LIFTING THE DOUBLE HIDDEN RISK: TOWARD A SELF-ORGANIZING REVIEW OF CLIMATE PHYSICAL RISKS FINANCIAL IMPACT MODELS *

ABSTRACT

Climate-related physical risks are increasingly recognized as "hidden risks" within financial systems. They are embedded in asset valuations, balance sheets, and macroeconomic projections, yet their true magnitude remains uncertain and often underestimated. However, beyond this first layer of invisibility lies a second hidden risk: the opacity and heterogeneity of the models used to assess these climate impacts. Actuaries, bankers, and researchers have developed a wide range of modeling approaches, each relying on different data, assumptions, and purposes making it difficult to determine which models are truly fit for financial decision-making.

This paper aims to address this "double hidden risk." After conducting a systematic literature review, we introduce a self-organizing methodology that classifies and visualizes the main families of models used to estimate the financial consequences of physical climate risks. Using Self-Organizing Maps (SOM) and Variational Autoencoders (VAE), we cluster the existing research landscape to reveal methodological proximities, overlaps, and blind spots. This auto-adaptive review provides a replicable framework to evaluate the transparency, comparability, and relevance of climate-finance models.

By mapping how the literature itself is structured, our approach contributes to lifting the second veil of opacity surrounding model quality and scope. It offers both academics and practitioners a data-driven tool to identify consistent methodologies, detect outliers, and benchmark analytical robustness across the field. Ultimately, the study advocates for a more transparent, cumulative, and accountable understanding of how physical climate risks propagate through the financial system.

Keywords Hidden Risks · Physical Climate Risk · Climate Finance · Financial Impact Modelling · Self-Organizing Maps · Auto-Adaptive Literature Review

1 Introduction

Human activity and pollution have already transformed the Earth's systems and altered the foundations of modern life [1]. Early systemic models such as *The Limits to Growth* [2] attempted to forecast how resource depletion, population growth, and industrial expansion could interact to produce large-scale environmental and economic consequences. Subsequent scientific work sought to quantify and predict the timing and magnitude of climate-related impacts on natural and human systems, both biological and economic [3]. Even before these models appear, Georgescu-Roegen had introduced a profound insight by applying the second law of thermodynamics to economic processes, arguing that every act of production and consumption inevitably increases entropy in the biosphere [4].

In the decades since, the intensification of global warming and the multiplication of extreme weather events have brought climate-related physical risks to the forefront of economic and financial research. These risks, floods, wildfires, droughts, heat waves, and rising sea levels, represent what can be described as a first form of *hidden risk*: they are embedded in the valuation of assets, the pricing of insurance contracts, and the balance sheets of firms and states, yet their true magnitude remains uncertain and often underestimated. As Battiston et al. (2017) and Bolton et al. (2020) have shown, such risks can propagate through financial networks, affecting asset prices, firm profitability, and sovereign

*: **. Title. Pages.... DOI:000000/11111.**

creditworthiness. They are not confined to the insurance sector but extend across banking, capital markets, and corporate finance (NGFS, 2020; [5]).

However, beyond this first layer of invisibility lies a second one: the uncertainty and opacity of the *models* used to quantify these physical impacts. The diversity of modeling practices ranging from actuarial damage functions to macro-financial stress tests, network contagion models, and asset-pricing frameworks creates a methodological fragmentation that prevents cumulative understanding. In other words, the way we measure and model climate risk introduces its own form of uncertainty: a *second hidden risk*. This conceptual blind spot has recently been highlighted by Higuera Roa et al. (2025) [6], who argue that assessing climate risks requires new approaches capable of integrating multiple layers of uncertainty, interactions, and cascading feedbacks. Yet, most financial risk models remain designed for stable, well-characterized systems, and struggle to capture the non-linearities inherent in climate processes.

This double hidden risk manifests both in the economy and in the scientific landscape itself. Because on one hand, we find financial markets internalizing climate hazards imperfectly, underpricing long-term risks and over-relying on historical volatility. And on the other hand, the proliferation of heterogeneous climate-finance models, each based on distinct assumptions, data sources, and methodological traditions, generates a second-order uncertainty about which models are "fit for purpose." This dual opacity raises a fundamental question: how can we ensure that the tools designed to reveal hidden risks do not, themselves, become sources of opacity?

A growing body of literature has therefore sought to quantify and model the financial implications of physical climate risks [5]. Over the past decade, the number of studies addressing the economic and financial impacts of climate-related hazards has increased markedly. In this paper, we review and analyze this emerging field through a dataset of 150 peer-reviewed articles extracted from Scopus and complementary sources.

To address the heterogeneity of methods, data, and perspectives, we employ unsupervised learning techniques. We use Self-Organizing Maps (SOM) [7] and Variational Autoencoders (VAE) [8] to classify and visualize the structure of this research landscape. Rather than conducting a purely thematic or systematic review, we propose an "self-organizing" literature review that dynamically organizes and clusters scientific contributions according to their methodological and conceptual proximity.

This approach draws inspiration from other disciplines that have already developed self-organizing classification frameworks, such as biology and its phylogenetic trees of life [9]. We argue that interdisciplinary research at the intersection of climate science and finance would greatly benefit from such an approach, providing both academics and practitioners with an intuitive and scalable way to explore a rapidly expanding body of knowledge.

Despite the growing awareness of climate-related financial risks, the field remains fragmented. Studies differ widely in scope, data sources, and methodologies from catastrophe modeling and econometric analysis to complex system simulations and network models [10]. This heterogeneity complicates cumulative understanding and hinders comparability. For instance, while actuaries model risk through expected loss functions and solvency constraints [11], macro-financial focus on systemic risk propagation and climate stress testing [12]. Regulatory frameworks such as the Task Force on Climate-related Financial Disclosures (TCFD, 2017) and the Network for Greening the Financial System (NGFS, 2020) encourage transparency but often rely on qualitative scenario narratives rather than quantitative comparability.

In this study, we argue that the challenge is not only to identify relevant models, but also to organize them meaningfully in a multidimensional space where their conceptual, methodological, and empirical proximities become visible. By projecting the corpus of scientific publications into a low-dimensional latent space, Self-Organizing Maps (Kohonen, 1982) and related unsupervised algorithms enable the detection of clusters of research that share common assumptions, models, or datasets. Such a framework directly addresses the second hidden risk, the methodological opacity, by revealing the structure, redundancy, and gaps within the existing literature.

This contribution therefore serves a dual purpose. First, it provides an automated and reproducible way to review the literature on the financial impacts of physical climate risks, helping researchers and policymakers navigate a complex, multi-disciplinary field. Second, it contributes methodologically to the emerging domain of computational bibliometrics and "auto-adaptive" literature reviews, extending machine learning applications beyond prediction into the meta-analysis of scientific models themselves. In doing so, we seek to lift the second veil of invisibility surrounding climate-finance research and contribute to a more transparent, structured, and cumulative understanding of how physical climate risks propagate through the financial system.

2 A glimpse at the vast landscape of Physical Risks Financial Impacts

Literature about Climate Physical Risks and Financial Impacts is trending. Since the Paris Agreements, the number of articles went from a few to a dozen every year.

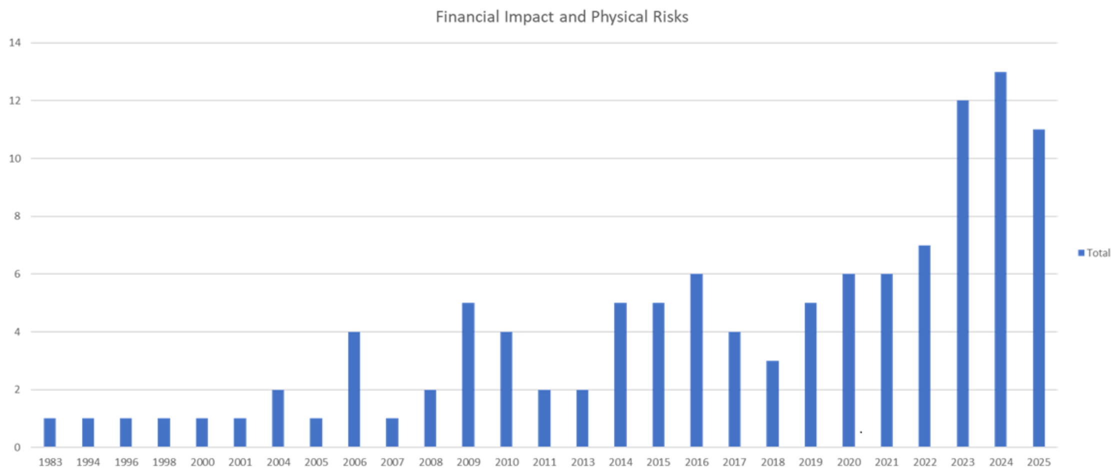

Figure 1: Number of articles found through Scopus, found using "FInancial Impact" AND "Physical Risk" in the Title, the Abstract or the keywords.

2.1 Physical Risks

Yet, despite the growing number of studies addressing the financial consequences of climate-related hazards, there is still no unified framework capable of linking the diversity of existing approaches. Current research is dispersed across disciplines, each adopting its own conceptual definitions, datasets, and analytical techniques. As a result, the field remains highly fragmented: economists focus on systemic propagation mechanisms, actuaries on local damage functions, and financial analysts on market valuation models. This lack of integration makes it difficult to compare results or to identify where complementarities and redundancies lie within the literature.

The central problem this paper addresses is therefore one of **organization and synthesis**: how can we construct a systematic and quantitative framework that captures the heterogeneity of studies assessing the financial impacts of physical climate risks? More specifically, how can we represent, in a single multidimensional space, the conceptual, methodological, and empirical proximities between these contributions? By answering this question through a self-organizing and data-driven approach, we aim to move beyond static typologies toward a dynamic, reproducible, and interpretable mapping of climate-finance research. What motivates us is figuring out how can each model address a specific need. In order to that, we first need to take a look at the rapidly changing landscape.

"Physical climate-related risks" or "physical climate risks", in this paper, is conceived as the economic costs and financial losses caused by the increasing frequency and severity of climate-related weather events (e.g., storms, floods, or heat waves) and by long-term changes in climatic patterns (e.g., ocean acidification, sea-level rise, or precipitation shifts) [13]. The Task Force on Climate-related Financial Disclosures (TCFD) distinguishes two main categories of physical risks: acute risks, arising from event-driven phenomena such as cyclones or wildfires, and chronic risks, associated with longer-term shifts in climate patterns such as rising temperatures or sea levels [14].

A wide range of models have been developed to quantify these risks depending on the underlying hazard. For instance, Palmer (1965) introduced the Meteorological Drought Index, which remains a benchmark in drought risk assessment [15]. In the case of cyclones, recent works [16] propose probabilistic frameworks linking wind-field modeling to insured losses. Sea-level rise projections are extensively studied by NOAA and related research groups, providing localized inundation scenarios for coastal planning (**[17]**). Wildfires have been addressed through both ecological and actuarial modeling (e.g., Abatzoglou and al., 2016 [18]), while earthquakes, although not directly climate-related, are often included in physical risk modeling frameworks because they share the same probabilistic and damage-function approaches used in catastrophe risk estimation.

These contributions form a landscape in which each hazard involves different timescales, data sources, and modeling assumptions. This fragmentation complicates the integration of physical risk modeling into consistent financial or actuarial frameworks and is precisely the issue that motivates the present study.

Table 1: Classification of climate-related hazards according to the European Commission [19]

Temperature-related	**Wind-related**	**Water-related**	**Solid mass-related**
		Chronic hazards	
Changing temperature (air, freshwater, marine water)	Changing wind patterns	Changing precipitation patterns and types (rain, hail, snow/ice)	Coastal erosion
Heat stress		Precipitation or hydrological variability	Soil degradation
Temperature variability		Ocean acidification	Soil erosion
Permafrost thawing		Saline intrusion	Solifluction
		Sea level rise	
		Water stress	
		Acute hazards	
Heat wave	Cyclone, hurricane, typhoon	Drought	Avalanche
Cold wave/frost	Storm (including blizzards, dust and sandstorms)	Heavy precipitation (rain, hail, snow/ice)	Landslide
Wildfire	Tornado	Flood (coastal, fluvial, pluvial, ground water)	Subsidence
		Glacial lake outburst	

2.2 Financial Impact and Financial Risk

The financial consequences of physical climate risks differ substantially depending on the perspective and objectives of the actors involved. While actuaries, bankers, and investors all attempt to quantify and manage the financial exposure arising from climate-related hazards, they operate under distinct frameworks of risk measurement, time horizons, and regulatory constraints.

2.2.1 Actuarial Research and Physical Climate Risks

In the insurance sector, the very notion of insurability is being redefined under the growing influence of climate change. As summarized by Charpentier (2008) [11], the increasing frequency and severity of natural hazards challenge the classical foundations of insurance, which rely on diversification, independence of risks, and the law of large numbers. When events become highly correlated across regions, the capacity of insurers to mutualize losses diminishes, giving rise to issues of moral hazard, adverse selection, and pricing instability. In this context, actuaries must adapt both their analytical tools and their understanding of what constitutes an insurable risk.

Actuarial modeling therefore approaches climate-related risks through the lens of loss distribution and expected damage. It relies on physical hazard models combined with exposure and vulnerability data to derive damage functions, which relate the intensity of an event (e.g., wind speed, flood depth, or temperature anomaly) to the expected monetary loss.

A typical formulation can be expressed as:

$$L = E \times V(I) = E \times \alpha\, I^{\beta} \tag{1}$$

where L is the expected loss, E the insured exposure (value of the asset or portfolio), and $V(I)$ a vulnerability function expressing the fraction of damage as a function of the hazard intensity I. The parameters α and β are empirically calibrated from past loss data or engineering studies.

For example, in flood risk assessment, β typically ranges between 1.2 and 2 depending on asset type, meaning that losses increase nonlinearly with flood depth. Similar relationships are used in windstorm models where damages scale approximately with the cube of wind speed, reflecting the kinetic energy of the event.

These functions are then integrated into probabilistic catastrophe models (CatModels) that combine hazard frequency, exposure, and vulnerability components to generate a full loss distribution:

$$Expected\ \ Annual\ \ Loss\ \ (EAL) = \int_0^{\infty} L(I)\, f(I)\, dI \tag{2}$$

where $f(I)$ denotes the probability density function of the hazard intensity.

Such models allow for the estimation of claim probabilities, solvency requirements, and reinsurance needs, but also for many other variables and parameters. Actuarial research increasingly integrates stochastic event sets and climate-adjusted hazard distributions to price policies under changing climatic conditions (Herweijer et al., 2009 [20]). The key objective for insurers is to comply with the Solvency II standards and to ensure that premiums adequately reflect a rising baseline of physical hazards.

2.2.2 Banking and Physical Climate Risks

In the banking sector, physical risks are primarily analyzed through their impact on credit risk parameters, the Probability of Default (PD), Loss Given Default (LGD), and Exposure at Default (EAD). A severe climate event such as flooding, drought, or heat waves can disrupt cash flows, reduce collateral values, or impair productive assets, thereby increasing a borrower's default probability.

Battiston and Monasterolo (2020) [12] have formalized how climate-related scenarios including both physical and transition channels can propagrate through balance sheets and modify default probabilities. Their framework explicitly models how shocks to asset values under different climate scenarios translate into increased credit risk, showing that investors' and firms' Probability of Default becomes a function of scenario-dependent exposures:

$$PD_i = f(X_i, H_t) \tag{3}$$

where PD_i denotes the probability of default of borrower i, X_i represents the financial and sectoral characteristics of the borrower, and H_t captures the physical hazard intensity at time t. A linear form of this dependence is:

$$\Delta PD_i = \gamma_0 + \gamma_1 \Delta H_t + \gamma_2 \Delta CF_i + \varepsilon_i \tag{4}$$

where ΔH_t represents the change in hazard intensity (e.g., flood frequency or temperature anomaly) and ΔCF_i the corresponding variation in cash flows. Positive γ_1 and γ_2 coefficients imply that worsening physical conditions and declining revenues both increase the default probability.

Within the Basel II/III Internal Ratings-Based framework, the Expected Loss (EL) from physical climate risks can then be expressed as:

$$EL = PD \times LGD \times EAD \tag{5}$$

Hence, a deterioration in any of these components directly amplifies credit losses. Supervisory authorities such as the European Central Bank (ECB) and the Network for Greening the Financial System (NGFS) increasingly incorporate these relationships into climate stress-testing exercises, evaluating how physical shocks affect portfolio resilience and systemic solvency.

By linking hazard scenarios to prudential regulation, this line of research connects climate science to financial stability analysis, aligning capital adequacy requirements with the long-term resilience of credit portfolios. Battiston's modeling approach thus provides a quantitative bridge between climate stress and default risk, allowing for the integration of scenario analysis into banks' internal risk assessment systems.

2.2.3 Financial Markets and Investment Research

From an investment perspective, the assessment of physical climate risks is increasingly embedded into investment decision-making and corporate valuation. Beyond macro-level stress tests, firm-specific exposure is now quantified using spatial and sectoral data to estimate potential losses under different hazard scenarios. Tools such as the web-based platform proposed by Andersson et al. (2017) [21] enable companies and investors to map vulnerabilities to floods, heat waves, or sea-level rise, translating them into financial risk indicators. These models often rely on *climate value-at-risk* (Climate VaR) measures, which express the potential loss in firm or portfolio value under an adverse climatic event:

$$ClimateVaR = \mathbb{E}[V] - V_p \tag{6}$$

where $\mathbb{E}[V]$ represents the expected firm or portfolio value and V_p its value under a specified physical climate scenario at probability level p. By quantifying this difference, investors can incorporate location-specific climate exposure directly into portfolio optimization and risk-adjusted performance metrics such as RoI or RoE. We believe this approach might help bridge scientific hazard modeling with financial valuation, fostering a more granular understanding of climate resilience in the business sector.

—

This tripartite framework is just an example, a glimpse to illustrate how chaotic our motivational question is. So let's put some order into it.

2.3 Physical Climate Risk Financial Impact

By targeting 3 different fields of research and 3 different types of stakeholder, we already categorize the literature. But in order to further help the decision-maker (i.e. the banker, the insurer, the company CFO...) to choose the right model, we need to go further.

We think that bridging the two concepts of physical risk and financial impact can take many forms. Over the past few years, several researchers have attempted to design models capable of linking environmental hazards with financial outcomes. Yet, the diversity of purposes, data sources, and disciplinary assumptions makes such integration difficult. The challenge lies precisely in connecting two distinct conceptual levels: the theoretical dimension, where climate risks and financial metrics are defined, and the empirical dimension, where observable proxies, such as drought events, wildfires, or asset devaluations, are measured and compared.

To clarify this relationship, we represent the overall research design using the *Libby Boxes* framework (Figure 2). Initially proposed in the field of accounting research [22], this approach illustrates how theoretical constructs can be linked to empirical proxies through explicit causal pathways. In our case, it serves as a visual bridge between the physical risk domain (left) and the financial impact domain (right), integrating control variables that account for firm-level heterogeneity and financial adjustments such as depreciations or provisions.

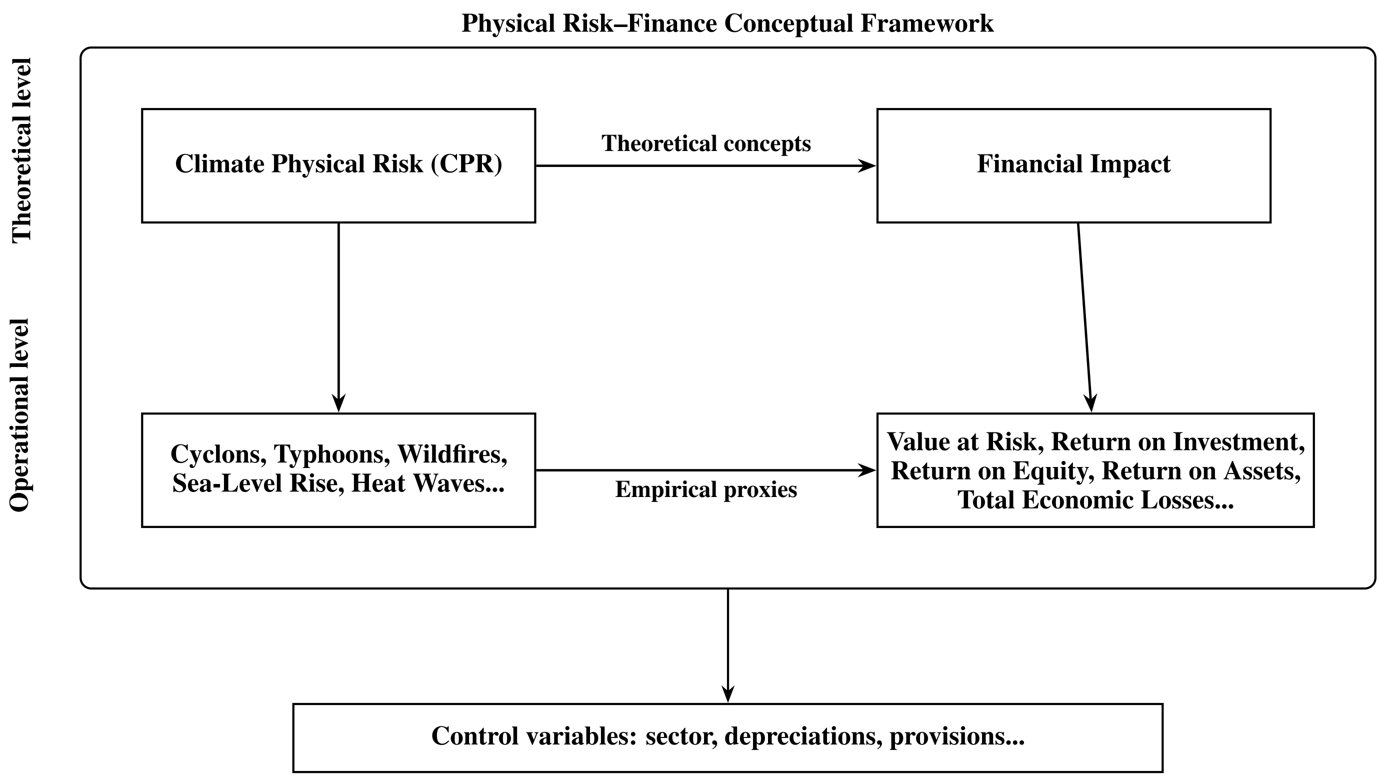

Figure 2: Libby Boxes representation of the conceptual bridge between physical climate risks and financial impacts. The upper boxes represent the theoretical constructs (Physical Risk, Financial Impact), while the lower boxes depict the empirical proxies used in applied models (e.g., tornadoes, cyclones, wildfires, and firm-level indicators such as profitability or asset value). Control variables such as depreciations or provisions are introduced to capture intermediate effects.

Libby boxes are useful as the show that each level, conceptual and empirical, can be connected analytically. The top layer captures theoretical reasoning, while the bottom layer represents measurable observations. This two-level reasoning underpins our effort to build an auto-adaptive literature review: by classifying papers according to their position in this conceptual space, we can identify which studies focus on the physical side, which on the financial side, and which attempt to connect the two.

In the next section, we show how this conceptual bridge can be operationalized using unsupervised learning techniques, particularly Self-Organizing Maps such as Kohonen Maps (Kohonen, 1982 [7]), to structure and cluster the existing literature on physical risk financial modeling.

3 Data and Methods

3.1 Data Collection and Framework

The initial dataset was constructed through a systematic search on Scopus using the keywords *"Physical Risk"* and *"Financial Impacts"* in the Titles, Abstracts and Keywords in the database gathering over 43 000 scientific journals. This query returned a large number (over 150) of peer-reviewed articles covering several disciplinary perspectives. Each paper was manually reviewed and classified according to a structured framework (see Appendix 6.3) including variables such as paper type, geographical scope, hazard category, data source, and methodological approach.

The resulting table allowed for a consistent coding of studies along the two main analytical axes of this research: Physical Risk and Financial Impact. These two dimensions guided the subsequent transformation of the qualitative framework into a quantitative dataset through the construction of a disjunctive matrix, later converted into a Burt table. This step enabled the representation of the literature corpus in an $n \times n$ matrix suitable for clustering analysis using self-organizing maps.

3.1.1 Data processing and variable coding

Each article included in the dataset was read and reprocessed manually following a systematic literature review protocol. For each study, a set of standardized variables was collected to ensure consistency across the corpus. These variables capture both the conceptual and empirical characteristics of each paper, including its treatment of physical climate risks and financial impacts.

To be retained in the final dataset, an article had to report at least one proxy of **financial impact** (e.g., losses, return, solvency, or credit indicators) and one proxy of **physical hazard** (e.g., flood, wildfire, drought, or heatwave intensity). This dual requirement ensured that all selected contributions explicitly linked a measurable physical risk to a quantifiable financial outcome.

After extraction, all entries were harmonized according to a structured nomenclature (see Table 2) to enable cross-comparison between disciplines and methodological approaches. This harmonization process relied on the framework presented in Appendix 6.3.

3.2 Method

Traditional literature reviews often rely on selective reading and thematic categorization, which can introduce subjectivity and limit reproducibility. In contrast, this work applies a multi-systematic and quantitative approach that enables the integration of both qualitative and categorical variables extracted from the reviewed studies. Each article identified through Scopus was encoded according to the framework described in Appendix 6.3, resulting in a structured dataset containing indicators for both Physical Risk and Financial Impact.

3.2.1 From qualitative variables to disjunctive and Burt tables

Because the dataset included heterogeneous variables (textual, binary, categorical, and numerical), we need to consider it as we would consider doing a multiple correspondence analysis (MCA). It was first transformed into a *complete disjunctive table Z*. Each modality of a qualitative variable was converted into a binary indicator (1 if present, 0 otherwise), allowing the representation of mixed data in matrix form:

Consider Z a *n*i* matrix with *n* the number of dimensions and *i* the number of individuals.

From this matrix, we computed the *Burt table B*, defined as:

$$B = Z^\top Z$$

Table 2: Variable nomenclature used for article coding and harmonization

Variable	Description
`Link`	DOI or permanent URL to the article source.
`Paper Name`	Full title of the publication.
`Paper Type`	Nature of the paper (empirical, literature review, methodological, policy report, etc.).
`Scope Covered`	Geographic or sectoral coverage of the study; harmonized variable unifies terminology across articles.
`Sector`	Main economic sector analyzed (insurance, banking, agriculture, real estate, etc.).
`Timeframe`	Period of analysis or data collection.
`Financial Impact Precision`	Level of financial granularity.
`Financial Proxy`	Type of indicator, the quantitative proxies used.
`Financial Data Source`	Origin of financial data (Bloomberg, ECB, balance sheets, surveys, etc.).
`Financial Sample Size`	Number of firms, countries, or assets included.
`Physical Risk Resolution`	Degree of spatial precision in the hazard modeling.
`Physical Risk Type`	Type of physical hazard (Specific, hybrid, compound perils or General) .
`Physical Risk Proxies`	quantitative proxies (e.g., intensity index, cumulative exposure).
`Physical Data Sources`	Data providers or databases used (NOAA, Copernicus, CMIP6, etc.).
`Physical Sample Size`	Number of observations.
`Control Variables`	Covariates used for econometric control and their data sources.
`Method`	Main analytical or statistical method used (regression, VAR, network model, machine learning, etc.).
`Projection (y/n) / Scenario`	Indicates whether the paper includes forward-looking projections (e.g., RCP, NGFS, or IPCC scenarios).
`Castle / Main Figure`	Key equations and graphical representations used to summarize the model.
`Summary / Key Findings and Results`	Concise synthesis of the study's main findings and empirical outcomes.
`Limitations and Assumptions`	Explicit limitations, simplifications, or boundary conditions acknowledged by the authors.
`Developing vs. Developed Countries`	Binary variable indicating whether the study focuses on developing economies, developed economies, or both.
`Damage Function (yes/no)`	Indicates whether a damage function or loss–intensity relationship is explicitly modeled.
`Transition Risk`	Whether the paper also considers transition risks in addition to physical ones.
`Accuracy`	Evaluation of the model's predictive performance, where applicable.
`Authors / Keywords / Journal`	Bibliographic metadata extracted from Scopus or the article itself.
`Nb Citation / Nb Ref / Year`	Citation count, number of references, and publication year.
`End-user`	Target audience or user of the model (insurer, regulator, bank, investor, policy-maker).

This symmetric $p \times p$ matrix summarizes the co-occurrence relationships between all variable modalities across the literature corpus. It serves as the input for the Kohonen self-organizing algorithm.

3.2.2 Kohonen Algorithm and literature clustering

We owe a lot to Kohonen's Self-Organizing Maps [7] allowing us to organize and treat both qualitative and quantitative data. We used the *Self-Organizing Map* (SOM) to classify and visualize articles within a two-dimensional grid. The SOM functions as a simplified neural network, composed of a single layer where each neuron corresponds to a node in the map. The algorithm iteratively projects the high-dimensional literature dataset onto this grid through competitive learning, following the steps below:

1. **Initialization.** Random initialization of the neuron weight vectors $\mathbf{w}_i(0)$ for all neurons i on the grid, where each $\mathbf{w}_i \in \mathbb{R}^n$ represents a prototype vector in the same n-dimensional feature space as the input data $\mathbf{x}$. Here, n denotes the number of attributes describing each article (e.g., hazard type, financial variable, method).
2. **Input selection.** At each iteration t, a new input vector $\mathbf{x}(t)$ (in this case, an article represented by its standardized attributes) is presented to the network. Each iteration corresponds to one learning step in the algorithm.
3. **Best Matching Unit (BMU).** The neuron whose weight vector is closest to the current input $\mathbf{x}(t)$ in Euclidean distance is identified as the BMU:
$$c = \arg\min_i \|\mathbf{x}(t) - \mathbf{w}_i(t)\|.$$
Here, c denotes the index of the best-matching neuron, and i indexes all neurons on the grid.
4. **Weight update.** The BMU and its neighboring neurons are updated according to the neighborhood function $h_{ci}(t)$, typically Gaussian:
$$\mathbf{w}_i(t+1) = \mathbf{w}_i(t) + \eta(t)\, h_{ci}(t)\, [\mathbf{x}(t) - \mathbf{w}_i(t)].$$
In this expression, $\eta(t)$ is the learning rate controlling the adjustment magnitude, and $h_{ci}(t) = \exp\left(-\frac{\|\mathbf{r}_c - \mathbf{r}_i\|^2}{2\sigma^2(t)}\right)$ is the neighborhood function, where $\mathbf{r}_c$ and $\mathbf{r}_i$ represent the two-dimensional coordinates of neurons c and i on the map grid, and $\sigma(t)$ defines the radius of influence.
5. **Neighborhood and learning decay.** Both the neighborhood radius $\sigma(t)$ and the learning rate $\eta(t)$ decrease over time according to exponential decay functions:
$$\eta(t) = \eta_0 e^{-t/\tau_1}, \quad \sigma(t) = \sigma_0 e^{-t/\tau_2}.$$
Here, η_0 and σ_0 denote the initial learning rate and neighborhood radius, while τ_1 and τ_2 are time constants controlling the rate of decay.
6. **Convergence.** The process is repeated for all input vectors until the map stabilizes. At convergence, neurons that are spatially close on the grid encode similar input patterns, thereby producing a topological and interpretable representation of the literature corpus in a two-dimensional space.

This process produces a two-dimensional topological representation of the literature corpus, where similar articles are located close to each other on the map. The resulting clusters correspond to coherent methodological or thematic families within the field of *Physical Climate Risk and Financial Impact*. This types of self-organizing maps have already been used in finance including for bankruptcy forecasting [23].

Finally, to complement this approach, we experimented with Variational Autoencoders (Kingma & Welling, 2013) as an alternative dimensionality reduction technique. Unlike in generative modeling, here the encoder network is used to project articles into a latent space, which can then be visualized and compared with the SOM projections to assess cluster stability.

4 Results

The models inside the articles can give us different information about how they adress the considered problem.

Table 3: Overview of a few models in the literature review

Research focus	**Hazard type**	**Financial indicator**	**Methodology**	**Main stakeholders**
Coastal real estate model	Flood and sea-level rise	Housing price or real estate discount	Linear regression (hedonic pricing or panel)	Investors and asset managers
Public finance and wildfires	Wildfire	Annual suppression expenses / budgetary cost	Time-series or linear regression; future cost scenarios	Public authorities, ministries of finance, forestry agencies
Agricultural exposure to drought	Drought	Agricultural income, margin, net profit per hectare	Linear or panel regression, sometimes simulated hydrological scenarios	Farms, agricultural enterprises, rural banks
CAR and natural catastrophes	Earthquake, hurricane, flood	Stock return, volatility, drawdown, market index	Event study, VAR model, or other linear econometrics	Investors and asset managers

4.1 Descriptive Analysis

Geographic and temporal coverage. The corpus exhibits a clear **North American bias**, with the majority of studies focusing on the United States and, to a lesser extent, Europe. Developing countries remain underrepresented despite their higher exposure to physical hazards. The timeframes covered by the studies range from short-term single-event analyses to long historical series spanning up to 123 years, reflecting the heterogeneity of data availability across hazards and sectors.

Sectoral distribution. As illustrated in Figure 3, financial markets account for roughly one third of all studies, followed by multisectoral analyses, energy, agriculture, and health. Infrastructure and transport remain marginal topics, indicating a research gap in asset-level and systemic interdependency modeling.

Figure 3: Sectoral distribution of studies in the literature sample

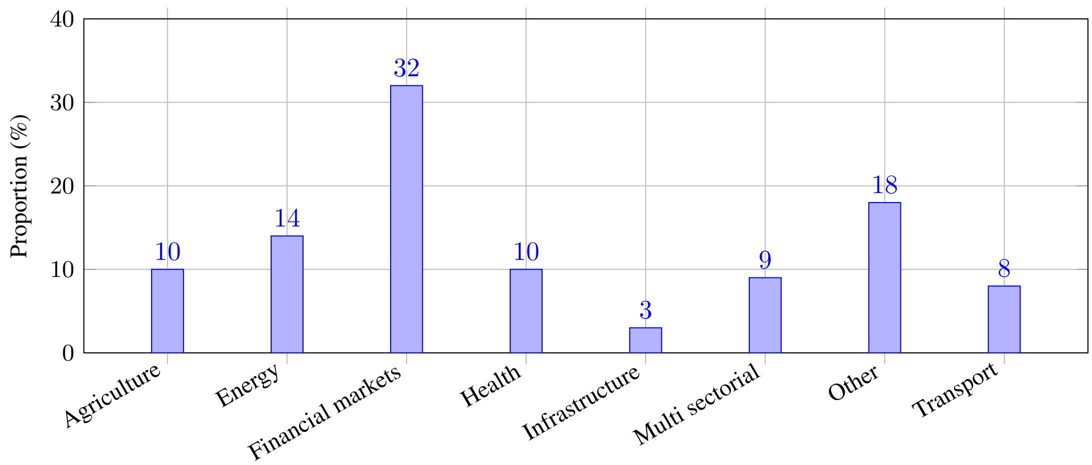

Scale and diversity of physical proxies. Table 4 summarizes the distribution of spatial and temporal scales used for physical risk proxies. The results show a wide diversity of modeling granularity: regional and local studies dominate, but asset-level and company-level analyses are gaining traction, reflecting growing interest in micro-level exposure assessment.

Table 4: Scale of physical proxies used in the reviewed studies

Scale (Physical proxy)	**Proportion**
Asset level	13%
Grid degrees	7%
Grid kilometers	15%
Local	20%
National	12%
Regional	27%
Subregional	5%
Company level	2%

Hazard typology. Most models address a **specific hazard** rather than aggregated climate indices. Multi-peril studies remain the largest group (31%), followed by drought (16%), wildfire (15%), earthquake (11%), and flood (10%). Heat waves and pandemics are still marginally represented, pointing to emerging but underdeveloped research areas.

Financial proxies and indicators. On the financial side, the diversity is equally marked (Table 6). Studies operate across scales ranging from individual firms to national aggregates, but few adopt portfolio-level or building-level perspectives. The most frequent financial variables are accounting cash flows (57%), followed by risk metrics (11%), credit risk indicators (10%), and market valuation measures (7%).

Table 5: Types of physical hazards covered in the literature

Hazard type	Proportion
Multi-peril	31%
Drought	16%
Wildfire	15%
Earthquake	11%
Flood	10%
Storm	7%
General (unspecified)	7%
Heat	2%
Pandemic	2%

Table 6: Scale of financial proxies used in the reviewed studies

Scale (Financial proxy)	Proportion
Asset level	15%
Building level	3%
Company level	7%
Individual level	17%
Local	10%
National	17%
Portfolio level	3%
Other	12%

Table 7: Main types of financial indicators identified in the literature

Financial indicator	Proportion
Accounting cash flow	57%
Credit risk	10%
Market valuation	7%
Risk metric	11%
Other	13%
Unspecified	2%

Summary. The field is methodologically rich but geographically and thematically uneven. Considering it is dominated by developed-country case studies, particularly in North America, and by hazard-specific analyses at the regional and local scales, we believe there is a need for harmonized frameworks, something that could help navigate this landscape.

4.2 Multiple Correspondence Analysis (MCA)

Beyond the classification of models presented in Table 3, we conducted an exploratory analysis of the literature dataset using a Multiple Correspondence Analysis (MCA). This method [24] allows the visualization of associations among qualitative variables such as hazard type, financial indicators, and stakeholder categories. In our case, each article is represented as an individual, described by several categorical attributes extracted from the harmonized framework (see Table 2).

The MCA projects this multidimensional qualitative dataset into a low-dimensional factorial space. The first two dimensions capture the largest amount of variance (here 9.1% and 8.6%), providing a simplified representation of how research topics relate to one another. Articles or keywords that appear close to each other in this space share similar profiles. For instance, they may rely on comparable types of hazards, financial variables, or methodological approaches.

The MCA visualization (Figure 4) reveals several patterns. We observe that *wildfires* and *droughts* tend to cluster together, reflecting their shared treatment within agricultural and ecological modeling frameworks. In contrast, *earthquakes* and *sea-level rise* form distinct poles, typically associated with real-estate valuation and infrastructure loss studies. Financially oriented keywords such as *impact*, *finance*, and *assess* occupy the center of the factorial plane, illustrating their transversal use across different hazard categories.

This initial map of the literature's structure shows thematic proximity and divergence among research contributions. we therefore confirm the heterogeneity observable in the field : while some hazards are studied through well-established

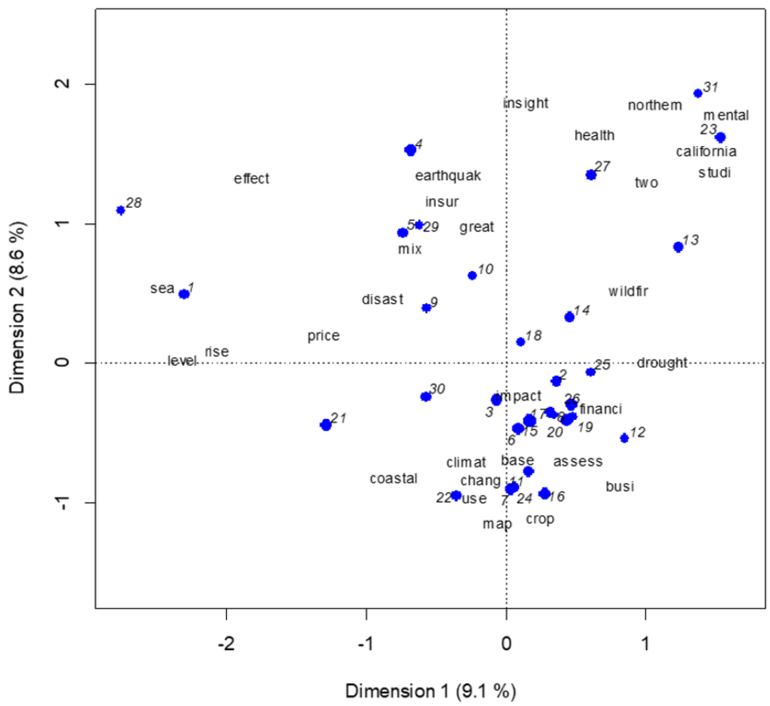

Figure 4: Example of Multiple Correspondence Analysis (MCA) projection of a 30 articles sample from taken from the literature dataset. The proximity between terms represents conceptual and methodological similarity across studies.

economic metrics, others remain isolated or poorly integrated into cross-sectoral analyses. The following subsection builds upon this insight using self-organizing maps to produce a more robust, data-driven clustering of the literature corpus.

4.3 Kohonen Maps

To complement the descriptive analysis, we applied a Kohonen Self-Organizing Map (SOM) to the harmonized dataset in order to cluster articles according to their methodological and thematic proximity. The SOM algorithm projects high-dimensional data onto a two-dimensional grid, preserving topological relationships among input vectors. In this representation, neighboring cells correspond to similar combinations of variables such as hazard type, financial impact proxy, or methodological approach (Kohonen, 1982).

Figure 5 shows the resulting SOM occupation grid. Each cell represents a cluster of articles sharing similar characteristics. The color gradient (from dark purple to yellow) indicates the number of observations per node: darker zones correspond to fewer articles, while lighter zones denote denser regions of methodological or thematic overlap. The central cells tend to capture more general studies dealing simultaneously with multiple hazard types and mixed financial indicators, whereas the corners are often associated with more specialized domains such as agricultural drought modeling, coastal flood valuation, or reinsurance-oriented risk analysis.

The map of representative individuals (Figure 6) illustrates examples of the dominant research orientations within each cluster. For instance, one cell gathers studies related to *wildfire budget risk planning* and *forest management finance*, while another includes works focusing on *seismic non-structural loss dynamics* or *asset-level climate tail risk*. Other clusters are associated with compound flood impact modeling, drought-related agricultural economics, or financial regulation of disaster risks.

This distribution suggests that the SOM produces coherent and interpretable groupings. Clusters (are supposed to) reflect subfields: physical modeling and catastrophe loss estimation on one side, financial valuation and regulatory assessment on the other. The presence of mixed clusters containing both economic and physical modeling perspectives indicates areas of methodological convergence across disciplines. It may reinforce the view that self-organizing maps are a powerful exploratory tool for revealing structure in a heterogeneous literature corpus, and for identifying underexplored intersections between the domains of climate science, insurance, and financial economics.

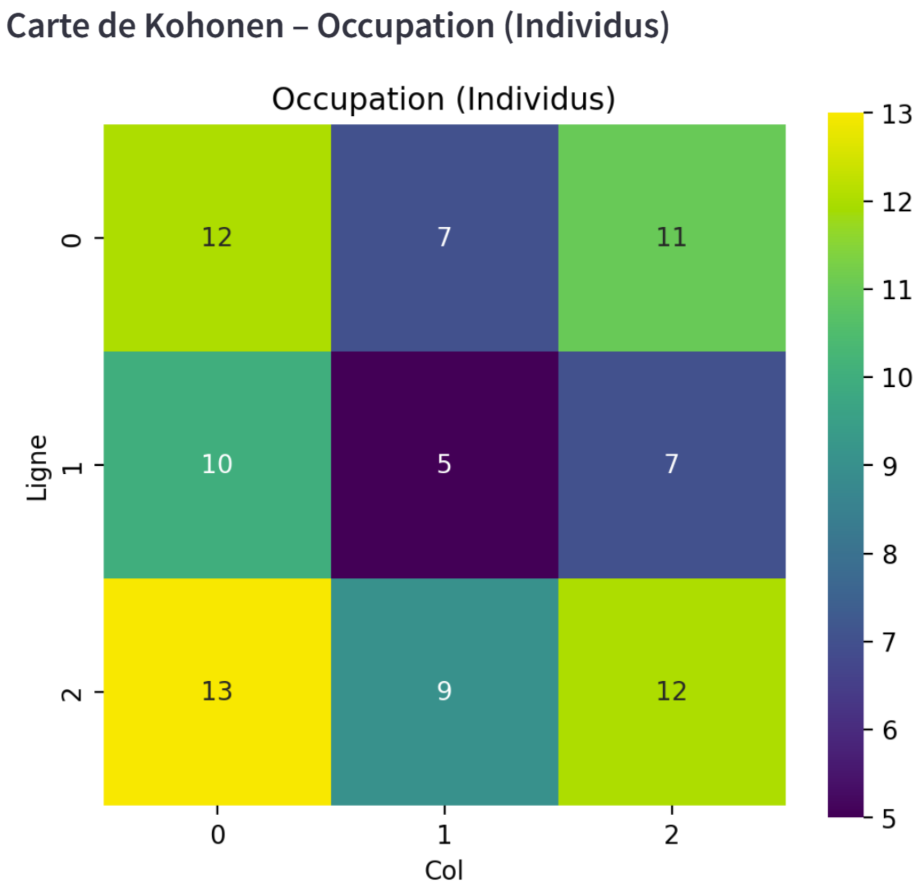

Figure 5: Kohonen Self-Organizing Map representing the occupation density of literature items across a 3×3 grid. The color scale indicates the number of articles clustered within each cell.

4.4 VAE Latent Projection Maps

While the self-organizing maps presented above already provide a robust two-dimensional representation of the literature corpus, we also explored the potential of Variational Autoencoders (VAEs) for latent-space projection. Unlike SOMs, VAEs rely on deep neural architectures capable of compressing high-dimensional data into a continuous latent manifold through probabilistic encoding and decoding (Kingma & Welling, 2013). The objective of this experiment was to assess whether the non-linear feature extraction capacity of VAEs could reveal additional structure within the literature dataset particularly in identifying continuous transitions between research domains rather than discrete clusters.

At this stage, the results are preliminary. The model successfully generated a latent representation of the articles, but the projection patterns remain too diffuse to yield interpretable clusters. This limitation is likely due to the relatively small sample size of the corpus and the mixed nature of the encoded features (binary and categorical variables). Future developments will involve optimizing the network architecture, regularization strength, and latent dimensionality to better capture meaningful proximity relationships between studies.

4.4.1 Interactive Clusterings

One of the most promising applications of the VAE approach is its ability to generate interactive latent-space maps. By labeling the encoded points and enabling zoom or filtering functions, users could dynamically explore subsets of the literature (e.g., all studies related to *drought* and *agriculture*), or all papers focusing on *financial regulation* and *risk pricing*. Such functionality would make the literature review adaptive, allowing researchers and practitioners to navigate through a living, continuously updated database rather than a static synthesis.

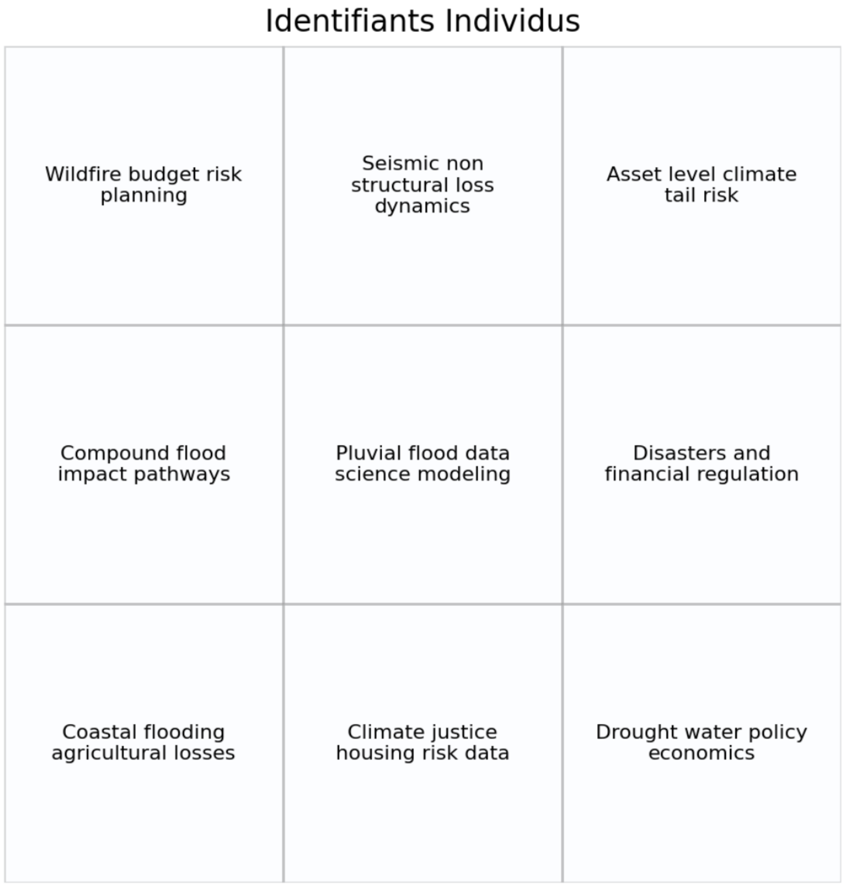

Figure 6: Example of article distribution by SOM cluster. Each cell lists representative studies, illustrating the thematic or methodological focus of each region of the map.

4.4.2 T-SNE and Fine-Tuning of the VAE Latent Space

To improve interpretability, future iterations will combine VAE projections with non-linear dimensionality reduction techniques such as t-SNE (t-Distributed Stochastic Neighbor Embedding). The idea is to reuse the latent space generated by the VAE and apply t-SNE to produce a more visually coherent 2-D map, while avoiding excessive dimensional compression that could distort proximity relationships. Our approach using VAE for representation learning, t-SNE for visualization, should make it possible to refine the clustering granularity while preserving the underlying topological structure identified by the SOM.

Although these experiments are still at an early stage, they point toward the next evolution of the auto-adaptive literature review: a continuous, machine-learned cartography of research, capable of self-updating as new publications emerge.

Our current goal here is to get something like in Figure 7 :

5 Discussion

The main difference with a regular systematic literature review here is that our self organizing maps can help us "navigate" through the models, clusterizing them and making them easier to read even when taking into account a high number of dimensions.

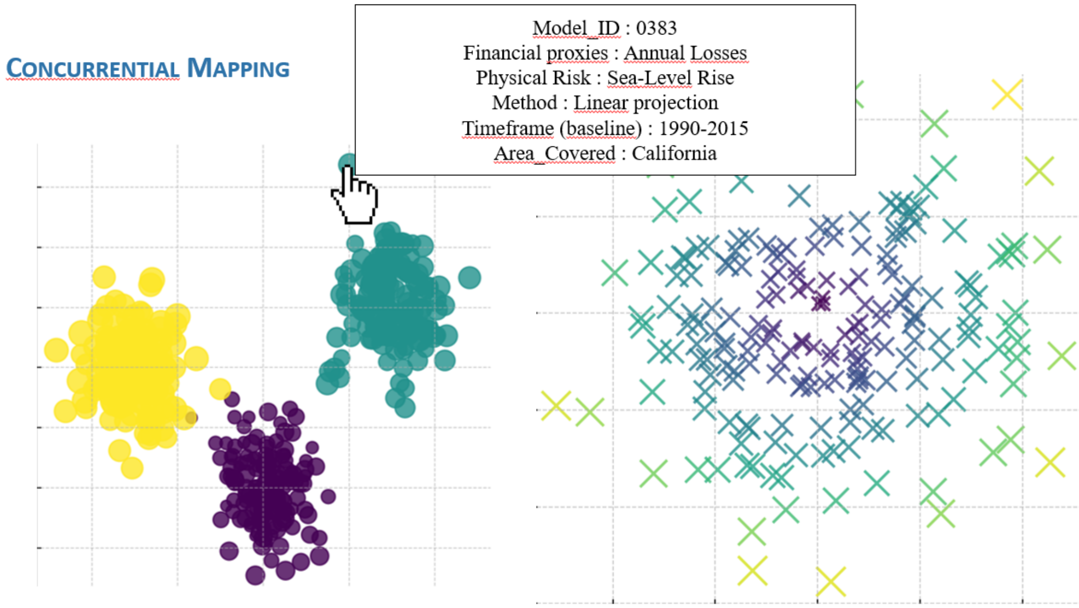

Figure 7: Objective

5.1 The insurance sector

The concept of physical climate risks in the insurance sector is directly linked to the issue of non-insurability. The main problem lies in the foundation of risk pooling and diversification, both of which are increasingly undermined by the systemic nature of climate hazards.

The concept of insurability is a cornerstone of actuarial science and becomes particularly challenged in the context of physical climate risks. Charpentier (2008) [11] reminds us that climate change fundamentally alters the traditional assumptions that make a risk insurable. Increasing frequency and severity of natural disasters generate spatial and temporal correlations among losses, thereby undermining the law of large numbers and reducing the effectiveness of diversification. When events are no longer independent, insurers face growing uncertainty in estimating premiums and reserves, which may lead to pricing that is either economically unviable or socially unacceptable.

Under such circumstances, certain areas or asset classes may enter what Charpentier describes as zones of "structural uninsurability," where private coverage becomes infeasible without public intervention or reinsurance support. This situation questions the capacity of private markets alone to manage systemic risks and highlights the crucial role of state-backed schemes and reinsurance institutions, such as France's Caisse Centrale de Réassurance (CCR), in maintaining risk-sharing mechanisms under climate stress.

Our empirical mapping confirms that insurance-related studies tend to form well-defined methodological clusters focused on damage functions, catastrophe modeling, and loss distribution. This concentration suggests that, while the insurance sector has developed robust quantitative tools to model risk ex-ante, it still faces structural limits in dealing with correlated, high-magnitude physical events.

5.2 The banking sector

For the banking sector, the main challenge identified by our analysis lies in the translation of physical hazards into credit and solvency risks. The SOM clusters show that banking-related research remains scattered, often bridging climate science and prudential regulation only indirectly through scenario analysis or macroeconomic stress testing. As emphasized by Battiston and Monasterolo (2020) [12], physical and transition shocks can propagate through balance sheets, affecting borrowers' cash flows, asset valuations, and collateral values.

Banks face a double asymmetry: limited data on firm-level exposure to physical risks, and uncertainty on how these exposures translate into credit parameters such as Probability of Default (PD) and Loss Given Default (LGD). Our results suggest that while progress has been made in modeling systemic contagion and portfolio stress testing, empirical applications linking hazard intensity to credit deterioration remain scarce. Developing this bridge between climate scenarios and internal risk models will be crucial for integrating physical risks into prudential frameworks and supervisory stress tests coordinated by the NGFS and the ECB.

5.3 C-level managers and corporate researchers

At the corporate level, the literature reflects growing concern with the operational and financial consequences of climate hazards for firms' assets and supply chains. However, the clustering analysis shows that corporate-focused studies are more heterogeneous than those in insurance or banking. They range from case-specific impact assessments (e.g., factory flood risk, agricultural yield loss) to cross-sectoral financial performance models linking Return on Assets (RoA) or market valuation to climate exposure.

This diversity reflects both the maturity of corporate climate disclosure and the fragmentation of analytical approaches. While many firms now recognize the materiality of physical risks, few possess the quantitative tools to measure their financial exposure internally. The literature thus points to a widening gap between corporate sustainability reporting often qualitative and narrative and the need for integrated financial modeling. Future work could leverage automated literature clustering to identify best practices in firm-level risk quantification and disclosure alignment with TCFD standards.

5.4 Academia and methodological development

For academic researchers, the clustering results reveal a rapidly expanding field but lacking standardization. Studies employ a variety of modeling paradigms from econometric regressions to catastrophe models, agent-based simulations, and network analysis. Yet few explicitly compare their predictive validity. The use of self-organizing maps and multiple correspondence analysis in this paper highlights the possibility of building meta-analytical tools to navigate this diversity and take a look "inside" the models.

From a methodological standpoint, academia has a key role to play in bridging disciplinary silos. By combining climate data science, actuarial modeling, and financial econometrics, researchers can contribute to the development of reproducible, open-source frameworks for assessing physical climate risk. The eventual goal should be to move toward cumulative, interoperable knowledge: a living, self-organizing literature where each new study finds its place within a continuously updated conceptual map of climate-finance research.

6 Conclusion

6.1 Main Results

This study represents a first attempt to address the *double hidden risk* in climate-finance research. The first hidden layer lies in the growing but still underestimated exposure of financial systems to physical climate hazards. The second, less visible one lies in the heterogeneity and opacity of the models designed to measure these impacts.

Through a systematic data collection and the use of unsupervised learning techniques, we proposed a self-organizing, or "auto-adaptive," framework capable of mapping and classifying this fragmented body of work. By applying Self-Organizing Maps (SOM) to a curated corpus of 150 peer-reviewed articles, we identified the main methodological families used to estimate the financial consequences of physical climate risks.

The resulting map reveals both structure and imbalance within the field. Insurance and actuarial research form a cohesive cluster built on well-established damage functions and loss-distribution models. Banking studies remain more scattered, reflecting the ongoing integration of physical hazards into credit-risk and solvency frameworks. Corporate-level studies are heterogeneous, often focused on case-specific impacts rather than generalized models. Academic work as a whole shows limited cross-fertilization between disciplines, indicating that model diversity has become both a strength and a source of opacity.

In short, our results suggest that while physical climate risks are increasingly modeled across sectors, the models themselves remain insufficiently connected. The proposed self-organizing review provides a first quantitative step toward lifting this second veil of invisibility, offering a reproducible structure through which future research can be continuously organized and compared.

6.2 Main Perspectives

The broader perspective emerging from this research points toward the automation of literature reviews and model clustering in interdisciplinary fields. A self-organizing review enables continuous updates and dynamic classification of research contributions as new papers are published. In future developments, integrating Variational Autoencoders (VAE) and interactive latent-space visualization could allow for adaptive, real-time exploration of scientific landscapes. Such a system would not only reduce the cognitive load of manual literature synthesis but also enhance transparency and reproducibility in how knowledge structures evolve over time.

This perspective aligns with the growing need for methodological coherence in climate finance: to connect fragmented knowledge, reveal hidden relationships between studies, and support cumulative scientific progress. The eventual objective is the creation of an open, evolving cartography of research, capable of informing both academic inquiry and practical decision-making in the face of accelerating climate risks.

6.3 Limits

As with any empirical review, several limitations should be acknowledged. First, *selection bias* may have affected the corpus construction: the reliance on Scopus and similar databases may exclude relevant gray literature, regional studies, or unpublished models. Second, *realization bias* arises from differences in data reporting, methodologies, and assumptions among studies, which may distort comparisons or cluster boundaries. Finally, *social desirability bias* may influence how certain results are framed, particularly in the corporate and policy-oriented literature, where the pressure to align with ESG or sustainability narratives can affect reporting transparency.

These limitations do not undermine the validity of the proposed framework but rather highlight the importance of iterative refinement and human oversight in any automated review process. The next step lies in scaling this approach to larger datasets, integrating new forms of unstructured data (e.g., full-text parsing and abstract embeddings), and reinforcing the dialogue between automated systems and expert judgment.

Acknowledgments

This is a work in progress. It was funded by the Institut Louis Bachelier's **Fondation Parc** and Fondation du Risque in Paris. We thanks Stéphane Voisin, Peter Tankov, the Institut des Actuaires for very insightful conversations and contributions to this work.

Appendix A — Literature Review Framework

The following framework was used to classify and code each article collected through the Scopus search (keywords: *Physical Risk* and *Financial Impacts*). It provides a structured view of the variables used to organize the corpus according to both the **Physical Risk** and **Financial Impact** dimensions. Each entry in the database was reviewed manually, and coded using the following categories: general information, financial and physical risk indicators, research design, and additional metadata (e.g., citations, keywords, end-user).

Table 8: Framework used to classify the articles included in the literature review

Category	**Subcategories / Variables**
General Information	Link; Paper name; Paper type; Scope / Covered area; Sector; Timeframe
Financial Impact	Financial impact precision; Financial impact type; Financial impact proxies; Financial data source; Financial sample size
Physical Impact	Physical risk precision or resolution; Physical risk type; Physical risk proxies; Physical data sources; Physical sample size
Research Design Information	Control variables; Control variable sources; Method; Method type; Projection (yes/no) ; Considered scenario (RCP / SSP) ; Equations; Castle (main figure); Summary; Key findings and results; Limitations and assumptions
Other Information	Developing vs. developed countries; Damage function (yes/no); Transition risk; Accuracy; Authors; Keywords; Journal; Number of citations; Number of references; Year; End-user

References

[1] Johan Rockström, Will Steffen, Kevin Noone, Åsa Persson, F. Stuart Chapin, Eric Lambin, Timothy M. Lenton, Marten Scheffer, Carl Folke, Hans Joachim Schellnhuber, Björn Nykvist, Cynthia A. de Wit, Terry Hughes, Sander van der Leeuw, Henning Rodhe, Sverker Sörlin, Peter K. Snyder, Robert Costanza, Uno Svedin, Malin Falkenmark, Louise Karlberg, Robert W. Corell, Victoria J. Fabry, James Hansen, Brian Walker, Diana Liverman, Katherine Richardson, Paul Crutzen, and Jonathan Foley. A safe operating space for humanity. *Nature*, 461(7263):472–475, 2009.

[2] Donella H. Meadows, Dennis L. Meadows, Jørgen Randers, and William W. Behrens. *The limits to growth*. Universe Books, New York, 1972.

[3] James Hansen, David Johnson, Andrew Lacis, Steve Lebedeff, Patrick Lee, David Rind, and Gary Russell. Climate impact of increasing atmospheric carbon dioxide. *Science*, 213(4511):957–966, 1981.

[4] Nicholas Georgescu-Roegen. *The entropy law and the economic process*. Harvard University Press, Cambridge, MA, 1971.

[5] Madison Condon. Climate services: The business of physical risk. *Boston University School of Law Research Paper*, 55:147–, 2023. Available at SSRN: https://ssrn.com/abstract=4396826.

[6] Oscar Higuera Roa, David N. Bresch, Jeroen C. J. H. Aerts, Thijs Tiggeloven, Reinhard Mechler, Stéphane Hallegatte, and Laurens M. Bouwer. Challenges and opportunities in climate risk assessment: Future directions for assessing complex climate risks. *Environmental Research Letters*, 20(5):053003, 2025.

[7] Teuvo Kohonen. Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43(1):59–69, 1982.

[8] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[9] Carl R. Woese, Otto Kandler, and Mark L. Wheelis. Towards a natural system of organisms: Proposal for the domains archaea, bacteria, and eucarya. *Proceedings of the National Academy of Sciences*, 87(12):4576–4579, 1990.

[10] Víctor Cárdenas. Financial climate risk: A review of recent advances and key challenges. *arXiv preprint arXiv:2404.07331*, 2024.

[11] Arthur Charpentier. Insurability of climate risks. *The Geneva Papers on Risk and Insurance: Issues and Practice*, 33(1):91–109, 2008.

[12] Stefano Battiston and Irene Monasterolo. On the dependence of investors' probability of default on climate transition scenarios. *SSRN Electronic Journal*, 2020.

[13] O. J. Oguntuase. Climate change, credit risk and financial stability. In *Banking and Finance*. IntechOpen, 2020.

[14] Task Force on Climate-related Financial Disclosures (TCFD). Recommendations of the task force on climate-related financial disclosures. Technical report, Financial Stability Board, Basel, Switzerland, 2017.

[15] Wayne C. Palmer. Meteorological drought. *Research Paper No. 45, U.S. Weather Bureau*, pages 1–58, 1965.

[16] Thibault Le Guenedal, Miguel Alvarado, Brian Mullen, and Alexander Ralston. A probabilistic model for tropical cyclone risk assessment: Application to insurance losses in the north atlantic basin. *Natural Hazards and Earth System Sciences*, 19(6):1337–1355, 2019.

[17] Allan Beltrán, David Maddison, and Robert J. R. Elliott. Is flood risk capitalized into property values? *Ecological Economics*, 146:668–685, 2018.

[18] John T. Abatzoglou and A. Park Williams. Impact of anthropogenic climate change on wildfire across western u.s. forests. *Proceedings of the National Academy of Sciences*, 113(42):11770–11775, 2016.

[19] European Commission. Commission delegated regulation (eu) 2021/2139 of 4 june 2021 supplementing regulation (eu) 2020/852 of the european parliament and of the council by establishing the technical screening criteria for determining the conditions under which an economic activity qualifies as contributing substantially to climate change mitigation or climate change adaptation and for determining whether that economic activity causes no significant harm to any of the other environmental objectives. Technical Report C(2021) 2800 final, European Commission, Brussels, 2021. Annex I, Appendix A – Classification of climate-related hazards.

[20] C. Herweijer, N. Ranger, and R. E. T. Ward. The insurance industry and climate change: Contribution to the global debate. *The Geneva Papers on Risk and Insurance: Issues and Practice*, 34(3):359–381, 2009.

[21] Donghyun Kim and Junbeom Lee. Development of a web-based tool for climate change risk assessment in the business sector. *Sustainability*, 8(10):1013, 2016.

[22] Robert Libby. *Accounting and human information processing: Theory and applications*. Prentice Hall, Englewood Cliffs, NJ, 1981.

[23] Philippe Dujardin, Frédéric Grelot, Jérôme Rouquette, and Darren Lumbroso. A multi-agent system for modeling decision-making processes in the financial sector exposed to flood risk. *Decision Support Systems*, 107:64–75, 2018.

[24] Ludovic Lebart, Alain Morineau, and Marie Piron. *Statistique exploratoire multidimensionnelle: Visualisation et inférence en fouille de données*. Dunod, Paris, 2000.