for heating will be. For example if the today average is $60°F$, we needs 5 degrees of heating ($HDD = 5$). On the contary if the today average is $70°F$, we do not need any heating, and the HDD is zero. Similarly a Cooling Degree Day (CDD) is defined as follows:

$$HDD = (T - 65°F)^+ .$$

The bigger the CDD is, the hotter the temperature is, and the larger the demand for cooling should be. The definition of a temperature index in those terms reflects the close relationship between the energy sector and the weather derivative market. Daily results are often cumulated to give a total on a given period, such as a week (Xmas-NewYearEve, sport event), a month (sales period, harvest period), or a quarter (summer holidays, opening season of a resort).

Weather derivatives differ in their characteristics, with different contract types such as Swap, Call and Put. The contract period can be the HDD season (November to March) or the CDD season (May to September).

The measurement authority, the reference index and payment conditions can vary as well. Contract schedule is generally made of three dates: transaction date (from two days to several months before the observation period), payment date (five days after the end of the period to get the data used in computing the payoff), adjustment date (three months after the payment date when final data are published and possible recording errors removed).

HDD swaps are contracts where both counterparties agree to exchange their climatic risk. It allows to smooth volatile cash flows. One counterparty accepts to pay the other when the index goes above a given level, and vice-versa if it goes below. The level is most of the time fixed so that expected positive and negative cashflows compensate exactly, and the current value of the swap is nil. For example, a beverage producer can enter a swap to get protection when summer time is too cold, and forego part of its benefits when it is too hot.

HDD calls are contracts where the buyer gets compensated when the climatic risk is above a predetermined level. The protection implies an upfront payment of a premium to the seller. For example an airline company can buy a call on the number of days where the average windspeed exceeds a given level. This protects against loss of revenues in case of delays or cancelations for a cost equal to the paid premium.

HDD puts are contracts where the buyer gets compensated when the climatic risk is below a predetermined level. Again we have an upfront payment of a premium to the seller. For example a ski resort can buy a put on the number of days where the average snowfall is lower than a given level. Loss of revenues because of a reduced number of skiers is diminished for a cost equal to the paid premium. Most of the time call and puts contracts are assorted with caps which bound the compensation payments but reduce the upfront payment.

2 The weather derivatives market

Most of the transactions are tailor-made and part of the OTC (Over The Counter) market. Usually OTC transactions are realized withing the ISDA standards (Master Agreement standards of the International Swap and Derivatives Association) which provide standardised contracts aimed at easing OTC transaction processes. Some go through specialized brokers but most of them are done without any intermediaries. Taylor-making is not surprising as these structures suit better the management of weather risk and the needs of the various players on this market. The organized markets are however rather successful, mainly because of the transparency, liquidity and security they offer. Among them, the most predominant one is the Chicago Mercantile Exchange (CME). The CME reported 1,041,439 trades in summer 2005 and winter 2005/6. Combined 2005/6 CME trades were up over 300% compared to $223,139$ trades in 2004/5 period.

2.1 Organized market: the Chicago Mercantile Exchange (CME)

CME weather derivatives are part of the CME Alternative Investment Products, with derivatives based on economic indices, housing price index or ethanol. Two major types of contracts are traded on the CME: futures and options on futures, both based upon a temperature-based index computed by Earthsat using weather data from the National Climate Data Center.

The first contracts appeared in 1999, first focusing on American cities. Now the CME offers contracts for 18 American cities (Las Vegas, Atlanta, Chicago, Cincinnati, New York, Dallas, Philadelphia, Portland, Tucson, Des Moines, Boston, Houston, Kansas City, Minneapolis, Sacramento, Salt Lake City, Detroit, Baltimore), together with 6 Canadian cities (Calgary, Edmonton, Montreal, Toronto, Vancouver, Winnipeg), 9 European cities (London, Paris, Amsterdam, Berlin, Essen, Stockholm, Rome, Madrid, Barcelona) and 2 Japanese cities (Tokyo, Osaka). Recently, the CME has experienced a fantastic growth, as shown in the figure below:

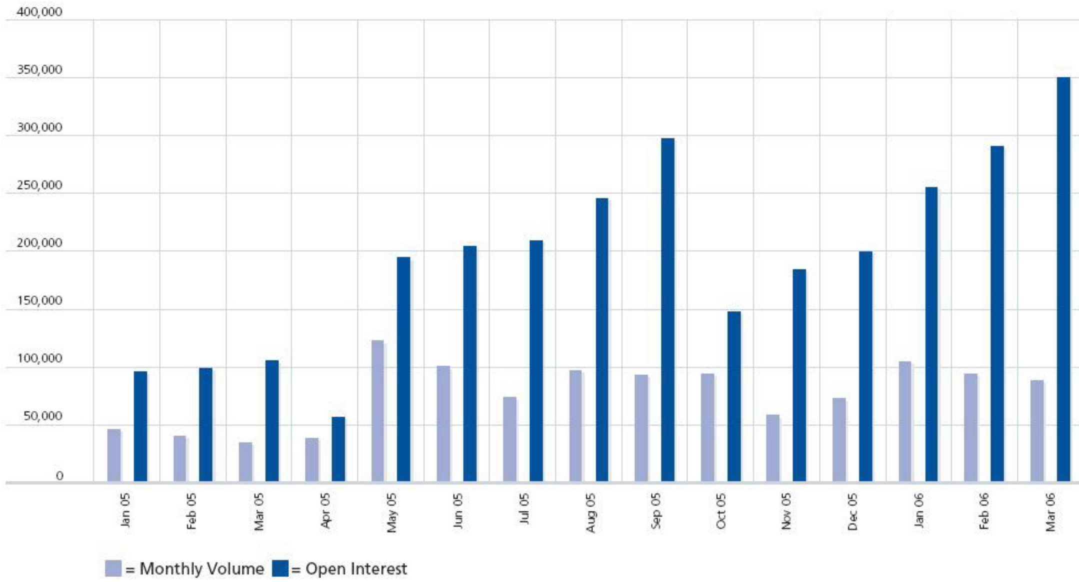

CME - Weather volume and open interest (From the CME Brochure "CME Weather Futures and Options")

As shown in the following figure,

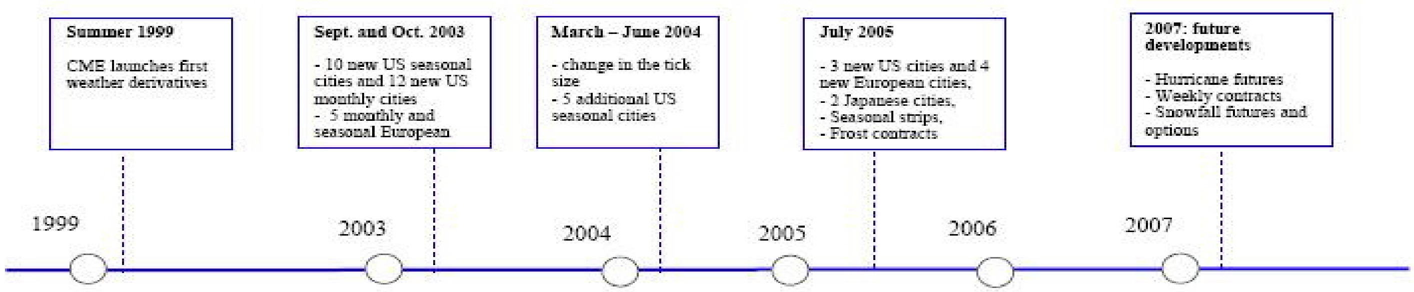

CME - Key historical facts

the history of the CME weather derivatives can be divided into two main parts: before 2003 and after. In 2003, everybody in the market thought the CME should close its weather activities because of an obvious lack of interest. But, surprisingly, the CME launched new contracts and modified the structure of the existing contracts, making the whole market more attractive.

The value of a CME future contract is simply expressed as

$$\text{Value of the index } \times \$20$$

The tick size, or monetary value of one index point, was first fixed at \$100 but then lowered down in 2004, as part of the desire to increase the overall flexibility of the market and its attractiveness. For Europe, the tick size is £20 and for Japan, ¥2500.

Today, there are three main types of contracts traded on the CME, depending on the considered period:
- the monthly contracts;
- the seasonal contracts, based upon a whole season, winter covering months between November and April, and summer covering the period between May and October;
- the seasonal strip contracts, where the market participants can choose a period from 2 to 6 consecutive months in a given season. This additional flexibility was introduced in September 2005 to focus on the real risk exposure, since the volatility can vary from one month to another.

There are three main types of indices, depending on the considered period:

- For the winter months, the monthly and seasonal contracts are based upon Cumulative Heating Degree Days, computed as the aggregated value of Daily Heating Degree Days over the considered period, $(65°F - T)^+$ in the US and $(18°C - T)^+$ elsewhere, where T is the average daily temperature defined as the average between the minimum observed temperature and the maximum observed temperature during the day.

- For the summer months, the monthly and seasonal contracts are based upon Cumulative Cooling Degree Days, computed as the aggregated value of Daily Cooling Degree Days over the considered period, $(T - 65°F)^+$ in the US. Elsewhere, contracts are based upon Cumulative Average Temperature.

- For Amsterdam, some specific contracts are available in winter (between November and March). They depend on the number of frost days, i.e., weekdays when frost is recorded.

The future developments of the CME are numerous. Since last March 2007, the CME has added hurricane futures and options on five U.S. defined areas - Gulf Coast, Florida, the Southern Atlantic Coast, the Northern Atlantic Coast, and the Eastern U.S, for the hurricane season beginning June 1. The underlying indexes are calculated by Carvill, a leading independent reinsurance intermediary in specialty reinsurance that tracks and calculates hurricane activity. These contracts provide an additional way to help addressing the needs of the insurance industry and other markets. In particular, insurers and others will be able to transfer their risk to the capital markets and thereby increase their capital in order to insure customers.

The CME is planning to offer weekly contracts, for the American cities only. These products enable market participants to hedge short-term intra-monthly weather risk. They will list two weekly contracts (Monday through Friday) and will be settled the following business day. The weekly contracts are different from the monthly and seasonal strip contracts essentially in the way that the weekly contract is an index of the week average temperature (not degree days).

Snowfall futures and options will be soon added to the existing products mainly based upon temperature. CME Snowfall futures will be geared to a CME Snowfall Index and will be offered initially on two U.S. cities - Boston and New York. These contracts will trade on a monthly basis from October through April.

2.2 European situation

Some attempts were made in Europe to launch an independent organized market for weather derivatives. In particular, the LIFFE launched three temperature-based European weather indices (Berlin, Paris, London) in July 2001, but this activity stopped later after the acquisition by Euronext.

In November 2005, Powernext, which is a European energy exchange based in France, and Meteo France launched the quotation of national temperature indices for 9 European countries (France, UK, Italy, Belgium, The Netherlands, Portugal, Spain and Switzerland). Three types of indices are provided: historical indices, observed indices and forecast (for Day D, $D+1$, $D+2$ and $D+3$) for average, minimum and maximum temperatures. The specificity of their indices relies upon the formula used to obtain a national aggregate. Various temperatures are weighted by population as to provide an (accurate) measure of risks related to the economic activity. Therefore:

$$\text{National index} = \frac{\sum n_i T_i}{\sum n_i},$$

where n_i is the population of region i and T_i is the relevant temperature in region i. The computed index can be seen as a gross estimate of the national risk. The objective is global and probably focused on energy producers. No contract is offered yet but these indices appear at the moment as decision tools and potential underlying for OTC weather transactions.

This initiative is further developed with the launch in June 2007 of Metnext, specialized in indices for weather risk management. Météo-France and Euronext, a subsidiary of NYSE Euronext, have teamed up to launch a joint venture named Metnext specializing in innovative solutions for index-based management of weather risk. Metnext, in which Météo-France holds a 65% interest and Euronext the remaining 35%, will initially offer two types of service for businesses, insurance companies, reinsurance companies, and banks. Indices tailored to the specific needs of individual firms in industry or other areas of business will enable them to anticipate, and thus manage, their exposure to weather risk more effectively. They will be based on historic data concerning the impact of variables such as temperature, rainfall, and the direction and wind speed on management parameters. Historical and forecasted indices will be available to cover forecast periods ranging from a few days to several weeks or several months, depending on the specific needs of the firm.

2.3 Electronic trading platforms

Since the beginning of the weather market in the late 90s, electronic trading platforms have always played an important role in the development of the market, especially Enron's platform in the early days. A lot of big market players have such a platform, and in particular Spectron is rather important in Europe.

These platforms typically act as service providers for the energy trading industry. They act as a neutral third party to trades and have no ownership function. Their role can be seen as a compromise between pure OTC transactions and organized markets: by putting together various participants on the weather risk marker, these platforms increase their opportunity to sign a deal that is still rather standardized but closer to their real exposure. Moreover, new structures of deals are typically agreed upon in these platforms: the first weekly swap or the first precipitation deal were signed on Spectron.

3 The World Bank program based upon weather derivatives

Another dimension of weather derivatives lies in the specificities of this widespread risk, with among the most exposed those relying on agriculture. Therefore, it seems natural to see the World Bank as one of the actors on this market. Indeed, according to the 2005 World Bank Report [9], "the World Bank can play a role in assisting countries in taking actions that effectively use limited government resources to facilitate market-based agricultural risk transfer".

3.1 General presentation

The role of the World Bank is crucial in developing countries. However, it is currently exposed to natural disasters without the benefit of ex-ante structures to finance losses. The main challenge seems to be the innovation in risk transfers, as to lead ways to solutions to many economic and social problems.

Among the various structures of the World Bank, the Commodity Risk Management Group (CRMG) has for main objective to deal with the agricultural risk in developing countries, where agricultural risk is defined as "negative outcomes stemming from imperfectly predictable biological, climatic and price variables". The economic and social consequences of this risk being so huge in some countries, it seems logical to seek for some form of protection against this risk. However, even if insurance companies were to write insurance policies against this risk, they would be typically too expensive for small farmers, and the compensation would take too long to be effective, partly because of the claim checking procedure. As far as the local government is concerned, its role is more to protect against catastrophic risk (low frequency - high severity).

The CMRG has developed several projects throughout the globe in order to deal with agricultural risk transfer using weather derivatives. Among these projects, the pilot program in India has been particularly successful.

3.2 From theory to practice: a pilot project for agricultural risk transfer in India

The Indian exposure towards agricultural risk was assessed in a national survey carried out in 1991. One sixth rural households had loans from rural financial institutions as to finance their agricultural activities, but only 35% of the actual credit needs were met through these channels. The World Bank estimated then that moneylenders, the reduction in farming inputs, overcapitalization and overdiversification of the activities have all led to a suboptimal asset allocation for Indian farmers. The challenge was therefore to innovate and find low-cost ways to reach farmers and help them to better manage their risk.

The first project of a total duration of 3 years was initiated in 2003 by the World Bank, based upon the idea of a large scaling-up and mainstreaming of weather insurance for small farmers. This project was run jointly by the CRMG, BASIX, an Indian microfinance institution created in 1996, and ICICI Lombard, a Mumbai based insurance company. Three different phases have to be distinguished in the launching and running process of the first weather insurance initiative in India.

The first year, 2003, was a pilot year, and the program was only targeting groundnut and castor farmers in the Andhra-Pradesh district of Mahbubnagar as indicated in the map below

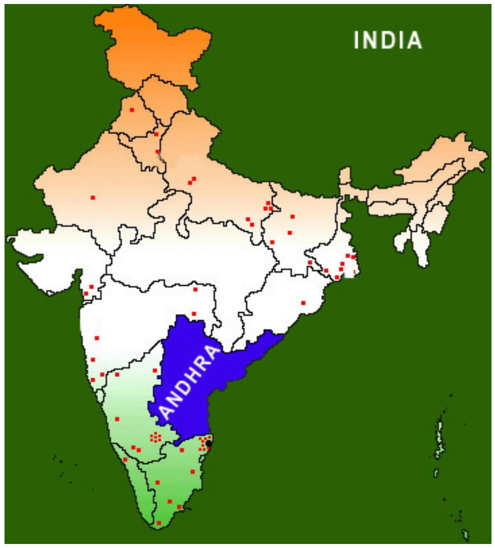

The design of the insurance scheme was made by ICICI Lombard, with the technical support of the CRMG and the consultation of BASIX, the objective being to protect the farmers from drought during the groundnut and castor growing season, which corresponds to Khariff (the monsoon season, between June and September). The entire portfolio sold by BASIX was insured by ICICI Lombard with a reinsurance contract from Swiss Re. The marketing and sales were made in 4 villages selected by a local bank, according to their current involvement in microfinance. Workshops and presentation meetings were arranged by BASIX. 230 farmers (154 producers of groundnut and 76 of castor) bought insurance against drought during Khariff, and most of them were considered as small farmers, with less than 2.5 acres of land. The insurance contract bought by the farmers had only the name of insurance. It was indeed a financial derivative based upon a reference weather index. The index based upon rainfall has been carefully established by farmers and biologists as to represent the real impact of rain on the growth of either the groundnut or the castor (two indices were considered depending on the crop). The different stages of the growth were also taken into account by weighting differently the different subperiods of the considered season. If the rainfall was not sufficient, then the contracts would be automatically triggered, and the farmers would automatically received their compensation. Because of this triggering mechansim, the costs of the overall scheme are lower than those of a traditional insurance programme, and the compensation recieved promptly by the farmers, helping them to survive until the next season.

The payout structure of the groundnut weather insurance policy was as follows:

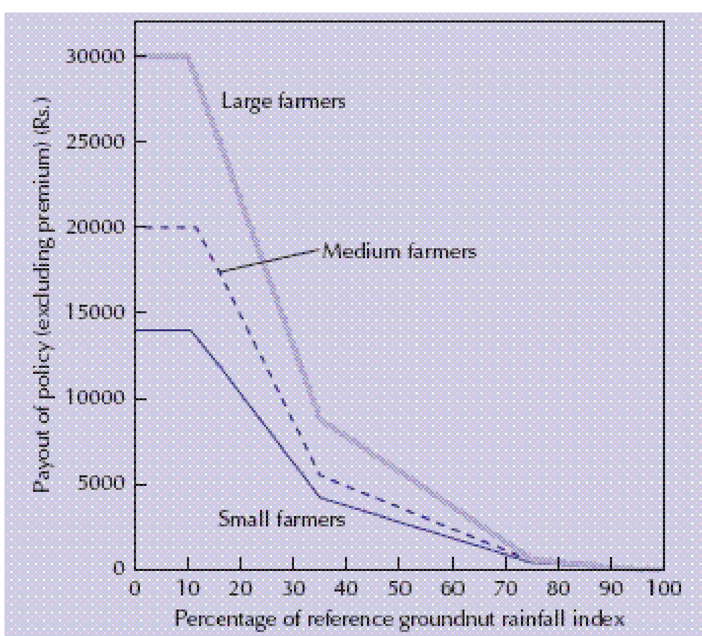

Different payoffs were considered depending on the size of the land, as to take into account scale effect.

The year 2004 witnessed an extension of the first experience to 4 new weather stations in Khamman and Antapur, in Andhra-Pradesh. Existing contracts were sligthly modified, following the feedback of farmers who bought them in 2003. In particular, more weight was put on the initial sowing period of groundnut in the computation of the rainfall index. New contracts were created for cotton farmers, targeting excess rainfall this time. 400 farmers bought these contracts from BASIX and 320 bought them directly from the insurance company ICICI Lombard, which did not buy reinsurance this time.

Further improvements were made in the third phase of the project in 2005: a total of 7685 contracts in 36 local in India were sold by BASIX, through the microfinance channel, but more generally more than 250000 farmers bought weather insurance in the country, directly from ICICI Lombard.

The overall project has been a major sucess, improving the financing condition of many small farmers in India. The scheme still exists. A long term question is certainly the study of the impact of these policies on the farmers behaviour.

3.3 Current and future projects

The CRMG has many other projects throughout the globe as shown in the following figure:

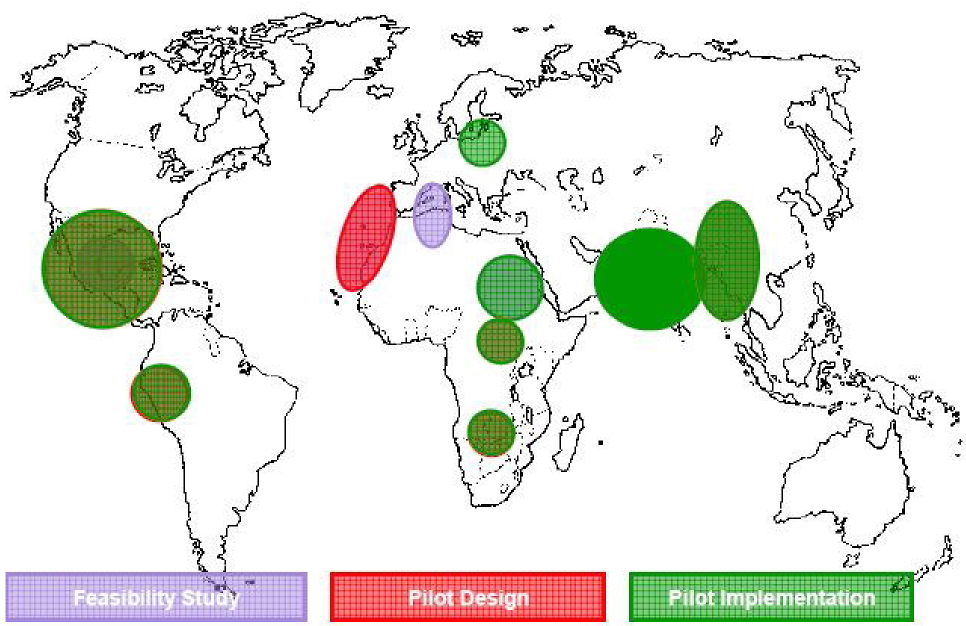

All these projects are based upon the same idea: the underwriting of insurance contracts dependent upon an index, i.e., weather derivatives, as to improve the overall efficiency of the risk management: cheaper product, less delay in receiving the compensation since the contracts are standardized and no claim check is required. The characterization of these indices is therefore essential for the success of these products. More precisely, these indices should be related to the agricultural risk, as this was the case in India. They are typically determined with the help of the local farmers and biologists to characterize the critical events. They involve typically:

- Rainfall measures, as to identify drought or excess rainfalls which can damage the culture, as in India for groundnut and castor;

- Some events: for instance, the so called "winterkill" index in Ukraine, counting the number of days when the temparture falls below $-16°C$. Temperatures below this level are killing the seeds in the soil and therefore can be seen as responsible for a poor harvest the consecutive season.

- Growing Degree Day, as to measure the growth and development of plants and insects during the "growing season":

$$\text{Daily Growing Degree Day} = (T - L)^+$$

The threshold L can vary during the life of the contract. The accumulated value is a good proxy to establish the development stages of a given grop or of an insect. As a consequence, a proper scheduling of pesticide or herbicide can be made.

4 Specific issues related to weather derivatives

Due to the specificities of the underlying risk, weather derivatives have also some specific issues. Some of them are common to all securities written on a non-financial underlying risk, as this is the case for insurance risk securities based on a catastrophic risk or on a mortality index.

4.1 Basis risk

When securitizing insurance risks, some basis risk is usually introduced as the new products usually depend on an index rather than on a single firm exposure. Basis risk simply translates the risk related to the spread that can exist between the agent exposure and the hedge proposed. Indeed, before a non-financial risk can be securitized, i.e., transfered to the capital markets, it must be in general standardized. This risk can theoretically create some problems and prevent the derivative market from developing.

As a consequence, the characterization and the computation of an appropriate index is an essential step. The reference index should be clearly identified, non manipulable, consensual, and widely available, with a long history. Usually this is not an issue with weather data that have been collected for many years by public and independent institutions. Still, the construction of an index from existing data requires transparency: in particular, the temperature used for the CME contracts is not a real temperature but a daily average of two extreme (real) temperatures. The computation rule should be explicit and made by an independent authority, here Earthsat. For the Powernext index, the computation rule is more complex but available for the subscribers; a database has been reconstructed for the index so that the different market participants should be able to relate their own risk exposure to the index. In general basis risk is more pronounced in the US than in Europe where distances are smaller, but local climate heterogeneity plays a key role in both geographic areas.
Note also that basis risk might not be such an issue for large energy producers, which are globally exposed towards weather conditions, with some concentration in large cities, densely populated. In their case, hedging with large cities contracts is a good approximation. The initiative of Metnext seems to fill in this gap in the European market.

The basis risk, even if clearly identified, can be a huge limit to the development of the market by small companies. How should a company based in Glasgow hedge its weather risk? Is the London contract representative of its risk? How is the risk "difference between London and Glasgow" understood? All these common sense remarks could explain partially the increased development of the CME market after the introduction of additional locations in 2003 but also its (still) relatively limited expansion. OTC transactions seem to be better suited for the management of this highly local risk.

Management of book of OTC weather derivatives exploits partially this imperfect link between measurement locations to achieve diversification in their positions. This geographical diversification supplements directional diversification to reduce risks. This requires a monitoring of the correlation between locations. Besides there often exist limits in terms of number of transactions and amount per transaction as well as types of trading and hedging strategies in order to keep everything under control.

4.2 Pricing issues

The literature on the pricing of weather derivatives is really important, especially on temperature derivatives. Two aspects are tackled: the understanding and modeling of the underlying distribution and the characterization of the pricing rule. The choice of the model and of the pricing rule is heavily dependent on the accounting classification of the weather derivative, either as an insurance product or as a financial derivative.

4.2.1 Underlying risk modelling

The statistical analysis of meteorological data is complex because of missing data, seasonality, nonstationarity (long term trend with wild short term blips), multidimension (different locations of the measurement units), and multiple interactions. Long term weather forecasts are notably imprecise and not always suitable for valuation and risk management.

US data are easy to get freely whereas European data are relatively expensive. US historical data are available through the National Climate Data Center (NCDC) which is part of the National Oceanic and Atmospheric Administration (NOAA). This constraints the choice of the measurement authorities. The optimal length of the database depends on the regularity of the data (trend, regular seasonality). Between 10 and 30 years is considered as the norm. The problem of the choice of the optimal sample size is well known to climatologits. It concerns the determination of the so-called Optimal Climate Normal (OCN) or optimal mean window of former years used to predict expected value for next year. The National Center for Environmental Prediction (NCEP) uses a methodology which measures deviations between the chosen OCN and a forecast made from a fixed period of 10 years. Most of the time specialised data vendors provides further servicing and control such as correcting human recording errors, harmonizing data in case of a change in the measurement device, desaisonalizing, and trend modelling (global warming, surface warming of the Pacific ocean known as El Nino phenomenon).

Statistical modelling of index dynamics, often after appropriate prewhitening (desaisonalising and detrending) can be broadly divided into parametric approaches and nonparametric approaches. Parametric approaches rely on standard time series tools such as Autoregressive-Moving Average (ARMA) models, Generalized Autoregressive Heteroscedastic (GARCH) models, and Markov switching models. These models allows introducing mean reversion phenomenon in the conditional mean or variance of the temperature data (Cao and Wei [8], Campbell and Diebold [7]) Parametric approaches may also rely on continuous time models (instead of a discrete time models) such as jump-diffusion models. The workhorse in that type of modeling is the Ornstein-Uhlenbeck process (Dischel [12], [13], Dornier and Queruel [14], Benth and Benth [4]) featuring mean reversion and stationary behaviour. Extensions based on fractional Brownian motion have also been considered to include long memory effects (Brody, Syroka and Zervos [4]). Nonparametric or semiparametric approaches rely on resampling methods by blocks of data (block bootstrap, subsampling) or by blocks of residuals (Roustant, Laurent, Bay and Carraro [20]). Other approaches try to incorporate weather forecasts to improve on the modelling. Jewson and Caballero [18] use medium-range forecasts (up to 11 days), while Yoo [21] uses seasonal forecasts based on ranking temperature being above, near or below normal.

Multivariate factor analysis also helps to identify and quantify factors relevant to explain differences between North versus South location measures as well as East versus West location measures. They can often be related to the presence of lakes, oceans (Pacific versus Atlantic), deserts or dominant wind directions (North).

4.2.2 Characterization of a pricing rule

Given the uncertainty and the flexibility in their accounting classification, but also their relative illiquidity, several pricing methods have been suggested for weather derivatives. They can be classified into three main categories: actuarial, financial and economic. In the following, we will briefly present these various approaches, focusing on the (forward) pricing rule of a weather derivative with a payoff F at a future time. Considering forward price allows us to simplify the problem in terms of interest rates and to focus on the pricing rule itself.

Actuarial method The first method based upon actuarial arguments uses the fair value, corrected by some margin as pricing rule. More precisely, denoting by $\mathbb{P}$ the statistical probability measure used as prior probability measure, the (forward) price of the derivative can be obtained as

$$\pi(F) = \mathbb{E}_{\mathbb{P}}(F) + \lambda \sigma_{\mathbb{P}}(F)$$

Different authors have studied the impact of the choice of the probability measure on the pricing. For instance, Jewson and Brix [17] compared two methods to compute $\mathbb{E}_{\mathbb{P}}(F)$: burn analysis, where the expected value is obtained as the average over historical data and index modeling, where a theoretical model for the relevant index is calibrated over past data, before simulations are conducted to find the expected value. For the price, both methods tend to be equivalent; for the sensitivities, the index modeling is better. Roustant, Laurent, Bay and Carraro [20] looked at the impact of a calibration error with index modeling: $\mathbb{P}$ is in fact $\mathbb{P}(\theta)$, with θ the model parameters. The impact is measured as $\pi_{\mathbb{P}(\widehat{\theta})}(F) - \pi_{\mathbb{P}(\theta)}(F)$.

Financial method The second method is more financial. Some papers assume the weather derivative market to be complete and therefore use the risk-neutral pricing rule:

$$\pi(F) = \mathbb{E}_{\mathbb{Q}}(F),$$

where $\mathbb{Q}$ is the unique risk-neutral probability measure. The underlying assumption is really strong and cannot be justified. A milder argument consists in assuming absence of arbitrage opportunities only. From the fundamental theorem of asset pricing, this is equivalent to the existence of at least one equivalent martingale measure

$\mathbb{P}^*$. Attainable contingent claims can be priced using one of these probability measures. However, writing

$$\pi(F) = \mathbb{E}_{\mathbb{P}^*}(F)$$

is equivalent to consider weather derivatives as attainable. Obviously, this is a strong underlying assumption, and the weather market is still very incomplete by essence.

In an incomplete market framework, there exist however many different methods to price a contingent claim, without creating any arbitrage opportunity. A rather standard approach involves utility maximization. To do so, we consider the preferences of each individual agent. This implies of course to make a distinction between *individual prices* and *equilibrium prices* that can be observed on the whole market.

Any individual wants to maximize the expected utility of her terminal wealth in this framework. The maximum price she is ready to pay for the weather derivative is therefore the price such that she is indifferent, from her utility point of view, between buying it or not buying it. For this reason, the price obtained by utility maximization techniques is called *indifference price* (many references exist on this subject, the seminal paper being that of Hodges and Neuberger [15]). Denoting by u the utility function of the agent we consider, and assuming that there is no interest rates (for the sake of simplicity), the indifference buyer price of F, $\pi^b(F)$, is determined as:

$$\mathbb{E}_{\mathbb{P}}\left(u\left(W_0 + F - \pi^b(F)\right)\right) = \mathbb{E}_{\mathbb{P}}\left(u\left(W_0\right)\right),$$

where W_0 is the initial wealth (which may be random). This price $\pi^b(F)$, which theoretically depends on the initial wealth and on the utility function, is not (necessarily) the price at which the transaction will take place. This gives an upper bound to the price the agent is ready to pay for the contract F. The agent will accept to buy the contract at any price below $\pi^b(F)$.

In the characterization of the indifference price, there is no question on the volume of the transaction. The potential buyer has two options: either buying 1 contract or not buying it. There is another possible approach, which consists of determining the price of the contract such that agreeing a little into the contract has a neutral effect on the expected utility of the agent. This notion of *fair price* was first introduced by Davis in [10] and [11] and corresponds to the zero marginal rate of substitution price. More precisely, the fair price p is determined such that:

$$\frac{\partial \mathbb{E}_{\mathbb{P}}\left(u\left(W_0 + \theta F - p\right)\right)}{\partial \theta}|_{\theta=0} = 0.$$

Economic approach The *transaction price* is an equilibrium price, either between the seller and the buyer only, or between the different players in the market. Note that a transaction will take place only if the indifference buyer price is higher than the indifference seller's price, which gives a lower bound to the price the seller is ready to accept for the contract. This is a necessary condition for a transaction, and more generally for a market equilibrium, which characterizes the situation where all agents in the market maximize their expected utility at the same time by exchanging their risks (such an equilibrium is also called Pareto-optimal).

Such an equilibrium approach has been adopted by different authors, such as Campbell and Diebold [7] or Horst and Mueller [16]. While the first authors look at how weather forecasts can influence the demand for weather derivatives, and hence their price, the latter consider the problem of pricing in an incomplete market, with a finite number of agents willing to exchange their weather risk exposure. The price of the contract is obtained as that of the Pareto-optimal equilibrium.

4.3 Design issues

As shown by various failed attempts of weather risk securitization, in particular the failed issued of a weather bond by Enron in 1999 (see for instance Barrieu and El Karoui [3] for a detailed study), the design of the new securities appears as an extremely important feature in the transaction. It may be the difference between success and failure. More precisely, as previously mentioned, the high level of illiquidity, deriving partly from the fact that the underlying asset is not traded on financial markets, makes these new products difficult to evaluate and to use. The characterization of their price is very interesting as it questions the logic of these contracts itself. Moreover, the determination of the contract structure is a problem in itself: on the one hand, the underlying market related to these risks is extremely illiquid, but on the other hand, the logic of these products itself is closer to that of an insurance policy. More precisely, even though weather derivatives have all the features of standard financial contracts, they are very different from the classical structures, as their underlying risk is related to a non-financial risk. Consequently the question of the product design, unusual in finance, is raised. Different papers have been looking at this question (see for instance Barrieu and El Karoui [3]). We briefly present here a simple approach to optimally design the cash flow structure of a weather derivative.

4.3.1 Simple framework

Let us consider two economic agents, henceforth called A and B, evolving in an uncertain universe modelled by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. At a fixed future date T, agent A is exposed to a weather risk for a random amount X. In order to reduce her exposure, A wants to issue a financial product, with a payoff F, and sell it to agent B against a forward price π at time T. Both agents are supposed to be risk-averse. For the purpose of this simple study, we assume that they are working with the same kind of choice criterion, an increasing exponential utility function u defined as:

$$u(x) = -\gamma \exp\left(-\frac{1}{\gamma}x\right)$$

with risk tolerance coefficient γ_A and γ_B, respectively.
Agent A objective is to choose the optimal structure for the weather derivative (F, π) as to maximize the expected utility of her final wealth, i.e., seeking:

$$\arg\max \mathbb{E}_{\mathbb{P}}\left[u_A\left(X - (F - \pi)\right)\right].$$

Her constraint is that agent B should have an interest to enter into this transaction. At least, the F-structure should not worsen agent B expected utility. To decide whether or not she should enter this weather transaction, agent B compares two expected utility levels, the first one corresponding to the case where she simply invests her initial wealth in a bank account and the second one to the situation where she enters the F-transaction. Thus, agent A is working under the constraint:

$$\mathbb{E}_{\mathbb{P}}\left[u_B\left((F - \pi) + W_0\right)\right] \geq \mathbb{E}_{\mathbb{P}}\left[u_B\left(W_0\right)\right],$$

where W_0 is the (forward) wealth of agent B before the F-transaction.

With the exponential utility functions, the problem to solve is:

$$\begin{aligned} &\min_{F,\pi} \gamma_A \mathbb{E}_{\mathbb{P}}\left[\exp\left(-\frac{1}{\gamma_A}\left(X - (F - \pi)\right)\right)\right] \\ \text{subject to} \quad &\mathbb{E}_{\mathbb{P}}\left[\exp\left(-\frac{1}{\gamma_B}\left((F - \pi) + W_0\right)\right)\right] \leq \mathbb{E}_{\mathbb{P}}\left[\exp\left(-\frac{1}{\gamma_B}W_0\right)\right]. \end{aligned} \tag{1}$$

Given the convexity of the program, the constraint is bounded at the optimum and the optimal pricing rule $\pi^*(F)$ of the financial product F is entirely determined by the buyer (agent B) as

$$\pi^*(F) = \gamma_B \ln \mathbb{E}_{\mathbb{P}}\left[\exp\left(-\frac{1}{\gamma_B}W_0\right)\right] - \gamma_B \ln \mathbb{E}_{\mathbb{P}}\left[\exp\left(-\frac{1}{\gamma_B}\left(F + W_0\right)\right)\right]$$

Agent B determines the minimal pricing rule, ensuring the existence of the transaction. The price $\pi^*(F)$ corresponds to the maximal amount agent B is ready to pay to enter the F-transaction and bear the associated risk, given her initial wealth W_0. In other words, $\pi^*(F)$ corresponds to the certainty equivalent of F for the utility function of agent B, or to the buyer indifference pricing rule.

4.3.2 Optimal structure

In the present simple framework, the optimal structure is given by the so-called Borch Theorem, presented below. In a quite general utility framework, Borch [5] obtained optimal exchange of risk, leading in many cases to familiar linear quota-sharing of total pooled losses.

Proposition 1 *The optimal structure of the weather derivative, given as solution of the optimization program* (1) *is given as a proportion of the initial exposure X, depending only on the risk tolerance coefficients of both agents:*

$$F^* = \frac{\gamma_B}{\gamma_A + \gamma_B} X \qquad \mathbb{P}\ a.s. \qquad (\textit{up to a constant}).$$

For the proof of this result, we refer the reader to Barrieu and El Karoui [2] for instance.
Note that agent A is transferring a part of her initial risk in this contract according to her relative tolerance. Moreover, if the issuer, agent A, has no exposure, no transaction will occur between both agents. In this sense, weather derivatives have a non-speculative underlying logic, and are there to help with the true risk transfer between the agents in the market.

4.3.3 Extensions

There are various ways in which this modelling approach can be extended. In particular, the buyer may also be exposed towards a weather risk. The potential diversification impact for both agents is taken into account since the optimal structure in this case becomes a transfer of the risk of each agent in proportion to their relative risk tolerance. Using obvious notation, the optimal structure of the contract becomes in this case:

$$F^* = \frac{\gamma_B}{\gamma_A + \gamma_B} X_A - \frac{\gamma_A}{\gamma_A + \gamma_B} X_B \qquad \mathbb{P}\ a.s.$$

Note that even if agent A, the seller, is not initially exposed towards a weather risk, there will still be a transaction between both agents due to the exposure of the buyer.

The diversification impact using the financial markets can also be easily taken into account in this framework. Assuming that both agents can choose optimally their financial investment in the market, as to maximize their expected utility, the optimization program (1) is transformed by referring to another probability measure $\widehat{\mathbb{Q}}$, instead of the prior probablity measure $\mathbb{P}$. $\widehat{\mathbb{Q}}$ is in fact the minimal entropy probability measure and enables to account for the diversification impact offered by optimal financial investment (for more details, please refer to Barrieu and El Karoui [2]).

Finally, more general decision criteria can be considered instead of the exponential utility. In particular, convex risk measures can be used as to assess the capital requirement needed by financial institutions of insurance companies when undertaking risk. When both agents have risk measures of the same family, then the proportional structure remains optimal as shown in Barrieu and El Karoui [3], otherwise some additional features maybe added to the structure (see for instance Jouini, Schachermayer and Touzi [19]).

References

[1] P. Barrieu and N. El Karoui, Reinsuring climatic risk using optimally designed weather bonds, Geneva Papers Risk and Insurance Theory. Vol. 27, 87-113 (2002).

[2] P. Barrieu and N. El Karoui, Inf-convolution of risk measures and optimal risk transfer, Finance and Stochastics 9, 269-298 (2005).

[3] P. Barrieu and N. El Karoui, Pricing, hedging and optimally designing derivatives via minimization of risk measures, in Volume on Indifference Pricing (ed: Rene Carmona), Princeton University Press (2007).

[4] F.E. Benth and J. Saltyte-Benth, The volatility of temperature and pricing of weather derivatives. To appear in Quantitative Finance (2007).

[5] K. Borch, Equilibrium in a reinsurance market, Econometrica 30 (1962), 424-444.

[6] D.C. Brody, J. Syroka and M. Zervos, Dynamical pricing of weather derivatives, Quantitative Finance, Vol.2, p.189-198 (2002).

[7] S. Campbell and F.X. Diebold, Weather forecasting for weather derivatives, Journal of the American Statistical Association, Vol. 100, No. 469 p.6 – 16 (2005).

[8] M. Cao and J. Wei, Weather derivatives valuation and market price of risk, The Journal of Futures Markets Vol. 24, No. 11, p. 1065-1089 (2004).

[9] Commodity Risk Management Group, Progress report on weather and price risk management work, Agriculture & Rural Development, The World Bank (http://www.itf-commrisk.org/) (2005).

[10] M. Davis, Option pricing in incomplete markets, in Mathematics of Derivative Securities, eds M.A.H. Dempster and S.R. Pliska, Cambridge University Press (1998).

[11] M. Davis, Pricing weather derivatives by marginal value, Quantitative Finance, Vol. 1, No 3, p. 305-308 (2001).

[12] R. Dischel, Black-Scholes won't do, Risk Magazine, E.P.R.M., October, p. 8-13 (1998).

[13] R. Dischel, Shaping history, Risk Magazine, E.P.R.M., October, p. 13-15 (1999).

[14] F. Dornie and M. Queruel, Caution to the wind, Risk Magazine, E.P.R.M., July, p. 30-32 (2000).

[15] S.D. Hodges and A. Neuberger, Optimal replication of contingent claims under transaction costs. Review of Futures Markets, vol. 8, 222-239 (1989).

[16] U. Horst and M. Mueller, On the spanning property of risk bonds priced by equilibrium. To appear in Mathematics of Operations Research (2007).

[17] S. Jewson and A. Brix, Weather derivative valuation: the meteorological, statistical, financial and mathematical foundations, Cambridge University Press (2005).

[18] S. Jewson and R. Caballero, The use of weather forecasts in the pricing of weather derivatives. Meteorol. App., 10, 367-376 (2003).

[19] E. Jouini, W. Schachermayer and N. Touzi, Optimal risk sharing for law invariant monetary utility functions. Mathematical Finance, to appear (2008)

[20] O. Roustant, J.-P. Laurent, X. Bay, L. Carraro, A bootstrap approach to the pricing of weather derivatives, Bulletin Français d'Actuariat (2005).

[21] S. Yoo, Weather derivatives and seasonal forecast, Working Paper, Cornell University (2003).