




POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE



Temperature Anomalies and Climate Physical Risk in Portfolio Construction

TESI DI LAUREA MAGISTRALE IN
MATHEMATICAL ENGINEERING - INGEGNERIA MATEMATICA

Carlo Bechi, 246853

Advisor:

Dr. Michele Azzone

Co-advisors:

Dr. Gabriele Sbaiz

Academic year:

2025-2026

Abstract: Driven by the increasing frequency and intensity of natural disasters and chronic climate threats, we investigate the impact of physical climate risk on global equity portfolios. By employing a panel regression analysis on sectoral returns, we provide statistical evidence that extreme temperature events exert a negative effect on most sectors. We introduce two novel metrics based on these temperature anomalies: *Climate Risk Exposure* and *Climate Exposure Volatility*, in order to measure the environmental vulnerability of a portfolio. Unlike available static country-level indices, these metrics incorporate the time varying probability of extreme events and their relations with firm-specific asset intensity. We integrate these measures into a multi-objective portfolio optimization framework. This approach extends the traditional mean-variance paradigm, allowing investors to construct portfolios that are resilient to physical climate shocks without sacrificing diversification. Finally, we conduct a backtesting analysis to show the practical benefits of incorporating these climate risk metrics into the investment process, evaluating how climate-aware strategies perform relative to traditional benchmarks.

Key-words: Climate Aware Portfolio Optimization, Physical Climate Risk, Temperature Anomalies, Climate Change

1. Introduction

Physical climate risk is becoming a relevant factor in the asset management industry, comprising both acute hazards and chronic long-term shifts (Bressan et al., 2024). Given the increasing frequency of climate-related phenomena¹ and the continued rise in greenhouse gas emissions, asset managers and financial institutions worldwide are seeking methods to incorporate these risks into investment decisions (Campiglio et al., 2023; Azzone et al., 2026b). Despite this growing attention, there remains a significant gap in effective quantitative metrics capable of capturing how corporate economic value is affected and, in turn, guiding the decisions of a rational investor (Bortolan et al., 2024; Katz and McIntosh, 2021). We propose two new metrics for the expected frequency and volatility of extreme climate events within a portfolio, based on temperature anomalies. We then integrate these metrics into a multi-objective optimization problem, extending the mean-variance framework to account for physical risk. To the best of our knowledge, this is the first attempt to offer a flexible, data-accessible

¹See <https://www.eea.europa.eu/en/analysis/indicators/economic-losses-from-climate-related>

approach for constructing optimal portfolios that explicitly consider physical climate risk.

Climate risk has emerged as a critical financial threat, broadly categorized into physical and transition risks (Bolton and Kacperczyk, 2023). Focusing on the former, the last decade has witnessed a sharp increase in both acute hazards, such as heatwaves, floods, and unprecedented wildfire activity, and chronic impacts like rising sea levels, with extreme events becoming up to 20 times more likely in certain regions (Jones et al., 2024). These environmental shifts disproportionately affect vulnerable systems and populations, often exacerbated by historical inequities. Despite global initiatives aimed at mitigation, the escalating frequency of these physical threats poses severe financial consequences for firms that remain ill-prepared for the resulting economic losses (Wang et al., 2023).

Classifying physical climate risk for a portfolio is a complex endeavor; consequently, very few contributions are currently available. While some country-level indices exist, such as the Notre Dame Global Adaptation Initiative (ND-GAIN),² they have significant limitations. The ND-GAIN index quantifies a country’s exposure, sensitivity, and adaptive capacity across six critical sectors (food, water, health, ecosystem services, human habitat, and infrastructure). However, it does not provide a multivariate time-series of extreme events but just a country level score. A timeseries framework would be essential for understanding the impacts on stock returns and the correlations between events, which are necessary conditions for determining whether portfolio diversification is possible. Some contributions that usually focus on a timeseries model do exist for a single risk e.g. wildfires for mortgages credit risk in the case of Kahn et al. (2024). They show that Mortgage-Backed Securities (MBS) can effectively mitigate physical climate risk for retail mortgages by minimizing the spatial concentration and correlation of the underlying assets. By leveraging these diversification metrics, the authors construct optimal, climate-resilient portfolios that minimize financial losses. We aim to apply a similar framework to global equity portfolios, addressing exposure to general climate-related phenomena, which we proxy using extreme temperature anomalies.

There is a large literature on incorporating climate risk in asset management. However, most of the contribution focus on sustainability (ESG) risk or transition risk and very few addresses physical risks (see Le Guenedal and Roncalli, 2022, for a survey). The literature on ESG portfolio selection generally follows two main approaches: incorporating ESG preferences directly into the utility function (Pedersen et al., 2021; Avramov et al., 2022) or applying explicit constraints to the optimization problem (De Spiegeleer et al., 2023; Morelli, 2024; Aprea and Sbaiz, 2025). Empirical studies utilizing the latter, specifically those minimizing variance or tracking error subject to ESG or carbon targets, outline that such constraints do not necessarily compromise risk-adjusted returns (Bolton et al., 2022; Azzone et al., 2026a).

To tackle the multi-objective problem that considers physical risk together with standard mean variance, we employ a Multi-Objective Particle Swarm Optimization (MOPSO) (see Manzoni et al., 2026, for example). Moore (1999) was the first to extend PSO (Kennedy and Eberhart, 1995) to multi-objective optimization, showing that swarm-based approaches are capable of effectively handling conflicting objectives and thereby establishing the basis for subsequent developments. Based on this paradigm, Coello et al. (2004) proposed a MOPSO algorithm that integrates Pareto dominance with an external archive to preserve solution diversity and enhance convergence. Their work also includes a sensitivity analysis, providing valuable insights into parameter tuning for improved efficiency and diversity. Durillo et al. (2009) performed a comparative study of existing MOPSO variants, identifying key challenges related to diversity maintenance. To address these issues, they introduced a MOPSO variant incorporating velocity constraint mechanisms, resulting in improved stability and diversity preservation. These findings highlight the necessity of adaptive control strategies in MOPSO algorithms. Furthermore, Kumar et al. (2020) analyzed the influence of critical parameters on MOPSO performance, showing that the inertia weight is essential for balancing exploration and exploitation. Finally, Reyes-Sierra et al. (2006) provided a comprehensive review of MOPSO developments, classifying algorithmic modifications according to leader selection, archive management, and velocity update strategies.

In this work, we bridge the gap between static country-level indicators and the need for dynamic, firm-specific physical risk assessment. Unlike existing literature that often relies on qualitative scores or focuses on specific asset classes like real estate, we provide a generalized framework applicable to global equity portfolios. First, we confirm the materiality of physical risk by empirically testing the impact of temperature anomalies on sector returns. Second, we model the time-varying probability of extreme events and define two novel metrics for climate physical risk that are based on physical risk correlations in different continents: the *Climate Risk Exposure* and the *Climate Exposure Volatility*. Finally, we integrate these metrics directly into a multi-objective portfolio optimization problem. This allows us to extend the standard mean-variance analysis to explicitly minimize the

²For further detail refer to: <https://gain.nd.edu/our-work/country-index/>

volatility induced by physical climate shocks, offering a viable tool for climate-resilient portfolio construction and climate risk diversification.

The remainder of this thesis is structured as follows. Section 2 details the methodology used to construct the time series of temperature anomalies and evaluates how these anomalies influence equity returns. Section 3 proposes two distinct risk metrics to quantify physical climate risk. Section 4 performs a portfolio optimization that incorporates these climate-sensitive measures. Finally, Section 5 conducts a backtesting analysis to evaluate how climate-aware strategies perform relative to traditional benchmarks.

2. Temperature Anomalies and Their Effects on Returns

This section details the methodology used to construct and model the time series of extreme temperature events. The goal is to estimate a time-varying probability of temperature anomalies and their correlations in different geographic regions, which serve as a fundamental input for the portfolio optimization framework. We choose to focus on temperature extremes as they are often the primary drivers of extreme climate events, such as wildfires, prolonged droughts, very intense heatwaves and glacial melt episodes. Nevertheless, we believe that our approach and methodology are designed to be flexible, allowing them to be extended to any other extreme climate event under consideration.

We use temperature data provided by Our World in Data ³, a scientific repository hosted by the University of Oxford. The platform sources information exclusively from official international institutions such as the IPCC, NASA, and the World Bank. We utilize monthly mean temperature data for the period spanning from January 1, 1940, to April 30, 2025. The analysis covers the six major continents: North America, South America, Europe, Africa, Asia, and Oceania. Antarctica is excluded from this study. This decision is based on the assumption that the percentage of revenues generated by companies within the MSCI World index in this continent is negligible, thus having no significant impact on the financial risk-return profile of the portfolio. We adopt a continent-level approach for temperature data because we have company revenue information available at this geographic scale. However, we believe that our analysis could be extended and rendered even more effective by utilizing temperature and revenue data at a national level. Such a detailed approach would allow for a more precise identification of the magnitude of climate risks to which the portfolio is exposed within each individual country.

2.1. Identification of Temperature Anomalies

To rigorously identify extreme temperature events, we implement a statistical detection method based on historical baselines. First, we establish a reference period (baseline) spanning from 1960 to 1990. For each continent k and each month $m \in \{1, \dots, 12\}$, we calculate the historical mean $\mu_{k,m}$ and the standard deviation $\sigma_{k,m}$ using only the observations within this 30-year window. This allows us to capture the seasonal cycle and the natural variability inherent to each specific month of the year and to each specific area of the globe.

To ensure comparability between different continents and months, the raw monthly temperature $T_{k,t}$ observed at time t is transformed into a standardized anomaly. The transformation is defined as

$$Z_{k,t} = \frac{T_{k,t} - \mu_{k,m}}{\sigma_{k,m}} , \quad (1)$$

where m is the month corresponding to time t . This normalization eliminates geographical bias by subtracting the historical mean and scaling temperatures by the specific historical variability of each region. This standardization allows for the consistent comparison of extreme events across diverse climatic zones, regardless of differences in baseline temperatures or local variances. Furthermore, removing seasonal bias standardizes the statistical distribution across the entire year; this ensures, for instance, that an anomaly in July is strictly comparable to one in December.

We define an indicator variable $B_{k,t}$ to formally identify the occurrence of a positive temperature anomaly at time t in continent k . An extreme event is recorded when the standardized temperature exceeds a critical threshold of 2. This threshold is statistically significant as, under the assumption of a normal distribution, it identifies events that lie in the extreme upper tail (approximately the top 2.28% of observations under Gaussianity assumption) and ensures that we focus on "extreme" events. The indicator variable is therefore constructed as follows:

³See <https://ourworldindata.org>

$$B_{k,t} = \begin{cases} 1 & \text{if } Z_{k,t} > 2 \\ 0 & \text{otherwise} \end{cases} . \quad (2)$$

To capture the dynamic evolution of climate risk over time, we estimate the probability $p_{k,t}$ that an anomaly occurs in continent k at time t . Given the binary nature of $B_{k,t}$, we employ a Logistic Regression model, which is the standard approach for modeling the probability of occurrence of discrete events.

To reflect the complexities of global warming patterns, the model incorporates a quadratic time trend. This allows us to account for potential non-linearities and accelerations in temperature increases that a simple linear trend might fail to capture. For each continent, the log-odds of the probability are modeled as

$$\text{logit}(p_{k,t}) = \ln \left(\frac{p_{k,t}}{1 - p_{k,t}} \right) = \beta_{0,k} + \beta_{1,k}t + \beta_{2,k}t^2 , \quad (3)$$

where:

- t represents the time index;
- $\beta_{1,k}$ captures the linear progression of the risk over time;
- $\beta_{2,k}$ accounts for the quadratic component, identifying whether the frequency of extreme events is accelerating or decelerating.

This specification ensures that our climate risk measures are forward-looking and capable of reflecting the intensifying nature of environmental threats, providing a more robust basis for the subsequent financial risk-return analysis.

As illustrated in Figure 1, the fitted probabilities $\hat{p}_{k,t}$ reveal a consistent trend across all analyzed continents. The probability of encountering an extreme temperature anomaly exhibits a significant increase over the sample period.




Figure 1: Probabilities of extreme temperatures over time in the period 1940-2020.

These pronounced upward trajectories provide empirical evidence that physical climate risk is non-stationary and has intensified in recent decades. In statistical terms, the traditional assumption of a constant risk distribution is invalidated by the clear intensification of temperature anomalies in recent decades. From a financial perspective, this systematic increase in $\hat{p}_{k,t}$ implies that the "tail risk" associated with climatic events is no

longer a static outlier but a dynamic variable, necessitating a dynamic approach to portfolio risk management.

The performance of the logistic regression models, evaluated through the AUROC and AUPR metrics, is summarized in Table 1. To ensure a robust assessment of the models' reliability, the regression parameters were estimated using a historical training period from 1940 to 2019. Subsequently, the performance metrics reported in the table were calculated over the extended horizon (1940-2025), thereby incorporating the out-of-sample test period (2020-2025) to verify the models' ability to project climate risk into the immediate future.

To evaluate the models, we primarily focus on the Area Under the Receiver Operating Characteristic curve (AUROC). This metric measures the model's discriminatory power: an AUROC of 0.5 suggests a predictive ability no better than a random guess, while a value of 1.0 represents perfect classification. As shown in the results, all continents exhibit exceptionally high AUROC values, ranging from 0.8392 for Europe to 0.9396 for Africa.

Furthermore, we evaluate the model performance using the Area Under the Precision-Recall curve (AUPR), which provides a more stringent evaluation in the presence of imbalanced datasets, a typical feature of rare climatic anomalies. While the AUROC focuses on the trade-off between successes and false alarms, the AUPR specifically evaluates the quality of positive predictions (Precision) against the model's ability to capture all real events (Recall).

In our analysis, the AUPR values for all continental models consistently and significantly outperform those of a baseline random classifier.

In conclusion, the high predictive scores across all geographical regions allow for the categorical rejection of a stationary climate hypothesis. This validates the use of these time-varying probabilities as robust inputs for the subsequent multi-objective portfolio optimization.

Table 1: Performance Metrics for Logistic Regressions

Continent	AUROC	AUPR	AUPR Random
Africa	0.9396	0.8031	0.2060
Asia	0.9125	0.7327	0.2225
Europe	0.8392	0.3779	0.0943
North America	0.8947	0.6100	0.1429
Oceania	0.8844	0.3216	0.0564
South America	0.8651	0.5211	0.1341

2.2. Correlation Analysis and Spatial Dependence

After estimating the time-varying probabilities of temperature anomalies for each continent, it is essential to analyze the spatial dependence between these events. The objective of this analysis is to determine the correlation structure of climatic shocks across different geographic areas. From a risk management perspective, understanding this dependency is crucial: we aim to quantify the risk that a global portfolio might be simultaneously affected by multiple extreme events occurring in different parts of the world, refraining from viewing global climate change as a series of isolated incidents rather than a coordinated threat to the entire investment universe. In the following, we focus on estimating this correlation structure and verifying its theoretical consistency. Specifically, we will examine the empirical correlations between the climate indicators and test whether the resulting correlation matrix satisfies the mathematical constraints inherent to a system of Bernoulli random variables. This verification is fundamental to ensuring that our model provides a robust and theoretically sound representation of global climatic interdependencies.

For each pair of continents (i, j) , we compute the Pearson correlation coefficient between their respective binary time series of anomalies, $B_{i,t}$ and $B_{j,t}$. We are able to visualize the result in the following empirical correlation matrix, Figure 2, that provides significant insights into the spatial dependence of temperature anomalies across continents.




Figure 2: Correlation matrix of extreme temperature time series.

A primary observation is that all estimated Pearson coefficients are positive. This uniform positivity reflects a global systemic trend in climate anomalies, suggesting that temperature extreme events are driven by a common underlying process, global warming, that tends to lift temperatures across the entire planet simultaneously. However, the magnitude of these correlations remains relatively low, with values ranging between 0 and 0.48. This evidence is particularly important for several reasons. The fact that correlations do not exceed 0.48 indicates that despite a common global trend, localized climatic drivers still play a dominant role in determining the occurrence of extreme events at the continental level. From a financial risk perspective, these "low-but-positive" correlations represent a double-edged sword. On one hand, the universal positivity confirms that no continent is truly immune or "negatively correlated" to the global warming trend. On the other hand, the moderate magnitude (below 0.48) suggests that a good degree of geographical diversification is still possible. In conclusion, the heatmap reveals a world where climate risk is globally interconnected but still characterized by significant regional specificities, necessitating a risk management strategy that accounts for both the common trend and the idiosyncratic volatility of each continent.

2.3. Fréchet-Hoeffding Bounds

While Pearson correlation is a standard measure for Gaussian random variables, its application to Bernoulli distributions requires careful consideration. Since our indicator variables $B_{k,t}$ are discrete and binary, the correlation coefficients are subject to mathematical bounds determined by the underlying probabilities of each event. Failure to account for these constraints could lead to a misinterpretation of the co-occurrence of extremes. By testing these theoretical limits, we ensure that our assessment of spatial dependence provides a reliable measure of the true systemic risk, accurately reflecting the probability that multiple geographic areas within a portfolio may be affected by climate shocks simultaneously.

Every pair of Bernoulli random variables has a maximum and minimum admissible values of correlation (Huber and Marić 2015), this constraint is formally expressed by the Fréchet-Hoeffding theorem:

Theorem (Fréchet-Hoeffding). For any two random variables X_1 and X_2 with cumulative distribution functions

F_1 and F_2 , with $U \sim \text{Unif}([0, 1])$ the following inequality holds:

$$\text{corr}(F_1^{-1}(U), F_2^{-1}(1 - U)) \leq \text{corr}(X_1, X_2) \leq \text{corr}(F_1^{-1}(U), F_2^{-1}(U)). \quad (4)$$

We prove the corollary that follows for the specific case of two Bernoulli variables, $X_1 \sim \text{Bern}(p_1)$ and $X_2 \sim \text{Bern}(p_2)$ with different probability parameters.

Corollary. The correlation coefficient ρ between two Bernoulli random variables $X_1 \sim \text{Bern}(p_1)$ and $X_2 \sim \text{Bern}(p_2)$ is bounded by $\rho_{\min} \leq \rho \leq \rho_{\max}$, where:

$$\rho_{\min} = \frac{\max(0, p_1 + p_2 - 1) - p_1 p_2}{\sqrt{p_1(1 - p_1)}\sqrt{p_2(1 - p_2)}} \quad (5)$$

$$\rho_{\max} = \frac{\min(p_1, p_2) - p_1 p_2}{\sqrt{p_1(1 - p_1)}\sqrt{p_2(1 - p_2)}}. \quad (6)$$

The result is derived by applying the Fréchet–Hoeffding theorem, which establishes the existence of upper and lower bounds for joint distribution functions.

To determine these limits, we analytically compute the covariance using the cumulative distribution functions of the two Bernoulli variables. By exploiting the properties of Bernoulli distributions, the terms $\text{Cov}(F_1^{-1}(u), F_2^{-1}(u))$ and $\text{Cov}(F_1^{-1}(u), F_2^{-1}(1 - u))$ can be solved in closed form. Finally, by dividing these analytical covariances by the product of the standard deviations, $\sqrt{p_1(1 - p_1)}\sqrt{p_2(1 - p_2)}$, we obtain the explicit expressions for $\rho_{\min}$ and $\rho_{\max}$. For the full proof, please refer to Appendix A.

These bounds imply that a perfect correlation ($\rho = 1$) is only attainable when the marginal probabilities are equal ($p_1 = p_2$). Conversely, a perfect anti-correlation ($\rho = -1$) can only be achieved if the probabilities are complementary, such that $p_1 = 1 - p_2$. If these conditions are not met, the theoretical maximum and minimum correlations will strictly lie within the interval $(-1, 1)$.

In our framework, the parameters $p_{k,t}$ are non-stationary, as they are derived from a time-dependent logistic regression. This time-varying nature makes the direct application of the Fréchet-Hoeffding formulas complex, as the bounds would technically shift at every time step t .

To provide an estimate of the correlation boundaries, we decide to use the mean value of the estimated parameters $\hat{p}_{k,t}$ over the training horizon (1940–2020) as the representative p_1 and p_2 for the boundary calculations:

$$\bar{p}_k = \frac{1}{T} \sum_{t=1}^T \hat{p}_{k,t}. \quad (7)$$

This approach allows us to verify if the observed empirical correlations between continents are consistent with the theoretical limits imposed by their average probability of occurrence.

The results of this analysis are summarized in Figure 3, which illustrates the empirical Pearson correlation for each pair of continents plotted against their respective Fréchet-Hoeffding boundaries. This visualization confirms that the empirical estimates fall strictly within the feasible theoretical range.




Figure 3: Correlation boundaries of extreme temperature time series.

It is well-established that when extending to a multivariate Bernoulli distribution involving more than two variables, additional constraints must be satisfied to define the range of admissible correlations (Huber and Marić, 2015, 2019). However, these higher-order constraints are generally not explicit and lack a closed-form expression, particularly for systems exceeding four variables. Since our system comprises six variables, one for each continent, applying such higher-order constraints becomes analytically intractable. Consequently, to the best of our knowledge, the pairwise Fréchet-Hoeffding bounds represent the most robust analytical tool available.

2.4. The Impact of Temperature Extremes on Equity Returns

We now investigate the empirical relationship between extreme temperature anomalies and financial performance across various economic sectors. The core objective is to determine whether climate physical risks, manifested through continental temperature extremes, exert a statistically significant influence on asset returns. We construct a linear regression model that includes climate risk measures among its covariates, similarly to the methodology proposed in Bortolan et al. (2024), where a linear regression analysis is employed to demonstrate that the deviation of extreme temperature event variance, calculated with respect to a historical reference period, exhibits a statistically significant negative relationship with equity earnings. In our specific case, we utilize temperature anomalies to assess the impact of climate extreme events on financial returns. The analysis conducted in this section is crucial, as it provides the empirical foundation for the subsequent sections. Specifically, it justifies the development of metrics designed to quantify the climate risk of a portfolio and the implementation of portfolio optimization strategies that explicitly incorporate climate-related measures.

To conduct a granular analysis, we construct sectoral return time series using the constituents of the MSCI World Index. We consider the companies belonging to the MSCI World universe, categorized by their respective sector. Specifically, we select only those firms provided with sufficient available data for the analysis and that were constituents of the MSCI World index in both 2020 and 2024. This selection criterion is essential to mitigate survivorship bias and ensure a consistent sample; by including only companies present at both the

start and end of the period, we avoid distorting the results with firms that might have entered or exited the index due to bankruptcy, mergers, or insufficient market capitalization during the timeframe of the study. We utilize total returns, which provide a comprehensive measure of performance by accounting not only for price changes but also for dividends, stock splits, and other corporate actions. This is crucial for ensuring that the observed stock fluctuations purely reflect market-driven price movements and economic reactions to external shocks.

For each sector s , the return $R_{s,t}$ is calculated as the market capitalization weighted average of the individual total returns of its constituent firms within the MSCI World.

The foundational step of our analysis involves extending the traditional Capital Asset Pricing Model (Sharpe, 1964). In its standard form, the CAPM explains the excess return of an asset through its sensitivity to the market portfolio. We augment this framework by introducing a climate-risk proxy to capture the potential impact of physical climate shocks, specifically, we explicitly incorporate into the model, as covariates, the time series of extreme temperature anomalies previously constructed for each continent. This approach allows us to directly quantify the sensitivity of asset returns to localized climatic extremes while controlling for traditional market factors.

For every sector s and every continent k , we estimate the following regression:

$$R_{s,k,t} - R_{f,t} = \alpha_{s,k} + \beta_{MKT,s,k}(R_{MKT,t} - R_{f,t}) + \gamma_{s,k}B_{k,t} + \epsilon_{s,k,t} \quad , \quad (8)$$

where:

- $R_{s,k,t} - R_{f,t}$ represents the Sectoral Excess Return, calculated by subtracting the monthly risk-free rate from the sectoral total return.
- $R_{MKT,t} - R_{f,t}$ is the Market Excess Return, derived from the performance of the overall MSCI World Index minus the risk-free rate;
- $B_{k,t}$ is the Temperature Anomaly Indicator for continent k , the time series described and constructed in the previous section;
- $\gamma_{s,k}$ is the parameter of interest, representing the *Climate Beta* or the sensitivity of the sector to extreme temperature events in a specific continent.

The model parameters are estimated over the 2004-2025 timeframe, which corresponds to the period for which financial return data are available. To ensure the robustness of our statistical inference, the regression incorporates HAC (Heteroskedasticity and Autocorrelation Consistent) robust standard errors (Newey and West, 1986). This approach is useful to correct for heteroskedasticity and temporal autocorrelation within the time series, ensuring that the estimated significance of the climate covariates is not inflated by the inherent properties of financial data. Furthermore, we exclude sectors regarding banks, financial and insurance companies from our analysis. This exclusion is justified by the fact that, without an in-depth analysis of their balance sheets and underlying portfolios, it is particularly challenging to determine their precise geographical exposure and the extent of their vulnerability to climate physical risks.

The results are presented in the following Table 2, specifically, only the *Climate Beta* that are statistically significant are reported.

Table 2: The table reports the estimated $\gamma_{s,k}$ coefficients and p-values for pairs reaching statistical significance ($p < 0.1$), displaying only the coefficients that are statistically significant.

Sector	Continent	Climate Beta	p-value
Applied Resources	North America	-0.0140	0.001***
	Oceania	-0.0110	0.049**
Automobiles & Auto Parts	Asia	-0.0138	0.052*
	Oceania	-0.0183	0.100*
Chemicals	Africa	-0.0064	0.017**
	Asia	-0.0076	0.003***
	North America	-0.0063	0.058*
	Oceania	-0.0119	0.001***
Consumer Goods Conglomerates	Africa	-0.0097	0.053*
	Asia	-0.0106	0.031**
	Oceania	-0.0094	0.067*
Cyclical Consumer Products	Asia	-0.0075	0.068*
	Europe	-0.0088	0.057*
	North America	-0.0083	0.069*
	Oceania	-0.0113	0.053*
Cyclical Consumer Services	Asia	-0.0109	0.013**
	Oceania	-0.0120	0.030**
Energy - Fossil Fuels	North America	-0.0153	0.063*
	Oceania	-0.0192	0.066*
Food & Beverages	Africa	-0.0058	0.039**
	North America	-0.0095	0.006***
	Oceania	-0.0115	0.010***
Food & Drug Retailing	North America	-0.0098	0.020**
	Oceania	-0.0135	0.015**
Industrial & Commercial Services	Africa	-0.0080	0.023**
	Asia	-0.0139	0.000***
	North America	-0.0096	0.001***
	Oceania	-0.0084	0.034**
	South America	0.0068	0.058*
Mineral Resources	Africa	-0.0136	0.028**
	Asia	-0.0148	0.011**
	North America	-0.0114	0.091*
Personal & Household Products & Services	North America	-0.0066	0.088*
Pharmaceuticals & Medical Research	North America	-0.0090	0.085*
Real Estate	Africa	-0.0077	0.070*
Technology Equipment	Asia	-0.0077	0.016**
	Europe	-0.0053	0.099*
Telecommunications Services	Oceania	-0.0133	0.008***
Transportation	North America	-0.0068	0.026**

*Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$*

The estimation of the model across all sector-continent pairs yields several noteworthy results. First, 17 out of 23 sectors exhibit returns that are significantly impacted by extreme temperature events in at least one continent. Second, a remarkably consistent pattern emerges among the statistically significant coefficients: the sign of $\gamma_{s,k}$ is universally negative, with only a single exception. This predominant negative relationship suggests that extreme temperature anomalies generally exert a downward pressure on asset returns.

Furthermore, we observe that sectors such as Chemicals, Consumer Goods Conglomerates, Cyclical Consumer Products, Food & Beverages, Industrial & Commercial Services, and Mineral Resources are negatively impacted by extreme temperature events in at least three different continents. This widespread sensitivity is consistent with economic expectations; indeed, these industries are characterized by a high intensity of tangible assets, which are inherently exposed to physical climate risks.

Specifically, the reliance on large-scale production plants and refineries for the Chemicals and Mineral Resources sectors and the reliance on extensive warehousing and distribution facilities for Consumer Goods Conglomerates, Cyclical Consumer Products, and Food & Beverages sectors creates significant points of vulnerability. Tangible assets are the primary targets of climatic shocks, as extreme events can lead to direct structural damage,

operational downtime, or the destruction of physical inventories. Consequently, the high concentration of fixed capital in these sectors acts as a direct transmission mechanism through which climatic anomalies translate into negative asset returns. These findings are perfectly aligned with the results presented by Bortolan et al. (2024). Moreover, it is important to note that the negative impact on asset returns is not necessarily caused by the direct physical damage to infrastructure that extreme temperatures may provoke. Specifically, severe temperatures can trigger significant disruptions in consumption patterns and a marked decline in labor productivity, which collectively undermine the operational efficiency of firms independently of any structural destruction (Starr, 2000; Graff Zivin and Neidell, 2014; Neidell et al., 2021).

While the initial findings offer a valuable preliminary perspective, they may be constrained by limited statistical power or omitted variable bias. To address these limitations and substantiate our results, we extend the analysis to estimate a specific *Climate Beta* for each sector in a panel dataset framework. Furthermore, to enhance the robustness of our estimates, we incorporate fixed effects to absorb systematic variations independent of temperature extremes.

In this framework, the data are organized by "stacking" the time series of different continents (k) for each specific sector (s). Instead of running separate regressions for each geographical area, this panel approach aligns the observations of extreme temperatures and returns vertically.

$$R_{s,k,t} - R_{f,t} = \alpha_s + \beta_{MKT,s}(R_{MKT,t} - R_{f,t}) + \gamma_s B_{k,t} + \mu_k + \delta_m + \epsilon_{s,k,t} \quad (9)$$

- γ_s is the unique climate coefficient for sector s , estimated across all continents;
- μ_k are Continent Fixed Effects, which control for time-invariant characteristics of each continent (e.g., geographic location or structural economic differences);
- δ_m are Month Fixed Effects, represented by 12 dummy variables (one for each month of the year), which are essential to control for seasonality in both climate patterns and economic activity.

This methodology offers the primary advantage of estimating a unique *Climate Beta* for each sector, which captures the common sensitivity of a sector's returns to extreme temperature shocks ($B_{k,t}$) across different regions, significantly increasing the statistical power of our test. Moreover, the inclusion of Continent Fixed Effects (μ_k) and Month Fixed Effects (δ_m) ensures to isolate the impact of extreme temperature fluctuations rather than reflecting unrelated geographical or temporal noise.

Table 3 presents the sectors characterized by a statistically significant γ_s . While most of these sectors were already introduced in Table 2, this updated selection also includes Retailers and Software & IT Services. Even under this specification, many sectors continue to exhibit statistically significant negative coefficients, with the exception of the Software & IT Services sector which has historically maintained minimal physical exposure prior to the recent data center expansion. These findings confirm that the negative impact of climate extremes remains persistent across diverse industries. Moreover, it is worth mentioning that the sectors included in Table 2 but omitted from Table 3 continue to display a negative *Climate Beta*, although these estimates are no longer statistically significant.

Table 3: The table reports the estimated *Climate Beta* (γ_s) and the respective p-values, displaying only the coefficients that are statistically significant.

Sector	Climate Beta	p-value
Chemicals	-0.0062	0.0014***
Energy - Fossil Fuels	-0.0089	0.0399**
Food & Beverages	-0.0050	0.0070***
Real Estate	-0.0055	0.0121**
Retailers	-0.0026	0.0811*
Software & IT Services	0.0026	0.0609*
Telecommunications Services	-0.0042	0.0718*
Transportation	-0.0032	0.0893*
<i>Significance levels: *$p < 0.1$, **$p < 0.05$, ***$p < 0.01$</i>		

As a final test of the relationship between sectoral returns and extreme climate events, we estimate a single

Global Climate Beta (Γ) by pooling all sectors and all continents into a single panel model

$$R_{s,k,t} - R_{f,t} = \alpha + \beta_{\text{MKT}}(R_{\text{MKT},t} - R_{f,t}) + \Gamma B_{k,t} + \text{FE}_{\text{month}} + \text{FE}_{\text{continent}} + \epsilon_{s,k,t}. \quad (10)$$

The results of this panel regression yield a coefficient Γ that is negative and highly statistically significant, with an estimated value of -0.003 and a p -value of 0.0096 . The fact that this result holds while including both continent and month fixed effects provides strong evidence that extreme temperature events represent a systematic physical risk factor. These climate shocks negatively affect firm valuations globally, justifying the inclusion of climate risk exposure in the portfolio optimization process.

3. Quantifying Portfolio Risk to Extreme Climate Events

This section outlines the construction of two distinct metrics designed to quantify physical climate risk within a financial portfolio. The objective is to measure the impact of extreme weather events through two specific dimensions: the *Climate Risk Exposure* (CRE), which captures the average expected exposure, and the *Climate Exposure Volatility* (CEV), which quantifies the instability and volatility of that exposure.

As previously mentioned and as a proxy for extreme climate events, this analysis utilizes the previously constructed time series of extreme temperatures for each continent. Temperature extremes are the fundamental drivers behind a vast array of physical disasters, including heatwaves, wildfires, droughts, intense storms, and hurricanes.

While this study focuses on temperature as a primary indicator, the framework is designed to be modular. The methodology can be easily extended to more specific climate hazards (e.g., flood risk, sea-level rise) or refined to a more detailed geographic scale, such as country-by-country analysis, provided the relevant Bernoulli time series and the geographical exposure of the firms are available.

3.1. Firms' Vulnerability

The sensitivity of a firm to physical climate risk depends on the geographic distribution of its revenues and the vulnerability of its physical assets and supply chain in those areas. Consequently, an industrial company with significant physical infrastructure is more vulnerable to a local climate disaster than a service-based firm. As in Azzone et al. (2026b), we define asset intensity of firm i as

$$AI_i = \frac{\text{Tangible Assets}_i}{\text{Revenue}_i}. \quad (11)$$

We believe this metric is highly informative, as it reflects the degree of a company's dependency on physical assets to generate revenue. When aggregating this metric by sector, our empirical results confirm our economic intuition: the Real Estate sector exhibits the highest average asset intensity (10.6), reflecting a business model where value is primarily linked to the ownership of physical structures and land. Conversely, the Retailers sector shows the lowest average asset intensity (0.6), consistent with operations characterized by high inventory turnover and a lower reliance on fixed tangible assets relative to total sales. This polarization highlights that physical vulnerability is not uniform across the economy but is strictly linked to the underlying asset composition.

To provide a geographical localization of the risk, we also define $S_{i,k}$ the revenue produced in continent k by firm i and by combining it with the asset intensity, we determine the importance of continent k for company i in term of physical risk. Furthermore, we define the climate normalized portfolio weight for continent k , denoted as $\alpha_k(\mathbf{w})$,

$$\alpha_k(\mathbf{w}) = \frac{\sum_{i=1}^N w_i \cdot AI_i \cdot S_{i,k}}{\sum_{k=1}^K \sum_{i=1}^N w_i \cdot AI_i \cdot S_{i,k}}, \quad (12)$$

where w_i is the portfolio weight of firm i and N is the number of stocks. By using these climate normalized weights we are able to capture the physical risk exposure of the portfolio to a given continent.

Unsurprisingly, calculating the climate normalized weights for the MSCI World index, Asia, North America, and Europe exhibit the highest climate normalized weights. In particular, Asia shows the most significant exposure; this is motivated by its role as the "world's factory," characterized by a dense concentration of manufacturing hubs and industrial facilities that inherently increase the physical asset intensity of the region.

3.2. Measuring Risk: Exposure and Volatility

Using the climate normalized weights and the temperature extreme time series $B_{k,t}$, we can define our two new physical risk metrics.

The first measure, *Climate Risk Exposure*, quantifies the expected average exposure of the portfolio to extreme climate events at a given time t . It represents the "baseline" physical risk, calculated as the expected value of the weighted sum of the time-varying Bernoulli variables $B_{k,t}$, using the climate normalized weights $\alpha_k(\mathbf{w})$. Unlike static indicators, this measure accounts for the non-stationarity of climate trends by utilizing the time-dependent probabilities $p_{k,t}$. It is defined as

$$\text{CRE}(w, t) = \mathbb{E} \left[\sum_{k=1}^K \alpha_k(w) B_{k,t} \right] = \sum_{k=1}^K \alpha_k(w) \cdot p_{k,t} \quad , \quad (13)$$

where $p_{k,t}$ is the probability of an extreme temperature event in continent k at time t . A higher $\text{CRE}(w, t)$ indicates that the portfolio is potentially highly affected by regions currently experiencing a high frequency of climate anomalies.

While the CRE provides a mean expectation of physical risk, it inherently lacks the capacity to account for the uncertainty or the potential for simultaneous disasters across different geographic areas. To address this limitation, we introduce the *Climate Exposure Volatility*, denoted as $\text{CEV}(\mathbf{w}, t)$. This metric quantifies the variance of the portfolio's climate exposure at time t , offering a rigorous measure of the stability of its risk profile. A high volatility suggests that the portfolio's risk is poorly diversified across climate zones. In such cases, the portfolio becomes highly susceptible to sudden, large-scale shocks where multiple geographic areas might be damaged concurrently. Formally, the function $\text{CEV}(w, t)$ is defined as the variance of the weighted sum of the Bernoulli variables $B_{k,t}$, utilizing the climate normalized weights $\alpha_k(\mathbf{w})$ previously derived

$$\text{CEV}(w, t) = \text{Var} \left(\sum_{k=1}^K \alpha_k(w) B_{k,t} \right) = \sum_{k=1}^K \alpha_k^2 \text{Var}(B_{k,t}) + 2 \sum_{k < j} \alpha_k \alpha_j \text{Cov}(B_{k,t}, B_{j,t}) \quad (14)$$

where $\text{Var}(B_{k,t}) = p_{k,t}(1-p_{k,t})$ is the variance of the Bernoulli process for continent k at time t and $\text{Cov}(B_{k,t}, B_{j,t})$ is the covariance between the events in continents k and j . Formally, the covariance is expressed as

$$\text{Cov}(B_{k,t}, B_{j,t}) = \rho_{k,j} \sqrt{\text{Var}(B_{k,t}) \text{Var}(B_{j,t})} = \rho_{k,j} \sqrt{p_{k,t}(1-p_{k,t}) p_{j,t}(1-p_{j,t})} \quad . \quad (15)$$

Substituting these expressions we can rewrite the *Climate Exposure Volatility* as

$$\text{CEV}(w, t) = \sum_{k=1}^K \alpha_k^2 p_{k,t}(1-p_{k,t}) + 2 \sum_{k < j} \alpha_k \alpha_j \rho_{k,j} \sqrt{p_{k,t}(1-p_{k,t}) p_{j,t}(1-p_{j,t})} \quad . \quad (16)$$

This formulation allows for a decomposition similar to the one found in Kahn et al. (2024), that decomposes the wildfire risk in the context of a portfolio of Mortgage-Backed Securities (MBS). The first term, $\sum_{k=1}^K \alpha_k^2 p_{k,t}(1-p_{k,t})$, is akin to a Herfindahl Index and measures the geographic dispersion of risk. If we assume uniform probabilities p_k , this term is minimized when the portfolio's climate exposure is equally distributed across all K regions (i.e., $\alpha_k = 1/K$ for all k). The second term, $2 \sum_{k < j} \alpha_k \alpha_j \text{Cov}(B_{k,t}, B_{j,t})$, accounts for the systemic component of climate risk, reflecting how the correlation of extreme events across different geographic areas amplifies the overall portfolio volatility, in other words it captures the systemic risk that arises when climate extreme events are geographically correlated: if extreme events in two continents (e.g., North America and Europe) are highly correlated, a portfolio with significant exposure to both will exhibit higher volatility. This metric effectively penalizes a lack of geographic diversification against synchronized climate shocks, providing a more comprehensive view of potential portfolio losses than a simple average exposure metric.

Furthermore, it is insightful to examine the behavior of this metric as the number of geographic regions K increases (e.g., shifting from a continental to a country-level allocation). Let us consider a simplified scenario where the portfolio's climate exposure is uniformly distributed across all regions, such that $\alpha_k = 1/K$ for all k , and assume a constant risk level $p_k = p$ and a uniform pairwise correlation $\rho_{k,j} = \rho$. Under these assumptions, the first term of Equation (16) becomes

$$\sum_{k=1}^K \left(\frac{1}{K} \right)^2 p(1-p) = \frac{1}{K} p(1-p) \quad . \quad (17)$$

Similarly to what was demonstrated in Kahn et al. (2024), as $K \rightarrow \infty$, this idiosyncratic component tends to zero, reflecting the benefit of granular geographic diversification. However, the second term (the systemic component) behaves differently. Since there are $K(K-1)/2$ pairs in the double summation, the correlation term simplifies to

$$2 \sum_{k < j} \frac{1}{K^2} \rho p(1-p) = 2 \frac{K(K-1)}{2} \frac{1}{K^2} \rho p(1-p) = \left(1 - \frac{1}{K}\right) \rho p(1-p) . \quad (18)$$

As $K \rightarrow \infty$, this term converges to the constant value $\rho p(1-p)$. This result highlights that while increasing geographic granularity can eliminate local risk dispersion, it cannot completely eliminate the systemic risk arising from the underlying correlation structure of climate events.




Figure 4: Evolution of Climate Exposure Volatility across three portfolio strategies.

The provided Figure 4 illustrates the evolution of CEV over time, comparing three distinct portfolio strategies to highlight different approaches to managing physical climate risk. Specifically, the analysis considers: i) the MSCI World Portfolio; ii) the Equal-Weighted Continent Portfolio, where the climate normalized weights are distributed equally across all continents; iii) the Inverse Risk-Weighted Portfolio, where geographic weights are assigned inversely proportional to the probability of extreme temperature events in each continent, thereby tilting the portfolio away from high-risk regions.

As clearly shown in the chart, the *Climate Exposure Volatility* is a function of time, reflecting the historical fluctuations in both the frequency and severity of extreme weather events. Notably, the risk measure exhibits a significant shift starting around 2010, where the volatility begins to trend downward. This phenomenon is explained by the statistical properties of the Bernoulli variables $B_{k,t}$. The variance of a Bernoulli distribution, given by $p_{k,t}(1-p_{k,t})$, reaches its maximum at $p_{k,t} = 0.5$. Beyond this threshold, as the probability of an extreme event continues to increase ($p > 0.5$), the variance paradoxically decreases. Since 2010, the parameters of several $B_{k,t}$ processes have exceeded 0.5 and continued to rise; consequently, the system reflects a higher degree of "certainty" regarding the occurrence of extreme events. This leads to a reduction in statistical volatility, even as the expected frequency (and thus the total risk) remains high.

Furthermore, the comparison reveals a clear hierarchy in risk mitigation. The Equal-Weighted Climate Portfolio consistently exhibits a lower CEV compared to the MSCI World baseline, demonstrating that simply diversifying exposure away from market-cap concentrations can stabilize the climate risk profile. Even more effective is the Optimal Risk-Weighted Portfolio, which achieves the lowest volatility among the three strategies. By tilting the

weights away from the most vulnerable regions, this approach significantly curtails the portfolio’s susceptibility to climatic shocks, proving that active geographic optimization is the most efficient tool for minimizing physical risk exposure.

To investigate the role and the impact of the correlation between continental extreme temperature events on the physical climate risk of a portfolio, we compute the Climate Exposure Volatility (CEV) of the MSCI World Index under three distinct scenarios. Specifically, we employ the empirical correlation estimated via the Pearson coefficient, alongside the theoretical maximum and minimum correlations derived from the Fréchet-Hoeffding bounds presented in the Figure 3.

The results, illustrated in Figure 5, show that the correlation structure plays a pivotal role in determining the overall risk level. It is evident from the findings that an increase in correlation leads to a significant rise in climate risk; higher synchronization of extreme events across geographic regions markedly amplifies the portfolio’s exposure volatility exposure, thereby reducing the benefits of geographical diversification.




Figure 5: Evolution of MSCI World Index Climate Exposure Volatility with three different correlation structures.

4. Climate-Aware Portfolio Optimization

In this section, we formalize the optimization problem designed to construct a climate-resilient portfolio. The investment universe is composed of $N=120$ stocks selected from the MSCI World Index. This reduction of the investment universe is necessary to maintain computational tractability, balancing the representativeness of the sample with the high-dimensional complexity of the optimization algorithm. To ensure the portfolio serves as a faithful proxy for global equity markets, the selection follows a stratified sampling approach based on market capitalization and location of the headquarter.

4.1. The Optimization Problem

The 120 assets are distributed across geographic macro-regions to reflect their relative importance in the global financial landscape. Each company is categorized based on the location of its corporate headquarter and the selection is partitioned according to the following weights:

- North America: 45% (54 assets);
- Europe: 25% (30 assets);
- Asia: 25% (30 assets);
- Oceania: 5% (6 assets).

This allocation strategy ensures that the optimization process operates on a diversified universe that mimics the capital distribution of the MSCI World. By anchoring the selection to headquarters' locations, we provide a realistic benchmark for assessing the trade-offs between financial performance and climate risk, while maintaining consistency with global indexing standards.

We seek the portfolio $w \in \mathbb{R}^{120}$ that solves the following optimization problem:

$$\begin{cases} \min_w & F(w) = [f_1(w), f_2(w), f_3(w)]^\top & (19a) \\ \text{s.t.} & \sum_{i=1}^{120} w_i = 1, & (19b) \\ & w_i \geq 0, \quad \forall i = 1, \dots, 120 & (19c) \end{cases}$$

where the three objective functions to be minimized are:

1. Minus the Expected Return: $f_1(w) = -w^\top r$;
2. Market Variance: $f_2(w) = w^\top \Sigma w$;
3. *Climate Exposure Volatility* (CEV): $f_3(w)$, as defined in Section 3.

Our optimization problem considers only portfolios where the total capital is fully invested in risky assets. Consequently, the model currently excludes the possibility of allocating a portion of the wealth to a risk-free rate or engaging in short-selling. However, this framework is inherently flexible and can be generalized to include a risk-free asset and short-selling opportunities by relaxing the constraints in (19b) and (19c).

4.2. Optimization Algorithm and Implementation

To solve this multi-objective optimization problem we employ a Multi-Objective Particle Swarm Optimization (MOPSO) algorithm (see Manzoni et al., 2026, for example). Unlike single-objective optimization, which seeks a unique global minimum, MOPSO aims to identify the Pareto Front. This front represents the set of all non-dominated solutions, that are portfolios for which it is impossible to improve one objective (e.g., increasing returns) without simultaneously worsening another (e.g., increasing risk or CEV).

The MOPSO is an iterative population-based algorithm where a set of candidate solutions, called particles, "fly" through the search space. Each particle represents a potential weight vector w and adjusts its trajectory based on its own historical best position and the best positions found by its neighbors. In a multi-objective context, the algorithm maintains an external repository, which is a specialized archive that stores the non-dominated solutions discovered during the search. The size of the population (n_{pop}) determines the density of the search, while the repository size (n_{rep}) controls the granularity and diversity of the final Pareto Front. By selecting leaders from this repository to guide the swarm, the algorithm ensures that the particles converge toward the global trade-off surface while maintaining a well-distributed set of investment options. For more details about the MOPSO and the pseudo-algorithm scheme, please refer to Appendix C.

In order to validate the algorithm's performance, we follow an incremental testing approach. The MOPSO is first applied to a classical Mean-Variance optimization framework (dual-objective). Subsequently, the third objective, *Climate Exposure Volatility* (CEV), is introduced to analyze how climate constraints shift the efficient frontier.

The estimation of the required input parameters is conducted as follows:

- Expected Returns: the vector r is estimated using the monthly total returns of the assets over the 5-year period preceding the optimization date;

- **Variance-Covariance Matrix:** the matrix Σ is initially estimated from daily returns over the same 5-year window and subsequently rescaled to a monthly frequency. To ensure the matrix is well-conditioned and to prevent the smallest eigenvalues from being too close to zero, we apply the Ledoit-Wolf shrinkage method (Ledoit and Wolf, 2004). By employing this shrinkage estimator, we are able to reduce the condition number of the covariance matrix, calculated over the entire available dataset, to below 1,000. Maintaining a condition number of this magnitude is a widely accepted rule of thumb in financial econometrics (Won et al., 2013) to ensure that the matrix is well-conditioned and easily invertible, thereby enhancing the numerical stability of the optimization process;
- **CEV Parameters:** for the calculation of the *Climate Exposure Volatility* (f_3), the Bernoulli parameters representing the probabilities of climate events are derived from the data of the last available month prior to the optimization.

Finally, to manage the computational complexity associated with an investment universe of $N=120$ assets, the CEV expressions are implemented using matrix-based operations rather than iterative loops. This vectorized approach significantly reduces the computational cost and accelerates the convergence of the MOPSO algorithm.

4.3. Algorithm Parameter Selection

To identify the optimal configuration for the MOPSO algorithm, we conduct an analysis focused on three fundamental hyperparameters: the number of iterations ($\text{iter} \in \{200, 300, 400\}$), the population size ($\text{npop} \in \{200, 500, 600\}$), and the repository size ($\text{nrep} \in \{100, 200, 400\}$).

The selection of these parameters is crucial because the model is intended for a monthly rebalancing strategy. Given that the algorithm will be executed at the beginning of each month throughout a backtesting period, we must find a robust trade-off between the quality of the resulting Pareto front and the computational time required for execution.

To evaluate the performance of each combination, we utilize three key indicators:

- **Execution Time:** the total time in seconds required for the algorithm to complete the specified iterations;
- **Hypervolume:** a comprehensive performance indicator that quantifies the volume of the objective space dominated by the Pareto front relative to a pre-defined reference point (Guerreiro et al., 2021). This metric serves as a dual measure of both convergence and diversity; a higher Hypervolume signifies that the approximation set is closer to the true Pareto-optimal frontier and achieves a superior distribution of solutions across the objective space;
- **Spacing:** this indicator quantifies the uniformity of the solutions' distribution along the Pareto frontier by measuring the standard deviation of the distances between adjacent non-dominated vectors (Schott, 1995). A Spacing value approaching zero indicates that the solutions are evenly spread throughout the objective space. Such a uniform distribution is critical in multi-objective optimization, as it ensures the investor is presented with a smooth and continuous spectrum of trade-offs between financial performance, traditional risk, and climate resilience.

The algorithm was run for every possible combination of the aforementioned parameters. To ensure the robustness and statistical reliability of the findings, each entry in the following tables represents the average value obtained from 10 independent runs of the optimization process. The tables of this performance analysis are categorized by the number of iterations and reported in the sections below.

Table 4: Algorithm Performance for 200 Iterations (iter = 200)

nrep	npop	Time (s)	Hypervolume	Spacing
100	200	19.04	3.01×10^{-6}	8.01×10^{-4}
100	500	50.36	4.47×10^{-6}	9.26×10^{-4}
100	600	58.54	4.92×10^{-6}	9.39×10^{-4}
200	200	30.16	8.85×10^{-6}	7.17×10^{-4}
200	500	74.62	1.08×10^{-5}	7.21×10^{-4}
200	600	86.25	1.03×10^{-5}	7.50×10^{-4}
400	200	47.52	2.04×10^{-5}	5.54×10^{-4}
400	500	115.17	2.29×10^{-5}	5.70×10^{-4}
400	600	140.62	2.48×10^{-5}	6.15×10^{-4}

Table 5: Algorithm Performance for 300 Iterations (iter = 300)

nrep	npop	Time (s)	Hypervolume	Spacing
100	200	27.64	4.10×10^{-6}	9.06×10^{-4}
100	500	71.91	4.99×10^{-6}	9.16×10^{-4}
100	600	115.21	5.41×10^{-6}	8.82×10^{-4}
200	200	42.96	9.48×10^{-6}	7.50×10^{-4}
200	500	107.04	1.11×10^{-5}	7.65×10^{-4}
200	600	160.93	1.10×10^{-5}	7.35×10^{-4}
400	200	73.74	1.91×10^{-5}	5.57×10^{-4}
400	500	217.39	3.37×10^{-5}	6.16×10^{-4}
400	600	257.48	3.17×10^{-5}	6.10×10^{-4}

Table 6: Algorithm Performance for 400 Iterations (iter = 400)

nrep	npop	Time (s)	Hypervolume	Spacing
100	200	42.27	3.68×10^{-6}	8.97×10^{-4}
100	500	95.20	4.38×10^{-6}	9.53×10^{-4}
100	600	123.88	5.03×10^{-6}	8.80×10^{-4}
200	200	58.95	9.39×10^{-6}	7.21×10^{-4}
200	500	149.82	1.15×10^{-5}	7.26×10^{-4}
200	600	186.80	1.18×10^{-5}	7.43×10^{-4}
400	200	100.14	2.40×10^{-5}	5.99×10^{-4}
400	500	243.93	4.03×10^{-5}	6.30×10^{-4}
400	600	287.35	4.70×10^{-5}	6.61×10^{-4}

Based on the results reported in Tables 3, 4, and 5, we can derive several key insights regarding the impact of hyperparameters on the optimization process.

Increasing the repository size (nrep) from 100 to 200 and 400 leads to a major improvement in the optimization quality. Specifically, by moving from 100 to 200, the Hypervolume increases by an order of magnitude, and the improvement remains substantial when further expanding the repository to 400. For this reason, we choose the configuration with nrep = 400, which yields clearly superior results while maintaining a computationally acceptable execution time.

Looking at the population size we see that a substantial gain in Hypervolume is observed when moving from 200 to 500 particles, whereas the improvement becomes less significant when further increasing the size to 600. Hence, we choose npop = 500.

The number of iterations significantly influences the solution quality. While the Spacing metric remains relatively stable and independent of this parameter, the Hypervolume shows a marked increase from iter = 200 to iter = 300, and a further, though less pronounced, improvement from iter = 300 to iter = 400. As expected, a higher number of iterations entails a increase in total execution time.

Given these considerations, we narrowed our selection to four candidate configurations:

1. $\text{iter} = 300, \text{nrep} = 400, \text{npop} \in \{500, 600\};$
2. $\text{iter} = 400, \text{nrep} = 400, \text{npop} \in \{500, 600\}.$

A direct comparison of the Pareto fronts generated by these four alternatives showed that they are remarkably similar, yielding nearly identical results in terms of objective space coverage. Among these options, the configuration with $\text{iter} = 300, \text{nrep} = 400,$ and $\text{npop} = 500$ is the one we deem most appropriate for our analysis. While performance is consistent across the four candidates, this setup yields the lowest Spacing value and is characterized by a lower computational cost, thus representing what we consider the optimal trade-off between efficiency and solution quality.

In Figure 6, we present the resulting three-dimensional Pareto frontier computed as of January 1, 2020. The visualization maps the complex interactions between the three conflicting objectives: the x-axis represents the Portfolio Variance, the y-axis shows the *Climate Exposure Volatility*, and the z-axis denotes the Expected Returns.

The quality of the frontier, characterized by a highly uniform distribution of points, its smooth and concave shape, and its broad extension across the objective space, confirms the effectiveness of the chosen optimization algorithm and the appropriate calibration of its parameters. Furthermore, the visualization clearly illustrates that the classic mean-variance trade-off remains fundamental, even with the integration of *Climate Exposure Volatility*. It is evident that to achieve a higher target for expected returns, one must necessarily accept an increase in portfolio risk, specifically in its variance.




Figure 6: Pareto Frontier for the 3-objective optimization problem. For a detailed visualization of the surface from multiple perspectives, please refer to Appendix B.1.

Beyond the objective space, we investigate the composition of the portfolios belonging to the frontier. Specifically, out of a total universe of 120 stocks, we observe that the average number of active positions is 22. This calculation is based on a threshold of 10^{-3} , where weights below this value are considered negligible and treated

as zero. A deeper analysis reveals that for portfolios with low variance, this number tends to be higher, reaching approximately 55 active weights. However, as we move along the frontier toward higher variance levels, the number of non-zero weights decreases. This trend further confirms that a reduction in diversification leads to an increase in risk and potentially to higher returns.

In contrast, when considering a standard mean-variance frontier that neglects *Climate Exposure Volatility*, the average number of non-zero weights drops to 14 (the charts resulting from the classic Mean-Variance optimization can be found in the Appendix B.2). This significant difference suggests that incorporating an additional risk dimension into the optimization framework emphasizes diversification as a primary tool for risk mitigation, preventing the optimizer from over-concentrating in assets that may appear financially optimal but remain highly vulnerable to climate shocks.

To quantify the structural differences between the allocations, we compute the continuous Jaccard distance between the various portfolios on the frontier using the following formulation:

$$d_J(x, y) = 1 - \frac{\sum_i \min(x_i, y_i)}{\sum_i \max(x_i, y_i)}. \quad (20)$$

The results are presented in the Heatmap in Figure 7, where portfolios are sorted by increasing order of variance. The Heatmap reveals two critical insights. First, portfolios with significantly different variances tend to have highly distinct allocations, as evidenced by the high Jaccard distance values. Second, the interference or influence of the third objective *Climate Exposure Volatility* is clearly visible, as it creates local variations in allocation patterns that deviate from the standard Mean-Variance progression. For the Heatmap of the standard mean-variance optimization please refer to Appendix B.2.




Figure 7: Jaccard Distance Heatmap of the Portfolios in the Frontier.

In order to further investigate the relationship between *Climate Exposure Volatility* and the portfolio's Expected Return, we perform a sensitivity analysis by slicing the 3D Pareto frontier into distinct variance intervals. Specifically, the total range of traditional financial variance is divided into several sub-intervals. For each interval, we select the particles (portfolios) falling within that specific range and project them onto a 2D scatter plot.

In these visualizations, the x-axis represents the CEV, while the y-axis represents the Expected Return. To provide a clearer trend analysis, we also plot an interpolated parabolic curve calculated at the mean variance value of the interval.

The results, as shown in Figure 8, reveal a clear and consistent pattern: an increase in Expected Return is systematically associated with an increase in CEV. This confirms that *Climate Exposure Volatility* acts as an additional dimension of risk. Just as traditional variance measures financial instability, CEV quantifies the portfolio's vulnerability to climate-related shocks.

This analysis demonstrates that CEV behaves much like traditional variance in relation to expected return. Investors seeking higher performance must be prepared to accept not only greater market volatility but also higher climate variance, highlighting the conflicting nature of green financial optimization.




Figure 8: Expected Return vs. *Climate Exposure Volatility*: 2D Cross-Sectional Pareto Projections.

Additionally, we construct a second Pareto frontier by applying the same methodology, replacing *Climate Exposure Volatility* (CEV) with *Climate Risk Exposure* (CRE). The results, shown in Appendix B.3, yield very similar insights, as the general shape of the frontier remains consistent despite the obvious differences in absolute values.

A comparable relationship is observed between Expected Return and CRE, mirroring the trade-offs identified in the Return-CEV space. Specifically, to achieve higher expected returns, the investor must be willing to accept higher levels of CRE, further identifying this metric as a distinct dimension of risk.

5. Backtesting Analysis of Climate-Aware Efficient Frontiers

In this section, we evaluate the potential benefits of selecting portfolios from the Pareto Frontier derived in Section 4, which explicitly incorporates climate risk. Specifically, the backtesting analysis covers the period from January 1, 2020, to April 30, 2025.

On the first business day of each month within this window, the Pareto Frontier is re-optimized following the methodology previously described: we simultaneously minimize the Market Variance and the *Climate Exposure Volatility*, while maximizing the Expected Return. The input parameters are updated monthly using a 5-year rolling-window approach. The Expected Returns are calculated as the mean of the monthly returns over the preceding five years, and the Covariance Matrix is estimated using daily returns from the same five-year look-back period. Furthermore, the temperature anomaly probabilities for each continent are updated using the most

recent available data, specifically those recorded in the month immediately preceding the rebalancing date. Following this monthly re-optimization on the first business day, the performance of the selected portfolios is evaluated at the end of the month. This cycle is repeated iteratively: at the start of each subsequent month, the portfolios are recalculated and rebalanced based on the newly updated Pareto Frontier, ensuring that the asset allocation remains aligned with the most recent market and climate data.

To provide a comprehensive benchmark, our analysis includes Market-Cap Weighted portfolio and the Equally-Weighted portfolio. Furthermore, we examine specific points on the Pareto Frontier, such as the minimum market variance portfolio and the maximum expected return portfolio.

To identify additional portfolios on the frontier that balance all three competing objectives, i.e. Expected Return, Variance, and CEV, we apply a Min-Max normalization to each objective. This ensures that all components are scaled within the $(0,1)$ range while preserving their relative ordering. We then combine these normalized objectives into a single scalar function using a convex combination governed by three positive parameters (a, b, c) such that $a + b + c = 1$. Since our optimization framework is structured as a minimization problem, we select the portfolio on the Pareto Frontier that yields the minimum value of this combined objective function. This approach allows us to explore different strategic tilts by varying the weights assigned to each objective. Specifically, we test the following weight combinations (a, b, c) :

- $(1/3, 1/3, 1/3)$: Balanced approach;
- $(0, 1/2, 1/2)$: Minimum risk (Financial and Climate);
- $(1/2, 1/4, 1/4)$: Return-oriented;
- $(1/4, 1/2, 1/4)$: Variance-oriented;
- $(1/4, 1/4, 1/2)$: Climate-oriented;
- $(1/2, 1/2, 0)$: Mean-Variance focus.

The results are presented in Figure 9, which illustrates the cumulative returns of the various portfolio strategies over the analyzed period.




Figure 9: Cumulative returns of the Pareto-optimized portfolios and benchmarks (January 2020 – April 2025).

To better evaluate the risk-adjusted performance of these portfolios, we measure several key financial metrics: the Compound Annual Growth Rate (CAGR), the Sharpe Ratio, the Sortino Ratio, the Maximum Drawdown and the average *Climate Exposure Volatility* (CEV) over time. In our analysis, we assume a risk-free rate of

zero for the calculation of both the Sharpe and Sortino ratios. Consequently, the Sharpe Ratio is calculated simply as the ratio of the expected return to the total volatility while the Sortino Ratio as the expected return over the downside volatility. The latter is computed only using the negative returns, since price increases over the expected return are actually beneficial to investors. Finally, the Maximum Drawdown represents the largest percentage drop from a peak to a trough, indicating the worst-case loss experienced during the period. These metrics are summarized in the Table 7 to allow for a direct comparison between the various portfolios taken into consideration.

Table 7: Performance Metrics Comparison (January 2020 – April 2025)

Strategy	CAGR	Sharpe Ratio	Sortino Ratio	Max Drawdown	CEV
Market-Cap Weights	15.75%	1.041	1.850	14.78%	0.0700
Equally-Weighted	17.49%	1.153	1.834	16.72%	0.0671
Min Market Variance	14.20%	1.282	2.324	11.47%	0.0660
Max Expected Return	3.32%	0.320	0.674	69.37%	0.1025
(1/3, 1/3, 1/3)	33.54%	1.153	2.787	31.54%	0.0610
(0, 1/2, 1/2)	14.65%	1.133	1.304	17.38%	0.0603
(1/2, 1/4, 1/4)	28.43%	0.920	2.077	46.10%	0.0618
(1/4, 1/2, 1/4)	30.57%	1.312	3.029	20.78%	0.0611
(1/4, 1/4, 1/2)	28.47%	1.045	2.382	30.57%	0.0605
(1/2, 1/2, 0)	33.35%	1.136	2.356	32.81%	0.0763

The analysis of the performance metrics provides several key insights into the behavior of the different strategies. The Maximum Expected Return portfolio exhibits the worst performance across all evaluated metrics. This shows that selecting a portfolio solely based on expected returns, without incorporating any risk measure, leads to extreme volatility and risk exposure. Indeed, it records the worst Maximum Drawdown of 69.37%, a Sharpe Ratio of only 0.320, which is significantly lower than all other strategies and the highest average CEV, also failing to address the volatility of extreme climate events.

The Equally-Weighted and Market-Cap Weighted portfolios, used as benchmarks, show similar trends, reflecting the general market performance over the last 5 years. In contrast, the Minimum Market Variance portfolio, while achieving a lower CAGR compared to the benchmarks, turns out to be the most effective in controlling market volatility. It records a Maximum Drawdown of 11.47%, the best among all observed portfolios, highlighting its superior risk-management capabilities.

Similarly, the (0, 1/2, 1/2) strategy, which ignores expected returns to assign equal weight to financial and climate variance, displays a profile very close to the Minimum Variance portfolio, achieving the second best Maximum Drawdown at 17.38% and the lowest average CEV.

When considering portfolios on the frontier that include Expected Return in their objective function, we observe the highest CAGR values. The (1/3, 1/3, 1/3) strategy, which weights all three objectives equally, achieves the highest overall CAGR (33.54%). The (1/2, 1/2, 0) portfolio shows very similar results; although it does not explicitly weight the *Climate Exposure Volatility* in the convex combination (achieving indeed the second worst average CEV), it still benefits from being selected from a Pareto Frontier that was originally constructed including this climate risk metric.

While these CAGR values (around 30%) may appear exceptionally high, it is important to note that this backtest does not account for transaction costs, bid-ask spreads, and in general rebalancing costs, which would realistically reduce net performance.

Finally, the remaining three portfolios show distinct behaviors. The (1/2, 1/4, 1/4) strategy records the second-worst Maximum Drawdown (46.10%), suggesting that its weights may over-prioritize returns at the expense of risk mitigation. Conversely, the (1/4, 1/2, 1/4) portfolio emerges as arguably the most balanced choice, as it yields the best results across all five metrics simultaneously, suggesting that this specific weight distribution provides one of the most effective compromise between the three competing objectives. In conclusion, the backtesting results underscore the fundamental trade-offs inherent in multi-objective portfolio optimization. While strategies focused exclusively on expected returns prove to be unsustainable due to extreme drawdowns, the integration of climate risk metrics does not necessarily come at the expense of financial performance. On the contrary, portfolios that balance market risk, climate exposure, and expected returns, most notably the (1/4, 1/2, 1/4) and (1/3, 1/3, 1/3) configurations outline superior resilience and efficiency. These findings suggest that incorporating climate risk into the optimization process acts as a robust diversifier, enabling the construction of portfolios that are better equipped to navigate the complexities of the modern financial and environmental landscape.

6. Conclusions

This study provides a flexible and robust quantitative foundation for climate-resilient portfolio construction, offering a viable tool for asset managers to navigate the escalating challenges of physical climate risk.

The work investigates the impact of physical climate risk on global equity portfolios, bridging the gap between theoretical climate science and practical investment management. Our empirical analysis confirms that extreme temperature anomalies represent a material financial threat, exerting a statistically significant negative pressure on asset returns across most economic sectors (see Table 2). This financial impact is deeply rooted in the non-stationary and intensifying nature of environmental hazards (see Figure 1); indeed, our results show that climate risk is not a static outlier but a dynamic variable that has significantly escalated in recent decades. Consequently, these findings invalidate the traditional assumption of a constant risk distribution, highlighting the urgent need for a dynamic approach to portfolio risk management that accounts for the evolving frequency and intensity of climate shocks.

The core contribution of this research lies in the introduction of two dynamic metrics: *Climate Risk Exposure* (CRE) and *Climate Exposure Volatility* (CEV) (see Section 3). Unlike existing static country-level indices, these measures capture the time-varying probability of extreme events and the systemic risk arising from their spatial correlation across continents. Furthermore, by incorporating the asset intensity of each firm, these metrics account for firm-specific vulnerability, reflecting the varying degrees of dependency on physical assets. This effectively distinguishes between capital-heavy sectors like Real Estate and more agile industries like Retailers, while considering the precise geographical distribution leveraging corporate revenues.

By integrating these measures into a Multi-Objective Particle Swarm Optimization (MOPSO) framework (see Section 4), we demonstrate that investors can explicitly account for climate resilience within the mean-variance paradigm, without necessarily sacrificing competitive Expected Return targets. The resulting 3D Pareto frontier (see Figure 6) illustrates the non-linear trade-offs between Expected Returns, Market Variance, and Climate Risk, proving that active geographic optimization is an efficient tool for minimizing vulnerability to physical shocks. Finally, our backtesting analysis (see Section 5) confirms the practical benefits of this approach, showing that climate-aware strategies can enhance portfolio resilience in a dynamic market environment, generating competitive performance compared to traditional benchmarks. In particular, the results suggest that integrating the *Climate Exposure Volatility* does not necessarily impair financial returns; rather, it provides an additional layer of risk management that helps identify more efficient portfolios along the Pareto Frontier, offering ways to reduce both overall volatility and maximum drawdown while maintaining a high target of Expected Return. The methodology proposed in this work is designed to be modular, offering several significant avenues for future research. First, extending the analysis from a continental scale to a national or sub-national level would allow for a more granular identification of localized climate risks. Such a shift would enable investors to manage the specific climatic threats of each country in detail, thereby unlocking the full potential of geographic diversification to mitigate idiosyncratic shocks.

Moreover, the framework is intrinsically adaptable to a broader spectrum of environmental phenomena. This approach could be extended to specific natural disasters, such as floods, sea-level rise, or wildfire activity, by constructing customized Bernoulli time series for each hazard. Furthermore, while this study explored triple-objective optimization, a promising development would involve a four-objective optimization framework. This would allow for the simultaneous consideration of Expected Return, Market Variance, *Climate Risk Exposure* (CRE) and *Climate Exposure Volatility* (CEV), offering an even more comprehensive mapping of the trade-offs between financial performance and various dimensions of climate risk. The joint inclusion of CRE and CEV as distinct objectives is particularly advantageous for a robust portfolio strategy. While CEV accounts for the cross-correlations between extreme events effectively guiding allocation toward geographically diversified areas with lower co-occurrence of shocks, relying on CEV alone could also lead to counterintuitive results. Specifically, since the variance of a Bernoulli process decreases as the probability p exceeds 0.5, an optimizer focused solely on minimizing CEV might inadvertently favor geographic areas with an extremely high risk of climate events. By integrating CRE as a separate objective, the framework ensures that the pursuit of lower correlation and variance does not come at the expense of an overall increase in the expected frequency of climate shocks, thereby balancing geographical diversification with absolute risk mitigation.

Finally, future studies could refine the statistical modeling of dependencies by allowing the correlation structure to be time-dependent. Since the probabilities of extreme events are non-stationary, modeling time-varying correlations would provide a more dynamic representation of how climatic interdependencies evolve over time. Such advancements would further strengthen the theoretical foundation of this study, providing asset managers with increasingly robust tools to navigate the escalating challenges of physical climate risk.

References

- Aprea, I. L. and Sbaiz, G. (2025). A neural network-particle swarm solver for sustainable portfolio optimization problems. *Decisions in Economics and Finance*, pages 1–40.
- Avramov, D., Cheng, S., Lioui, A., and Tarelli, A. (2022). Sustainable investing with ESG rating uncertainty. *Journal of Financial Economics*, 145(2):642–664.
- Azzone, M., Barucci, E., and Stocco, D. (2026a). Asset management with an ESG mandate. *Journal of Banking & Finance*, page 107640.
- Azzone, M., Ghesini, M., Stocco, D., and Viola, L. (2026b). Physical climate risk in asset management. *Corporate Social Responsibility and Environmental Management*, 33(1):1077–1094.
- Bolton, P. and Kacperczyk, M. (2023). Global pricing of carbon-transition risk. *The Journal of Finance*, 78(6):3677–3754.
- Bolton, P., Kacperczyk, M., and Samama, F. (2022). Net-zero carbon portfolio alignment. *Financial Analysts Journal*, 78(2):19–33.
- Bortolan, L., Dey, A., and Taschini, L. (2024). Volatile temperatures and their effects on equity returns and firm performance. Technical report, CESifo Working Paper.
- Bressan, G., Đuranović, A., Monasterolo, I., and Battiston, S. (2024). Asset-level assessment of climate physical risk matters for adaptation finance. *Nature Communications*, 15(1):5371.
- Campiglio, E., Daumas, L., Monnin, P., and von Jagow, A. (2023). Climate-related risks in financial assets. *Journal of Economic Surveys*, 37(3):950–992.
- Coello, C. A. C., Pulido, G. T., and Lechuga, M. S. (2004). Handling multiple objectives with particle swarm optimization. *IEEE Transactions on evolutionary computation*, 8(3):256–279.
- De Spiegeleer, J., Höcht, S., Jakubowski, D., Reyners, S., and Schoutens, W. (2023). ESG: A new dimension in portfolio allocation. *Journal of Sustainable Finance & Investment*, 13(2):827–867.
- Durillo, J. J., García-Nieto, J., Nebro, A. J., Coello, C. A. C., Luna, F., and Alba, E. (2009). Multi-objective particle swarm optimizers: An experimental comparison. In *International conference on evolutionary multi-criterion optimization*, pages 495–509. Springer.
- Graff Zivin, J. and Neidell, M. (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 32(1):1–26.
- Guerreiro, A. P., Fonseca, C. M., and Paquete, L. (2021). The hypervolume indicator: Computational problems and algorithms. *ACM Computing Surveys (CSUR)*, 54(6):1–42.
- Huber, M. and Marić, N. (2015). Multivariate distributions with fixed marginals and correlations. *Journal of Applied Probability*, 52(2):602–608.
- Huber, M. and Marić, N. (2019). Admissible Bernoulli correlations. *Journal of Statistical Distributions and Applications*, 6(1):2.
- Jones, M. W., Kelley, D. I., Burton, C. A., Di Giuseppe, F., Barbosa, M. L. F., Brambleby, E., Hartley, A. J., Lombardi, A., Mataveli, G., McNorton, J. R., et al. (2024). State of wildfires 2023–2024. *Earth System Science Data*, 16(8):3601–3685.
- Kahn, M. E., Ouazad, A., and Yönder, E. (2024). Adaptation using financial markets: Climate risk diversification through securitization. Technical report, National Bureau of Economic Research.
- Katz, D. and McIntosh, L. (2021). Corporate governance update: “Materiality” in America and Abroad. In *Harvard Law School Forum on Corporate Governance*. The President and Fellows of Harvard College Cambridge, MA.
- Kennedy, J. and Eberhart, R. (1995). Particle swarm optimization. In *Proceedings of ICNN’95-international conference on neural networks*, volume 4, pages 1942–1948. IEEE.

- Kumar, D., Kumar, V., et al. (2020). Impact of controlling parameters on the performance of MOPSO algorithm. *Procedia Computer Science*, 167:2132–2139.
- Le Guenedal, T. and Roncalli, T. (2022). Portfolio construction with climate risk measures. *Available at SSRN 3999971*.
- Ledoit, O. and Wolf, M. (2004). Honey, i shrunk the sample covariance matrix. *The Journal of Portfolio Management*, 30(4):110–119.
- Manzoni, L., Puzzoli, L., and Sbaiz, G. (2026). Sustainable multi-objective optimization problem for environmental impact and electricity production analysis of dams. *Energy Systems*, pages 1–30.
- Moore, J. (1999). Application of particle swarm to multiobjective optimization. *Technical report*.
- Morelli, G. (2024). Responsible investing and portfolio selection: a Shapley-CVaR approach. *Annals of Operations Research*, 342(3):1991–2019.
- Neidell, M., Graff Zivin, J., Sheahan, M., Willwerth, J., Fant, C., Sarofim, M., and Martinich, J. (2021). Temperature and work: Time allocated to work under varying climate and labor market conditions. *PloS one*, 16(8):e0254224.
- Newey, W. K. and West, K. D. (1986). A simple, positive semi-definite, heteroskedasticity and autocorrelation-consistent covariance matrix. *Technical Working Paper*.
- Pedersen, L. H., Fitzgibbons, S., and Pomorski, L. (2021). Responsible investing: The ESG-efficient frontier. *Journal of Financial Economics*, 142(2):572–597.
- Reyes-Sierra, M., Coello, C. C., et al. (2006). Multi-objective particle swarm optimizers: A survey of the state-of-the-art. *International journal of computational intelligence research*, 2(3):287–308.
- Schott, J. R. (1995). Fault tolerant design using single and multicriteria genetic algorithm optimization. Technical report, AIR FORCE INST OF TECH WRIGHT-PATTERSON AFB OH.
- Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance*, 19(3):425–442.
- Starr, M. (2000). The effects of weather on retail sales. *Available at SSRN 221728*.
- Wang, F., Harindintwali, J. D., Wei, K., Shan, Y., Mi, Z., Costello, M. J., Grunwald, S., Feng, Z., Wang, F., Guo, Y., et al. (2023). Climate change: Strategies for mitigation and adaptation. *The Innovation Geoscience*, 1(1):100015–1.
- Won, J.-H., Lim, J., Kim, S.-J., and Rajaratnam, B. (2013). Condition-number-regularized covariance estimation. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 75(3):427–450.

A. Proof of Correlation Bounds for Bernoulli Variables

In this section of the Appendix we provide the full proof of the corollary regarding the Correlation Bounds for two Bernoulli Random variables.

Corollary. The correlation coefficient ρ between two Bernoulli random variables $X_1 \sim \text{Bern}(p_1)$ and $X_2 \sim \text{Bern}(p_2)$ is bounded by $\rho_{\min} \leq \rho \leq \rho_{\max}$, where

$$\rho_{\min} = \frac{\max(0, p_1 + p_2 - 1) - p_1 p_2}{\sqrt{p_1(1-p_1)}\sqrt{p_2(1-p_2)}} \quad (21)$$

$$\rho_{\max} = \frac{\min(p_1, p_2) - p_1 p_2}{\sqrt{p_1(1-p_1)}\sqrt{p_2(1-p_2)}}. \quad (22)$$

Proof. To derive the explicit forms of $\rho_{\min}$ and $\rho_{\max}$, we start from the Fréchet–Hoeffding bounds

$$\text{corr}(F_1^{-1}(U), F_2^{-1}(1-U)) \leq \text{corr}(X_1, X_2) \leq \text{corr}(F_1^{-1}(U), F_2^{-1}(U)), \quad (23)$$

with $U \sim \text{Unif}([0, 1])$, and we apply it to two Bernoulli Random variable $X_1 \sim \text{Bern}(p_1)$ and $X_2 \sim \text{Bern}(p_2)$ with respective cumulative distribution function F_1^{-1} and F_2^{-1} .

For a Bernoulli variable $X \sim \text{Bern}(p)$, the quantile function is

$$F^{-1}(u) = \begin{cases} 0 & \text{if } 0 \leq u < 1 - p \\ 1 & \text{if } 1 - p \leq u \leq 1. \end{cases} \quad (24)$$

First of all, to estimate the upper bound ρ_{max} , we consider the product $F_1^{-1}(u)F_2^{-1}(u)$ which is equal to 1 if and only if

$$u \geq 1 - p_1 \quad \text{and} \quad u \geq 1 - p_2 \implies u \geq \max(1 - p_1, 1 - p_2) = 1 - \min(p_1, p_2) \quad (25)$$

and integrating over $[0, 1]$

$$E[F_1^{-1}(U)F_2^{-1}(U)] = \int_{1-\min(p_1, p_2)}^1 1 \, du = \min(p_1, p_2). \quad (26)$$

Thus, the $\text{Cov}(F_1^{-1}(U), F_2^{-1}(U))$ is

$$\begin{aligned} \text{Cov}(F_1^{-1}(U), F_2^{-1}(U)) &= E[F_1^{-1}(U)F_2^{-1}(U)] - E[F_1^{-1}(U)] \cdot E[F_2^{-1}(U)] \\ &= E[F_1^{-1}(U)F_2^{-1}(U)] - E[X_1]E[X_2] \\ &= \min(p_1, p_2) - p_1p_2. \end{aligned} \quad (27)$$

Dividing by the product of standard deviations $\sigma_1\sigma_2 = \sqrt{p_1(1-p_1)}\sqrt{p_2(1-p_2)}$, we obtain $\text{corr}(F_1^{-1}(U), F_2^{-1}(U))$, leading to

$$\rho_{max} = \frac{\min(p_1, p_2) - p_1p_2}{\sqrt{p_1(1-p_1)}\sqrt{p_2(1-p_2)}}. \quad (28)$$

Now, to calculate the lower bound ρ_{min} , we examine the product $F_1^{-1}(u)F_2^{-1}(1-u)$ that is 1 if and only if

$$u \geq 1 - p_1 \quad \text{and} \quad 1 - u \geq 1 - p_2 \implies 1 - p_1 \leq u \leq p_2. \quad (29)$$

This interval has a non-zero length only if $1 - p_1 \leq p_2$, which is equivalent to $p_1 + p_2 - 1 \geq 0$. The length of the interval is $\max(0, p_2 - (1 - p_1)) = \max(0, p_1 + p_2 - 1)$.

Integrating over $[0, 1]$

$$E[F_1^{-1}(U)F_2^{-1}(1-U)] = \int_{1-p_1}^{p_2} 1 \, du = \max(0, p_1 + p_2 - 1) \quad (30)$$

and thus, the $\text{Cov}(F_1^{-1}(U), F_2^{-1}(U))$ is

$$\begin{aligned} \text{Cov}(F_1^{-1}(U), F_2^{-1}(1-U)) &= E[F_1^{-1}(U)F_2^{-1}(1-U)] - E[F_1^{-1}(U)] \cdot E[F_2^{-1}(1-U)] \\ &= E[F_1^{-1}(U)F_2^{-1}(1-U)] - E[X_1]E[X_2] \\ &= \max(0, p_1 + p_2 - 1) - p_1p_2. \end{aligned} \quad (31)$$


Dividing by the product of standard deviations $\sigma_1\sigma_2 = \sqrt{p_1(1-p_1)}\sqrt{p_2(1-p_2)}$, we obtain $\text{corr}(F_1^{-1}(U), F_2^{-1}(1-U))$, leading to

$$\rho_{min} = \frac{\max(0, p_1 + p_2 - 1) - p_1p_2}{\sqrt{p_1(1-p_1)}\sqrt{p_2(1-p_2)}}. \quad (32)$$


B. Multi-objective optimization

B.1. The 3D Pareto Frontier: Exploring More Perspectives


In this section of the Appendix, additional viewpoints of the Pareto frontier, calculated as of January 1, 2020, are presented to better visualize the multidimensional nature of the optimal solutions. Following the methodology described in Section 4, this frontier is constructed within a multi-objective optimization framework designed to simultaneously minimize both Market Variance and *Climate Exposure Volatility* (CEV), while maximizing the Expected Return. By observing these diverse perspectives, it is possible to appreciate the complex trade-offs between traditional financial risk, climate-related vulnerability, and potential growth across the set of optimal portfolios.



(a) Perspective 1



(b) Perspective 2




(c) Perspective 3

B.2. Mean-Variance Optimization: Standard Efficient Frontier


In this section, we present the results of the standard Mean-Variance optimization, illustrating the classical Efficient Frontier and the inherent trade-off between Expected Return and Market Variance. This framework serves as a fundamental baseline for the multidimensional analysis discussed in the previous sections.

The resulting frontier in Figure 11a is remarkably broad, spanning a wide range of return and variance values. It clearly exhibits the classical risk-return trade-off: as investors seek higher expected returns, they must inevitably accept a higher level of financial variance, representing the increased uncertainty of those returns. Within this bi-dimensional space, any point on the frontier represents the maximum possible return for a given level of risk, or conversely, the minimum risk for a target level of return.

To further investigate the composition of these optimal solutions, we analyze the Jaccard Distance Heatmap in Figure 11b, where portfolios are ordered by increasing levels of variance. The heatmap reveals a strong correlation between financial characteristics and portfolio concentration: as the distance between two portfolios' variance (and consequently their expected return) increases, their compositions become significantly more distinct. Specifically, the Jaccard distance shows that portfolios at opposite ends of the frontier, such as the Minimum Variance and the Maximum Return portfolios, share very few assets in common. This confirms that moving along the frontier does not merely involve a re-weighting of the same assets, but a structural shift in the portfolio's selection to accommodate the escalating risk-return requirements.



(a) Mean-Variance Pareto Frontier



(b) Mean-Variance Jaccard Distance Heatmap

B.3. Triple-Objective Optimization: Integrating Climate Risk Exposure


In this section of the Appendix we present the results of the multi-objective optimization performed as of January 1, 2020, where we utilize the *Climate Risk Exposure* (CRE) instead of the *Climate Exposure Volatility* (CEV) to evaluate the portfolio's climate-related risk. By adopting CRE as our objective, we aim to measure the absolute magnitude of the portfolio's vulnerability to extreme temperature anomalies, providing an alternative perspective on environmental risk.

The optimization framework simultaneously considers three competing objectives: Expected Return, Market Variance, and *Climate Risk Exposure* (CRE). In line with the methodology described in Section 4, Expected Returns and the Financial Covariance Matrix are estimated using a 5-year rolling window of historical data. Furthermore, the probabilities of extreme temperature anomalies for each continent used in the CRE are sourced from the latest available data prior to the optimization date, ensuring that the risk assessment reflects the most current environmental trends.


The analysis is supported by three key visualizations. The Figure 12a displays the 3D Pareto Frontier, which appears dense, well distributed and convex. This surface, similarly to the one in Figure 6, illustrates the fundamental trade-offs in a multidimensional space, where any improvement in climate resilience or financial stability necessitates a compromise in expected return.

To analyze the structural composition of these optimal solutions, the Figure 12b presents a Jaccard Distance Heatmap, where portfolios are ordered by increasing Market Variance. The heatmap clearly highlights that portfolios with different objective values possess distinct asset compositions; as the distance between portfolios along the frontier increases, their constituent weights diverge significantly, confirming that the optimization effectively explores diverse regions of the asset space.

Finally, we investigate the specific relationship between *Climate Risk Exposure* and Expected Return by slicing the 3D frontier into distinct variance intervals and projecting the points onto 2D scatter plots. The results can be observed in Figure 13, where we observe a consistent pattern: an increase in Expected Return is systematically associated with an increase in *Climate Risk Exposure*. This confirms that CRE acts as an additional dimension of risk, analogous to Traditional Variance. Much like financial instability, CRE quantifies a portfolio's vulnerability to environmental shocks, demonstrating that investors seeking higher performance must accept not only greater market volatility but also heightened climate risk. This highlights how *Climate Risk Exposure* represents a high-quality dimension of climate risk, serving as a robust alternative to *Climate Exposure Volatility* in quantifying the portfolio's overall climate-related vulnerability.



(a) 3D Pareto Frontier integrating CRE



(b) Jaccard Distance Heatmap for the CRE Optimization




Figure 13: Expected Return vs. *Climate Risk Exposure*: 2D Cross-Sectional Pareto Projections.

C. Multi-Objective Particle Swarm Optimization (MOPSO)

The objective of this appendix section is to delve into the details of the Multi-Objective Particle Swarm Optimization (MOPSO), the optimization algorithm employed to compute the Pareto fronts for the portfolio optimization problems discussed in this work. Specifically, this section provides comprehensive information regarding the algorithm's parameters and their functional roles, the specific values utilized in our simulations (see Table 8), and a pseudocode representation illustrating the algorithm's execution flow (see Algorithm 1). Unlike single-objective optimization, which seeks a unique global minimum, MOPSO aims to identify the Pareto Front. This front represents the set of all non-dominated solutions (portfolios) for which it is impossible to

improve one objective (e.g., increasing returns) without simultaneously worsening another (e.g., increasing risk).

MOPSO is an iterative, population-based algorithm where a set of candidate solutions, termed particles, "fly" through the multi-dimensional search space. Each particle represents a potential weight vector w and adjusts its trajectory based on two primary influences: its own historical best position and the best positions discovered by its neighbors retrieved from an external repository. This external repository acts as a specialized archive to store non-dominated solutions found during the search.

The performance of the MOPSO algorithm in identifying the optimal portfolio frontier is closely tied to its parameters, which balance the search between exploring new asset allocations and refining existing high-performing ones.

The scale of the search is defined by the population size (n_{pop}) and the number of iterations ($iter$). While n_{pop} represents the total number of candidate particles (portfolios) "flying" through the search space simultaneously, the iteration count determines the duration of the optimization process. Specifically, a higher n_{pop} enhances the exploration of the search space, allowing the algorithm to evaluate a broader range of objective function values. On the other hand, a larger number of iterations ($iter$) is crucial for the refinement of these solutions, enabling the algorithm's empirical Pareto Front to converge more closely toward the true, theoretical Pareto Front of the optimization problem. To manage the non-dominated solutions discovered, the repository size (n_{rep}) limits the number of portfolios stored in the final archive, ensuring that the resulting Pareto Front remains manageable yet representative of the entire trade-off spectrum.

To guarantee a well-distributed frontier, avoiding clusters of portfolios with nearly identical risk-return profiles, the algorithm employs a grid system (n_{grid}) and an inflation rate (α). The grid system works by partitioning the objective space into hypercubes (cells). This discretization allows the algorithm to monitor the density of solutions across the frontier. By tracking how many portfolios fall into each cell, the MOPSO can steer the swarm toward less-populated regions, ensuring that the final output is not just a collection of optimal points, but a uniform representation of the entire risk-return trade-off.

In this context, the inflation rate (α) is used to slightly expand the boundaries of the grid beyond the actual minimum and maximum objective values found in the repository (s_{min} and s_{max}). Specifically, it calculates an offset $dc = (s_{max} - s_{min}) \cdot \alpha$, which is subtracted from the minimum and added to the maximum values. Moreover, the selection pressure (β) and deletion pressure (γ) parameters are used to favor solutions in less-populated areas of the grid. This mechanism is crucial for portfolio optimization, as it forces the swarm to identify a diverse range of investment options, from minimum-variance portfolios to those targeting maximum returns.

The core of the MOPSO optimization process lies in the velocity update rule, which determines how each particle, representing a specific portfolio allocation, adjusts its movement through the search space. The direction and magnitude of a particle's displacement are governed by a stochastic velocity equation that balances three fundamental forces: inertia, social influence, and cognitive memory. Specifically, the velocity update combines the particle's previous velocity, modulated by the inertia weight ω , with two acceleration components directed toward one of the global best positions found by the swarm ($gbest_{pos}$) and the individual particle's own historical best position ($pbest_{pos}$).

The inertia component, ω , provides the necessary momentum to maintain the particle's previous direction, ensuring stability in the search and preventing erratic changes that could lead to sub-optimal convergences. The social component, weighted by the coefficient c_1 and a random variable, pulls the particle toward a leader selected from the non-dominated repository. This selection, influenced by the pressure parameters β and γ , ensures that the swarm gravitates toward the most promising and less explored regions of the current Pareto Front. Simultaneously, the cognitive component, weighted by c_2 , attracts the particle back to its personal best configuration. The interaction of these vectors, each scaled by random stochastic factors, defines both the trajectory and the intensity of the search. By modulating these weights, the algorithm effectively balances global exploration with local refinement, allowing the particles to adaptively shift the portfolio weights until the empirical frontier converges toward the theoretical optimum.

Finally, the mutation rate (μ) introduces random variations in the positions of the particles, a vital step to maintain genetic diversity within the swarm and prevent the algorithm from converging prematurely on local sub-optimal frontiers.

Table 8: MOPSO Algorithm Parameters and Configuration

Parameter	Description	Value
n_{pop}	Population size	500
$iter$	Number of iterations	300
n_{rep}	External repository size	400
n_{grid}	Number of grids per dimension	10
α	Grid inflation rate	0.1
β	Selection pressure (global best selection)	1
γ	Deletion pressure (repository maintenance)	1
μ	Mutation rate	2
ω	Inertia weight	0.8
c_1	Personal learning coefficient	1.5
c_2	Global learning coefficient	1.5
lb	Lower bound of asset weights	0
ub	Upper bound of asset weights	1

```

1: Initialization:
2:  $pos \leftarrow$  list of  $npop$  particles random initialized
3: ensure particles' components are within bounds  $(lb, ub)$ 
4: ensure particles' components sum up to 1
5:  $vel \leftarrow$  list of particle velocities initialized to 0
6:  $objs \leftarrow$  list of objective values of the particles
7:  $pbest \leftarrow objs$ ,  $pbest\_pos \leftarrow pos$ 
8:  $rep, rep\_obj \leftarrow$  non-dominated particles from  $\{pos, objs\}$ 
9:  $grid \leftarrow create\_grid(rep\_obj, ngrid, \alpha)$ 
10: Main Loop:
11: for  $t = 0$  to  $iter - 1$  do
12:    $pm \leftarrow (1 - t/(iter - 1))^{1/\mu}$  (Mutation probability)
13:   for  $i = 0$  to  $npop - 1$  do
14:     Update Position:
15:      $gbest\_pos \leftarrow rep[select\_gbest(\beta)]$ 
16:      $vel[i] \leftarrow \omega vel[i] + c_1 r_1 (gbest\_pos - pos[i]) + c_2 r_2 (pbest\_pos[i] - pos[i])$ 
17:      $pos[i] \leftarrow pos[i] + vel[i]$ 
18:     ensure components of  $pos[i]$  are within bounds  $(ub, lb)$ 
19:     ensure components of  $pos[i]$  sum up to 1
20:      $objs[i] \leftarrow cal\_obj(pos[i])$  (update the objective values of the particle)
21:     Mutation:
22:     if  $rand() < pm$  then
23:        $new\_pos \leftarrow mutate(pos[i])$ 
24:        $new\_obj \leftarrow cal\_obj(new\_pos)$ 
25:       if  $new\_obj$  dominates  $objs[i]$  or  $(objs[i]$  not dominates  $new\_obj$  and  $rand() < 0.5$ ) then
26:          $pos[i] \leftarrow new\_pos$ ,  $objs[i] \leftarrow new\_obj$ 
27:         ensure components of  $pos[i]$  are within bounds  $(ub, lb)$ 
28:         ensure components of  $pos[i]$  sum up to 1
29:       end if
30:     end if
31:     Update Personal Best:
32:     if  $objs[i]$  dominates  $pbest[i]$  or  $(pbest[i]$  not dominates  $objs[i]$  and  $rand() < 0.5$ ) then
33:        $pbest[i] \leftarrow objs[i]$ ,  $pbest\_pos[i] \leftarrow pos[i]$ 
34:     end if
35:   end for
36:   Update Repository and Grid:
37:    $rep.extend(pos)$ ,  $rep\_obj.extend(objs)$ 
38:   remove dominated solutions from  $rep$  and  $rep\_obj$ 
39:    $grid \leftarrow create\_grid(rep\_obj, ngrid, \alpha)$ 
40:   while  $len(rep) > nrep$  do
41:      $del\_ind \leftarrow select\_deletion(\gamma)$ 
42:      $rep.pop(del\_ind)$ ,  $rep\_obj.pop(del\_ind)$ 
43:   end while
44: end for
45: Result Collection:
46: return  $result \leftarrow \{rep, rep\_obj\}$ 

```

Abstract in lingua italiana

A fronte della crescente frequenza e intensità dei disastri naturali e delle minacce climatiche croniche, questo studio analizza l'impatto del rischio climatico fisico sui portafogli azionari globali. Sfruttando tecniche di panel regression sui rendimenti settoriali, forniamo evidenza statistica di come eventi di temperature estreme abbiano un effetto negativo sulla maggior parte dei settori. Per l'esposizione al rischio climatico di un portafoglio introduciamo due nuove metriche basate sulle anomalie di temperatura: la *Climate Risk Exposure* e la *Climate Exposure Volatility*. A differenza degli indici statici attualmente disponibili a livello nazionale, queste metriche catturano la probabilità variabile nel tempo degli eventi estremi e la loro relazione con gli asset tangibili aziendali. Integriamo queste misure in una ottimizzazione di portafoglio multi obiettivo, che estende la classica frontiera media-varianza, permettendo agli investitori di costruire portafogli che sono resilienti agli shock climatici senza sacrificare la diversificazione. Infine, attraverso un'analisi di backtesting, mostriamo i benefici pratici nell'incorporare queste metriche di rischio climatico nel processo di investimento, valutando come strategie consapevoli del rischio climatico si comportino rispetto ai benchmark tradizionali.

Parole chiave: Ottimizzazione di Portafoglio Consapevole dei Rischi Climatici, Rischio Climatico Fisico, Anomalie di Temperatura, Cambiamento Climatico

Acknowledgements

I would like to express my deepest gratitude to my supervisor, Michele Azzone, and my co-supervisor, Gabriele Sbaiz, for giving me the opportunity to carry out this work.

I am sincerely grateful for their constant guidance and support throughout this journey. I deeply appreciate the patience they showed and the advice they provided at every stage of the research. Their availability for any clarification and their insightful suggestions were fundamental to my growth and to the completion of this work.



POLITECNICO
MILANO 1863