

Pricing Rainfall Based Futures Using Genetic Programming

Sam Cramer¹(✉), Michael Kampouridis¹, Alex A. Freitas¹,
and Antonis K. Alexandridis²

¹ School of Computing, University of Kent, Canterbury, UK
`sc649@kent.ac.uk`

² Kent Business School, University of Kent, Canterbury, UK

Abstract. Rainfall derivatives are in their infancy since starting trading on the Chicago Mercantile Exchange (CME) since 2011. Being a relatively new class of financial instruments there is no generally recognised pricing framework used within the literature. In this paper, we propose a novel framework for pricing contracts using Genetic Programming (GP). Our novel framework requires generating a risk-neutral density of our rainfall predictions generated by GP supported by Markov chain Monte Carlo and Esscher transform. Moreover, instead of having a single rainfall model for all contracts, we propose having a separate rainfall model for each contract. We compare our novel framework with and without our proposed contract-specific models for pricing against the pricing performance of the two most commonly used methods, namely Markov chain extended with rainfall prediction (MCRP), and burn analysis (BA) across contracts available on the CME. Our goal is twofold, (i) to show that by improving the predictive accuracy of the rainfall process, the accuracy of pricing also increases. (ii) contract-specific models can further improve the pricing accuracy. Results show that both of the above goals are met, as GP is capable of pricing rainfall futures contracts closer to the CME than MCRP and BA. This shows that our novel framework for using GP is successful, which is a significant step forward in pricing rainfall derivatives.

Keywords: Rainfall derivatives · Derivative pricing · Gibbs sampler · Genetic programming

1 Introduction

Rainfall derivatives fall under the umbrella concept of weather derivatives, which are similar to regular derivatives defined as contracts between two or more parties, whose value is dependent upon the underlying asset. In the case of weather derivatives, the underlying asset is a weather type, such as temperature or rainfall. The main difference between normal derivatives and weather derivatives is that weather is not tradable. Hence, typical pricing methods that exist in the literature for other derivatives are not suitable for weather derivatives.

In this problem domain the underlying asset is the accumulated rainfall over a given period, which is why it is crucial to predict rainfall as accurately as possible to reduce potential mispricing. Contracts based on the rainfall index are decisive for farmers and other users whose income is directly or indirectly affected by the rain. A lack or too much rainfall is capable of destroying a farmer's crops, hence their income. Thus, rainfall derivatives are a method for reducing the risk posed by adverse or uncertain weather circumstances. Moreover, they are a better alternative than insurance, because it can be hard to prove that the rainfall has had an impact unless it is destructive, such as severe floods or drought. Similar contracts exist for other weather variables, such as temperature.

Within the literature rainfall derivatives is split into two main parts. Firstly, predicting the level of rainfall over a specified time and secondly, pricing the derivatives based on different contract periods/length. Both aspects carry their own unique problems, with the former being a very hard time series to predict accurately, due to its volatility and redundancy of its reoccurring pattern. The latter part of rainfall derivatives constitutes an incomplete market¹. This means the standard pricing models such as the Black-Scholes model are incapable of pricing rainfall derivatives, because of the violation of the assumptions of the model; namely no arbitrage pricing. This paper focuses on pricing rainfall derivatives based on the predicted level of rainfall.

In order to predict the level of rainfall for rainfall derivatives, Markov-chain extended with rainfall prediction (MCRP) [1] and spatial-temporal rainfall (STR) models [2] have been used. More recently, Genetic Programming (GP) has been applied as an alternative predictive technique [3–5]. By predicting the underlying variable of rainfall, this increases the accuracy of pricing, which is crucial because contracts are priced ahead of time—up to a year ahead. Having the best possible predictive method reduces uncertainty in the market and boosts confidence in rainfall derivative pricing.

There is little literature on rainfall derivatives, due to being quite a new concept and rainfall being very difficult to accurately measure. The pricing techniques that have been applied so far are indifference pricing [6] and arbitrage free approach [7]. Both work in slightly different ways, with indifference pricing assuming the investor has a utility function which is a function of risk. The arbitrage free approach on the other hand is a method to change the measure of the underlying asset using the Esscher transform, taking the user from the real world to the risk neutral world through a probabilistic shift. Prior to contracts trading on the Chicago Mercantile Exchange (CME), indifference pricing was the initial technique [6,8], since contracts began trading in 2010, the arbitrage free approach has now become the standard pricing technique [7,9,10].

This paper derives contract prices for U.S.A. cities using the most recent pricing technique of the arbitrage free approach. As mentioned previously, we use the arbitrage free approach since the contracts have started to trade and are able to compare contract prices between different techniques against actual contract

¹ In incomplete markets, the derivative can not be replicated via cash and the underlying asset; this is because one can not store, hold or trade weather variables.

prices quoted on the CME. In order to derive contract prices, we require a technique that can maximise the accuracy of the prediction process before pricing. We use GP for three reasons, (i) GP has been shown to outperform the standard approach of MCRP in [3–5]. (ii) There is a correlation between predictive accuracy and pricing accuracy [11, 12], meaning that GP should improve pricing performance. (iii) We are able to define contract specific equations, rather than a single model representing an entire year.

In this paper we use GP to predict the level of rainfall for a selection of cities from around U.S.A, which has not been done before to the best of our knowledge. Moreover, we create a novel framework for calculating the derivative prices using GP to estimate the underlying variable of rainfall. In order to do so using the arbitrage free approach, we need to generate a probability density function (PDF) out of the deterministic equations generated by GP. We develop a range of strategies in order to create a PDF, that represents the rainfall process before translating these into the risk-neutral world using the Esscher transform. The strategies consist of taking a sample of predictions across the evolution period over several runs. Thus, the computational overhead of GP is reduced.

One potential issue using a sampling approach is not having sufficient samples to generate a PDF that can adequately reflect rainfall. To cope with this issue we use Markov chain Monte Carlo (MCMC) in order to estimate the population density from a sample of our top predictive models generated from GP. The use of MCMC helps to generate a PDF that we can use with confidence.

Additionally, we show the limitation of pricing over the year using a single model. The current framework within MCRP only allows for one model to replicate rainfall over the year. On the other hand, GP is capable of producing more flexible equations that can be specifically evolved for a certain contract period. Therefore, we hypothesise that having contract-specific models will further improve the pricing accuracy, over a single GP equation for an entire year.

Hence, the contribution of this paper is a novel framework for estimating the prices under risk-neutral conditions for GP. This allows us to explore the importance of having a more suitable and robust rainfall prediction method for deriving contract prices for rainfall derivatives. Using the proposed framework with and without contract-specific models in this paper, the prices can be compared against real market data and we can evaluate how well GP is able to price rainfall derivatives.

The paper is laid out as follows. In Sect. 2 we introduce the data used in our experimentation. In Sect. 3, we give an overview of how the Esscher transform is used to derive contract prices. In Sect. 4, we introduce the methodology proposed in this paper to translate the predictions generated by GP into rainfall derivative prices. In Sect. 5, we outline the experimental setup used in this work. In Sect. 6 we show the results for both the rainfall prediction process and the pricing steps. Finally, we conclude in Sect. 7.

2 Data

We focus on the data sets that have been used for pricing derivatives on the CME. Rainfall contracts are available for 10 cities in U.S.A.: Chicago, Dallas, Des Moines, Detroit, Jacksonville, Kansas, Los Angeles, New York, Portland and Raleigh. The contracts traded for these cities are futures and options for monthly or seasonal rainfall indices.

For the monthly index, the contract is defined by the accumulated amount of rainfall within each calendar month. Whereas, the seasonal contracts are defined by the accumulated amount of rainfall between two and eight consecutive months. The contracts themselves are only available between the months March through till October. The notional value of one rainfall contract is 50 USD per 0.1 index point, where 1 index point equates to 1 in. of rainfall.

To maximise the predictive performance of GP before pricing, we use the data transformation proposed in [3] to smooth the problem landscape. As an example, Fig. 1(a) shows the original daily rainfall data for Detroit, and Fig. 1(b) shows the result of the data transformation. By making the problem landscape of rainfall prediction a simpler problem, we can enhance the rainfall prediction accuracy, which is the first of the two steps to pricing.

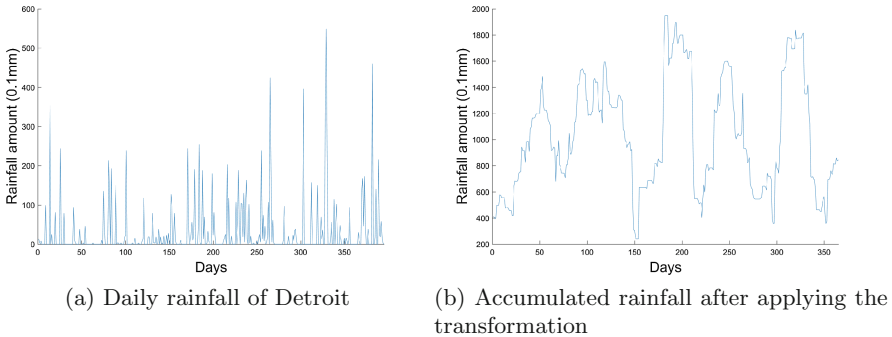


Fig. 1. A comparison between the daily rainfall time series of Detroit and after applying the data transformation from [3].

3 Pricing

One of the key characteristics of the weather derivative market is the nature of the incomplete market, whereby the underlying weather indices are non-tradable assets and can not be replicated by other risk factors. In other words, it is impossible to construct a riskless hedge portfolio containing the weather derivative. The standard approach is to price a futures contract $F(t; \tau_1, \tau_2)$ at time t with accumulation period $[\tau_1, \tau_2]$, by calculating the risk-neutral expectation Q of the rainfall index $I(\tau_1, \tau_2)$ with accumulation period $[\tau_1, \tau_2]$ based on the information set F_t available at time t . Therefore, the underlying variable is required to

calculate the index over an accumulation period. We can express the price of a futures contract by the following:

$$F(t; \tau_1, \tau_2) = \exp^Q [I(\tau_1, \tau_2)|F_t] = \exp^Q \left[\sum_{\tau=\tau_2}^{\tau_1} R_\tau |F_t \right]. \quad (1)$$

Our rainfall estimates $I(\tau_1, \tau_2)$ is considered the expected price under the canonical measure P , but are within the ‘risky’ world. Therefore, we require $Q \sim P$ such that all tradable assets in the market are martingales after discounting. Since the market is incomplete, there will exist many different martingales (Q), where it is impossible to find a unique risk-neutral measure Q [13, 14], such that Q is equivalent to the physical measure P . In order to calculate the arbitrage price under risk-neutral conditions, we require an equivalent martingale measure where $Q = Q_\theta$ using the Esscher transform [15], where θ represents the market price of risk (MPR). The MPR is the additional return or risk premium expected by investors for being exposed to undertaking the futures contract. When pricing with Black Scholes and similar pricing models, the unique equivalent martingale measure is obtained by changing the drift in the Brownian motion. The Esscher transform has been widely used across financial applications [16], more recently across rainfall derivative pricing [7, 10]. To use the Esscher transform we require estimating the type of distribution for our predictions. We then apply a constant MPR to transform our distribution to find the expected price under the risk-neutral measure Q_θ , where θ is calibrated to the market data. The transformation of probability density $f(x)$ of a random variable X to a new probability density $f(x; \theta)$ with parameter θ is the Esscher transform, given by:

$$f(x; \theta) = \frac{\exp(\theta x) f(x)}{\int_{-\infty}^{\infty} \exp(\theta x) f(x) dx} \quad (2)$$

4 Adapting Genetic Programming to the Esscher Transform

The motivation behind using Genetic Programming (GP) is threefold. Firstly, GP has not been applied to the pricing of rainfall derivatives, whilst it has been shown to improve the prediction against the currently used methods. Secondly, it has been noted [11, 12] that improving the prediction of the underlying variable leads to more accurate pricing. Therefore, by applying GP to the pricing domain, we aim to improve the pricing performance, which should help boost confidence in trading contracts. Finally, GP provides a flexible platform to predict rainfall contracts on a contract by contract basis. Therefore, we further tailor GP for maximising the pricing performance.

Before pricing, we need to perform an intermediate step in order for GP to calculate risk-neutral prices using the Esscher transform. One of the key aspects

of the Esscher transform is the probabilistic shift under $P \sim Q$ in order to find a unique equivalent martingale close to the predicted level of rainfall. This requires constructing a probability density function (PDF) out of the predictions generated. The current approaches in rainfall derivatives are stochastic processes that simulate unique rainfall pathways on each iteration. Despite GP being a stochastic algorithm, the output is a deterministic model and can not be used to estimate the expected index of rainfall. However, GP does generate many different equations to describe the rainfall process over the evolution process.

In order to recreate the outcome of a stochastic process (e.g. MCRP), we create a subset of rainfall equations generated from GP for every run of GP we perform. By building a large enough sample using many subsets of different rainfall equations, we can form a PDF of the expected level of rainfall for each day. The PDF that is generated can be manipulated to price under the risk-neutral density using the Esscher transform. Based on the nature of rainfall, we do expect to generate a non-gaussian distribution similar to the underlying data, which is assumed to follow either a gamma or mixed-exponential distribution [1].

4.1 Strategy for Prediction Selection

To generate a PDF requires having several different observations for the same time point, but we require a sufficient number in order to determine what form the distribution takes. Using MCRP, one would typically run the chain 10,000 times in order to generate sufficient samples, which for GP is unfeasible given the computational cost to run GP 10,000 times if we were to take the final prediction from each run. To reduce the overhead, we propose taking a sample of the best solutions from the final generation. Not only will this reduce the computational cost, but is a simple method to extract the required information. One concern is that taking too many samples from the final generation may reduce the fit of a distribution, by poor predictions being selected, which will heavily skew results. Thus, we must find the best possible balance.

We present a sample of results for a contract of March using various strategies in Fig. 2. Figures 2(a-c) show the PDF of choosing between 1, 5 and 10 individuals per GP run over 50 runs and Figs. 2(d-f) show the outcome over 100 runs instead. Therefore, Fig. 2(a) contains a total of 50 samples, whereas, Fig. 2(f) contains a total of 1000 samples. We choose the sample sizes based on avoiding longer runs of GP and reducing the risk of selecting too many extreme values that may exist from poor fitting solutions. We noticed that samples of 25 or larger posed a risk of selecting extreme values. What can be seen from the figures is the nongaussianity of the predictions, which we would expect and we do not wish to exceed 100 runs of GP, due to computational overhead. In some cases we do witness that GP seems to find a modal value for the prediction with a fairly narrow distribution. From a pricing perspective this is a positive sign as GP is able to determine what it believes to be the expected outcome.

Based on the PDF's generated, we notice that in several cases no clear distribution can be easily identified, shown in Fig. 3. This is sign that not enough

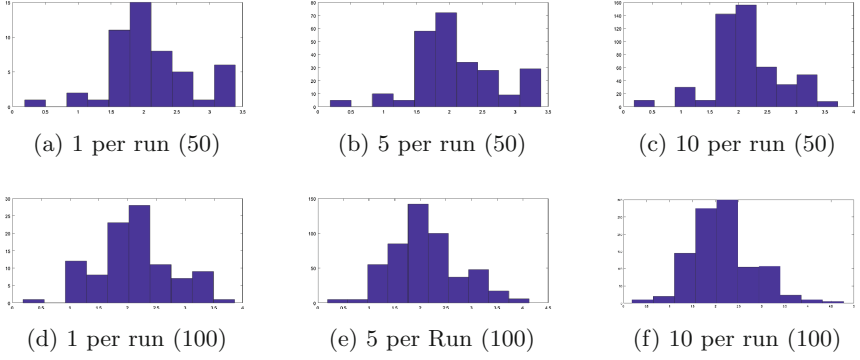


Fig. 2. The probability densities generated from GP for different strategies are shown for Detroit for the contract period of March (01/03/2011–31/03/2011). Values in brackets represents the number of GP runs, with the number per run showing how many samples are chosen to form the PDF at each run.




Fig. 3. Situations where a clear density can not be identified for contracts traded for Detroit in June (a), Jacksonville in June (b) and New York in April (c). Values in brackets represents the number of GP runs, with the number per run showing how many samples are chosen to form the PDF at each run.

samples have been generated and we would anticipate that generating more samples would lead to a clear distribution. However, we are attempting to reduce the computational overhead and necessity of running 10,000 separate GP runs. Hence, we employ Markov chain Monte Carlo (MCMC) to estimate the true parameters for the distribution we expect and can replicate the missing samples required to generate a PDF of our rainfall predictions.

4.2 Markov Chain Monte Carlo with Gibbs Sampling

The first key ingredient in Bayesian inference is the observation whose values are initially uncertain and described through a PDF. Another critical aspect is the previous belief about values of the parameter of interest, before observing the data. Bayesian theory is based on Bayes' Theorem, which allows new evidence to update beliefs through probabilities. Consider a random sample $x = (x_1, \dots, x_n)$ and the parameter of interest $\theta \in \Theta$ with Θ being the parameter space.

The likelihood function of θ is defined as: $f(x_1, \dots, x_n | \theta)$, the prior distribution $p(\theta)$ is the PDF before the observation of the value x . The inference is then based on the probability distribution of θ after observing the value of x , upon which information becomes available. We can then obtain the posterior distribution:

$$p(\theta | x_1, \dots, x_n) = \frac{\prod_{i=1}^n f(x_i | \theta) p(\theta)}{\int \prod_{i=1}^n f(x_i | \theta) p(\theta) d\theta} \propto \prod_{i=1}^n f(x_i | \theta) p(\theta). \quad (3)$$

In order to estimate the posterior distribution, using Eq. 3, we can use MCMC simulation when the posterior distribution is available. We can draw new samples of parameter $\theta = (\theta_1, \dots, \theta_p)$ directly from the joint posterior $p(\theta | x_1, \dots, x_n)$. We can estimate the joint posterior using Gibbs sampler, which is one type of MCMC to estimate the posterior. Gibbs sampling begins with an initialised vector of $\theta^0 = (\theta_1^0, \dots, \theta_p^0)$. At each iteration t , each component θ_j^t is sampled from the conditional distribution given all the other components of θ to generate a new vector of $\theta^t = (\theta_1^t, \dots, \theta_p^t)$. The sampling step of θ follows as:

$$\begin{aligned} \theta_1^t &\sim p(\theta_1 | \theta_2^{t-1}, \theta_3^{t-1}, \dots, \theta_p^{t-1}, x_1, \dots, x_n) \\ \theta_2^t &\sim p(\theta_2 | \theta_1^{t-1}, \theta_3^{t-1}, \dots, \theta_p^{t-1}, x_1, \dots, x_n) \\ &\vdots \end{aligned}$$

The sampling steps end once the last iteration has been reached, with sufficient iterations to achieve convergence. The predictive rainfall r_t of days of interest t follows an independent reparameterised Gamma distribution in the form of the mean and standard deviation of the initial rainfall predictions:

$$f(r_t | \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} r_t^{\alpha-1} e^{-\beta r_t}, \quad \alpha = \frac{\mu^2}{\sigma^2}, \quad \beta = \frac{\mu}{\sigma^2}. \quad (4)$$

Hence the parameters of interest of the likelihood distribution in Eq. 4 are the mean and the standard deviation parameters. The prior probability distributions are the same for both parameters of interest, note that they both have vague priors, the Uniform priors $U(0, 1)$. In order to estimate the posterior of the parameters of interest, we use JAGS [17], which is an iterative MCMC simulation method, using a Gibbs sampler described previously. We run a total of 50,000 iterations including 10,000 iterations being the burn-in period.

Figure 4 shows the density plots of the Markov chains for both the shape and the rate parameters of the Gamma distribution obtained by using JAGS. Note that each posterior density, all simulated Markov chains converge to stationary, shown by the clear peak. Hence, the number of iterations and burn-in period used are sufficient to achieve convergence of the Markov chains.




Fig. 4. The estimated kernel density of the predicted rainfall amounts and the density plots of our posterior estimates for the shape and rate parameter.

Figure 4 also shows the estimated kernel density of the predicted rainfall amounts by using the posterior means of the shape and the rate parameters in the Gamma distribution. We can see that the estimated density is representable of the target density of the rainfall prediction. This shows that using MCMC has assisted GP in creating a PDF that can be used for pricing, whilst minimising the overhead.

4.3 Pricing Using Estimated Densities

Now that the densities of our rainfall prediction under P have been estimated, we can apply the Esscher transform to shift our density to discover the expected price under the risk-neutral measure Q_θ . To do so requires one final step, which is to estimate the distribution generated and describe it under a non-gaussian distribution using Normal-Inverse Gaussian (NIG) [18]. NIG is a four parameter distribution, suitable for semi-heavy tails and skewness, which can be observed from our rainfall predictions. Applying the NIG will give us the expected price under Q , whereas, our previous steps gave us the expected price under P . We require the price under Q since this is the risk-neutral level. Moreover, NIG gives us a flexible framework to price under and maintains the statistical properties after changing the MPR (θ).

NIG has been used for several applications of risk-neutral modelling across a variety of financial problems. The NIG has a PDF in the closed form of:

$$f(x|\alpha, \beta, \mu, \delta) = \frac{\alpha \delta \exp(\delta \sqrt{\alpha^2 - \beta^2} + \beta(x - \mu))}{\pi \sqrt{\delta^2 + (x - \mu)^2}} K_1 \left(\alpha \sqrt{\delta^2 + (x - \mu)^2} \right), \quad (5)$$

where K_1 denotes the modified Bessel function of the second kind. α controls the steepness, μ the location, β the skewness and δ is the scaling of the distribution. By using the Esscher transform, we manage to keep the shape of the distribution and we modify the NIG under the Esscher transform with MPR (θ) becoming $\text{NIG}(\alpha, \beta + \theta, \mu, \delta)$. In order to estimate the parameters for NIG, we use an optimisation algorithm of Expectation-Maximisation (EM).

Table 1. The optimal configuration of GP/GA found by iRace. Parameters with a * are used by both the GP-part and GA-part of GP from [5].

GP/GA parameters			
Max depth of tree	9	Elitism*	2%
Population size*	1000	Number of generations*	90
Crossover*	96%	ERC negative low	-412.30
Mutation*	46%	ERC negative high	-201.63
Primitive	42%	ERC positive low	96.87
Terminal/Node bias	25%	ERC positive high	382.84
Set	Value		
Functions	ADD, SUB, MUL, DIV, POW, SQRT, LOG		
Terminals	11 r_t periods $\{r_{t-1}, r_{t-2}, \dots, r_{t-11}\}$, 10 r_y periods $\{r_{y-1}, r_{y-2}, \dots, r_{y-10}\}$, ERC, Constants in the range $[-4,4]$		

5 Experimental Setup

The purpose of our experiments is to compare the GP against the existing MCRP and burn analysis (BA) against contract prices as quoted by the Chicago Mercentile Exchange (CME). We use BA as it is the most frequently used benchmark in financial applications, where the expected prices under P are calculated based on the cost and payout of the same contract in the previous year. It computes the expected outcome over the accumulation period $I(\tau_1, \tau_2)$ with an additional risk premium that may occur. Therefore, $Q = P$ and the MPR is zero. BA can not price contracts on a daily basis, but acts as a reasonable benchmark.

For this paper we use the GP outlined in [5], which is a hybrid GP that decomposed the problem of rainfall prediction into smaller subproblems assisted by a Genetic Algorithm (GA). For details of this technique, hereafter referred to as GP/GA, see [3]. GP/GA has been shown to outperform the current approach of MCRP on European data sets. We tune the parameters for MCRP on the historical data and tune GP/GA using a package called iRace [19] on the training set. We present the optimal parameters along with the terminals and functions for GP/GA in Table 1.

Evaluating the predictive performance of the algorithms at predicting rainfall is crucial before we look at the pricing performance. We expect that the better the rainfall prediction the better the pricing performance. Therefore, we run the GP/GA algorithm and compare its rainfall predictive performance against MCRP, on all 10 data sets specified in Sect. 2. This is the first time GP has been applied to the data of U.S.A. cities that are used for rainfall derivatives.

We train GP/GA from 01/Jan/2001–31/Dec/2010 before testing on the unseen test set (01/Jan/2011–31/Dec/2011). Recall that Sect. 4.1 discussed the effect of different strategies. Based on previous experiments we discovered that

100 runs, saving the best 10 predictions per run, gave us the best performance in order to estimate a PDF that we can price under. Therefore, we report the average predictive performance over the 100 runs for each city. Additionally, we run MCRP 10,000 times using the mixed-exponential distribution to estimate the level of rainfall on a particular day.

6 Results

The performance of rainfall prediction of GP/GA and MCRP is shown in Table 2 based on the average RMSE performance on the testing set for each city.

Table 2. The average RMSE performance in tenths of mm for each of our approaches across each city. The best performance is shown in bold.

City	GP/GA	MCRP
Chicago	776.19	703.05
Dallas	524.05	1543.45
Des Moines	526.55	825.22
Detroit	653.42	498.86
Jacksonville	667.10	1098.33
Kansas	513.59	891.70
Los Angeles	321.97	941.00
New York	1087.66	987.21
Portland	530.04	636.07
Raleigh	561.43	829.17

GP/GA achieved the lowest RMSE over seven data sets compared to MCRP, which outperformed GP/GA three times. This is a very good result, showing that GP/GA predicts the underlying weather variable of rainfall more accurately than the most commonly used approach in the literature. This is essential to avoid problems of mispricing, which decrease the confidence within the price of derived contracts.

In order to determine whether the above results are statistically significant, we compare the two approaches by using the Wilcoxon signed-rank test. The null hypothesis that there is no significant difference between both approaches is rejected with a p value of 0.0195, which is less than the 5% level. Therefore, GP/GA provides more accurate rainfall predictions.

We now turn to pricing the contracts using GP/GA with MCRP and BA as benchmarks. Due to the limited availability of complete data from Bloomberg, we present the results of pricing for 3 of the above cities, Detroit, Jacksonville and New York over the monthly contracts for periods March to October. Based on the general performance of the algorithms from Table 2, we would expect MCRP to price more accurately in two out of these three cities. Table 3 shows

Table 3. The prices derived from GP/GA, MCRP and BA are shown with the comparison from the quoted prices on the CME for monthly contracts from March - October.

	Contract period	CME	GP/GA	MCRP	BA
Detroit	March	2.30	2.90	2.57	2.38
	April	2.70	2.99	3.05	2.88
	May	4.10	4.18	3.52	3.40
	June	3.50	3.23	3.46	3.55
	July	3.60	3.32	3.69	3.24
	August	3.00	3.06	3.60	3.31
	September	3.00	3.21	3.09	3.33
	October	2.40	3.44	2.91	2.53
Jacksonville	March	3.70	5.32	3.99	3.87
	April	2.40	2.55	2.87	2.84
	May	2.80	2.67	3.18	2.50
	June	7.50	6.39	5.72	6.16
	July	7.00	6.91	7.57	6.32
	August	7.00	4.10	8.29	6.44
	September	8.10	4.72	6.90	8.02
	October	2.60	2.92	4.19	4.03
New York	March	4.20	3.59	3.75	4.26
	April	4.40	3.68	3.96	4.30
	May	3.20	3.71	4.54	3.76
	June	5.00	4.37	4.69	4.42
	July	4.50	2.87	4.85	4.93
	August	4.30	4.46	4.53	4.52
	September	4.20	4.95	3.99	3.88
	October	4.60	6.06	3.98	3.90

the prices for each contract against the actual contract price traded on the CME. In order to generate the prices, we follow the methodology outlined in Sect. 4. Our goal is to price as close to the CME as possible.

Table 3 shows that GP/GA prices closer to the CME 9 times, while MCRP priced closer 7 times and BA priced closer 8 times. According to the Friedman test, at the 5% significance level there was no significant difference between approaches with a p value of 0.5818. This important result shows that GP/GA is able to price rainfall derivatives comparably to those on the CME. Interestingly, these findings support the hypothesis that better prediction leads to better pricing, which is shown in Table 2, where MCRP outperformed GP/GA for Detroit and New York for rainfall prediction. For each of those two cities MCRP priced closer to the CME than GP/GA. Similarly, GP/GA priced closer to the CME

than MCRP for Jacksonville. For the purpose of this paper we leave the MPR constant at 0, but it would be possible to shift our predictions in line with the contract prices quoted by the CME.

Please note these prices were calculated based on all available information up to 2011 (31/12/2010) and are the result of the historical average payoff. Usually with pricing, the price changes over time depending on whether more or less rainfall is expected and trader behaviour gets nearer to maturity. Unfortunately we are unable to track the prices due to the lack of data available.

One of the drawbacks to the current procedure of having a single rainfall equation to explain a year of rainfall is the difficulty on predicting the chaotic nature of rainfall time-series data. This creates models that can not capture the dynamics of rainfall over time, but instead capture the general trend over the year. Therefore, we propose moving away from a one-size-fits-all model and have contract-specific models.

To improve the pricing under Q , we propose building individual rainfall models for each contract. From the experimentation, we noticed that having separate models should increase the predictive accuracy, as we focus on a smaller subset of data. We split the training set into 3 month partitions, where the first two months are the months prior to the contract of interest. For example, predicting a contract in March would consist of the data from 01/Jan–31/Mar. We perform this partition for every year to train our model before testing on the same partition. We provide the results of RMSE for GP (GP/GA-P) and MCRP (MCRP-P) under this partition set up in Table 4, along with the new prices of GP. The prices for MCRP remain unchanged, because the model is still the same. This is one of the weaknesses of MCRP is that it can not be used to develop dynamic models for pricing, whereas GP can take full advantage of dynamic modelling.

Using the results from Table 4, we observe better predictive accuracy when partitioning the data into 3 month segments and the behaviour of rainfall is better explained. By partitioning the data we outperform MCRP in 16 cases in terms of rainfall prediction (66.67%), which shows that having a dynamic model is better and in no case do we perform worse than GP predicting over an entire year. Based on the Wilcoxon signed rank test, GP/GA statistically outperforms MCRP at the 5% significance level with a p value of 0.0278, showing that partitioning the data does statistically lead to better predictive accuracy. Moreover, this reflects in more accurate pricing: in 19 cases we price closer to the CME than GP/GA before partitioning. Table 5 shows the results of the Friedman test and Holm post-hoc test for the results in Table 4.

Table 5 shows that partitioning the data has a really positive effect on the accuracy of our pricing, whereby GP/GA-P achieves an average rank of 1.95, compared with BA, MCRP and GP/GA (2.40, 2.73 and 2.92 respectively). Note that partitioning the time series leads to a significant increase against GP/GA for pricing with a p value of 0.0101 at the 5% significance level. GP/GA-P does not outperform MCRP at the 5% significance level for the problem of pricing at the CME, but it does so at the 10% significance level.

Table 4. The predictive accuracy of GP/GA-P and MCRP-P, along with the prices for GP/GA and GP/GA-P and CME. The bold values represents a superior performance in either rainfall prediction or pricing performance.

City	Contract	Rainfall prediction		Contract prices		CME
		GP/GA-P	MCRP-P	GP/GA-P	GP/GA	
Detroit	March	540.62	575.27	2.24	2.90	2.30
	April	606.06	586.26	2.93	2.99	2.70
	May	562.86	607.70	4.18	4.18	4.10
	June	533.23	538.05	3.68	3.23	3.50
	July	648.17	484.01	3.58	3.32	3.60
	August	738.62	540.95	3.05	3.06	3.00
	September	791.66	857.00	2.86	3.21	3.00
	October	593.26	827.28	3.71	3.44	2.40
Jacksonville	March	459.87	1349.03	3.47	5.32	3.70
	April	361.95	1397.80	2.32	2.55	2.40
	May	570.78	1270.63	2.70	2.67	2.80
	June	535.54	1051.15	7.94	6.39	7.50
	July	571.43	660.87	7.13	6.91	7.00
	August	606.88	563.55	5.23	4.10	7.00
	September	772.97	760.28	6.05	4.72	8.10
	October	538.14	975.60	2.86	2.92	2.60
New York	March	611.27	343.46	3.52	3.59	4.20
	April	512.40	430.12	4.22	3.68	4.40
	May	462.04	658.33	2.69	3.71	3.20
	June	483.19	955.26	5.14	4.37	5.00
	July	1214.59	1510.32	5.21	2.87	4.50
	August	1885.99	1933.50	4.50	4.46	4.30
	September	1921.82	1718.51	4.61	4.95	4.20
	October	768.11	1458.09	4.05	6.06	4.60

One of the issues of having contract specific equations is the increase in complexity by having multiple equations to explain the rainfall process over the year. For each city we would have eight models explaining the data rather than one model. Future work should look at condensing the concept of GP/GA-P into a single model.

To summarise, GP/GA performs very well for predicting the underlying weather variable of rainfall and for pricing against MCRP. Moreover, we have shown the benefits of having a separate GP for each contract by partitioning the time series into shorter time frames. We witness a statistical improvement over GP/GA without partitioning and MCRP and outperform (but not statistically) BA, which should be the more accurate approach for the initial pricing

Table 5. The mean ranks of the four approaches, the Friedman test statistic and the respective p values for the Holm post-hoc test (using the best method (GP/GA-P) as the control method). Significant results are shown in bold.

Friedman p -value	0.0467		
Approach	Ranks	p value	Holm score
GP/GA-P	1.95	-	
BA	2.40	0.2404	0.0500
MCRP	2.73	0.0386	0.0250
GP/GA	2.92	0.0101	0.0167

of contracts. This is a very important step within pricing, being able to predict the level of rainfall better than current approaches and to price appropriately as well, with the data sets we have available to us.

7 Conclusion

This paper introduces a novel approach for dealing with the pricing of rainfall derivatives. Our novelty is the proposed creation of probabilistic models generated from Genetic Programming (GP), with the assistance of Markov chain Monte Carlo (MCMC). By developing this approach, we are able to price rainfall derivatives using the Esscher transform, a popular technique for calculating risk premiums. The motivation for this paper comes from the work of [3–5] where GP was used to predict the rainfall time series with a range of alternative approaches across European cities. However, the work did not present information on pricing. In this paper we show the effect of pricing under GP with the assistance of MCMC in order to create a probabilistic density function, which we are able to price under risk-neutral conditions.

We evaluate the performance of rainfall prediction on U.S.A. cities using a decomposition based Genetic Programming (GP/GA) [5]. We find sufficient evidence that the algorithm has a superior predictive power than the most currently used approach within the literature of rainfall derivatives. Based on the hypothesis that better prediction leads to better pricing [12], we would expect our model to perform better for pricing in those cities where GP/GA outperformed the standard approach. We find that there is evidence to suggest that the hypothesis is true, in the contracts for cities we had available to us. In an attempt to increase the pricing performance under GP/GA we proposed generating contract-specific models, rather than a one-size-fits-all approach. We found this to significantly increase predictive accuracy and subsequently the pricing performance.

Future work will include looking at GP to produce stochastic equations to describe the rainfall time series, this would replace the sampling strategy and MCMC before applying the risk-neutral densities. More analysis on the proposed contract specific approach in an attempt to maximise the performance. Finally,

upon the availability of rainfall derivative data, we can understand the pricing dynamics and how the prices change over time, which includes studying the effect that the market price of risk has on rainfall derivative contracts.

References

1. Wilks, D.S.: Multisite generalization of a daily stochastic precipitation generation model. *J. Hydrol.* **210**, 178–191 (1998)
2. Rodriguez-Iturbe, I., Cox, D.R., Isham, V.: Some models for rainfall based on stochastic point processes. *Proc. R. Soc. Lond. A: Math. Phys. Eng. Sci.* **410**(1839), 269–288 (1987)
3. Cramer, S., Kampouridis, M., Freitas, A.A., Alexandridis, A.: Predicting rainfall in the context of rainfall derivatives using genetic programming. In: 2015 IEEE Symposium Series on Computational Intelligence for Financial Engineering and Economics, pp. 711–718, December 2015
4. Cramer, S., Kampouridis, M., Freitas, A.A.: Feature engineering for improving financial derivatives-based rainfall prediction. In: Proceedings of 2016 IEEE Congress on Evolutionary Computation, Vancouver. IEEE Press, July 2016
5. Cramer, S., Kampouridis, M., Freitas, A.: A genetic decomposition algorithm for predicting rainfall within financial weather derivatives. In: Proceedings of the Genetic and Evolutionary Computation Conference 2016, GECCO 2016, pp. 885–892. ACM, New York (2016)
6. Carmona, R., Diko, P.: Pricing precipitation based derivatives. *Int. J. Theor. Appl. Financ.* **08**(07), 959–988 (2005)
7. Cabrera, B.L., Odening, M., Ritter, M.: Pricing rainfall futures at the CME. *J. Bank. Financ.* **37**(11), 4286–4298 (2013)
8. Leobacher, G., Ngare, P.: On modelling and pricing rainfall derivatives with seasonality. *Appl. Math. Financ.* **18**(1), 71–91 (2011)
9. Ritter, M., Mußhoff, O., Odening, M.: Minimizing geographical basis risk of weather derivatives using a multi-site rainfall model. *Comput. Econ.* **44**(1), 67–86 (2014)
10. Noven, R.C., Veraart, A.E.D., Gandy, A.: A lévy-driven rainfall model with applications to futures pricing. *Adv. Stat. Anal.* **99**(4), 403–432 (2015)
11. Jewson, S., Ziehmann, C., Brix, A.: *Weather Derivative Valuation*. Cambridge University Press, Cambridge (2010)
12. Alexandridis, A., Zapranis, A.: *Weather Derivatives: Modeling and Pricing Weather-Related Risk*. Springer, New York (2013)
13. Jensen, B., Nielsen, J.: Pricing by no arbitrage. In: Cox, D., Hinkley, D., Barndorff-Nielsen, O. (eds.) *Time Series Models: In Econometrics Finance and Other Fields*. Chapman & Hall/CRC/Taylor & Francis, New York (1996)
14. Benth, F.E., Benth, J.: *Modelling and Pricing Derivatives on Precipitation*, chap. 8, pp. 179–195. World Scientific (2012)
15. Bingham, N., Kiesel, R.: *Risk-Neutral Valuation: Pricing and Hedging of Financial Derivatives*. Springer Finance Textbooks. Springer, Heidelberg (2004)
16. Gerber, H., Shiu, E.S.W.: Option pricing by Esscher transforms. *Insur. Math. Econ.* **16**(3), 287 (1995)
17. Plummer, M.: JAGS: a program for analysis of Bayesian graphical models using Gibbs sampling. In: Proceedings of the 3rd International Workshop on Distributed Statistical Computing (2003)

18. Barndorff-Nielsen, O.E.: Normal inverse Gaussian distributions and stochastic volatility modelling. *Scand. J. Stat.* **24**(1), 1–13 (1997)
19. López-Ibáñez, M., Dubois-Lacoste, J., Stützle, T., Birattari, M.: The irace package: iterated racing for automatic algorithm configuration. Technical report, IRIDIA, Université Libre de Bruxelles, Belgium (2011)

Applications of Evolutionary Computation
20th European Conference, EvoApplications 2017,
Amsterdam, The Netherlands, April 19-21, 2017,
Proceedings, Part I
Squillero, G.; Sim, K. (Eds.)
2017, XXIV, 905 p. 268 illus., Softcover
ISBN: 978-3-319-55848-6