Pricing a Basket Derivative on Temperature and Precipitation: A Copula-Based Approach with an Application to the Kansas Winter Wheat Harvest

Jason Carranza Randolph Russell

April 2026

Abstract

We construct and price a multi-index weather derivative to hedge temperature and rainfall risk for winter wheat farmers in Wichita, Kansas during the harvest period from June 10 to July 10. Temperature is modeled via the modified Ornstein–Uhlenbeck framework of Alaton et al. (2002), while the seasonal rainfall index is fitted with a Weibull distribution following the index value simulation methodology from Mußhoff et al. (2006). To capture dependence between the cumulative average temperature (CAT) index and the cumulative rainfall (RF) index, we use a Frank copula, selected among several copula candidates by AIC and the K-S test, following the copula-based pricing framework of Bressan & Romagnoli (2021). Using 50,000 joint Monte Carlo scenarios, we price a weighted basket call option that pays in bad harvest conditions, specifically cool and wet weather. A standard independence copula underprices the option compared to our Frank copula, emphasizing the importance of incorporating the negative dependence between temperature and rainfall in the hedge.

Contents

1 Introduction 4

2 Data and Index Construction 5
- 2.1 Data Source . . . 5
- 2.2 Seasonal Index Construction . . . 6
- 2.3 Dependence Between CAT and RF . . . 7

3 Temperature Modeling 8
- 3.1 The Alaton Mean Structure . . . 8
- 3.2 Stochastic Component: Mean Reversion and Volatility . . . 9
  - 3.2.1 Estimation of Monthly Volatilities . . . 9
  - 3.2.2 Estimation of the Mean-Reversion Parameter . . . 10
- 3.3 CAT Simulation . . . 11

4 Rainfall Modeling 12
- 4.1 Daily Precipitation Models versus Index Value Simulation . . . 12
- 4.2 Candidate Distributions . . . 12

5 Copula Framework and Calibration 14
- 5.1 Motivation . . . 14
- 5.2 Pseudo-Observations . . . 14
- 5.3 Candidate Copulas . . . 15
- 5.4 Copula Selection . . . 15
- 5.5 Joint Simulation . . . 16

6 Basket Option Pricing 17
- 6.1 Contract Design . . . 17
- 6.2 Theoretical Pricing Formula . . . 18
- 6.3 Monte Carlo Results . . . 18
- 6.4 The Value of Modeling Dependence . . . 20

7 Limitations and Extensions 20

8 Conclusion 21

A Varying ω_1 and ω_2 24

B Different Strike Levels **24**

1 Introduction

Weather is an important factor for crop yields in agriculture. Poor weather conditions during important stages of crop growth can cause poor harvests, negatively affecting farmers' profits. Although traditional crop insurance offered by the USDA is available, it has issues that weather derivatives can alleviate. In the United States, taxpayers subsidize roughly 63 percent of crop insurance premiums offered by the USDA (Neuhofer & Rosch, 2019), and the programs are susceptible to adverse selection and moral hazard: risky producers are more likely to purchase coverage, and insured producers may alter their behavior once protected (Mußhoff et al., 2006).

Weather derivatives offer an alternative hedging mechanism to insurance. Since weather derivatives have payouts determined by weather indices, they eliminate the moral hazard present in indemnity-based insurance. Weather derivatives were first used in the late 1990s to hedge temperature risk (Alaton et al., 2002; Bressan & Romagnoli, 2021). Since then, weather derivatives have expanded to different indices such as snowfall, rainfall and wind speed.

Most of the academic literature on weather derivatives focuses on temperature-based products, using the Ornstein–Uhlenbeck model introduced by Alaton et al. (2002) and later expanded upon by Benth et al. (2007, 2008). Rainfall derivatives have received less attention, despite rainfall's potentially greater importance in agriculture (Mußhoff et al., 2006). Agricultural producers, however, are not exposed to just one weather factor, but to many all at once. Excessive rainfall and colder temperatures during the wheat harvest period can significantly reduce yield and quality (Smith & Goodwin, 2004), with excess rainfall being the more damaging of the two. Excess rainfall delays drying and promotes fungal diseases, and colder temperatures slow maturation and further delay drying. These risk factors are not independent; in Kansas, wetter weather often coincides with colder temperatures in the winter wheat harvest period.

A derivative written solely on temperature or rainfall ignores this dependence. The copula framework is a natural solution as it allows the distributions of temperature and rainfall to be modeled separately while capturing their dependence through a copula function (Nelsen, 2006; Bressan & Romagnoli, 2021).

In this paper, we construct and price a bivariate basket weather derivative tied to the winter wheat harvest window in Wichita, Kansas. Specifically, we:

(i) model daily temperature using the Alaton Ornstein–Uhlenbeck model and simulate the cumulative average temperature (CAT) index;

(ii) fit the cumulative rainfall (RF) index with a Weibull distribution using the index value simulation approach;

(iii) calibrate a Frank copula to capture the negative dependence between CAT and RF;

(iv) price a weighted basket call option via Monte Carlo simulation.

The remainder of this paper is organized as follows. Section 2 describes the data and index construction. Section 3 presents the temperature model. Section 4 discusses the rainfall marginal. Section 5 introduces the copula framework and calibration. Section 6 defines the basket option and reports the Monte Carlo price. Section 7 discusses limitations and extensions, and Section 8 concludes.

2 Data and Index Construction

2.1 Data Source

Daily weather observations were obtained from the National Oceanic and Atmospheric Administration (NOAA) for the Dwight D. Eisenhower National Airport station in Wichita, Kansas (station ID: USW00003928). The dataset spans from January 1, 1956 to April 10, 2026, comprising 25,668 daily records. The variables accessed were:

- daily maximum temperature, TMAX (°F),
- daily minimum temperature, TMIN (°F),
- daily total precipitation, PRCP (inches).

Daily average temperature was constructed as

$$\mathrm{TAVG}_d = \frac{\mathrm{TMAX}_d + \mathrm{TMIN}_d}{2}\,. \tag{1}$$

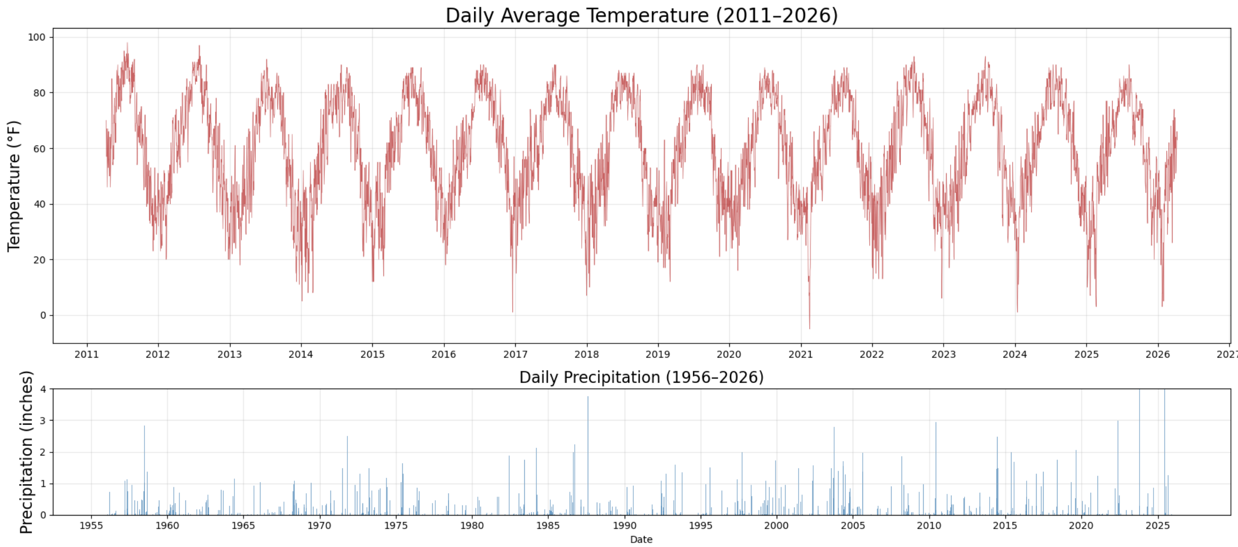

Figure 1: Daily average temperature and rainfall in Wichita, Kansas.

2.2 Seasonal Index Construction

We target the 31-day winter wheat harvest window from June 10 to July 10. Denote this set of calendar days by $\mathcal{W}$. For each year y in the sample, we compute two seasonal indices:

$$\mathrm{CAT}_y = \sum_{d \in \mathcal{W}} \mathrm{TAVG}_{y,d}\,, \qquad \mathrm{RF}_y = \sum_{d \in \mathcal{W}} \mathrm{PRCP}_{y,d}\,. \tag{2}$$

The CAT index aggregates temperature exposure over the harvest window, while the RF index aggregates total precipitation. Across 70 usable years (1956–2025), the key summary statistics are reported in Table 1.

Table 1: Summary statistics for the seasonal CAT and RF indices, 1956–2025.

Statistic	CAT (°F·days)	RF (inches)
Mean	2448.67	4.76
Standard deviation	87.69	2.60
Minimum	2306 (1995)	0.28 (1974)
Maximum	2665 (1990)	10.77 (2004)
Skewness	0.317	0.525

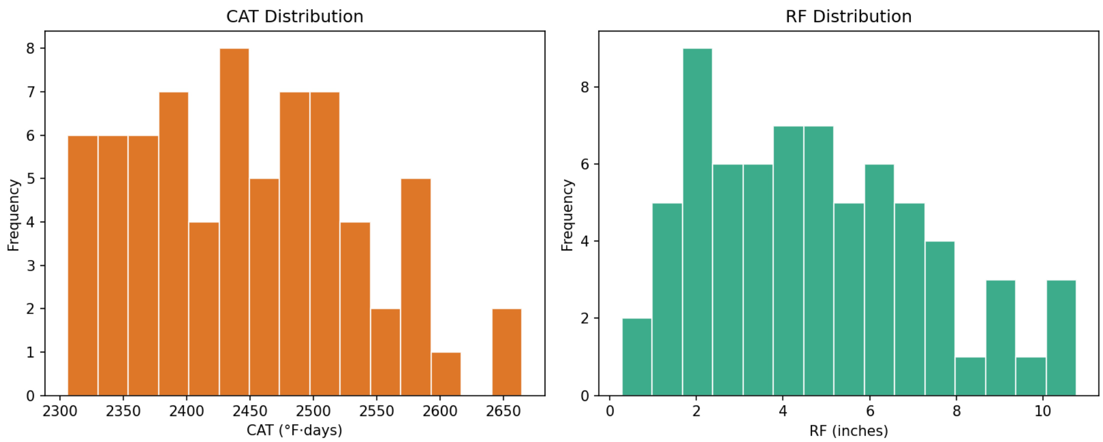

Figure 2: Historical distributions of CAT and RF during the harvest window.

2.3 Dependence Between CAT and RF

Two measures of association between the indices are computed:

$$\hat{\rho}(\mathrm{CAT}, \mathrm{RF}) = -0.5212\,, \qquad \hat{\tau}(\mathrm{CAT}, \mathrm{RF}) = -0.3929\,. \tag{3}$$

The negative dependence confirms that cooler seasonal windows tend to be wetter, which is precisely what we want to hedge against.

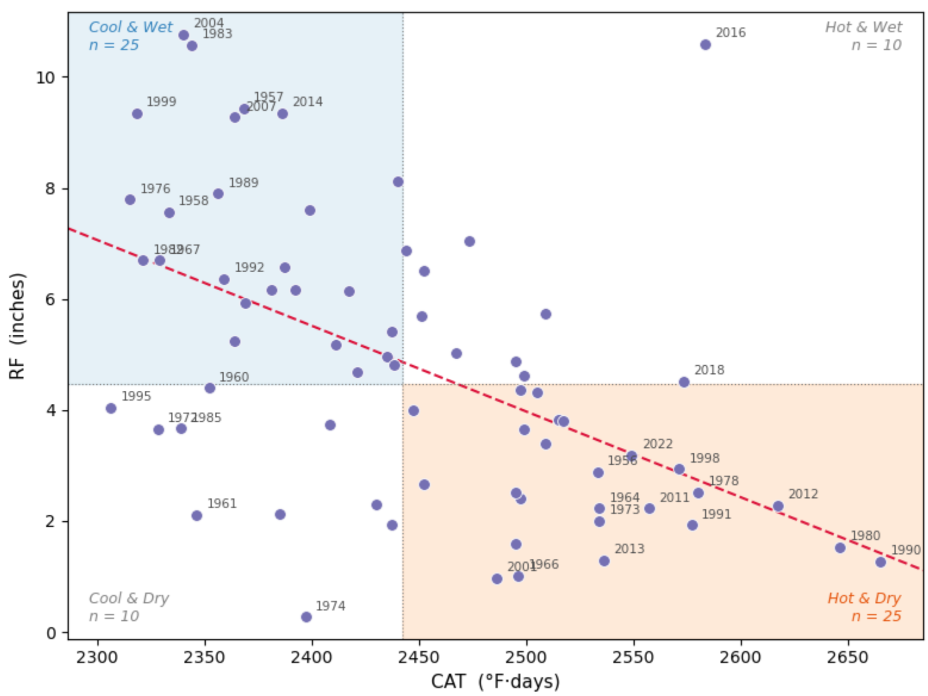

Figure 3: Scatterplot of CAT vs RF in the harvest period exhibiting negative dependence.

3 Temperature Modeling

3.1 The Alaton Mean Structure

Following Alaton et al. (2002), we model the deterministic seasonal mean temperature as

$$T_t^m = A + Bt + C\sin(\omega t + \varphi)\,, \tag{4}$$

where t is measured in days from the beginning of the dataset and $\omega = 2\pi/365.25$ is the angular frequency adjusted for leap years.

To estimate the parameters, we rewrite the sinusoidal component using the amplitude–phase decomposition:

$$C\sin(\omega t + \varphi) = a_3 \sin(\omega t) + a_4 \cos(\omega t)\,, \tag{5}$$

which renders the model linear in the unknown parameters $\boldsymbol{\xi} = (a_1, a_2, a_3, a_4)^\top$. Defining the observation vector $\mathbf{b} \in \mathbb{R}^n$ as the daily average temperatures and the design matrix $\mathbf{A} \in \mathbb{R}^{n\times 4}$ with rows

$$\mathbf{A}_{i,\cdot} = \big(1,\ i,\ \sin(\omega i),\ \cos(\omega i)\big), \qquad i = 1, \ldots, n = 25{,}668,$$

the least-squares problem

$$\hat{\boldsymbol{\xi}} = \arg\min_{\boldsymbol{\xi}} \|\mathbf{A}\boldsymbol{\xi} - \mathbf{b}\|_2^2 \tag{6}$$

yields the estimates

$$A = 55.8498\,, \quad B = 0.000120\,, \quad C = 24.6936\,, \quad \varphi = -2.0137\,. \tag{7}$$

The small positive trend B is consistent with the assumption of increasing temperatures. The amplitude $C \approx 25\,^\circ\mathrm{F}$ implies a roughly $50\,^\circ\mathrm{F}$ difference between the hottest and coldest seasons in Kansas.

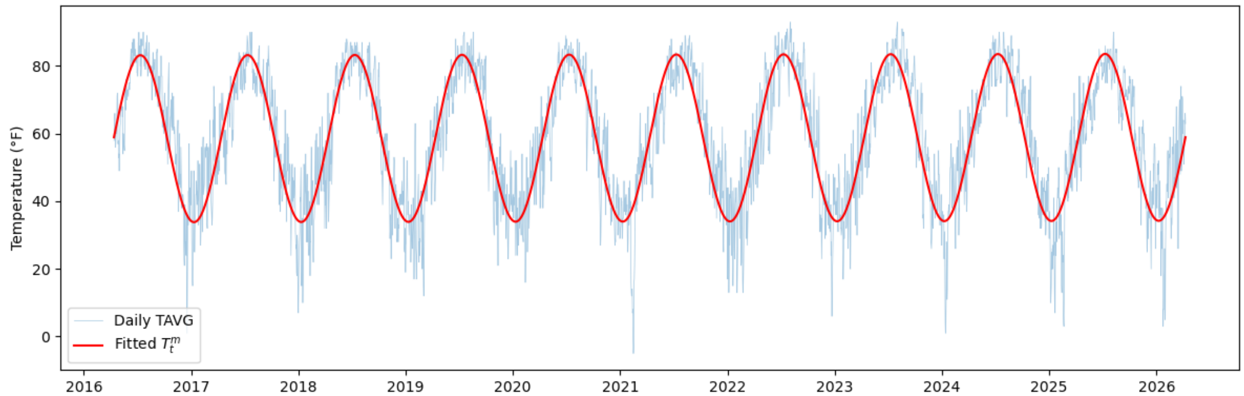

Figure 4: Historical temperature (2016-2026) with fitted seasonal mean overlaid.

3.2 Stochastic Component: Mean Reversion and Volatility

The stochastic component of temperature is modeled as a mean-reverting Ornstein–Uhlenbeck process. The naive SDE

$$dT_t = a(T_t^m - T_t)\,dt + \sigma_t\,dW_t$$

does not revert to the mean T_t^m in the long run (Dornier & Queruel, 2000). Following Alaton et al. (2002), the corrected SDE is

$$dT_t = \left[\frac{dT_t^m}{dt} + a(T_t^m - T_t)\right] dt + \sigma_t\,dW_t\,, \tag{8}$$

where

$$\frac{dT_t^m}{dt} = B + \omega C\cos(\omega t + \varphi)\,.$$

The solution to (8), starting at $T_s = x$, is

$$T_t = (x - T_s^m)\,e^{-a(t-s)} + T_t^m + \int_s^t e^{-a(t-\tau)}\sigma_\tau\,dW_\tau\,. \tag{9}$$

3.2.1 Estimation of Monthly Volatilities

The diffusion coefficient σ_t is assumed to be a piecewise constant function, taking a distinct value σ_μ for each calendar month μ. The quadratic variation estimator is

$$\hat{\sigma}_\mu^2 = \frac{1}{N_\mu}\sum_{j=0}^{N_\mu-1}(T_{j+1} - T_j)^2\,, \tag{10}$$

where N_μ is the total number of days in month μ across all years. The volatility estimates for the harvest window months are:

$$\hat{\sigma}_{\text{Jun}} = 4.5324\,, \qquad \hat{\sigma}_{\text{Jul}} = 4.1392\,. \tag{11}$$

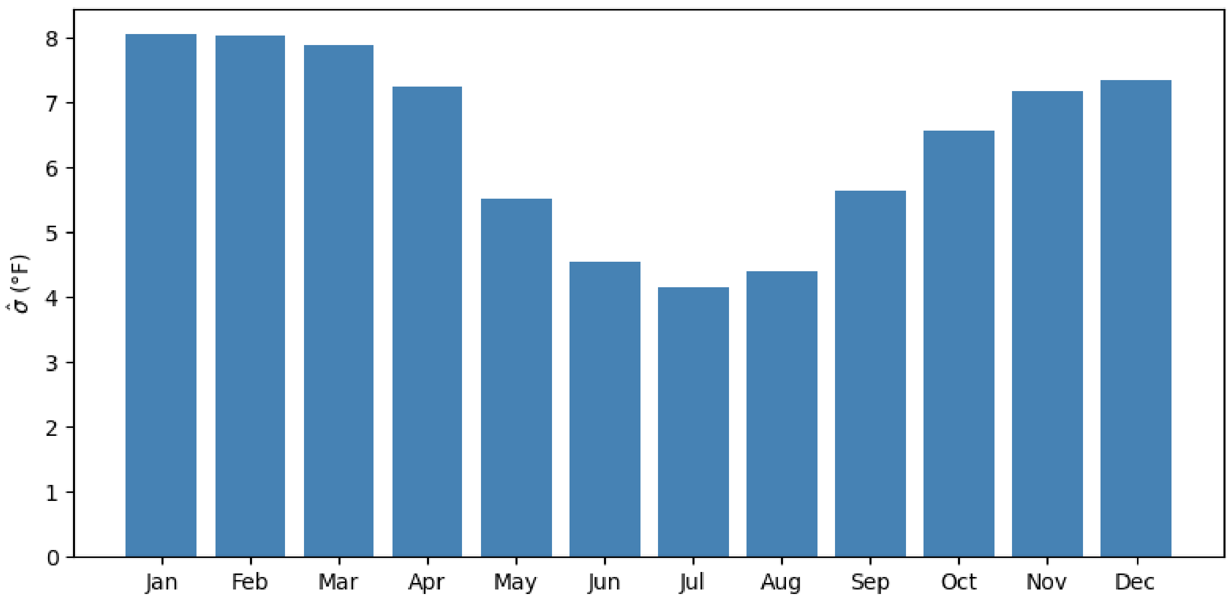

Figure 5: Bar chart of monthly temperature volatility estimates $\hat{\sigma}_\mu$.

3.2.2 Estimation of the Mean-Reversion Parameter

The mean-reversion speed a is estimated using the martingale estimation function method of Bibby & Sørensen (1995). Define

$$G_n(a) = \sum_{i=1}^{n} \frac{\dot{b}(T_{i-1}; a)}{\sigma_{i-1}^2} \big\{T_i - \mathbb{E}[T_i \mid T_{i-1}]\big\}\,, \tag{12}$$

where $b(T_t; a) = dT_t^m/dt + a(T_t^m - T_t)$ is the drift and $\dot{b}(T_t; a) = T_t^m - T_t$ is its derivative with respect to a. Using the conditional expectation

$$\mathbb{E}[T_i \mid T_{i-1}] = (T_{i-1} - T_{i-1}^m)\, e^{-a} + T_i^m\,,$$

and setting $G_n(\hat{a}_n) = 0$, we obtain the closed-form estimator

$$\hat{a}_n = -\log\left(\frac{\sum_{i=1}^{n} Y_{i-1}(T_i - T_i^m)}{\sum_{i=1}^{n} Y_{i-1}(T_{i-1} - T_{i-1}^m)}\right), \qquad Y_{i-1} \equiv \frac{T_{i-1}^m - T_{i-1}}{\sigma_{i-1}^2}\,. \tag{13}$$

Inserting the numerical values yields

$$\hat{a} = 0.3132\,. \tag{14}$$

3.3 CAT Simulation

With all parameters estimated, daily temperatures over the 31-day contract window are simulated using the Euler–Maruyama discretization of the Ornstein–Uhlenbeck process:

$$T_{i+1} = (T_i - T_i^m)\, e^{-a} + T_{i+1}^m + \sigma_i \sqrt{\frac{1 - e^{-2a}}{2a}}\, \varepsilon_i\,, \qquad \varepsilon_i \sim \mathcal{N}(0, 1)\,. \tag{15}$$

Using $N = 50{,}000$ simulated paths for the future 2026 contract window, the simulated CAT distribution has mean and variance

$$\mu_{\mathrm{CAT,sim}} = 2472.41\,, \qquad \sigma_{\mathrm{CAT,sim}} = 76.25\,, \tag{16}$$

which are close to the historical values: $\mu_{\mathrm{CAT,hist}} = 2448.67$ and $\sigma_{\mathrm{CAT,hist}} = 87.69$. As in Alaton, we treat the simulated CAT distribution as normal:

$$\mathrm{CAT} \sim \mathcal{N}(2472.41,\ 76.25^2)\,. \tag{17}$$

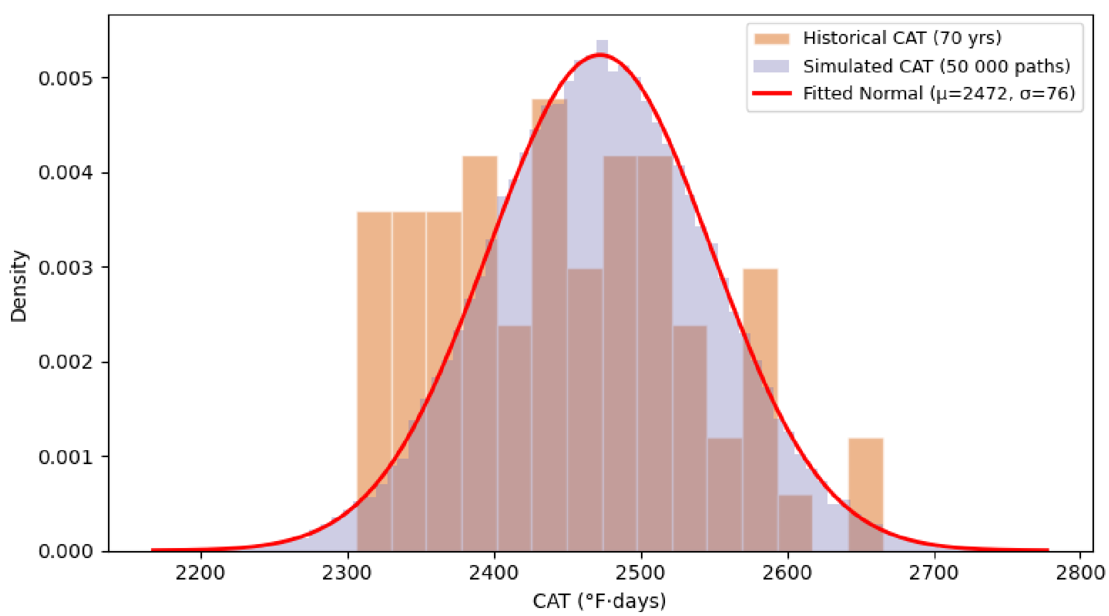

Figure 6: Histogram of simulated CAT distribution with fitted normal density, overlaid on the historical CAT histogram.

4 Rainfall Modeling

4.1 Daily Precipitation Models versus Index Value Simulation

A standard approach to modeling daily precipitation uses a first-order Markov chain for rainfall occurrence combined with a mixed exponential distribution for rainfall amounts (Mußhoff et al., 2006). The Bernoulli random variable $X_t \in \{0, 1\}$ indicates whether rainfall occurs in a given day t. X_t depends on the transition probabilities q_t^{01} (dry to wet) and q_t^{11} (wet to wet), and the rainfall amount y_t (assuming $X_t = 1$) is

$$f(y_t \mid X_t = 1) = \frac{\alpha_t}{\beta_t} \exp\left(\frac{-y_t}{\beta_t}\right) + \frac{1 - \alpha_t}{\gamma_t} \exp\left(\frac{-y_t}{\gamma_t}\right), \tag{18}$$

where α_t is the probability that a rainy day is "light" and $0 < \beta_t < \gamma_t$ represent the intensity of rainfall given the event.

However, a problem with daily simulations is the underestimation of the variance of rainfall over multi-week periods, leading to underpricing of rainfall derivatives by as much as 16 to 26 percent for put options Mußhoff et al. (2006). To avoid this issue, we instead follow the index value simulation approach. We fit a distribution to the RF index values instead of modeling daily rainfall.

4.2 Candidate Distributions

Two distributions were considered for the 70 years of RF observations:

1. **Gamma distribution** with shape k and scale θ:

$$f(x; k, \theta) = \frac{x^{k-1}\, e^{-x/\theta}}{\theta^k\, \Gamma(k)}, \qquad x \geq 0\,. \tag{19}$$

2. **Weibull distribution** with shape c and scale λ:

$$f(x; c, \lambda) = \frac{c}{\lambda} \left(\frac{x}{\lambda}\right)^{c-1} \exp\left[-\left(\frac{x}{\lambda}\right)^c\right], \qquad x \geq 0\,. \tag{20}$$

Using maximum likelihood estimation, we obtain the parameter values and goodness-of-fit metrics reported in Table 2.

Table 2: Fitted parameters and goodness-of-fit metrics for the RF index.

Metric	Gamma	Weibull
Parameters	$k = 2.8914$, $\theta = 1.6474$	$c = 1.9253$, $\lambda = 5.3720$
AIC	329.28	326.91
KS statistic	0.0670	0.0637
KS p-value	0.8909	0.9216

Both distributions fit the data well, as evidenced by the large KS p-values, but the Weibull distribution is selected due to its lower AIC value.

$$\mathrm{RF} \sim \mathrm{Weibull}(c = 1.9253,\ \lambda = 5.3720)\,. \tag{21}$$

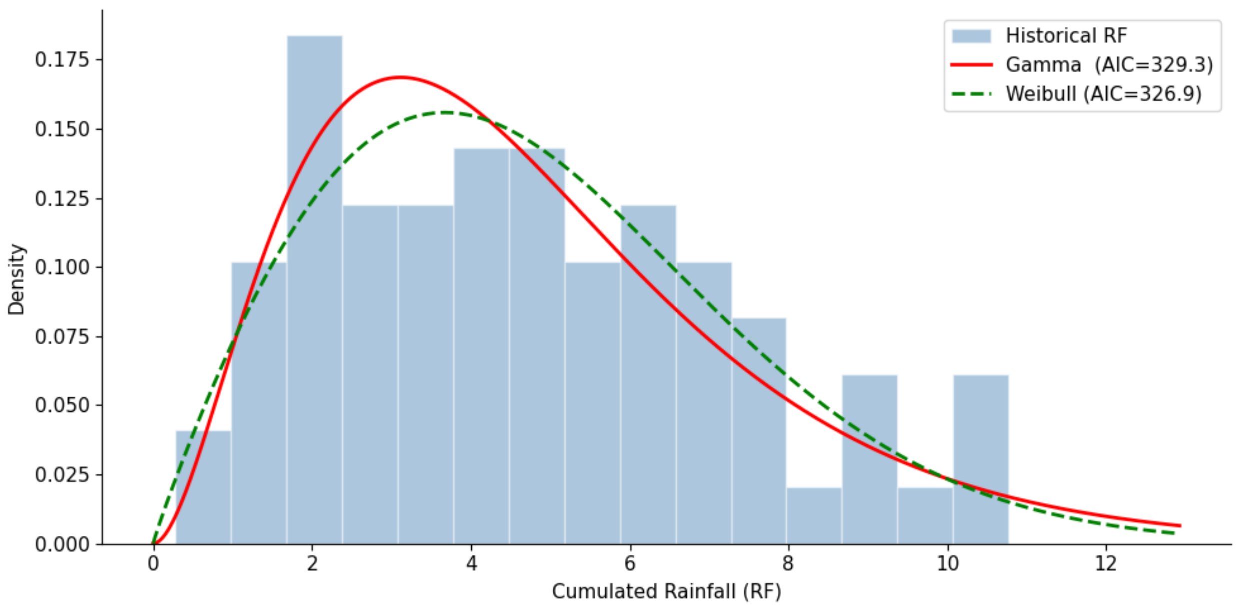

Figure 7: Histogram of historical RF index with fitted Weibull and Gamma densities. We choose Weibull.

5 Copula Framework and Calibration

5.1 Motivation

We estimated the normal CAT distribution and the Weibull RF distribution separately:

$$\text{CAT} \sim \mathcal{N}(2472.41, 76.25^2), \qquad \text{RF} \sim \text{Weibull}(1.9253, 5.3720).$$

The indices are not independent, however, as seen in their negative dependence (3). In order to take this dependence into account for more accurate pricing of the joint option, we use a copula. Following the approach of Bressan & Romagnoli (2021), we model the marginal distributions and the dependence structure separately. By Sklar's theorem, the joint distribution of CAT and RF can be written as

$$F_{\text{CAT,RF}}(x, y) = C(F_{\text{CAT}}(x), F_{\text{RF}}(y)), \tag{22}$$

where C is a copula. This allows the marginals to be combined with a dependence structure calibrated from the paired empirical data.

5.2 Pseudo-Observations

Let $(\text{CAT}_i, \text{RF}_i)$, $i = 1, \ldots, n$, denote the paired historical harvest-window observations. We transform the historical index pairs to pseudo-observations using two methods: first, using the fitted marginal CDFs (parametric method):

$$u_i = F_{\text{CAT}}(\text{CAT}_i), \qquad v_i = F_{\text{RF}}(\text{RF}_i). \tag{23}$$

and second, using a rank-based method:

$$u_i = \frac{\text{rank}(\text{CAT}_i)}{n+1}, \qquad v_i = \frac{\text{rank}(\text{RF}_i)}{n+1}. \tag{24}$$

The transformed observations (u_i, v_i) are approximately uniform on $[0, 1]^2$ in each coordinate, so any remaining structure in the transformed data is interpreted as dependence between CAT and RF rather than marginal behavior. We then fit candidate copulas using both methods by maximum likelihood.

5.3 Candidate Copulas

Following Bressan & Romagnoli (2021), we choose the Clayton, Gumbel, and Frank copulas as candidates. Because the empirical dependence is negative, standard Clayton and Gumbel copulas (which only model positive dependence) are not suitable. They are appropriately modified by evaluating at $(1-u, v)$, which is counterintuitively referred to as a 90°-rotation.

1. **Rotated Clayton (90°):** Captures lower-tail dependence in the original variables. The parameter satisfies $\theta = 2|\hat{\tau}|/(1-|\hat{\tau}|)$.

2. **Rotated Gumbel (90°):** Captures upper-tail dependence. In this case $\theta = 1/(1-|\hat{\tau}|)$.

3. **Frank copula:** A symmetric copula with no tail dependence, defined by

$$C(u, v; \theta) = -\frac{1}{\theta} \ln\left(1 + \frac{(e^{-\theta u} - 1)(e^{-\theta v} - 1)}{e^{-\theta} - 1}\right), \tag{25}$$

   where $\theta \in \mathbb{R} \setminus \{0\}$ is related to Kendall's τ through the Debye function:

$$\hat{\tau} = 1 - \frac{4}{\theta}\big[1 - D_1(\theta)\big], \qquad D_1(\theta) = \frac{1}{\theta}\int_0^\theta \frac{t}{e^t - 1}\, dt\,. \tag{26}$$

   The inversion from $\hat{\tau}$ to θ is solved numerically via Brent's method.

5.4 Copula Selection

For each candidate copula, the log-likelihood is evaluated on the pseudo-observations and the AIC is computed. The results are presented in Table 3.

Table 3: Copula selection results. Both methods agree on Frank as the best copula.

	Parametric		Rank-based	
Family	$\hat{\theta}$	AIC	$\hat{\theta}$	AIC
Rot. Clayton	0.6603	−12.63	0.8418	−14.60
Rot. Gumbel	1.4218	−19.97	1.5329	−17.08
Frank	−3.5863	**−22.93**	−4.0459	**−21.72**

Under both the parametric and rank-based approaches, the **Frank copula** achieves the lowest AIC. We proceed using parametric pseudo-observations with

$$C = \text{Frank}(\hat{\theta} = -3.586)\,. \tag{27}$$

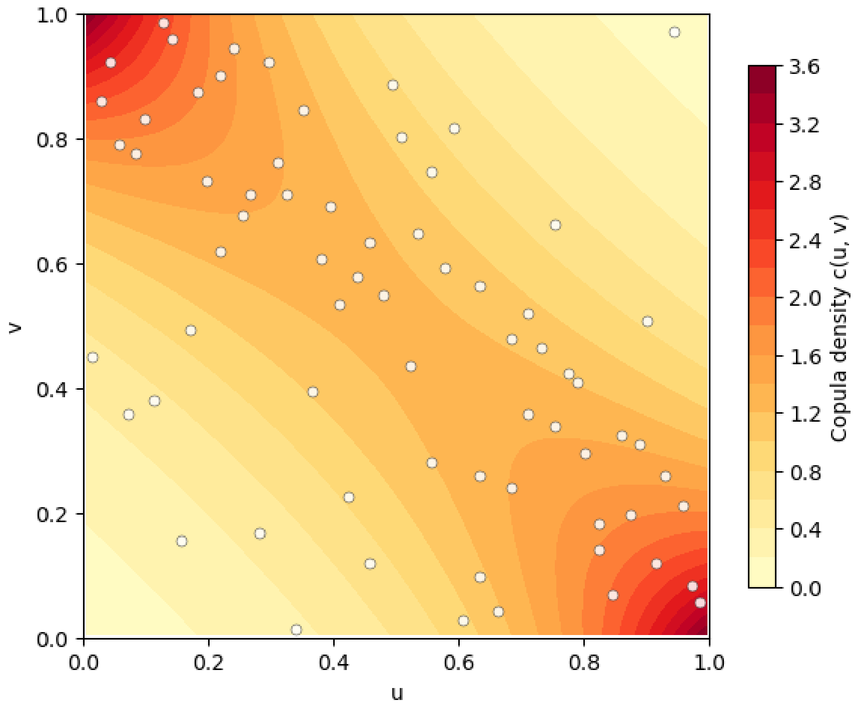

Figure 8: Density of the fitted Frank copula ($\theta = -3.586$) with rank-based pseudo-observations overlaid. Higher density along the off-diagonal corners reflects the negative dependence between CAT and RF.

5.5 Joint Simulation

Using the fitted Frank copula and the two marginal distributions (17)–(21), we generate 50,000 correlated $(\mathrm{CAT}, \mathrm{RF})$ scenarios via the following algorithm:

1. Draw two independent uniform random variables u and w from $\mathcal{U}(0, 1)$.

2. Compute the conditional inverse of the Frank copula to obtain a dependent pair (u, v):

$$v = -\frac{1}{\theta} \ln\left(1 + \frac{w(e^{-\theta} - 1)}{w + (1 - w)\, e^{-\theta u}}\right).$$

3. Apply inverse CDF transforms to obtain physical-scale values:

$$\mathrm{CAT} = F_{\mathrm{CAT}}^{-1}(u) = \mu_{\mathrm{CAT}} + \sigma_{\mathrm{CAT}}\, \Phi^{-1}(u)\,, \qquad \mathrm{RF} = F_{\mathrm{RF}}^{-1}(v)\,.$$

Validation confirms that the simulated samples reproduce both the target marginal statistics and the target Kendall's $\hat{\tau}$.

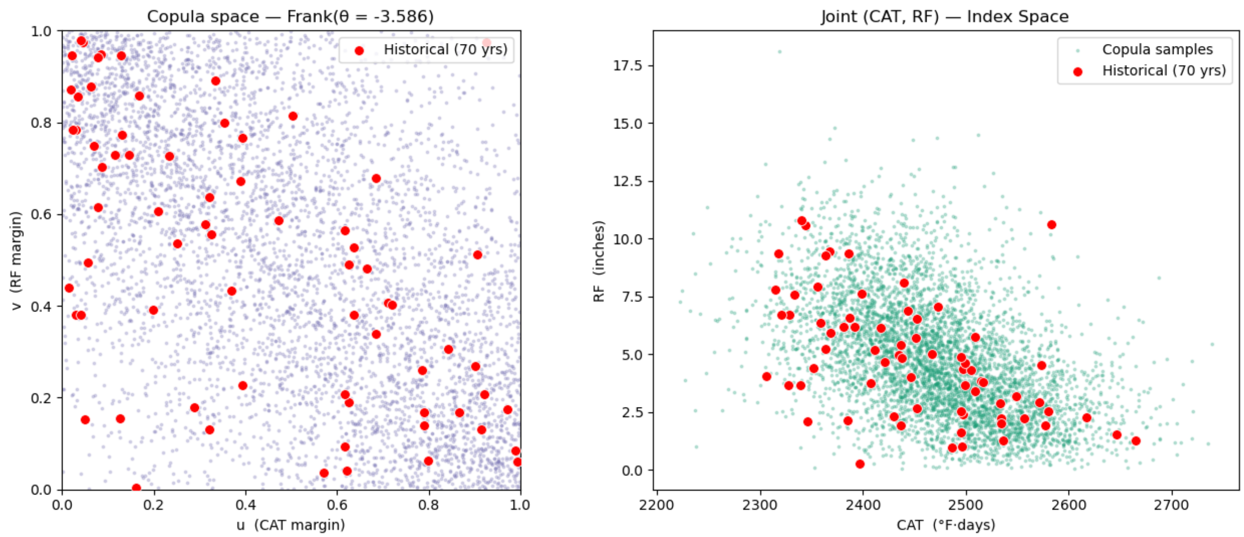

Figure 9: Scatter plots of 5,000 of the 50,000 simulated (CAT, RF) pairs drawn from the fitted Frank copula, with the 70 historical observations overlaid in red. (Left) Copula space showing uniform marginals linked by the estimated Frank parameter θ. (Right) Index space, with CAT in °F $\cdot$ days and RF in inches, obtained by applying the inverse marginal transforms. The clustering of historical points within the simulated cloud shows the suitability of the joint dependence structure.

6 Basket Option Pricing

6.1 Contract Design

Our derivative is designed to pay the holder when conditions are bad for the wheat harvest: when there is a lot of rain and temperatures are cool. Following the basket option framework described in Bressan & Romagnoli (2021), we construct the payoff in three steps:

1. **Standardize** both indices to z-scores using their model-implied moments:

$$\mathrm{CAT}_{\mathrm{std}} = \frac{\mathrm{CAT} - \mu_{\mathrm{CAT}}}{\sigma_{\mathrm{CAT}}}, \qquad \mathrm{RF}_{\mathrm{std}} = \frac{\mathrm{RF} - \mu_{\mathrm{RF}}}{\sigma_{\mathrm{RF}}}.$$

2. **Reverse** the CAT index so that high values correspond to cold conditions:

$$\mathrm{CAT}_{\mathrm{rev}} = -\,\mathrm{CAT}_{\mathrm{std}}.$$

3. **Form a weighted basket**, tilted toward rainfall ($\omega_1 < \omega_2$):

$$S = \omega_1\,\mathrm{CAT}_{\mathrm{rev}} + \omega_2\,\mathrm{RF}_{\mathrm{std}} \tag{28}$$

Under this construction, $S = 0$ corresponds to a "normal" summer, and large positive values of S indicate anomalously cool and wet conditions. The call option payoff is

$$\Pi = \max\bigl(S - K^B,\ 0\bigr) \tag{29}$$

where $K^B = \omega_1 K_{CAT} + \omega_2 K_{RF}$, and K_{CAT} and K_{RF} are the scaled strikes for temperature and rainfall, respectively.

6.2 Theoretical Pricing Formula

Under the assumption of a zero market price of weather risk and a constant risk-free rate r, the call option price at time t is

$$C_t = e^{-r\tau}\, \mathbb{E}[\Pi] = e^{-r\tau} \int_{K^B}^{\infty} (x - K^B)\, f_S(x)\, dx\,, \tag{30}$$

where f_S is the density of the basket index S. The CDF of S can be expressed via a C-convolution through the copula (Bressan & Romagnoli, 2021):

$$F_S(z) = \int_0^1 D_1\, C_{\bar{\omega}_1\mathrm{CAT},\, \bar{\omega}_2\mathrm{RF}}\Bigl(\omega,\ F_{\bar{\omega}_2\mathrm{RF}}\bigl(z - F^{-1}_{\bar{\omega}_1\mathrm{CAT}}(\omega)\bigr)\Bigr)\, d\omega\,. \tag{31}$$

Here, we evaluate the expectation in (30) via Monte Carlo simulation rather than numerical integration of (31).

6.3 Monte Carlo Results

The option price is computed as

$$\hat{C}_t = e^{-r\tau}\, \frac{1}{N} \sum_{i=1}^{N} \max\bigl(S^{(i)} - K^B,\ 0\bigr)\,. \tag{32}$$

Our simulation uses $r = 0.05$, $\tau = 9/12$ (nine months to the start of the contract window), and $N = 50{,}000$ scenarios. This τ is a reasonable choice, as 9 months is the approximate growing time of winter wheat. A farmer would likely buy an option like this for hedging, soon after his wheat has been planted. In this example, we further use $\omega_1 = 0.2, \omega_2 = 0.8$ and $K^B = 0.42$ (resulting from $K_{CAT} = 0.1$ and $K_{RF} = 0.5$, meaning the temperature index is 0.1 standard deviations below average and rainfall is 0.5 standard deviations above average). The standard error is

$$\mathrm{SE} = e^{-r\tau}\, \frac{\hat{\sigma}_\Pi}{\sqrt{N}}\,.$$

The resulting price is:

$$\boxed{\hat{C}_t = 0.2032\,, \qquad 95\%\text{ CI}:\ [0.1994,\ 0.2070]}\,. \tag{33}$$

The price is reported in standardized units. A notional value ν (dollars per unit of S) converts the price to a real premium: Premium $= \nu \times \hat{C}_t$. The farmer chooses ν to match their estimated per-unit crop loss from one standard deviation of adverse weather.

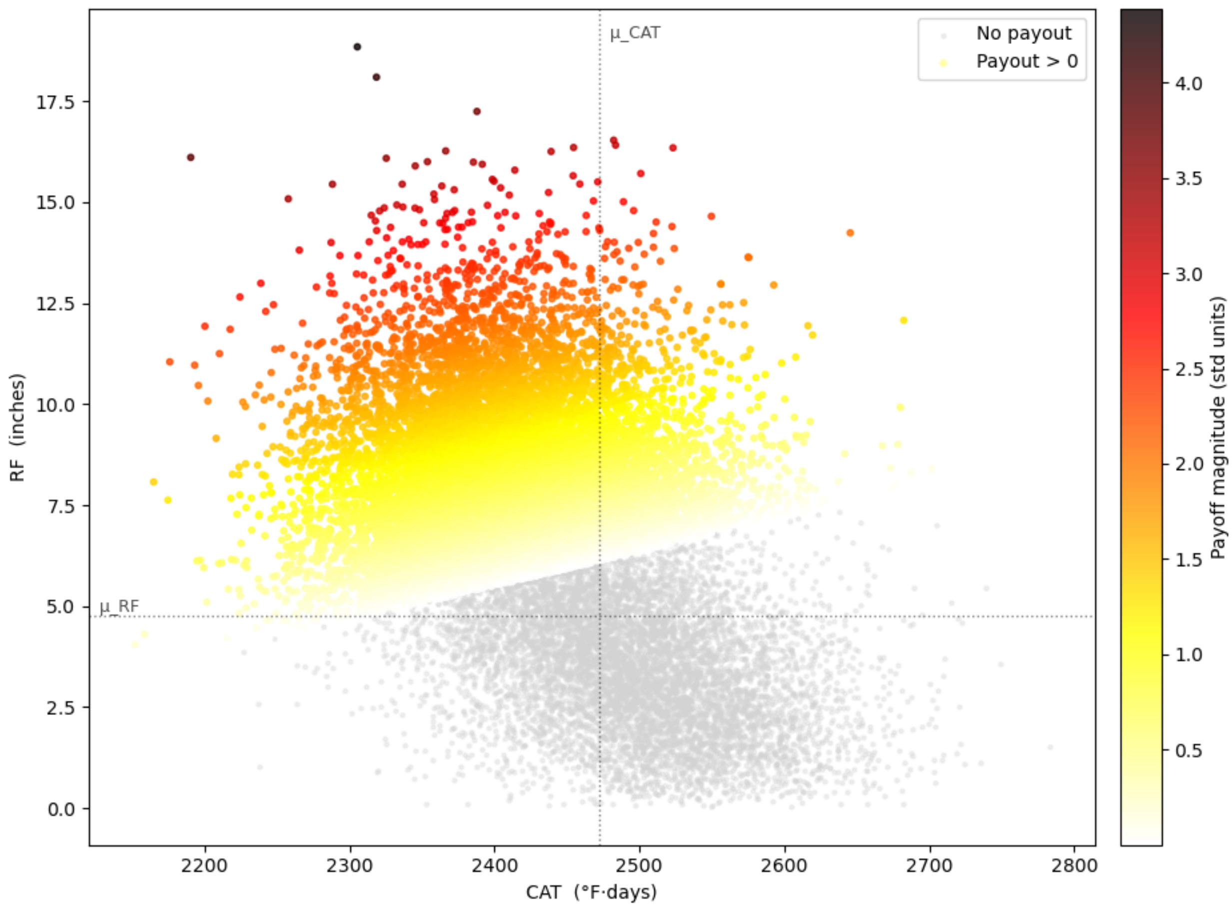

Figure 10: Scatterplot of CAT vs RF colored by payoff, showing exercise region in the upper-left (cool and wet).

6.4 The Value of Modeling Dependence

A product based only on rainfall would miss the fact that cool and wet conditions tend to arrive together. Under an independence copula $C(u, v) = uv$ (which assumes no dependence), the same option price was calculated to be roughly 0.1703, which, when scaled by notional value, would lead to a substantial underpricing, a concern emphasized by Bressan & Romagnoli (2021). The negative dependence captured by the Frank copula ensures that the hedge reflects the realistic co-movement of the two risk factors.

7 Limitations and Extensions

Several limitations of our analysis should be noted:

1. **Market price of risk.** As is standard in the weather derivatives literature, we assume the market price of weather risk is zero. This cannot be avoided because there is no liquid traded market from which to calibrate this risk.

2. **Basis risk.** The Wichita airport station is assumed to be a proxy for weather conditions faced by Wichita winter wheat farms. Rainfall, however, is highly geographically varied. Even nearby farmers may observe different weather patterns.

3. **Basket weights and strike.** The weights $(\omega_1, \omega_2) = (0.2, 0.8)$ and the strike $K^B = 0.42$ are chosen to be economically sensible but are not calibrated to empirical yield-loss data. Realistically, these need to be calibrated on regional yield-loss data.

4. **Normality of CAT distribution.** The Alaton model assumes the Gaussian distribution, which may not hold in the empirical data.

5. **Copula choice.** While the Frank copula fits best among the three families tested, more flexible copulas (e.g., the Survival BB8 copula considered by Bressan & Romagnoli (2021)) may improve the fit, particularly in capturing tail dependence.

Natural extensions include calibrating the basket weights to yield/loss data, testing additional copula families, correlating weather conditions with dollar losses to determine an appropriate tick size, and incorporating spatial weather data from multiple stations to mitigate basis risk.

8 Conclusion

We constructed and priced a multi-index weather derivative to hedge the joint temperature-rainfall risk faced by Kansas winter wheat farmers during the harvest window of June 10 to July 10. Using seventy years of daily weather data from Wichita, Kansas, we built seasonal CAT and RF indices and found strong negative dependence ($\hat{\tau} = -0.39$), confirming that cooler harvest windows tend to coincide with wetter conditions. This is precisely the combination that threatens winter wheat yields.

On the modeling side, the Alaton Ornstein-Uhlenbeck framework provided a good fit to the seasonal temperature dynamics. For rainfall, the Weibull distribution was selected and the index value simulation approach was used to avoid the variance underestimation associated with daily precipitation models. The Frank copula, chosen among three candidates by AIC under both parametric and rank-based estimation, captured the negative dependence between the two indices.

We then gave a framework and example of the resulting call option pricing. Using 50,000 joint Monte Carlo scenarios, we priced a weighted basket call option, tilted 80% toward rainfall and 20% toward temperature, at a price of 0.2032 standardized units, convertible to a dollar premium using a notional value calibrated to the farmer's crop exposure. A comparison with the independence copula, which yielded a price of roughly 0.1703, emphasized the importance of modeling dependence: ignoring the co-movement of temperature and rainfall leads to underpricing of the derivative. Taken together, these results demonstrate the feasibility of copula-based multi-index weather derivatives for managing the compound climate risks inherent in farming.

References

Alaton, P., Djehiche, B., & Stillberger, D. (2002). On modelling and pricing weather derivatives. *Applied Mathematical Finance*, 9(1), 1–20.

Benth, F. E., & Benth, J. Š. (2005). Stochastic modelling of temperature variations with a view towards weather derivatives. *Applied Mathematical Finance*, 12(1), 53–85.

Benth, F. E., Benth, J. Š., & Koekebakker, S. (2007). Putting a price on temperature. *Scandinavian Journal of Statistics*, 34, 746–767.

Benth, F. E., Benth, J. Š., & Koekebakker, S. (2008). *Stochastic Modeling of Electricity and Related Markets*. World Scientific, Singapore.

Bibby, B. M., & Sørensen, M. (1995). Martingale estimation functions for discretely observed diffusion processes. *Bernoulli*, 1(1/2), 17–39.

Bokusheva, R. (2018). Using copulas for rating weather index insurance contracts. *Journal of Applied Statistics*, 45(13), 2328–2356.

Bressan, G. M., & Romagnoli, S. (2021). Climate risks and weather derivatives: A copula-based pricing model. *Journal of Financial Stability*, 54, 100877.

Cong, R.-G., & Brady, M. (2012). The interdependence between rainfall and temperature: Copula analysis. *The Scientific World Journal*, 2012.

Dornier, F., & Queruel, M. (2000). Caution to the wind. *Weather Risk Special Report*, Energy & Power Risk Management.

Mußhoff, O., Odening, M., & Xu, W. (2006). Modeling and pricing rain risk. Contributed paper, International Association of Agricultural Economists Conference, Gold Coast, Australia.

Nelsen, R. B. (2006). *An Introduction to Copulas* (2nd ed.). Springer, New York.

Neuhofer, Z. T., & Rosch, S. (2019). Federal crop insurance: Specialty crops. Congressional Research Service, CRS Report R45459. Available at: `https://www.congress.gov/crs-product/R45459`.

Smith, V. H., & Goodwin, B. K. (2004). Crop insurance, moral hazard, and agricultural chemical use. *American Journal of Agricultural Economics*, 78(2), 428–438.

Zhu, W., Tan, K., Porth, L., & Wang, C. (2018). Spatial dependence and aggregation in weather risk hedging: A Lévy subordinated hierarchical Archimedean copulas approach. *ASTIN Bulletin*, 48(2), 779–815.

A Varying ω_1 and ω_2

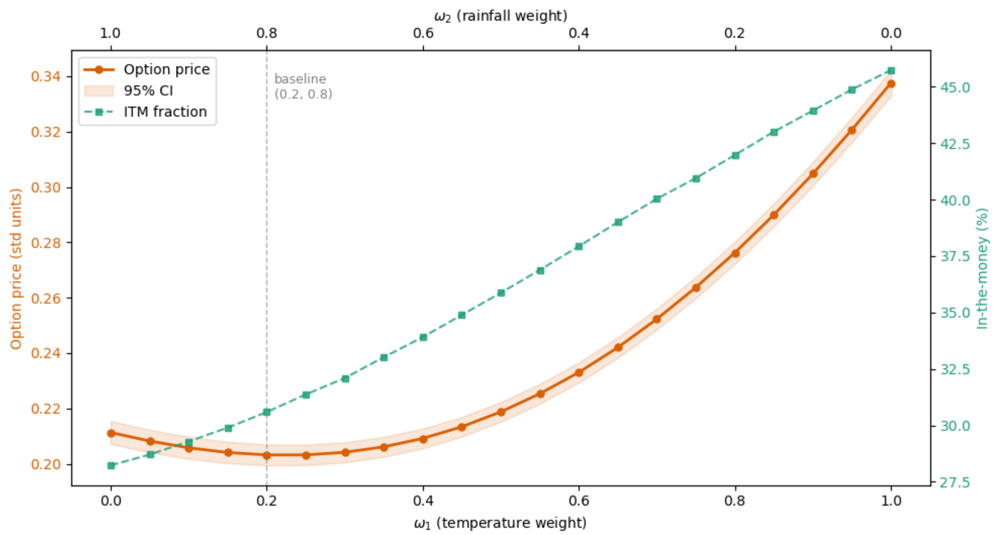

Figure 11: Sensitivity of the basket option price to the weight allocation (ω_1, ω_2) with $\omega_1 + \omega_2 = 1$. The individual strike components $K_{\text{CAT}} = 0.1$ and $K_{\text{RF}} = 0.5$ are held fixed, so the composite strike K^B varies with the weights. The price is nearly flat for rainfall-dominated allocations and rises as temperature weight increases, resulting from their respective distributions and the choice of strikes.

B Different Strike Levels

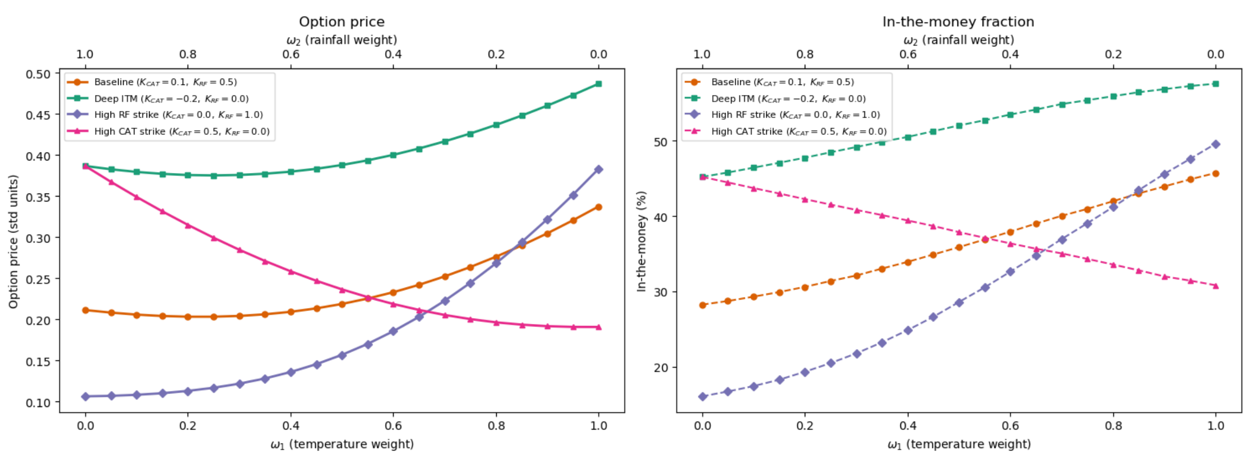

Figure 12: Sensitivity under four strike configurations. The deep ITM curve is expensive because its low composite strike is easily exceeded. The high RF strike curve rises steeply with ω_1: when rainfall is heavily weighted, the large $K_{\text{RF}} = 1.0$ suppresses payouts. The high CAT strike curve shows the opposite pattern, as expected. The baseline sits between these extremes due to its moderate strikes on both indices.