Institute of Actuaries of Australia

Weather Derivatives: Pricing and Risk Management Applications

Jon Tindall B.Sc(Hons.), B.E., Grad. Dip. AppFin.

Presented to the Institute of Actuaries of Australia
Financial Services Forum 2006
11 and 12 May, 2006

The Institute of Actuaries of Australia
Level 7 Challis House 4 Martin Place
Sydney NSW Australia 2000
Telephone: +61 2 9233 3466 Facsimile: +61 2 9233 3446
Email: actuaries@actuaries.asn.au Website: www.actuaries.asn.au

Abstract:

Key words: Derivatives, weather, HDD, CDD, risk, temperature, rainfall, Ornstein-Uhlenbeck, stochastic, mean-reversion.

The primary objective of this paper is in the discussion of suitable methods for the modelling of weather variables and the pricing of their respective derivative contracts. This work attempts to bring together much of the current thinking in terms of the pricing of weather derivative contracts. In addition to the theoretical overview provided, an empirical investigation is undertaken using historical data from the Sydney region. Both temperature and rainfall dynamics are investigated with some case studies undertaken to highlight the practical applications of these financial contracts.

“Draft Version Only”

10 May 2006

Contents

1: Introduction

1.1 Background

Nearly the entire industrialised world is in some way affected by variations in weather patterns, be they seemingly random or predictable. These fluctuations invariably have financial impacts on the producers and consumers of the products that are affected, thus leading to the desire to minimise as much of this risk as possible. Weather derivative contracts are a relatively new innovation in financial engineering that have been receiving significant attention in recent years as the industrial world continues to realise the magnitude of risk management applications that these contracts posses. Whilst both the market and literature has been dominated by temperature-based contracts from their beginnings in the 1990's, recently a more diverse range of weather variables have been utilised, such as rainfall, snowfall, wind speed and barometric pressure.

The weather derivatives market was largely born out of the deregulation and privatisation of many energy industries amongst developed nations during the late 1980's and early 90's, most notably in the US. These newborn utility companies were now being funded by private investors who were more scrupulous in regard to their investments than most governments have the ability to be. In particular these new investors required more certainty in their revenue stream and hence looked towards derivatives in order to facilitate this.

The first genuine transaction of a weather derivative contract occurred in 1996 when Entergy-Koch and Enron completed a HDD swap for the winter of 1997 in Milwaukee, WI. From these beginnings the market for these contracts has been growing steadily with the realisation that nearly all industries can benefit from these risk management tools. The US Department of Commerce estimates that 70% of all US companies and as much as 22% of its total GDP is significantly impacted by the weather conditions.

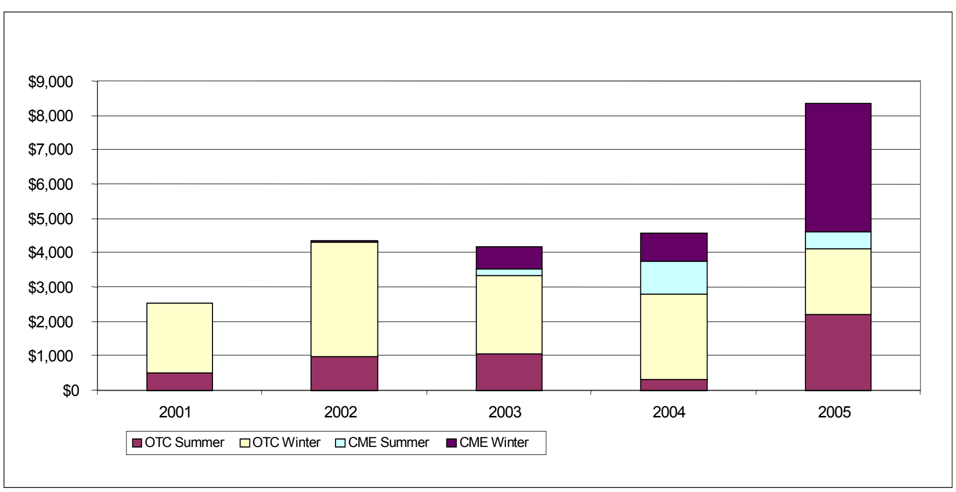

<u>Figure 1.1</u>: Total Notional Value [Source: PwC November 2005 Survey for WRMA]

Today the weather derivatives market size has reached around $US 8.4 billion (PwC[1] 2005) in total notional value. Initially weather derivatives were restricted to a mature OTC (Over-the-counter) market but in recent times significant attempts have been made to standardise the contracts and provide an electronic platform on which the instruments can be formally transacted. Trading on the Chicago Mercantile Exchange (CME) began on September 22, 1999, and now represents around 55% of the total global annual turnover in the weather derivative market (PwC 2005). The CME primarily offers temperature-based contracts over the major European and US cities.

The failure of US company Enron has been a much-publicised event in terms of the accounting irregularities that were involved in its demise, but, more significant were the implications that this event would have on the rapidly expanding but still immature weather derivative market. Some authors disagree on the long-term impacts of the failure as it is seen that many of the practices at Enron were unethical at best and that the collapse of the company was as a result of natural market pressures. Adding to this, the many employees of Enron (that aren't in prison) have now been redistributed amongst other smaller organisations and some argue that this newfound diversification has actually enhanced the growth of the market, providing a greater level of transparency[2]. Either way, it is evident from Figure 1.1 that the market size remained stagnant for a period after the collapse of Enron late in 2001, however, in 2005 the market has shown promising signs of returning to the growth rates that were being experienced during the late 1990's.

Enron began its life as a gas and oil pipeline manager and then soon became a trader of weather derivatives in order to manage the risk inherent in oil and gas supply contracts. Enron was to a large extent the creator of the weather derivatives market where they 'made-markets'[3] on a large number of contracts on the CME. Many of the current procedures for the modelling of weather risk can be attributed to the teams who used to work at this once leading energy company and the 'release' of this expertise to the general community has stimulated new research into these risk management tools.

1.2 Contract Types

Active markets exist in the trading of options, futures and swaps over a variety of underlying weather variables. Often these derivative contracts have special, tailor-made features that are introduced in order to properly match the hedging requirements of the client. Typically these specific features are designed to limit the payout in some way and thus make the product more affordable to the consumer. They include cap's and option barriers such as 'up-and-in' barriers that can be applied to the contract in a variety of ways.

As has previously been alluded to, the majority of the weather derivatives market is based around two weather variables, namely: temperature and rainfall. We will describe the typical features of these contracts as well as an indication of their use in managing weather related risk.

[1] PwC compile a market survey each year for the Weather Risk Management Association (WRMA)
[2] For an account of the impact of the Enron failure on weather derivative markets see Nicholls [2002].
[3] Market 'makers' are traders that offer buy and sell quotes simultaneously over a particular security.

1.2.1 Temperature

Heating Degree Days (HDD) and Cooling Degree Days (CDD) contracts were responsible for over 85% of the weather derivative transactions made throughout 2005. The primary purpose of these contracts is to allow organisations to hedge against excessively high or low temperature distributions over a pre-specified period of time.

Firstly, we define the average daily temperature as the arithmetic average of the maximum and minimum temperatures recorded in a 24hr period. In other words:

$$T_i = \frac{T_{\max} + T_{\min}}{2} \tag{1.1}$$

Naturally this average temperature is more reliable to model than either the maximum or minimum temperature are by themselves however information is thus lost and this measure would not provide a satisfactory hedge for those seeking protection from extreme daily temperatures.. A HDD is simply the number of degrees the days average temperature was below some reference level, $\overline{\mathrm{T}}$, generally set at 18˚. Hence the HDD's for the month are given as:

$$HDD = \sum_{month} \max\{0, (\overline{T} - T_i)\} \tag{1.2}$$

The specific terminology derives from the fact that 'heating' is generally required for temperatures that are below the reference level (here set at 18˚) thus requiring the expenditure of energy. The converse is also true for CDD's with the consumer's general requirement for 'cooling' energy above the reference temperature. From here an option contract can then be set with a payoff and exercise price based on these measures. Hence a call option on the monthly HDD index would have a payoff of the form:

$$V_t = \max\{0, (HDD - K)\} \cdot tick \tag{1.3}$$

Where K is the exercise price of the option and the 'tick' provides the conversion from degrees to \$ (i.e. its units are \$ per degree) and generally ranges anywhere from \$1,000 to over \$1,000,000 per degree.

A CDD contract is the reverse of the HDD in terms of its payoff. A CDD is just the number of degrees a particular days average temperature was above some pre-determined reference level. As before the accumulated CDD's for a month are given as:

$$CDD = \sum_{month} \max(0, (T_i - \overline{T})) \tag{1.4}$$

Other variables over which contracts are common include the monthly or daily average temperature (Asian type option) as well as monthly and yearly cumulative temperatures. The maximum daily temperature provides a significant hedge to crop farmers against severe temperatures that can have a devastating impact on a harvest's quality and size. Take as an example the production of wine who's yield and quality are highly dependent on temperatures during the two months prior to harvesting. If excessively high temperatures are encountered it can affect the physical vine growth as well as the quality of the grape produced including its flavour and odour. Hence a viticulturist could easily utilise a weather derivative over the maximum daily temperature experienced at a weather station that was in reasonably close proximity to the vineyard. This contract would compensate the farmer for reductions in yield and quality that result from the variations in the temperature.

1.2.2 Rainfall

Rainfall contracts have received significantly less focus, in terms of both the literature and the amount of trading activity, when compared to temperature based derivatives. Primarily this was due to the fact that the weather derivatives market was largely created from the needs of energy and utility companies who's major exposure was to temperature related weather variations. Adding to this, rainfall has proven significantly more difficult to model accurately, particularly where geographically small areas are concerned. This is primarily due to the discrete nature of rainfall and as we shall see in Section 4, two geographically close recording stations can produce widely different precipitation readings, something that is not encountered with temperature based statistics.

These difficulties manifest as geographical 'basis' risk in rainfall derivative contracts in that the situation of the risk must be relatively close to the measuring station for an effective hedge to be possible. This makes it increasingly hard to find the two (minimum) willing counterparties that all derivative contracts require. Not only are you required to find someone who will benefit and someone who will lose-out if rain falls but you also have to create an underlying rainfall index (i.e. weather station) that will protect both users. This has proven difficult to achieve in practice. As can be seen in Figure 1.2, the value of rainfall-based contracts has been increasing rapidly in recent years primarily due to the activities of hydroelectricity providers and consumers with the ongoing privatisation of water supply arrangements in many developed economies.

1.2.3 Other Variables

Whilst temperature and rainfall contracts account for the majority of the traded derivatives, other important weather variables are also becoming realised for their hedging opportunities. Interest in these new products is purely demand driven and is generally initiated by a client demanding protection from unfavourable weather conditions. In fact, any variable that possesses a reliable measurement practice can potentially be used as a basis for a derivative contract to provide protection against unfavourable financial outcomes.

Wind speed contracts have been one of the major growth points for the weather derivatives market, accentuated by a significant global 'push' towards wind based power generation and other renewable sources of energy. Wind-generating power companies are at the mercy of the winds in terms of the amount of electricity that they can produce and hence require some type of protection against lower than expected wind speeds. The Merrill Lynch - Global Commodities Wind Power Indices (WPIs) are produced as a reference by which wind farmers and other producers can hedge their exposure to the variability of wind speeds.

The chart that follows shows how the 'mix' of transactions in the global weather derivative market has been changing in recent times. Note that the large volume changes during 2005, evident in Figure 1.1, means that although some categories appear to have dropped in overall proportion, in fact all segments have increased their nominal turnover in 2005 when compared with the previous year.

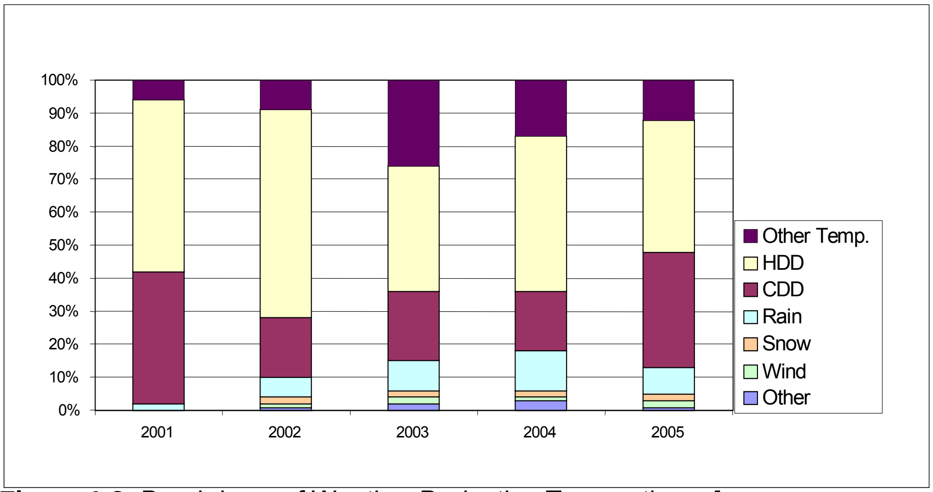

Figure 1.2: Breakdown of Weather Derivative Transactions. [Source: PwC November 2005 Survey for WRMA]

Others variable of interest include Sunlight hours, cloud cover, snow melt, humidity and atmospheric pressure. All of these variables can in some way impact upon the earnings of producers and hence the cost to consumers.

1.3 Weather Derivative Markets

There are two primary markets for exchange traded (standardised) weather derivative contracts that cover US, European and some Canadian cities. The two major exchanges that offer an automatic and standardised market for the trading of weather derivatives are:

CME: The Chicago Mercantile Exchange has the largest weather derivatives market in the world. The exchange offers both futures and option contracts over a range of US and European cities. In 2005 the CME traded 4.25 billion of notional value on its exchange which accounts for almost 56% of the total traded volume worldwide.

Liffe: London International Financial Futures and Options Exchange. In July of 2001, Liffe launched a series of contracts based on indices related to the daily average temperatures in London, Paris and Berlin. This organisation has recently been working with several technology companies to create an internet based trading platform for European weather derivatives.(Risk, March 2000). In 2004 Liffe suspended its trading of weather derivative contracts due to a lack of turnover and significant structural issues.

The acceptance of weather derivatives has been slower in Europe than in the US market. One of the main reasons for this is the lack of a standardised weather-recording framework that exists across the Atlantic. This is as a result of the many individual countries that make up Europe who have vastly different levels of development, which has meant that it has historically been very difficult to obtain consistent, reliable data. The recent expansion of the European Union (EU) and the creation of a single European currency should greatly facilitate the propagation of unified procedures for data capture and analysis.

More recently Asia, and in particular Japan, have become a source of weather derivative demand as their energy markets gradually became deregulated. The first official transaction in Japan was during 1999 between Mitsui Marine and a local sporting goods manufacturer that consisted of an option over the snow depth recorded for the following winter season. Natsource Japan, a large energy-based broker, is one of the major promoters of weather derivatives in Japan and has created a measure, the Japan Weather Derivatives Index (JWDI), around which market participants can design weather risk management products. As well as this the company has created an interbank electronic exchange, the Japanese Weather Exchange (JWX), on which the large financial organisations can transfer weather risk both within Japan as well as to the rest of the world.

The first weather derivative contract in Australia occurred in march of 1998 between United Energy Marketing and Utilicorp[4], a US based energy utility. The contract called for a payout if the temperature rose above 35°C in Melbourne or 33°C in Sydney for 5 days or more during the summer months. As it turns out Sydney reached this level on 5 days and Melbourne on 6 days and hence the contract exercised at roughly 8 times the initial one-off premium paid. Currently, in Australia, the formal market is practically non-existent. Whilst a partnership was formed in 2001 between Aquilla energy and Macquarie[5] Bank to provide trading and consulting services in the weather derivatives market, the initiative appears to have faded and there are few genuine transactions that have been recorded and made public in the past 2 years.

1.4 Market Participants

In the early years of the weather derivative market the major transactions were generally between large energy companies and large financial institutions. The overall complexities of the risk inherent in these products required a significant amount of research in order to properly price and was beyond the expertise of most of the participants. The financial impact of weather was so great on these energy and utility companies that is was economical to spend large amounts of money endeavouring to reduce this uncertainty.

To achieve this they utilised temperature-based derivatives in an attempt to smooth the uncertainty in their financial performance that was attributable to temperature variations. As an example, consider a company who supplies gas to consumers for heating during the winter months. Clearly the company will see reduced profits if temperatures are higher than usual during the winter period in question and, conversely, will experience higher profits if the region experiences lower than average temperatures. Hence this business might seek to reduce the overall variability of its expected profits buy purchasing a put option on the HDD index for the particular months in question.

[4] Energy and Power Risk Management – June, 1998
[5] 'Aquila To Enter Australian/New Zealand Weather Derivative Market' – Insurance Journal - August 20, 2001

Figure 1.3 shows the relationship between the maximum power load delivered and the average temperature for the New England region (source: McIntyre [2001]) . This chart clearly reveals the large correlation between these variables and shows graphically the hedging opportunities that are available to energy producers and consumers.

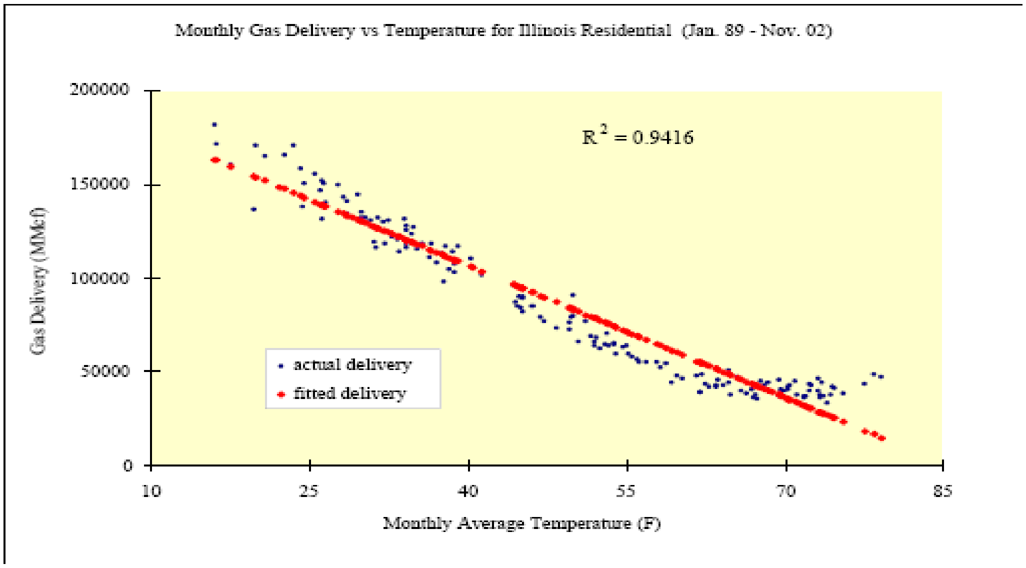

Figure 1.3: Power Load vs Temperature
(source: McIntyre [2001])

1.4.1 Weather Derivatives and Insurance

There are many similarities that exist between traditional insurance policies and weather derivative contracts in terms of their risk management applications. Both have the ability to recoup losses incurred by adverse weather conditions across both commercial and personal lines of business. Having said this, there are several important differences in terms of the coverage offered and the payouts received by the purchaser of these two forms of financial protection. Some of the most important points of difference include:

- Identifiable Loss: Insurance contracts are more often designed for protection from the extreme weather events such as a storm or flood where an identifiable loss has occurred. This loss is required as evidence that a claim can be made against the insurer. The payout of a weather derivative on the other hand is determined by reference to an index whose composition is transparent, such as temperature and rainfall. In this sense the purchaser of the protection does not require proof and significant savings can be made on legal fees required to defend appropriate payments.

- Moral Risk: The moral risk that is inherent in all insurance contracts can be nearly entirely removed as the reference is made to an index that is (hopefully) out of the control of both the counterparties. This will again act to reduce the cost of covering the associated risk and hence make weather derivatives a more affordable option to primary producers.

- Basis Risk: There are very little differences between the loss incurred and the final payout in an insurance contract. In other words, an insurance contract reduces the amount of basis risk present in a financial hedge by tailoring the payout to match the client's overall financial impact. Weather derivatives on the other hand are generic by there very nature and the returns from the contract will never properly match the exposure faced by the client (unless the client is situated in the weather station!). This is the most significant restriction to the expansion of weather derivative contracts.

The insurance industry has contributed to much of the growth in recent years. Weather derivative contracts have allowed many insurers and reinsurers to offer a wider range of products to their clients as well as being able to assist in managing the significant weather related risk that a general insurer will often retain. These now include non-catastrophic weather insurance contracts that can protect crops and their associated income streams from what are relatively predictable weather extremes.

Large reinsurers such as Swiss Re (New Markets division) and Hannover Re are currently using weather derivatives as part of their broader risk management programs and more recently in Asia companies such as Westpac Bank and Element Re have entered the market. Whilst these large reinsurers are well suited due to the generally large size of weather derivative contracts, theoretically they should be more appealing to geographically more concentrated insurers who are less easily able to obtain 'natural' diversification. Brokers such as Willis, AON and Marsh are actively promoting weather derivative products to their clients however these contracts tend to be highly tailored to meet the specific needs of the consumer and hence generally require specific pricing approaches.

1.4.2 Other Users

There are also other, more technical, users of these new derivative contracts that rely on their specific return profiles to achieve a better mean-variance outcome for their portfolio. In fact, weather derivative assets are beginning to be seen as attractive for a well-balanced portfolio as they have very small correlations with traditional asset classes such as bonds and equities thus creating a unique diversification tool. On top of all of this, hedge fund interest during 2004 and 2005 provided a new source of expansion that added to the growth that the market has experienced during recent years. These funds are typically attracted to inefficient markets, such as the market for weather derivatives, where their pseudo-arbitrage strategies can make significant excess returns over well-diversified portfolios.

More and more financial institutions are promoting these derivative products to their clients and, as will be shown throughout this paper, the breadth of risk management applications appear to have no bounds. These include:

- Construction companies – hedging against temperature, rainfall, snowfall;
- Drink manufactures - temperature;
- Farmers – all weather variables;
- Event Organisers – mainly precipitation;
- Tour Operators – all weather variables.

1.5 Literature

As has already been alluded to, the study of weather derivatives is fairly recent pursuit and the body of literature represents this. General pricing methodologies do not exist and the majority of pricing is completed via simulation of historical data, more commonly known as 'burning cost' analysis.

Cambell and Diebold (2002) laid out a framework for the modelling of temperature that has since been extended by several authors as well as tested empirically across different regions. Alaton[2002] provides a closed form approximation for the pricing of CDD options, based around a normal approximation, that can equally be applied to HDD based contracts. On the practical front, Garcia et. al. [2001] gives an excellent overview of how these contracts can be applied to practical situations, in particular a soft drink manufacturer.

Rainfall derivatives have received even less attention than the temperature contracts with much of the research being kept 'in-house' by the financial institutions that invest heavily in their development. Cao et al [2004] provide a comparison of three methods for the modelling of rainfall distributions than can be used as the basis for the pricing of derivative contracts over these statistics. Moreno[2002] provides an alternative approach as well as undertaking an analysis of basis risk in rainfall contracts in London.

1.6 Overview

The general aim of this paper is to overview a variety of approaches to the modelling and pricing of weather related risks, particularly temperature and rainfall based contracts. Throughout the paper the numerical methods outlined are highlighted by an empirical study on several weather stations based in Sydney. Data from the Australian Bureau of Meteorology (BOM) is used during this paper to provide practical examples of the numerical processes that are discussed.

Section 2 gives an outline of the pricing frameworks that are generally used by practitioners in the weather derivatives market. An overview of mean reverting dynamics and their association with the Uhlenbeck-Ornstein process is given and solutions to these stochastic equations are derived. A variation on the Black-Scholes framework proposed by Jewson and Zervos [2003] is discussed along with the limitations of the general BS approach for the pricing weather derivatives. A short revision of the common numrical techniques ('Burn' analysis, Monte Carlo simulations) is also provided to refresh their use before pricing options via these methods in the following sections.

Temperature-based derivatives, where the majority of the current literature resides, are covered in section 3. A variety of methods in the modelling and pricing of temperature based derivatives are discussed in detail. As well as this several case studies are included to assist in illuminating the usefulness of these derivatives in managing risk. The modelling of temperature roughly follows the method used in Benth(2005) and Alaton(2002) , however it is applied to the historical data from Sydney, Australia.

Section 4 investigates the modelling and pricing of rainfall derivatives. A similar analysis to Moreno (2002) is undertaken on two geographically close weather stations to determine if similar relationships can be found. Again several case studies are provided indicating the areas into which these relatively new financial products might find themselves. The final section of this paper attempts to summarise the main points covered with a view as to how the weather derivatives market might best move forward from these relatively primitive foundations.

2: Pricing Principles

The theory in terms of weather derivative pricing is still extremely sparse with no widely satisfactory formula that can be used by the growing number of practitioners in the marketplace. The current pricing methodologies can be broadly put into two groups; analytical solutions or numerical solutions. Whilst standard equity options have the famous Black-Scholes equation to provide practitioners with a reliable pricing basis, their weather-based counterparts have no agreeable equivalent to the BS framework and generally require a numerical approach.

A quick revision of standard option pricing theory is helpful when investigating its extensions later. Arithmetic Brownian motion is commonly represented by the following stochastic differential equation:

$$dX_t = \mu.dt + \sigma.dW_t \tag{2.1}$$

Whilst this might be suitable for modelling biological processes, Geometric Brownian motion (gBm) is the process that is generally used to model financial variables such as stock and commodity prices. Its necessity arises out of the fact that a log function does not permit negative values, essential when modelling asset prices. gBm is described by the following stochastic differential equation:

$$\frac{\partial X_t}{X_t} = \mu.dt + \sigma.dW_t \tag{2.2}$$

This can be reduced to arithmetic Brownian Motion via the substitution into the above equation of $y = F(X_t) = \log X_t$ and the use of Ito's formula, to give:

$$dF = (\mu. - \frac{1}{2}\sigma^2)dt + \sigma.dW_t$$

which with an initial condition, X_0, has a solution given by:

$$F(X_t) = \log x_0 + (\mu - \frac{1}{2}\sigma^2)(t - t_0) + \sigma W_{t-t_0} \tag{2.3}$$

A solution for the process X_t is then found by reversing the substitution (exponentiating), Hence we arrive at our distribution for the initial gBm process:

$$X_t = X_0.e^{[(\mu - \frac{1}{2}\sigma^2)(t-t_0) + \sigma W_{t-t_0}]} \tag{2.4}$$

From here Black and Scholes use their now famous hedging method to derive a partial differential equation (i.e. no longer a stochastic equation) for the dynamics of the option price that is based on the Brownian motion of equation (2.2). The Black-Scholes p.d.e is:

$$\frac{\partial V_t}{\partial t} = rV - ry.\frac{\partial V}{\partial y} - \frac{1}{2}\sigma^2 y^2 \frac{\partial^2 V}{\partial y^2} \tag{2.5}$$

Where V_t represents the payoff of the option contract. This differential equation can then be solved to obtain the explicit Black-Scholes formula. Alternatively a martingale approach can be adopted via equation (2.4), i.e. seeking a solution to the equation:

$$V_t = \mathbf{E}_Q[X_0 . e^{[(\mu - \frac{1}{2}\sigma^2)(t-t_0) + \sigma W_{t-t_0}]} \mid F_t] \qquad (2.6)$$

Where Q represents the risk-neutral martingale measure.

2.1 Mean Reversion

The key difference to modelling weather-based variables when compared with a general financial variable, such as a stock price, is that most weather components exhibit some degree of mean reversion. A mean reverting process is one in which the drift component of the stochastic differential equation (equation 2.1) always acts in a direction that opposes the current displacement from the mean process, in much the same way as a spring acts on a weight.

This concept is utilised frequently in the modelling of interest rates, which like many weather variables are at least partially mean reverting in that do not rise or fall without bound. The Vasicek(1979) model of forward interest rates is based on this mean reverting approach as well as other well known interest rate model such as the Hull-White model and the Cox-Ignersoll-Ross model.

The mean reversion component is deterministic and is an extension of the drift term, μ.

$$\frac{dX_t}{dt} \propto -(X_t - \overline{X}) \qquad (2.7)$$

where $\overline{X}$ represents the mean process. In this way the drift will always act in a way so as to bring the process closer to its mean. A proportionality constant is now required, called the mean-reversion parameter, which is a measure of the restoration force acting on the process dynamics. It is akin to the spring constant k, for all those who can still recall high school physics classes.

$$\frac{dX_t}{dt} = -\gamma.(X_t - \overline{X}) \qquad (2.8)$$

Substituting this into the standard Brownian motion dynamics, equation (2.1), we get:

$$dX_t = \gamma(\overline{X} - X_t).dt + \sigma.dW_t \qquad (2.9)$$

Note that this process will no longer be gBm and that negative values are permitted by equation (2.9). The probability that negative values will occur depends on the mean process level as well as the strength of the mean reversion. For a strongly reversionary process who's mean is significantly above zero it is highly unlikely that the process would ever go negative. The Ornstein-Uhlenbeck process is the name given to this modified Brownian motion. It has since been shown in Dornier and Queruel [2002] that the process given in equation (2.8) does not actually revert to its mean when temperature is used as the variable. This is due to the fact that the mean process that the equation is reverting to, $\overline{X}$, is not constant. To overcome this an extra term (the time derivative of the mean process) is required to be added to the drift component of the stochastic equation above. Hence we get:

$$dX_t = \left[\gamma(\overline{X} - X) + \frac{d\overline{X_t}}{dt} \right] dt + \sigma . dW_t \tag{2.10}$$

This representation is now mean reverting in the long-run, in other words: $\mathbf{E}[X_t] = \overline{X}$. Another advantage of including this term is that equation (2.10) can now be solved by the traditional integrating factor method. Multiplying through by $e^{\gamma s}$ we obtain:

$$e^{\gamma s} dX_s - e^{\gamma s} \gamma(\overline{X} - X)dt + e^{\gamma s}.d\overline{X_s} = e^{\gamma s} \sigma . dW_s \tag{2.11}$$

Now the left hand side of the above expression is just the differential of a product, i.e:

$$e^{\gamma s} dX_s - e^{\gamma s} \gamma(\overline{X} - X)dt + e^{\gamma s}.d\overline{X_s} = d[e^{\gamma s}.(X_s - \overline{X}_s)] \tag{2.12}$$

$$d[e^{\gamma s}.(X_s - \overline{X}_s)] = \sigma . \int_0^t e^{\gamma s} dW_s \tag{2.13}$$

Note that this would not have been possible without the extra term being added to equation (2.10) to make it properly mean reverting. After rearranging this expression we obtain the solution to the stochastic process:

$$X_t = \overline{X}_t + (X_0 - \overline{X_o}).e^{-\gamma \Delta t} + \int_s^t e^{-\gamma \Delta t}.\sigma_\tau dW_\tau \tag{2.14}$$

This equation will then become the basis of the Monte Carlo simulations that are undertaken in section 3.

2.2 Black-Scholes formulation

The seminal paper of Black and Scholes [1977] provided an analytical framework for the pricing of contingent claims and in particular options, however the Black Scholes (BS) formula relies on some fairly stringent assumptions. Most importantly it is assumed that the underlying process is driven by gBm as given above by equation (2.1). Most empirical studies show that in fact asset returns are strongly leptokurtic (More concentrated in the middle and 'fat-tailed'), such as Fama (1965).

As well as this, with respect to weather derivatives, there is generally no underlying process that is actively traded (i.e. people don't trade degrees ….yet) and as such the BS framework has no method of hedging the derivative in order to derive an analytical solution. For these reasons a standard BS approach is not applicable and other more indirect methods must be pursued.

2.2.1 Asian Options

Many common weather derivative contracts have as their underlying index, an average of some statistic over a period of time. For example, many temperature-based derivatives will have a payoff that is determined by reference to the average temperature over a week or a month. These types of options are generally referred to as 'Asian' options and are common in both fixed-interest and equity derivative markets. The averaging can be done on a variety of time frames however it is generally the daily temperature that is averaged. For example a call option on the monthly average temperature would have a payoff of the form:

$$V_t = \max\{0, (\frac{\sum_{i=1}^{n} T_i}{n} - K)\} \cdot tick \quad (2.15)$$

where there are n days in the month. Analogies to the BS partial differential equation can be derived for the varieties of 'Asian' options that exist[6] however, as outlined above, we still have no underlying asset with which to perform the required hedge. These types of contracts are particularly appropriate for rainfall-based derivatives whose discreetness can be smoothed by averaging over a larger period.

2.2.2 Alternative Black Scholes

Several authors have decided to proceed with the BS framework ignore the problems of the underlying assumptions or have attempted to alter the BS approach to accommodate these deficiencies. The most serious assumption that must be relaxed is that there is no underlying asset to base the derivative price around. Jewson et. al [2003] has suggested that an alternative BS formula can be derived akin to the derivation of an option on a futures contract. To overcome the fact that there does not exist an underlying, traded asset, Jewson creates a hypothetical forward weather index that is used as the underlying asset which the BS framework can use as a hedge to enable the pricing of the contingent claim.

The theory is based on the Black (76) model where a futures contract is used as the hedge when deriving the partial differential equation that governs its motion. To illuminate, let us assume that the futures price process is governed by the 'cash-and-carry' relationship:

$$Y_t = X_t . e^{r(T-t)}$$

By using Ito's formula when substituting this relationship into the standard gBm form equation (2.1) we get the altered stochastic differential equation:

$$dY_t = y[(\mu - r)dt + \sigma dW_t] \quad (2.16)$$

Following the same hedging procedure as used in deriving the Black Scholes partial differential equation we arrive at the following relationship for the dynamics of options on a futures contract:

$$\frac{dV_t}{dt} = rV - \frac{1}{2}\sigma^2 y^2 \frac{d^2V}{dy^2} \quad (2.17)$$

[6] These include fixed strike and floating strike amongst others. See Buchen [2002] for derivations.

If we compare this relation with equation (2.5) we can see that the term, $-ry.\frac{\partial V}{\partial y}$ is removed from the right hand side of the equation. This is the same equation one would get if calculating the price of an option on a dividend paying stock where the dividend yield was equal to the risk free rate. Hence by using notation similar to Buchen [2002] we can write the option price over a futures contract as:

$$\begin{aligned} V(y,t) &= BS(ye^{-r\tau}, t, r, \sigma) \\ &= e^{-r\tau}.BS(y,t,0,\sigma) \end{aligned} \tag{2.18}$$

where $BS(x,t,r,\sigma)$ represents the 'standard' Black Scholes pricing formula. This formula can then be used, after appropriate modelling of the futures price process, to calculate the premium applicable to a range of weather based option contracts.

2.3 'Burn' Analysis

This is a typical actuarial approach adopted to price a contingency where no assumptions are required to be made as to the nature of the process on which the contingency relies. A typical 'burn' analysis seeks to answer the question: "What would be the return from the contract had I purchased it each year for the last x years?" Generally, an arithmetic average is then taken of the results.

For example, to price a February HDD call with exercise of 100, simply find what the financial return would have been for each of the February months in the historical data set, with the appropriate indexing of the exercise value in order to standardise the temperatures over time[7]. This approach will be used in the pricing of both temperature and rainfall based derivatives in sections 3 and 4 over a range of exercise levels.

2.4 Monte Carlo Simulations

Monte Carlo simulations differ from the 'burn' analysis approach above in that they require assumptions to be made as to the dynamics of the underlying variable. Essentially they involve running a series of simulations based on a statistically derived model and then calculating what the expected return is from all of these simulations.

Mathematically speaking we find a solution to:

$$\mathbf{E}[f(X_t)] = \frac{1}{N}.\sum_{i=1}^{N} f(\overline{X}(t,\psi_i)) \tag{2.19}$$

i.e. the arithmetic average of the simulation outcomes. Here t represents time and ψ_i is the series of calculation points.

The concept of mean reversion discussed earlier has important consequences for the selection of an appropriate starting point for the simulation process. If the time period of

[7] If using data sets over 100 years then a quadratic parameterisation of the long-term trend will be required to properly estimate the effective strikes. See section 3.3.1 for more discussion.

interest is in the short-term then it is necessary to begin the simulations from the present day values as any current deviation from the mean will have an impact on the process in the short-term. Consider for example a one-day option over the average daily temperature that was being valued at both 1 day and 1 month before the contract period. For the valuation close to the contract period it is important to take into account the current displacement from the mean as this value will have an impact on the temperature on the following days. Hence the simulation should be started on the actual raw value. If however we are 1 month from the contract period then we can begin the simulation on the mean temperature process as the actual displacement from the mean will not have an appreciable influence on the temperatures in a month's time. An example of pricing via a Monte Carlo simulation is shown in section 3 when the valuing of temperature based derivatives are investigated.

3: Modelling Temperature

Compared with rainfall, the analysis of temperature has received significant attention from the literature in recent years. Primarily this is due to the fact that the majority of traded contracts in the weather derivatives market are temperature based, generally related to energy supply or demand. More recently the search for a reliable statistical model for temperature dynamics has been intensified by the need to provide sound evidence for the impacts of human interaction on the planet. The temperature model that is developed in the next sections follows a combination of approaches used by Benth et al [2002] and Alaton [2002] and has been applied to a range of weather station recordings.

3.1 Data

In order to highlight much of the discussion made herein, reference will be made to an analysis undertaken using temperature data from the Sydney region for the period 1856-2005. The data was obtained from the Australian Bureau of Meteorology (BOM). The following subsections of the data were used in this investigation:

- Sydney Airport Jan 1940 – Dec 2005
- Observatory Hill, Sydney Jan 1940 - Dec 2005
- Prospect Dam Jan 1965 - Dec 2005

The majority of the analysis in this paper is in respect of the Sydney Airport weather station which consists of 66 complete years of minimum and maximum temperature as well as precipitation readings. Please refer to the appendix for a discussion of the treatment of missing values in these data sets.

3.2 Temperature Distributions

The following histogram shows the distribution of temperatures at Sydney Airport for the period Jan 1940 through to Dec. 2005

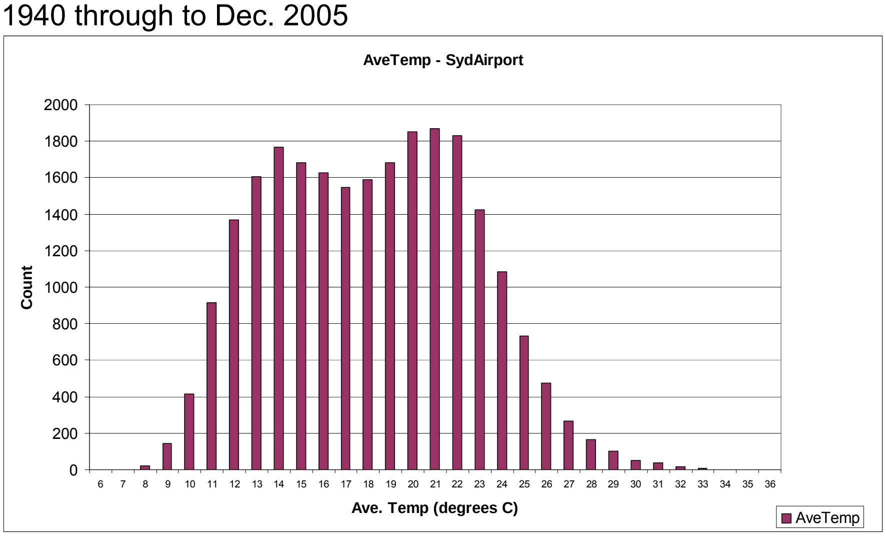

Figure 3.1: Temperature Histogram – Sydney Airport

As figure 3.1 shows, the distribution is bimodal, reflecting peaks for both the summer and winter months. This feature is common amongst most temperature distributions throughout the world and Appendix B shows a range of these. It is interesting to note that southern hemisphere locations tend to have a negatively skewed distribution where as northern hemisphere locations are usually positively skewed. If one were to restrict attention to just summer months or just winter months then the following patterns are revealed:

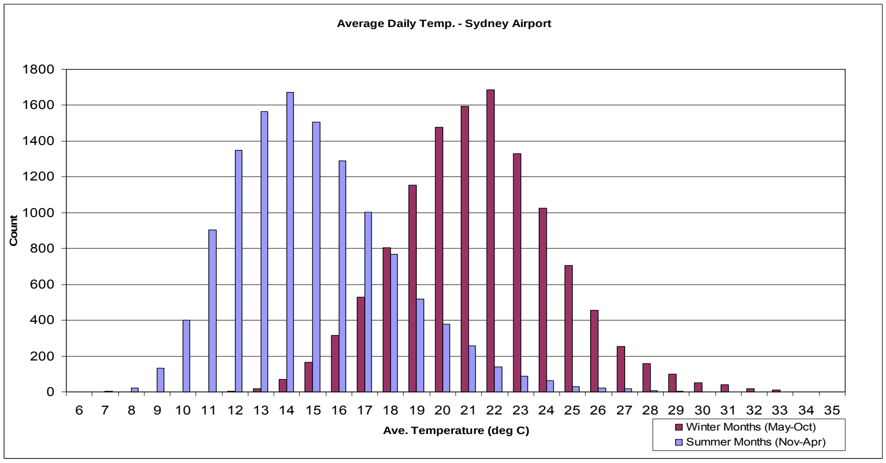

Figure 3.2: Temperature Histogram – Sydney Airport

The initial part of the modelling process is centred on removing both the long-term and seasonal trends from the data (the deterministic part) and from there a stochastic model is fitted to the residuals. The first part of this process is referred to as 'detrending' and several authors have proposed methods for achieving this, for example Benth [2005] and Alaton [2002].

3.3 De-trending

De-trending is the term used to describe the removal of a deterministic approximation to several of the more predictable patterns in the weather. Specifically it refers to the removal of seasonal patterns as well as linear (or quadratic) terms that represent long term heating or cooling of the earths atmosphere as well as the added impacts caused by human interaction.

The decision to de-trend the data can be time consuming and costly so there have been attempts made to quantify the necessity to carry out the procedure. Jewson [2004] provides a decision rule for the introduction of linear trend but finds that it does not significantly beat a rule to always de-trend or to always not de-trend. He found that it beat a no trend rule when the time series was relatively long and conversely beat a linear trend when the sample was short.

3.3.1 Long-term Trends

Analysis of most long-term temperature data reveals a slight positive trend that represents the gradual warming of the globe that has been occurring since the last ice age. Study of longer time periods shows that the trend is more significant in latter years and suggests that a quadratic term should be used to model these long-term environmental changes.

$$T_{Linear} = a + b.t + c.t^2$$

However, for shorter time periods (<80 years) it is reasonable to assume a linear form.

$$T_{Linear} = a + b.t \tag{3.1}$$

This is generally acceptable given that the majority of the weather derivative contracts that are entered into have periods of a year or less. This means that to make projections over time periods where the quadratic term becomes significant are rare in practice and for the purposes of this investigation a linear trend will be assumed.

3.3.2 Seasonal Variation

Seasonal variation is by far the most dominant term in the overall temperature variation. The periodic nature of temperature that results from the seasonal variation can generally be described by a sinusoidal function and for the purposes of this analysis we will assume it to be represented via a truncated Fourier series of the form:

$$T_{Seasonal} = \varepsilon\alpha_0 + \sum_i \alpha_i .Sin(\gamma t + \phi) + \sum_i \beta_i Cos(\lambda t + \theta) \tag{3.2}$$

Here the coefficient α_0 is not required as this effect would be captured in the linear trend already discussed. The remaining parameters are to be estimated from the data. Combining these two effects, equations (3.1) and (3.2), we now have an equation for the mean temperature process given by:

$$\overline{\mathrm{T}} = \mathrm{T}_{\mathrm{linear}} + \mathrm{T}_{\mathrm{seasonal}}$$

hence,

$$\overline{T} = a + b.t + \sum_i \alpha_i .Sin(\gamma t + \phi) + \sum_i \beta_i Cos(\lambda t + \theta) \tag{3.3}$$

For the purposes of this investigation a first order Fourier series will be used so that only one α and one β need to be estimated from the data. This will entail a total of 6 parameters to be estimated in equation (3.3) above.

3.3.3 Parameter Estimation

A Least-squares estimation algorithm is used to determine the values of the parameters given in equation (3.3). In relation to the ‘speeds’ of the seasonal processes (γ and λ), there are two possible approaches to determining suitable parameter values. Most authors choose to constrain the two values to $\frac{2\pi}{365}$ so that a Doppler effect is not encountered for projections over longer time periods. Alternatively the ‘speeds’ could be allowed to vary, independently of each other, so that a better fit to the raw data is achieved. When the two methods were undertaken on the Sydney Airport data series the following deviances were obtained:

Fixed γ and λ: 147,614
Variable γ and λ: 143,255

Whilst this shows that the removal of the constraint improves the model fit (as would be expected) the magnitude of the improvement does not warrant the loss of tractability that the model suffers as can be seen in the following graph that shows the resultant parameter values when the wave speeds are allowed to vary.

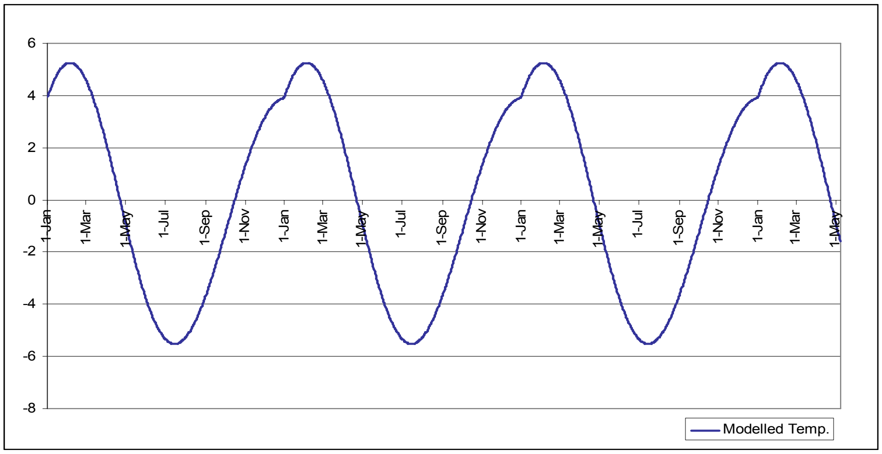

Figure 3.3: Over fitting of raw data when wave speeds allowed to vary.

The clear 'kink' in the graph at around the new year is hard to justify in terms of common experience. For this reason we will constrain the individual wave speeds to the value $\frac{2\pi}{365}$ so that we have the superposition of waves with equal frequencies. The following table summarises the results of the parameter estimation process for the three weather stations involved:

	Syd. Airport	Observatory	Prospect
a	16.925	17.434	17.26
b	$6.30*10^{-5}$	$5.16*10^{-5}$	$4.91*10^{-5}$
α	5.14	4.91	5.194
β	0.69	-0.20	0.986
φ	1.097	1.25	1.100
θ	0.97	1.10	0.675

Table 3.1: Parameter estimation results.

In order to estimate the mean-reversion parameter a measure of the autocorrelation must be established. To achieve this, today's temperature is regressed against yesterday's temperature (i.e. an AR(1) model) and then the following relationship yields the parameter:

$$\gamma = e^{-\kappa} \tag{3.4}$$

where κ is the regression parameter. This was undertaken for each of the weather stations yielding the following parameters:

	Syd. Airport	Observatory	Prospect
γ	0.555	0.570	0.561

Now that all of the parameter values have been estimated it is then possible to simulate typical paths that represent the dynamics of average temperature throughout the year. The following graph shows the actual and modelled temperatures for the last 10 years.

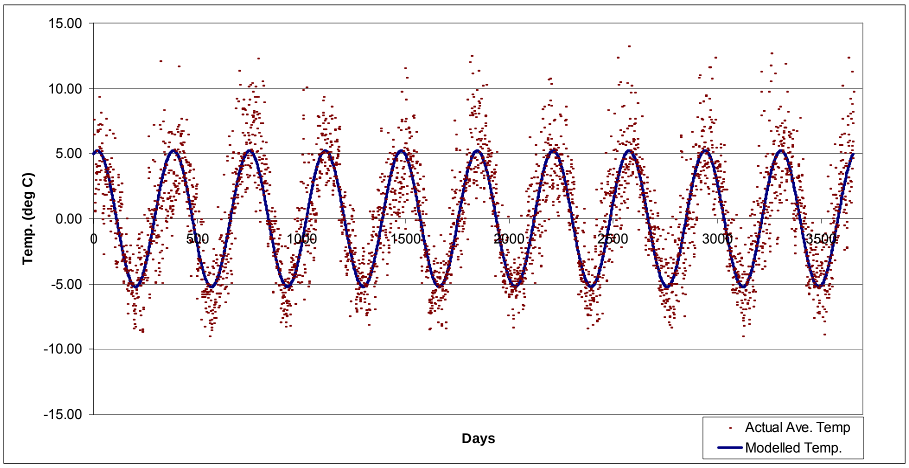

Figure 3.3: Modelled Seasonal Temperature – Sydney Airport

The increased volatility of temperatures in the summer months is evident from the graph but what is less obvious is how volatility varies throughout the rest of the year.

3.4 Patterns of Volatility

Most pricing models for weather derivative contracts are derived on the basis of a constant volatility over the term of the contingency. On a closer inspection of figure …… it is evident that the volatility is not uniform over the year and there are in fact distinct seasonal patterns of volatility. Note the following pattern for the seasonal variation of the volatility of temperature at Sydney Airport:

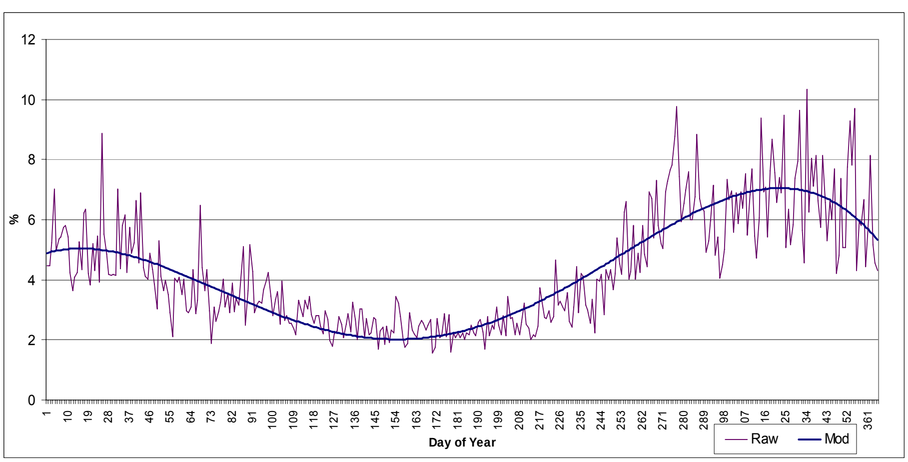

Figure 3.4: Seasonal Volatility Pattern – Sydney Airport

A degree-4 polynomial has been fitted to the data and is shown in figure 3.4 above. Other authors suggest modelling volatility as a Fourier serries (Benth [2005]), as a separate Ornstein-Uhlenbeck process (Bhowan[2003]) or as a piecewise construction of constant monthly volatilities (Alaton [2002]).

Once all of the aforementioned factors are removed from the raw data we are left with the final residuals which if the model assumptions are correct should approximate a normal distribution. The following chart shows the distribution of the residuals:

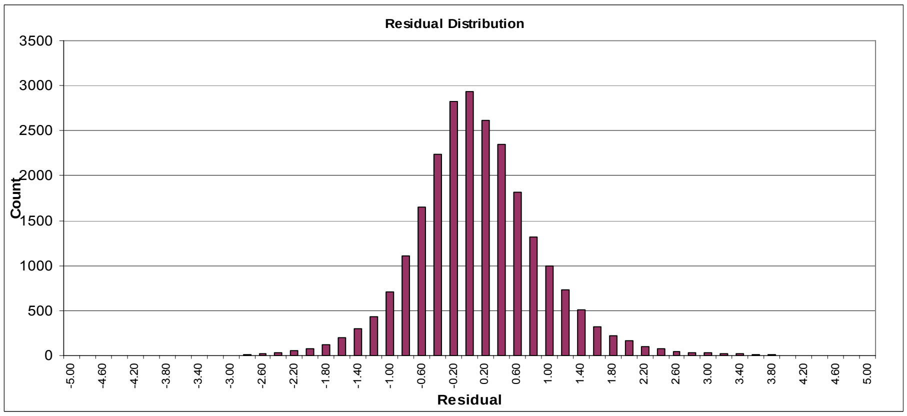

Figure 3.5: Residual Distribution (Z_i)

As can be seen from figure 3.4, the residuals are a good approximation to a normal distribution at a high level although they appear to be leptokurtic and slightly right-skewed.

3.5 Stochastic representation

Recall from equation (2.8) that the stochastic differential equation that we are using to model temperatures is:

$$dT_t = \left[\gamma(\overline{T} - T_t) + \frac{d\overline{T}}{dt} \right] dt + \sigma.dW_t \tag{3.5}$$

This has a solution for an initial conditions of, T_0 & $\overline{T_0}$, that is given by:

$$T_t = \overline{T}_t + (T_0 - \overline{T_o}).e^{-\gamma\Delta t} + \int_s^t e^{-\gamma\Delta t}.\sigma_\tau dW_\tau \tag{3.6}$$

In order to undertake a simulation of the stochastic equation (3.3) we need to obtain a discrete approximation to the dynamics so that daily readings can be predicted. A Euler approximation to equation (3.5) is:

$$T_{t+1} - T_t = \gamma\,(\overline{T} - T_y).+\frac{d\overline{T}_t}{dt} + \sigma.\mathrm{Z} \tag{3.7}$$

Where $\mathbf{Z} \sim \mathbf{N}(0,1)$.

This equation can be used to simulate sample trajectories of the stochastic process and will enable the undertaking of Monte Carlo type analysis.

3.6 Pricing Option Contracts

Section 2 outlined the common approaches that are currently employed in the pricing of weather derivative contracts. Later in this section a numerical comparison of these methods will be undertaken with reference to the weather stations outlined in section (3.1). Firstly, we will look at an approximation that can be used to price both HDD and CDD option contracts.

3.6.1 Normal Approximation

Alaton [2002] proposes an approach for deriving a closed form approximation to the price of an option over CDD or HDD contracts. The catch is that generally the CDD price will only work in summer months and the HDD in winter months, however, these periods would generally represent the exposure periods for the majority of the traded weather contracts.

The basics behind the approach are to assume that the max() function has no effect on the resulting distribution of HDD's or CDD's. If for example the 18 degree limit was used as the basis of a CDD contract in Sydney during January it would be a rare day that has an average temperature below this value. For the record, out of the 2045 January days since 1940 only 11 have recorded an average temperature below 18 degrees (in today's terms), hence this may not be as restrictive an assumption as it may first appear.

Now consider the measurement of CDD's as given earlier by equation (1.3) but with the reference level set at 18˚C, i.e.:

$$CDD_n = \sum_n \max\{0, (T_i - 18°)\} \tag{3.8}$$

If we make the assumption that the average daily teperature will always be greater than or equal to the reference temperature (18˚C) then we can rewrite this relation as:

$$CDD_n = \sum_n T_i - 18n \tag{3.9}$$

Now we have removed the complication of the max() function and we are free to extend the assumption of normally distributed CDD's. The relationship in equation (3.9) above is a linear combination of a Gaussian process which itself will be Gaussian and a normal approximation can therefore be legitimised for relatively large values of n.

We can then calculate the moments of this via:

$$\begin{aligned}\mathbf{E}[CDD|F_{\mathbf{t}}] &= \mathbf{E}[\sum_{i=1}^{n} T_i - 18n] \\ &= \sum_{i=1}^{n} \mathbf{E}[T_i] - 18n\end{aligned}$$

and

$$\mathbf{V}[CDD \mid F_t] = \sum \mathbf{V}[T] + 2\sum\sum Cov[T_{i,} T_j]$$

The value of a call option over this CDD distribution can be given as:

$$c(t) = e^{-r(t_n - t)}\mathbf{E}[\max(CDD_n - K, 0) \mid F_t] \tag{3.10}$$

where K is the strike price of the call option and $t_n > t$. Proceeding under the assumption of normality, Alaton shows that:

$$c(t) = e^{-r(t_n - t)} \int_K^\infty (x - K).f_{CDD_n}(x).dx$$

which on evaluation of the integral gives:

$$c(t) = e^{-r(t_n - t)}[(\mu_n - K)\mathbf{\Phi}(-a_n) + \frac{\sigma_n}{\sqrt{2\pi}} e^{\frac{-a^2{}_n}{2}}] \tag{3.11}$$

Similarly for a put option over the CDD index we can write:

$$p(t) = e^{-r(t_n - t)}\mathbf{E}[\max(K - CDD_n, 0) \mid F_t] \tag{3.12}$$

$$p(t) = e^{-r(t_n - t)} \int_0^K (K - x).f_{CDD_n}(x).dx$$

which as before yields:

$$p(t) = e^{-r(t_n - t)}[(K - \mu_n)[\mathbf{\Phi}(a_n) - \mathbf{\Phi}(\frac{\mu_n}{\sigma_n})] + \frac{\sigma_n}{\sqrt{2\pi}}(e^{\frac{-a^2{}_n}{2}} - e^{\frac{-\mu_n{}^2}{2\sigma_n{}^2}})] \tag{3.13}$$

where $\mathbf{\Phi}$ represents the cumulative (standard) normal distribution and r is the risk-free investment rate. This is a powerful approximate method for the pricing of weather based option contracts particularly in geographic areas where the effect of the max() function is minimal.

3.6.2 An Example

To highlight much of the discussion so far we will price a CDD option for the month of January in Sydney. The CDD option is chosen as it would be popular to power generating organisations seeking to smooth their returns due to the variability of electricity demand during summer months. We will use 3 methods to calculate the option price: Normal approximation, ‘Burn’ analysis and Mote Carlo simulations.

The specifics of the options to be priced are:

Period:	January
Measure:	Cumulative CDD
Exercise Prices:	170 / 180 / 190 / 200 CDD’s
Tick::	$100,000 /CDD
Location:	Sydney Airport (Kingsford Smith)

The pricing is undertaken for the three weather stations investigated throughout this section.

Burn Analysis

A burn analysis is undertaken using 66 years of data from the Sydney Airport weather station. To begin with, it is necessary to calculate the present value of the temperatures by adding to the raw temperatures the linear and quadratic trends that were calculated in section 3.2.1. As the linear trends will not affect the arithmetic average it is only necessary to inflate the daily average temperature rather than the raw maximum and minimum values. These present value temperatures are then used to calculate the expected payoff for this contract had it been purchased every January for the past 66 years.

Monte Carlo Simulation

The following chart shows a typical sample simulation for the month of January at Sydney Airport that is used to determine the expected payoff from the CDD option. The simulation was achieved via the use of equation (3.7) along with a random number generator.

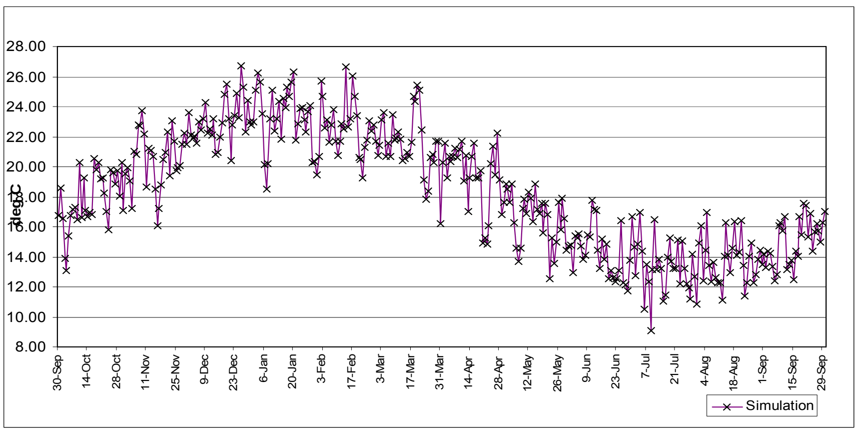

Figure 3.6: Simulated Sample ‘Path.

Results

The tables that follow show the results obtained for the Sydney Airport weather station:

Sydney Airport
January

Method	Exercise (CDD) 170	180	190	200
'Burn' Analysis	$473,306	$268,763	$137,865	$59,582
Monte Carlo Simulation	$489,044	$230,479	$90,848	$13,990
Normal Approximation	$463,670	$288,627	$167,993	$91,012

Table 3.2: Sample Option Prices

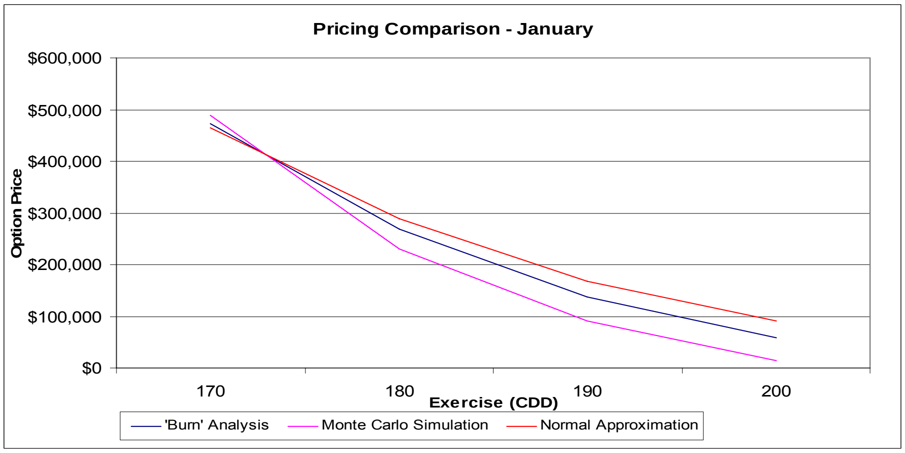

Figure 3.7: Comparison of option pricing methodologies

3.7 Risk Management Applications

Beverage Manufacturers

In May of 2000 Corney and Barrow, an operator of a chain of London wine bars, entered in to a weather derivative contract in order to hedge against variability in its summer sales. The company had found that sales were significantly reduced when lower than normal average temperatures were experienced and sought to reduce the dependency of their financial returns on the weather. The contract called for the payment of £1500 for every degree the average daily temperature was below 24 capped at a maximum of £15,000 per day and £100,000 for the entire summer. The derivative only applied to Thursdays and Fridays, which represented the majority of business for the inner city wine bars and hence provided an effective hedge.

An illustrative case study was undertaken by Garcia et. al. [2001] that highlights the effects that temperature can have on juice and milk sales in Switerland, UK, Germany and France. Whilst little relationship was found in the milk sales, the juice showed a high positive correlation with the average daily temperature. The authors investigate the strength of the relationship across the various countries as well as seeking an ideal level for a cap, that is, a level of tempreture above which further increases have little impact on sales (or possibly even a negative impact).

Construction Industry

There are a variety of ways in which the prevailing temperature can have an adverse impact on general construction activities and this has lead to an increasing interest in weather derivatives by industry participants.

Whilst abnormally hot temperatures can cause significant delays it is generally cold temperatures that result in the most significant cost impacts. Some of the weather hazards that most building and construction companies are exposed to include:

- Concrete: Frost causes water that exists in all concrete structures to freeze which in turn can lead to the member cracking under the pressure of the expanding water. Also concrete requires particular temperatures during the setting process if it is to obtain its ‘ideal’ strength and prevalently low temperatures can cause the delay of the construction.

- Stop Work: Workers that are not in air conditioned areas have maximum and minimum temperature that they are allowed to work at. Also at temperatures near the extremes workers become less efficient i.e. less work get completed in hot temperatures.

- Snow: Snowfall that is unexpected in either its occurrence or its magnitude can cause extensive delays to a construction site as well as the added chance of collapse of any temporary structures that are required during construction.

4: Rainfall Derivatives

The market for rainfall derivatives has been slow to gain momentum after first having been neglected in favour of the more popular temperature-based derivatives. The body of literature remains narrow and there is presently a lack of a unified pricing approach to give the newborn market participants confidence. As can be seen in figure 1.2 the market for rainfall contracts is now approximately 10% of the total market size and represented around $US 840 million of notional value in 2005.

Whilst the demand for rainfall protection is high there are several technical difficulties that are restraining the further growth of these contracts. ‘Basis risk’ (discussed in the following section) and a lack of reliable statistical models have meant that most banks and other large financial institutions are reluctant to quote prices over these complex and highly uncertain derivative products. Despite this there have been several recent examples of businesses utilising the risk management power of these contracts.

During 2002 a German golf course operator entered into a rainfall based contract with the large French bank, Société Generale. In recent years the operator had seen significantly lower earnings due to a much higher than average rainfall and wished to hedge their exposure to the weather through a derivative contract. The contract called for the payment of a fixed sum for every day over 50 that received more than 1mm of rain during the months of May through September.

4.1 Data

Data was analysed from the same weather stations as used in section 3 for the temperature analysis. Most weather stations in the data series supplied by the Australian BOM have recorded precipitation values significantly earlier than temperature recordings and as such we have recordings as early as 1858 for the Observatory Hill weather station and 1887 for Prospect Dam. The data sets used in the analysis were:

- Sydney Airport Jan 1929 - Dec 2005
- Observatory Hill, Sydney Jan 1858 - Dec 2005
- Prospect Dam Jan 1887 - Dec 2005

A similar analysis is performed on the Sydney data to that undertaken by Moreno [2002] for rainfall dynamics in London, UK where the local discrepancies between two geographically close[8] weather stations in that city were analysed with a view towards pricing rainfall derivatives. The proposed methodology of Cao et al [2004] is empirically evaluated with reference to weather statistics.

As before, Appendix A contains a discussion on the treatment of missing values in the data series used in this investigation. In general the measuring of precipitation values over the historical data set tends to be more reliable that the temperature based measurements with few significant gaps in the series.

[8] The two stations were Heathrow Airport and St James Park, a distance of 32 km apart.

4.2 Local Correlation

Temperatures diffuse quickly to produce homogenous areas with nearly uniform levels of the desired statistic. Rainfall, on the other hand, is discrete by its very nature and this provides difficulties when attempting to produce reliable models of its variability. This local variability manifests itself as 'basis' risk when considering the use of derivatives to hedge rainfall exposure and is one of the primary reasons that these types of weather derivatives have been slow to take up compared with temperature derivatives. Often the derivative contracts are based on the average of several nearby weather stations in order to reduce much of the spatial basis risk that is inherent in precipitation derivatives.

Unlike temperature statistics, rainfall shows a much higher degree of local variation (volatility) due to the discrete nature of rainfall. Hence two weather stations in close proximity to each other often record vastly different rainfall measures on the same day. This provides a considerable issue in terms of their suitability for hedging rainfall exposures as it is highly dependant on the location of the measuring station.

To highlight this occurrence I will briefly analyse the correlation between two sites in Sydney, Sydney Airport and Observatory Hill, which are a distance of 11.2 km (7 miles) apart. We will also make a comparison with Prospect Dam weather station which is a distance of around 34km west from Observatory Hill.

4.2.1 Variation

On average the Observatory Hill station received 140.7mm more rain per year than Sydney Airport (1229.9mm and 1089.2mm) and only 1 in every 4 years does Sydney Airport record a higher cumulative rainfall for the year. The graph that follows depicts the correlation of the yearly cumulative rainfall from the two weather stations.

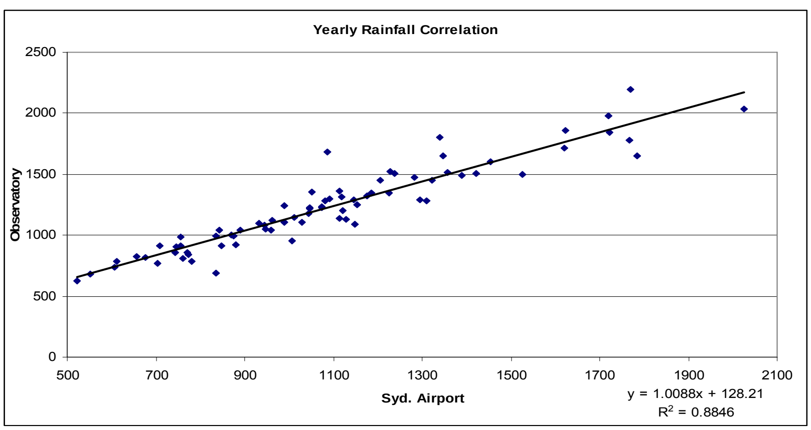

Figure 4.1: Annual Rainfall Correlation – Syd. Airport vs Observatory

The monthly rainfall correlations are more important than the yearly relationships as derivative contracts of the latter are rare. Figure 4.2 that follows shows these monthly correlations for the two weather stations.

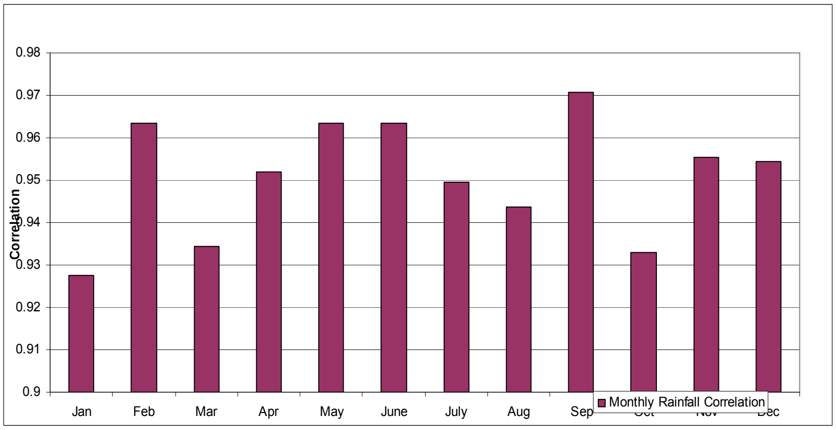

Figure 4.2: Monthly Rainfall Correlation – Syd. Airport vs Observatory

These correlations are in a relatively tight band with all monthly correlations being greater than 0.92. which shows that the basis risk involved with precipitation derivatives in Sydney is not as pronounced as that found by Moreno [2002] for London based weather stations. The daily correlation between the two sites is much weaker that the yearly correlations and often differences of 100% are encountered. A correlation of 0.89 was calculated for the daily average rainfall between the two sites

4.3 Modelling Rainfall

As has previously been alluded to, the statistical modelling of rainfall provides challenges that are not inherent in the modelling of temperature. Most of this difficulty arises from the discreetness of rain when compared with temperature distribution which requires different measures in order to properly model them. Wilks[1998] derived a model of precipitation comprised of two components; a Markov chain to capture the frequency and a mixed exponential distribution to represent the magnitude of the rainfall process.

Cao el al [2004] extend this by deriving and comparing three proposed models: a gamma distribution, a mixed exponential distribution and a kernal density approach.

Kubo and Kobaysahi [2002] argue that for the pricing of daily rainfall options it is sufficient to use a model based on a frequency assumption for that particular day of the year rather than considering the autocorrelation inherent in rainfall patterns, as in Moreno [2002].

4.3.1 Linear Trends

Due to the long-term trends that were found in temperature data it would be reasonable to expect that there might be some long-term, discernable trends in the rainfall patterns. A regression was performed of the yearly cumulative rainfall at Sydney Airport since 1940 to determine if any statistically significant trend could be established. Figure 4.2 displays the results. A slope of 0.055 with a standard error of 1.952 rejected any assumption of a linear trend in yearly cumulative rainfall at Sydney Airport.

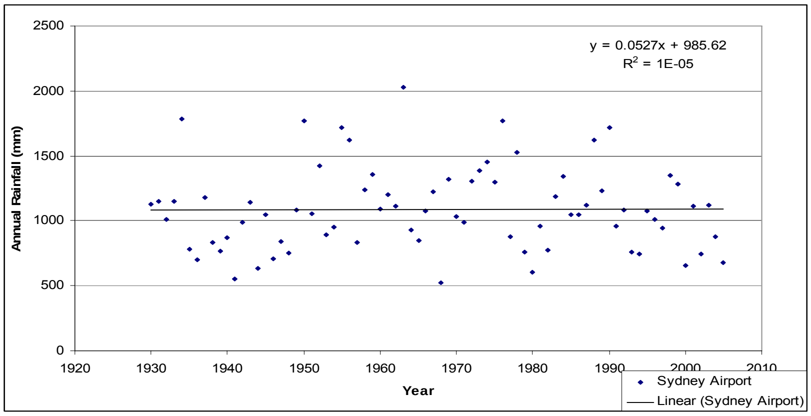

Figure 4.3: Annual Cumulative rainfall – Sydney Airport.

This means that the clear linear trends uncovered in the temperature data in section 3 do not transfer into a noticeable trend in rainfall patterns. Even when a linear regression was undertaken on the entire Observatory Hill data set (148 complete years) there were still no clear evidence of any linear pattern. The slope parameter was estimated at 0.2565 with a standard error of 0.655.

4.3.2 Frequency Process

The frequency of the rain process is simply the Boolean answer to the question: “Did any precipitation enter the measuring devise during the 24hr period?” Hence the amount of rain has no effect of the frequency process that we propose as will be seen in section 4.3.3. Later, however, we will allow the magnitude process to depend on the outcome of the frequency process described here.

4.3.2.1 Seasonal

Daily frequency resides in a 20%-50% band that has a mild seasonal dependence. Frequency tends to be higher in the summer months than in the winter months but noise appears to be the dominant process. Figure 4.3 shows the yearly distribution of daily rainfall frequency with a fitted degree 4 polynomial that will be used for subsequent projections.

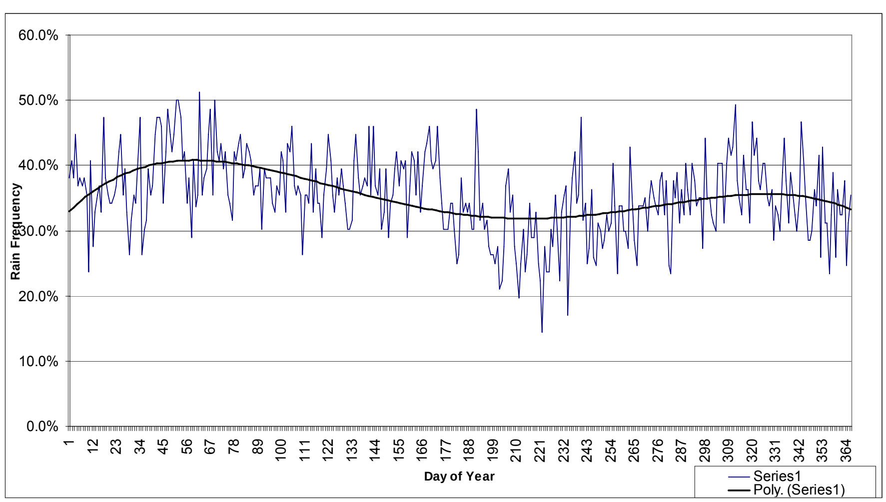

Figure 4.4: Seasonal Rainfall Frequency – Syd. Airport

4.3.2.2 Persistence

It is evident from every day weather observation that the chance of it raining on a particular day is strongly related to wether it rained on the previous day. This leads us to attempt to model the length of a rain period rather than just the frequency of rain.

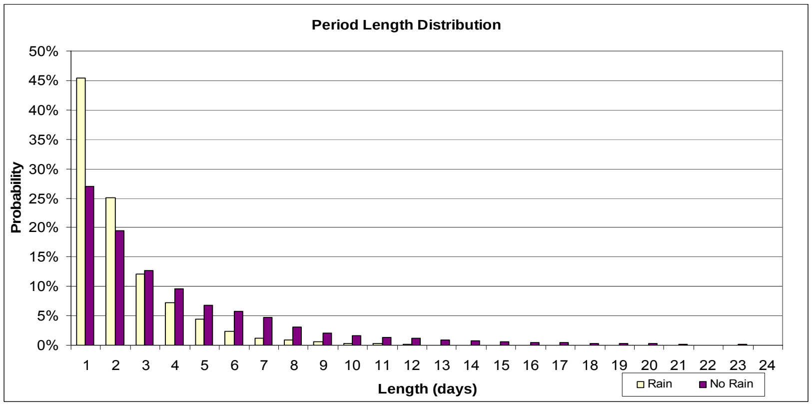

Figure 4.5: Rain/no Rain Persistence Distribution-Sydney Airport

We could use the above distribution of rain and no-rain periods to complete simulations over a particular period in the future. However as Moreno identifies and as can be seen in figure 4.3, the frequency of rain is not constant throughout the year and hence it would be unreasonable to assume that the lengths of rain/no-rain periods are constant either. This means that we would be required to produce separate models for all of the individual periods where there are significantly different frequencies.

4.3.2.3 Modelling

We propose a Markov chain for the modelling of the frequency of rain as outlined in Cao et al [2004]. Moreno's approach is to model via a Bernoulli distribution that was dependant on its history. Hence the probabilities were of the form:

$$P_t = P[X_t = 1 | X_{t-1}, X_{t-2}, \ldots\ldots, X_{t-n}] \tag{4.1}$$

The author then proceeded to determine the value of k in the above equation that produced the most reliable results and found that k = 1 gave the smallest relative error. So the process depends only on its last value and hence is equivalent to a Markov chain. Seasonality that is present in the frequency, as evident from figure 4.3, can be introduced through the variation of transition probabilities throughout the year. We propose fitting a degree 4 polynomial to each of the distribution depicted in figure 4.6.

The difference now is that we need to estimate the frequencies conditional on wether or not it rained on the previous day, an extension of the analysis carried out in section 4.2.2.1.The following graph shows the seasonal distribution for the two conditions along with their respective fitted distributions.

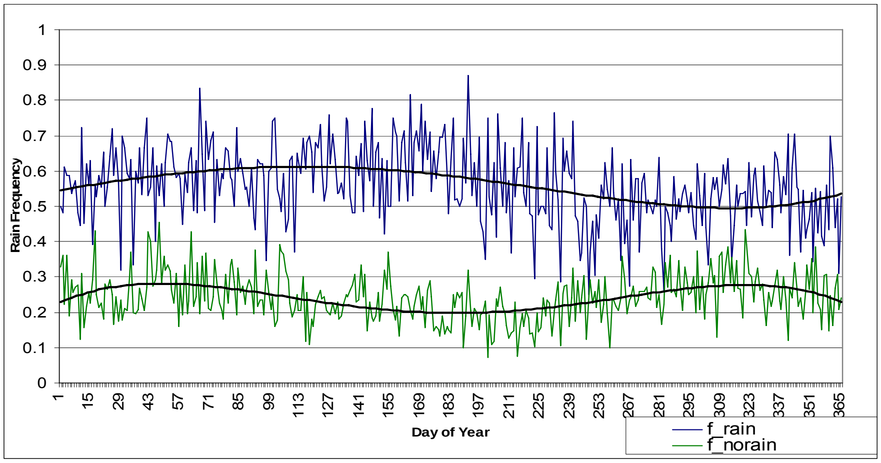

Figure 4.6: Seasonal Rainfall Frequency – Syd. Airport

This Markov model of the frequency process can then be used (after determining the parameter values) to predict the occurrence of rain on a daily basis over any time period out into the future. The table that follows shows the transition probabilities obtained for the Sydney Airport weather station:

	Rain	No rain
Rain	0.55	0.45
No rain	0.28	0.72

Table 4.1 Transition probabilities – Syd. Airport

4.3.3 Magnitude Process

The second part of the modelling process involves predicting the magnitude of the rainfall conditional on the fact that it rains on that day. We represent this size process by Y_t. Here we adopt a combination of procedures utilised by several authors where we segment the size distributions depending on the occurrence of rainfall on both the preceding day as well as the following day. This is facilitated by the fact that the frequency process is first used to predict the occurrence of rain over the entire contract period allowing the magnitude of the process to depend on the future frequency state.

The following table shows the basic statistics of the four segments

		Rain t+1	
Rain t-1		No	Yes
No	Average	3.74	7.20
	StdDev	6.21	12.23
	Max	79.0	182.1
	Min	0.1	0.2
Yes	Average	6.43	13.89
	StdDev	11.22	21.46
	Max	132.6	216.2
	Min	0.2	0.2

Table 4.2: Basic statistics of the four segments – each is be fitted with its own gamma distribution – (units: ˚C)

Gaussian distributions were fitted to the segments via maximum likelihood estimation of the two parameters (α and β) that make up the distribution given by the form:

$$f(x) = \frac{x^{\alpha-1}.e^{\frac{-x}{\beta}}}{\beta^{\alpha}\Gamma(\alpha)} \tag{4.2}$$

The chart below shows the differences in the parameters of the four gamma functions fitted to the segmented distributions. Table 4.2 and figure 4.7 show the different statistical properties of the four segments and underscores the importance of this classification. Possibly, the N/R/R and the R/R/N segments could be combined if one wished to simplify the model as their parameters are relatively close

Gamma distribution parameter values

	NRN	RRN	NRR	RRR
alpha	0.58	0.45	0.47	0.49
beta	19.32	12.09	13.07	6.39

Table 4.3: Parameter values of gamma distribution - estimated by least squares.

These four distributions are shown graphically in the following picture:

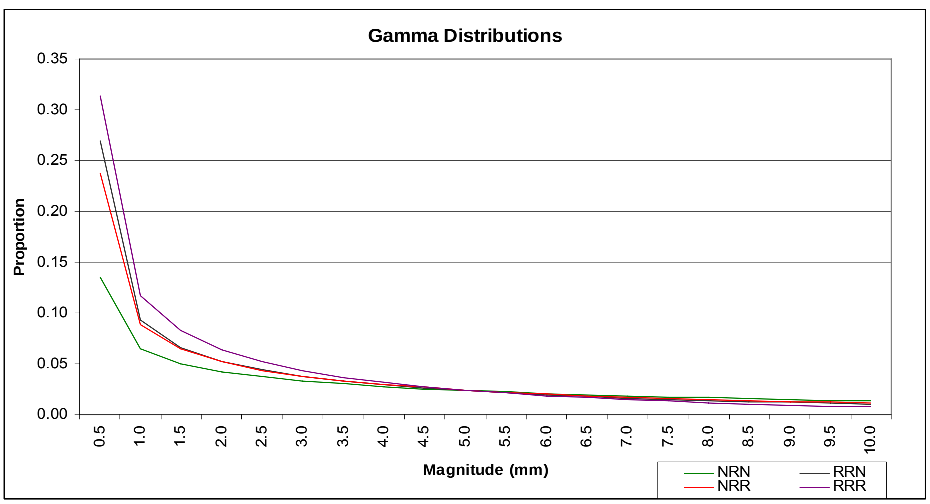

Figure 4.7: Gamma Distributions – Sydney Airport

Figure 4.8 below shows a detailed section of the actual vs expected chart for the no rain / rain / no rain segment. It is evident from the graph that the model fit is satisfactory, even well into the tail of the distribution. Similar results were obtained for the three other segments that represent the range of possible situations for the magnitude process. These four gamma distributions now become the basis for the magnitude process that is simulated in the following section.

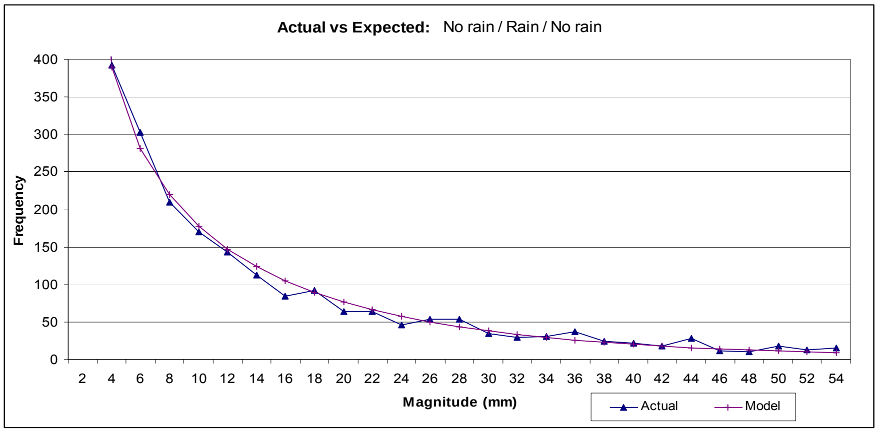

Figure 4.8: Model Fit - Gamma Distribution – No rain / rain / no rain

4.3.4 Simulation

Once parameters have been estimated for both the frequency process, X_t and the magnitude process, Y_t we are free to simulate sample precipitation values and compare them with historical results.

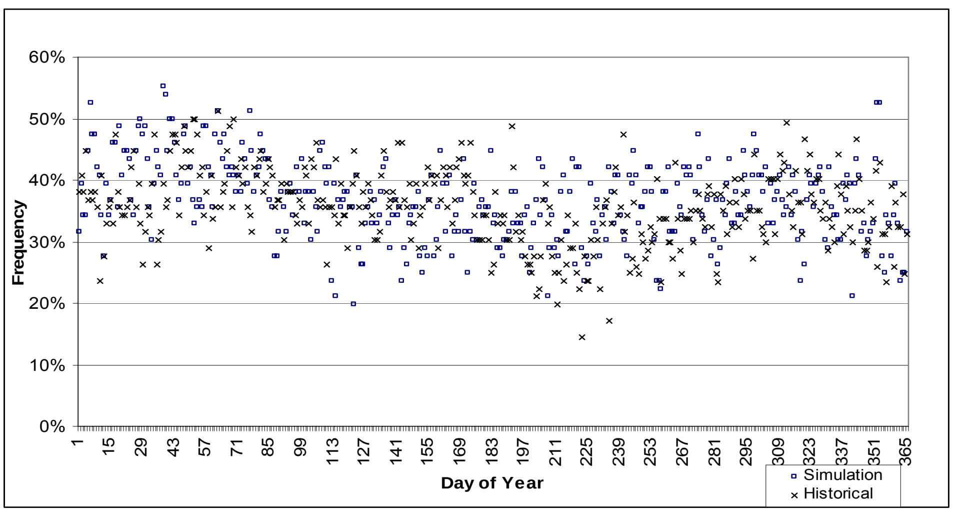

Figure 4.9: Sample simulated seasonal frequency vs Historical – 76 years.

Simulated over 76 years the following histograms show the distribution of the magnitude of rainy days as predicted by the model against the actual values recorded at Sydney Airport. model for the months of February, March and April.

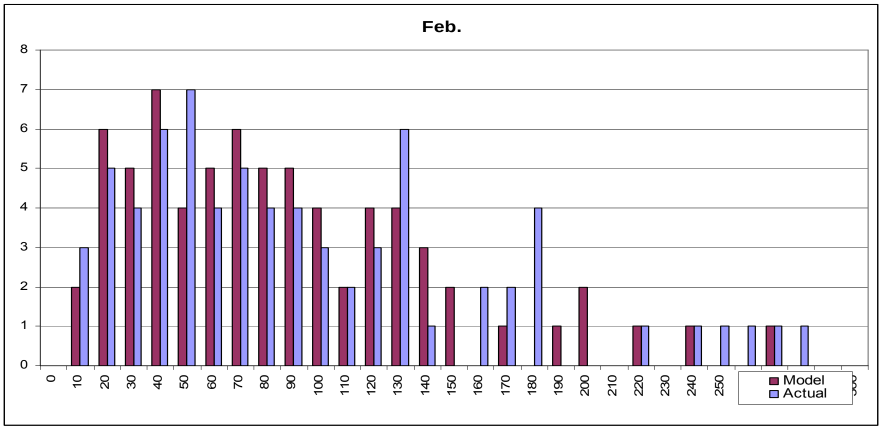

Figure 4.10: Sample simulated rainfall (mm) – Sydney Airport – February - 76 years.

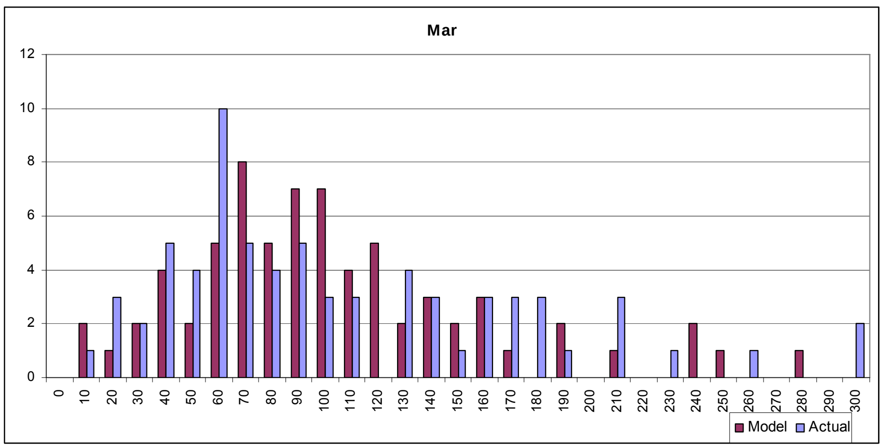

Figure 4.11: Sample simulated rainfall (mm) – Sydney Airport – March - 76 years.

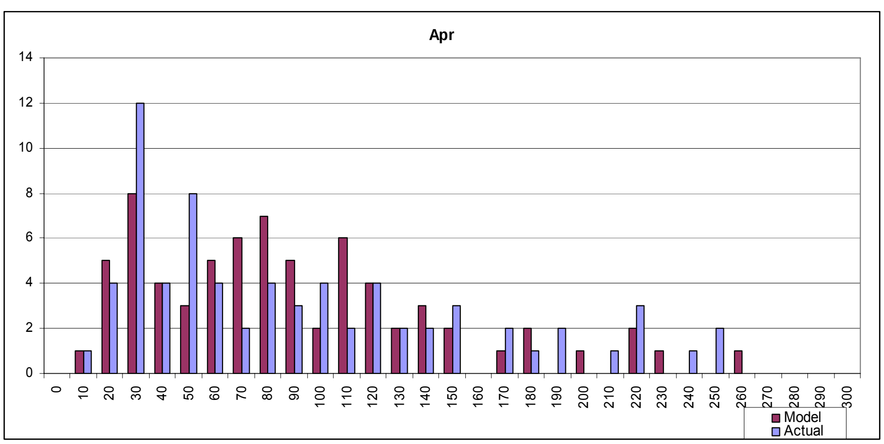

Figure 4.12: Sample simulated rainfall (mm) – Sydney Airport – April - 76 years.

Visually, the fit appears to be acceptable although the simulated rainfall frequency appears to be a little thin at the higher magnitudes. Figures 4.10-4.12 display these results.

4.4 Pricing Example

To illustrate these principles we will calculate the price of a series of options using a burning cost approach. This will be undertaken for the nearby weather stations, namely Sydney Airport and Observatory Hill, so that the extent of the geographic basis risk can be appreciated.

The specifics of the options to be priced are:

Period:	February
Measure:	Cumulative Monthly Rainfall
Exercise Levels:	150 / 170 / 190 / 210 mm
Tick::	$10,000 /mm
Location:	Sydney Airport & Observatory Hill

The burning cost analysis used similar data sets for both of the weather stations, i.e. 1939 – 2005, even though data for the Observatory goes back significantly further than 1939.

Burning' Cost Summary

	Exercise Level												
	120	130	140	150	160	170	180	190	200	210	220	230	240
Sydney Airport	$357,632	$328,750	$302,908	$277,908	$253,908	$233,987	$216,724	$202,250	$187,776	$173,303	$159,461	$146,303	$133,171
Observatory Hill	$444,645	$406,566	$371,921	$342,447	$315,987	$293,987	$274,092	$254,658	$236,237	$219,263	$203,671	$189,197	$175,671

Table 4.4: Option prices – Sydney Airport vs Observatory Hill – 1939-2005

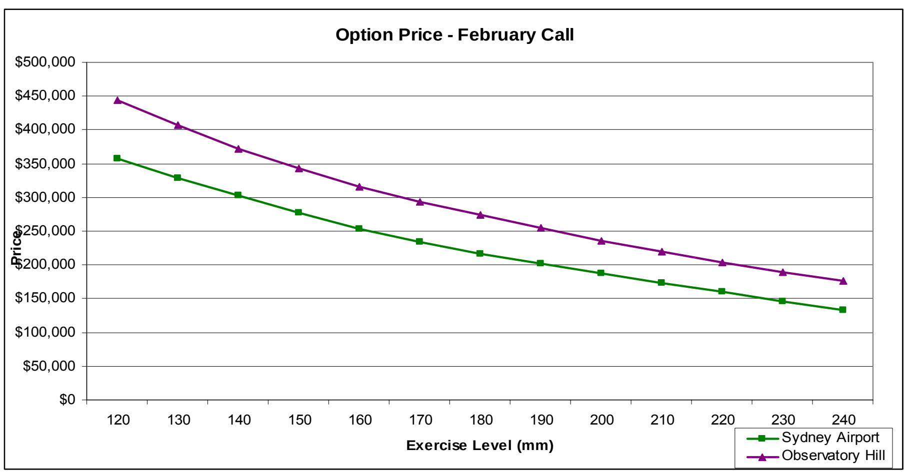

Figure 4.13: Option prices – Sydney Airport vs Observatory Hill – 1939-2005

This graph clearly shows that the cost of the Observatory Hill option is consistently greater that the similar option at Sydney Airport in what appears to be a fairly constant proportional difference.

4.5 Risk Management

Application 1: Hydro-Power Generation.

In 1999 the finance committee of Hydro-Quebec recommended the "adoption of a weather risk management approach using weather derivatives"[9] so that they could better manage the volatility of their earnings. Management recognised that the variation in runoff was the most significant risk to the future earnings of Hydro-Quebec as over 90% of their energy supply was sourced from hydropower.

To manage this risk the company entered into a derivative contract with Enron that consisted of a combination of a put and a call option, often referred to as 'long straddle' when the strategy is used in equity or interest rate markets. As can be seen in the following chart, the final payoff depended on 2 dimensions: the Rainfall and the Gas price as show in the following distribution:

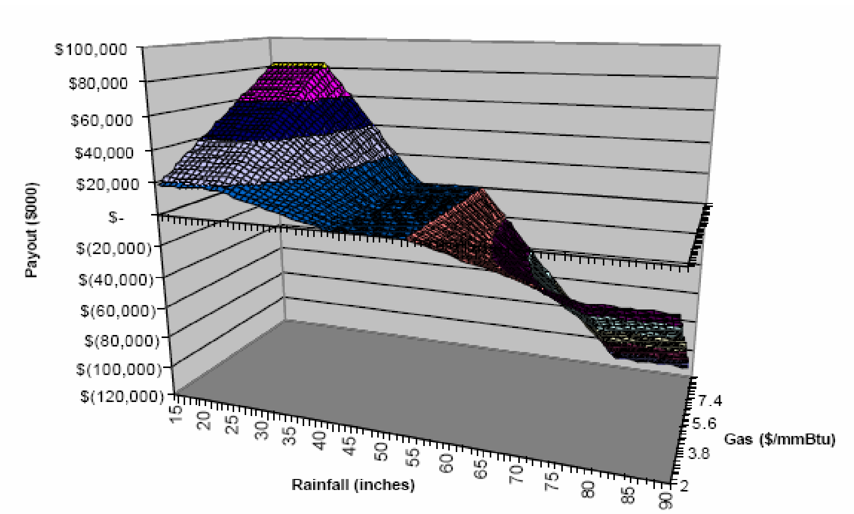

Figure 4.14: Pay-off structure - Hydro-Quebec Rainfall protection.

Whilst this is a more complicated option structure than investigated in the previous sections it serves to illustrate how these derivative contracts are structured in practice.

Application 2: Crop Protection

Traditional crop insurance has proven difficult to implement in a sustainable way and insurance companies covering these risks generally have very few options to transfer these risks onto reinsurers or other financial institutions. Some of the obstacles to a sustainable insurance solution include:

- High Administrative costs – heavily underwritten;
- Asymmetric Information – requires significant research and expertise;
- Moral Hazard – i.e. during economic downturns.
- Loss assessments need to be undertaken for each individual farm.

9 Hydro-Quebec's 1999 Annual Report

Weather derivative contracts overcome many of these difficulties by providing a transparent measure that can be used to obtain reliable prices that are available to the entire marketplace.

Developing Countries Risk Transfer:

Recently there have been many efforts made to attempt to use weather derivatives, or more precisely, indexed weather insurance to overcome many of the agricultural pressures that inflict developing countries. The World Band Commodity Risk Management group are currently investigating these areas and have developed the BASIX indexed insurance system which aims to provide a more effective way of insuring the income of some of the worlds poorest people. During 2003 and 2004 the pilot of the project entered 6 villages in the western region of the African country.

Essential index based insurance contracts are called such because their payout is not defined by a contract (as in a traditional insurance policy) but rather is linked to the value of a particular index. This means that the moral risk is nearly completely removed as the underlying index is calculated in a transparent and efficient manner. On top of this the BASIX system incorporates high-level insurance against extreme weather shocks which are reinsured on to the international market.

More recently, on the 6th of March this year, a derivative transaction was finalised that is being credited as the first humanitarian based derivative transaction. A leading insurance broker played a large part in negotiating a derivative contract between the United Nations World Food Program and a large French reinsurer, AXA Re. The contract called for the payment of $7.1 million in the event of an extreme drought, which was defined by strict measurement of well-established indexes.

Application 3: Event Insurance

This is a very recent development in the sphere of precipitation-based derivatives. Whilst event insurance has been available for some time now there has generally not been offered rainfall payouts or if they had the insurer (or indeed reinsurer) was generally unable to properly price these risks or transfer them to the broader financial markets.

Take as an example the operators of the Sydney Cricket Ground who offer a refund to patrons of a one-day cricket match where there are less than a certain number of overs bowled during the game. To hedge against the exposure to rain they could purchase a call options over a daily rainfall index so that, in general, the contract pays out on those days that refunds are given to the spectators. The challenge here would be to design the contract so that it matches their exposure as closely as possible, i.e. to minimise the basis risk between the derivative contract and the economic exposure they face. As a further example consider the operators of the Sydney Royal Easter Show that is held every year over a 2 week period during Autumn. Ticket sales for the show during the period are heavily reliant on the prevailing weather conditions and the operators could seek protection to both adverse temperature and rainfall outcomes that would see a reduction in the revenue received from the event.

Most event organisers would concede that weather conditions are the most unpredictable aspect in their attempt to bring together a large event and nearly all events would be in some way financially exposed to extremes in most (if not all) of the weather variables. As such, this field offers great potential as an avenue of expansion for these derivative contracts.

5: Where to from here?

5.1 Research

The level of academic literature is still sparse and much work needs to be done in order to enable the market to expand. It is essential for the growth of any new market that a transparent and reliable approach to pricing is put in place so that the participants can feel secure in their transactions. Thus, for the weather derivatives market to expand, a unified approach to pricing needs to be established for use by the growing number of practitioners in the marketplace.

The current approaches to pricing are highly data intensive and the models adopted are generally kept in-house by those organisations that spend significant resources in developing them. Affordable, standardised data warehouses are required so that there are no information barriers between the market participants. Europe faces the biggest challenges in term of achieving a central data collection methodology due to the vast range of constituent countries that make up the EU. There currently exists no rules governing the operation of weather stations throughout the EU and an individual station can change is method of recording or even shut down altogether without any formal notification. Despite many of these difficulties progress is being achieved, such as the formation of WeatherXchange, a joint venture between the UK Met office and a large broker to provide up-to-date historical information well as forecasting services.

During 2002 Euronext, a large European electronic exchange, released a series of temperature-based indices covering a variety of regions in France. Whilst raw historical data is not available the transparency with which these indices are calculated should provide a solid foundation for those seeking to transact temperature derivatives. In a recent development the Swedish Meteorological and Hydrological Institute and Energy-Koch Trading have teamed up to launch a Nordic Precipitation Index. This index is based on 17 stations, 9 in Norway, and 8 in Sweden. This index provides a reference so that companies with large precipitation exposures in this region can purchase protection

5.2 New Markets

The weather derivatives market in Australia is practically non-existent. Being a continent that relies so heavily on primary industries, Australia has much to gain out of a greater understanding of its financial exposure to the weather and ways in which these risks can be effectively managed. Most derivative products provide their purchasers with price protection (eg. cattle or wool futures) but weather derivatives are nearly unique in being able to provide protection from adverse quantity, even if the price paid or received is unchanged.

Whilst its uses in the agricultural sphere are potentially enormous it is first necessary to overcome the lingering issue of the sizeable basis risk that exists, particularly for rainfall-based contracts. These difficulties will be overcome through a combination of clever product design as well as improved knowledge of the dynamics of temperature and rainfall statistics. There is also an obvious lack of counterparties to many of the contracts that would be popular in Australia apart from the usual risk bearers and speculators such as insurers and large banks. Whilst in an urban environment it might be possible to find two natural counterparties for most of the rural weather risk it would be generally one-way with risk speculators required to take the other side.

5.3 New Interest

The potential for expansion of the weather derivative market globally is truly massive. More and more industries are realising the extent to which their returns are governed by variations in the weather and as has been highlighted throughout this paper there have been pioneers in many fields who have dived into this new risk management tools. Weather derivatives are unique in as far as derivative contracts are concerned as they combine aspects of both the financial and scientific worlds and the scientific accuracy of weather forecasters is proving significant in the approaches to pricing weather derivatives going forward.

Interest via weather-based index insurance is shaping up to be one of the major growth areas in the next 10 years from both an economic as well as humanitarian perspective. Overall, the convergence of the properties of insurance and derivative products is enabling a wider range of risks to be offered for coverage by risk ‘takers’ as well as facilitating a more efficient vehicle to transfer these risks to global financial markets.

References

Alaton, P. [2002]. *On Modelling and Pricing Weather Derivatives.* Applied Mathematical Finance, Volume 9 – Issue 1.

Benth, F.E. and Saltyte-Benth, J. [2005] *The Volatility of Temperature and Pricing of Weather Derivatives.* Centre of Mathematics for Applications, Oslo March 2005.

Black, F. and Scholes, M. [1973] *The pricing of Options and Corporate Liabilities*. Journal of Political Economy 81:637-654

Buchen, P. [2002]. *Options, Derivatives and Exotics.* 2002 Course Notes - School of Mathematics and Statistics, University of Sydney.

Cambell, S.D and Diebold, F.X, [2002] *Weather Forecasting for Weather Derivatives*. Wharton, The Rodney White Centre for Financial Research.

Cao, M., Li, A. and Wei, J. [2004] *Precipitation Modelling and Contract Valuation: A Frontier in Weather Derivatives*. The Journal of Alternative Investments – Fall 2004 (pp. 93-99)

Dornier and Queruel [2000]. *Pricing Weather Derivatives by Marginal Value*. Quantitative Finance – Volume 1.

Dunis C.,L. and Karalis V. [2003]. *Weather Derivatives Pricing and Filling Analysis for Missing Temperature Data*, Liverpool Business School and CIBEF.

Fama, E. F. [1965] *The Behaviour of Stock-Market Prices*. Journal of Business, vol. 38, pp. 34-105.

Garcia, A.F. and Sturzenegger, F. [2001] *Hedging Corporate Revenues with Weather Derivatives: A Case Study*. Universit´e de Lausanne Ecole des Hautes Etudes Commerciales. HEC - 2001

Jewson, S. and Zervos, M. [2003] *The Black-Scholes equation for Weather Derivatives.* RMS and Kings College – London. August 2003.

Jewson, S. and Penzer, J. [2004] *Weather Derivative Pricing and a Preliminary Investigation into a Decision Rule for De-trending.* London School of Economics.

Kubo, O. and Kobaysahi, Y. [2002] *Stochastic Modelling of Daily Rainfall for Pricing Weather Derivatives*. Hitachi Ltd.

McIntyre, R. [2001] *Weather Risk Management – Janruary 2001.* Enron Europe Finance & Trading Limited publication.

Moreno, M. [2002] Rain Risk, Speedwall Weather Derivatives Limited article.

Mraoua, M. and Bari, D. [2005] *Temperature Stochastic Modelling and Weather Derivatives Pricing: Empirical Study with Moroccan Data.* OCP Group / Maroc Meteo - October 2005.

Nicholls, M. [2002] *Silver lining for weather market*. Environmental Finance June 2002. http://www.environmental-finance.com/2002/0206jun/enron.htm

Price Waterhouse Coopers [2005]. *Weather Derivatives Survey – 2005.* Published on WRMA website. http://www.wrma.org/wrma

Swishchuk, A. [2005]. *Mean-Reverting Models in Financial and Energy Markets.* Colloquium, Dept. of Mathematics and Statistics, U of C.

Weil, Gotshal & Manges [2006]. *Participates in World's First Humanitarian Weather Derivative Transaction; Provides Novel Pro Bono Representation to United Nations World Food Programme* - Weil Gotshal Press Release - Monday, 6 March 2006.

Websites

http:// www.wrma.org Weather Risk Management Association, WRMA

http://www.climetrix.com Climetrix – An initiative of RMS. Pages: Market News, Market Overview, Weather Data

http://www.weatherderivs.com/ - Speedwall Weather Systems – Speedwall Derivatives.

Appendix A: Missing Values

The treatment of missing values is one of the most crucial steps in the data 'cleansing' process and as such a brief discussion of the typical approaches is warranted. It is common in any large time series that there will be missing values that can arise for a wide range of reasons. In respect of the data obtained from the Australian Bureau of Meteorology there are minimal gaps the majority of which appear during the periods of the great wars.

Temperature

The methods for filling missing data in temperature based time series are far more developed then for other variables in the main part due to the homogeneity in its geographical distribution. In other words, this uniformity means that information from near-by weather stations can be used to estimate values based on a historical relationship between the two stations.

Some of the common methods that are used as a basis for filling missing values include:

- Fallback methodology
- Naïve approach
- Expectation Maximization (EM) algorithm
- Data Augmentation (DA) algorithm
- Neural Networks Regression (NNR) models
- Principal Component Analysis (PCA).

The first two of these are well established and are generally used as benchmarks for the measuring of the other four. Dunis and Karalis [2003] undertook a comparison of the various methods for filling missing data from Philadelphia International weather station ('Fallback' of – Allentown, Bethlehem.) and concluded that the PCA method offered the most reliable results when tested in terms of a number of performance measures including:

- Mean squared error
- Root mean squared error
- Theil's Inequality Coefficient
- Mean absolute percentage error
- Mean error

Based on the results of Dunis and Karalis [2003] the PCA method is selected to fill missing temperature values contained in the BOM data.

Rainfall

The approach to filling missing data points in the precipitation data is different to that used for temperature data and at the present time there exists very little research on appropriate methods for correcting rainfall based time series data. Due to the discrete nature of rainfall, the measurements at near-by stations are much more loosely linked than with temperature readings and as such can not be as effective in relating to the missing values of the primary weather station. Here we use a more naïve approach than that used for temperature time series data.

As before we will rely on near-by weather stations to supplement the missing values except here we look to find 2 such backup weather stations to remove some of the impact of the discreetness of rainfall. If the backup stations have the required entries then the average of the two stations is recorded for the primary weather station.

For example, the following sample:

Date	Primary Station	Backup 1	Backup 2
1/02/2005	0	0	0
2/02/2005		0	0
3/02/2005		0	5.5
4/02/2005	0	0	0
5/02/2005		10.5	15.3
6/02/2005	1.5	5	2.3

after the application of the above procedure becomes:

Date	Primary Station
1/02/2005	0
2/02/2005	0
3/02/2005	2.75
4/02/2005	0
5/02/2005	12.9
6/02/2005	1.5

Appendix B: Temperature Distributions

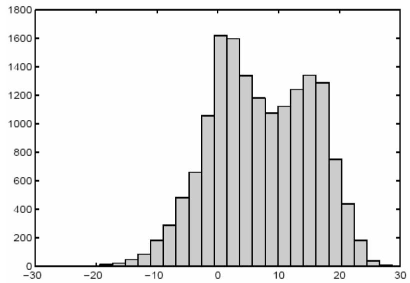

Figure B1: Histogram of daily average temperatures from Stockholm. (source: Benth et al [2005]) - (°C)

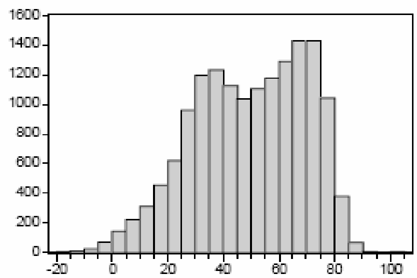

Figure B2: Histogram of daily average temperatures from Chicago. (source: Cambell and Diebold [2002]) - (°F)

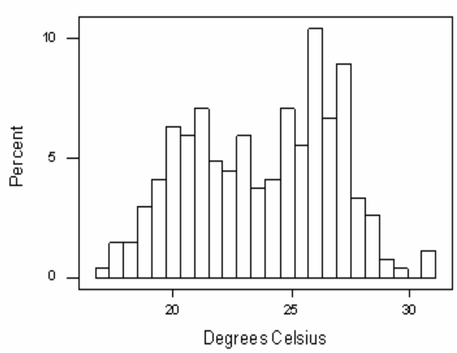

Figure B3: Histogram of daily average temperatures from St Loius. (1845-1978). (source: *Journal of Statistics Education* Volume 9, Number 1) - (deg C)