PARC
Paris Agreement Research Commons

CLIMATE RELATED FINANCIAL RISKS COOKBOOK

A Practical Guide to Comparing and Selecting Models That Assess the Financial Impacts of Climate Hazards



Institut des ACTUAIRES

2026

TABLE OF CONTENTS

FOREWORD 3

AUTHORS AND CONTRIBUTORS 4

EXECUTIVE SUMMARY 5

KEY OBJECTIVES 7

KEY FINDINGS 10

READER'S GUIDE 17

PART 1: THE MENU – UNDERSTANDING PHYSICAL CLIMATE RISK AND FINANCIAL IMPACT 19

PART 2: THE RECIPES – FAMILIES OF PHYSICAL RISK FINANCIAL IMPACT MODELS 44

PART 3: THE INGREDIENTS – DESIGN CHOICES AND DATA 236

PART 4: MODEL CLUSTERING 341

PART 5: THE CLIMATE FINANCIAL RISK HUB, REFERENCES AND APPENDIX 360

APPENDIX 1 – QUESTIONS TO ASK CLIMATE DATA AND MODEL PROVIDERS 390

APPENDIX 2 - MODEL CATEGORY DESCRIPTIONS 401

TABLE OF CONTENTS 402

LIST OF FIGURES 404

LIST OF TABLES 408

LIST OF PROVIDER BOXES 411

FOREWORD

This publication was produced by the **PARC (Paris Agreement Research Commons) Foundation**, hosted by the **Institut Louis Bachelier.** It continues the Alignment Cookbook series previously published by the Foundation. Following the publication of the latest Alignment Cookbook (2024), this new volume focuses on models used to assess the financial impacts of physical climate risks for banks, insurers and asset managers.

This work was made possible with the support of the **Institut des Actuaires (French Institute of Actuaries)**.

About the PARC Foundation

The **PARC (Paris Agreement Research Commons) Foundation**, established within the Fondation du Risque and hosted by the Institut Louis Bachelier, supports interdisciplinary academic research addressing the key challenges of the Paris Agreement. Bringing together researchers, regulators, financial institutions and other economic stakeholders, the Foundation develops independent research, methodologies and analytical tools to support the alignment of financial systems with sustainability objectives and contribute to the achievement of the Paris Agreement goals.

About the Institut des Actuaires

The **Institut des Actuaires (French Institute of Actuaries)** is the professional body responsible for organising and representing the actuarial profession in France. It promotes high standards of actuarial practice, ensuring compliance with professional standards and ethical principles while contributing to long-term financial and economic stability. As a member of the International Actuarial Association (IAA) and the European Actuarial Consultative Group, the Institut des Actuaires works to strengthen the recognition and reputation of the actuarial profession in France and internationally. In a context of growing risks and evolving prudential regulations, it plays a key role in upholding professional standards and supporting the actuarial profession in addressing major economic and societal challenges.

Contact

parc@institutlouisbachelier.org

www.parc-research.org

Suggested citation

Institut Louis Bachelier et al. (2026). Climate Related Financial Risks Cookbook: A Practical Guide to Comparing and Selecting Models That Assess the Financial Impacts of Climate Hazards.

AUTHORS AND CONTRIBUTORS

Lead author: P. Vigneron (Institut Louis Bachelier)

Authors: S. Voisin (Institut Louis Bachelier, PARC Foundation)

The authors would like to express their sincere gratitude to the **Institut des actuaires**, particularly **F. Picard**, **L. Barry**, **A. Charpentier**, **O. Lopez** and **M. Thomas**, for their valuable support and contributions throughout this project.

The authors are also grateful to the experts and financial institutions who contributed to the development of this report through bilateral consultations, interviews and reviews. In particular, they would like to thank the Institutional **Investors Group on Climate Change** (IIGCC), the **United Nations Environment Programme Finance Initiative** (UNEP FI), the **ACPR**, the **Banque de France**, and all the representatives from insurance companies, banks, asset owners and asset managers for sharing their expertise and perspectives.

Special thanks are also extended to **P. Tankov** (ENSAE, Institut Polytechnique de Paris and Institut Louis Bachelier), **J. Raynaud** (Institut Louis Bachelier Fellow), **C. Gardes** (IMF), **A. Mandel** (Université Paris 1 Panthéon-Sorbonne), **T. Barreau**, **N. Toroman**, **M. Rais** and **E. Dutreuil** (Institut Louis Bachelier) for their valuable insights and contributions to this research.

Finally, the authors would like to thank all the **model developers and data providers** who generously shared information on their methodologies and contributed to the review and validation of the case studies presented in this report.

The authors used large language models (LLMs) to assist with document formatting, proofreading, and the creation of illustrations. All AI-assisted outputs were reviewed and edited by the lead author, who takes full responsibility for the content of this work.

EXECUTIVE SUMMARY

Ten years ago, the Task Force on Climate-related Financial Disclosures (TCFD) was launched by the Financial Stability Board (FSB) in the wake of the Paris Agreement. It helped bring climate risks to the attention of financial institutions, businesses and policymakers, while establishing the first common frameworks for identifying, assessing and disclosing their financial implications. The TCFD completed its mandate in 2023, but the challenge it set out to address remains. Recognizing climate risk is one thing. Measuring its financial consequences with sufficient rigor to inform decisions is another.

This Cookbook builds on the foundations laid a decade ago and takes up that next challenge. Over the past ten years, a growing range of models has emerged to translate physical climate hazards into financial impacts. They differ in their methods, assumptions, data requirements and outputs, often substantially. The resulting landscape is rich but fragmented, making it difficult for users to understand what each model can reliably deliver, how models compare, and which are fit for a given purpose.

Meanwhile, the reality of climate change has become increasingly apparent through physical impacts whose financial consequences are ever more tangible. What was once largely perceived as a long-term concern is increasingly reflected in insurers' claims, banks' loan losses, asset impairments and disruptions to economic activity. From floods and wildfires to droughts and heatwaves, physical climate hazards are no longer distant scenarios. They are already affecting asset values, business operations and financial stability, making the need to measure and anticipate their financial consequences more pressing than ever.

This growing financial materiality also reinforces the distinction, central to the TCFD framework, between physical and transition risks. While transition risks arise from the policy, technological and market responses to decarbonization, physical risks stem from changes in the climate system itself: acute events such as floods, cyclones, wildfires, droughts and heatwaves that can directly damage assets, and chronic changes such as rising temperatures, sea-level rise and water stress that can erode their value more gradually. Measuring physical climate risk requires combining three components: hazard, exposure and vulnerability.

$$\text{Physical Climate Risk} = \text{Hazard} \times \text{Exposure} \times \text{Vulnerability}$$

As Physical Climate Risk equals Hazard times Exposure times Vulnerability:

- hazard is the physical event itself
- exposure is the assets or economic activity within its reach
- vulnerability is the degree of damage sustained per unit of hazard intensity.

Climate risks are physical by nature. Hazards do not always have a financial impact. A flood is a hydrological event, a wildfire a combustion process, a drought a climatic anomaly. When they happen during the life of an economic process, they turn into physical climate risks. Financial institutions must then translate physical observations into estimates of loss, impairment, default or repricing, through a chain that runs from climate signal to physical hazard, to physical damage, to operational disruption, to economic consequence, and finally to financial impact. Each link in that chain introduces its own uncertainty, its own proxy, and its own modelling choice. This translation problem, more than any single hazard or dataset, is the central object of this cookbook.

Asset-level vulnerability assessment sits inside the hazard-to-damage link of this chain, since physical damage is precisely "hazard intensity" filtered through an asset's vulnerability, instead of forming a separate step of its own. Part 3.2 unpacks that modelling choice in full, covering how exposure and vulnerability assumptions, location, construction type, adaptation measures already in place, are built into a given model.

A large and rapidly expanding market of catastrophe models in which we identified over 100 data providers, climate scenarios, geospatial datasets, vulnerability functions and financial translation tools has emerged to perform this translation. Over 200 models were reviewed to build this cookbook.

Many of these models remain partially proprietary and operate as black boxes. Two models built on identical climate scenarios can produce radically different financial outputs, not because one is wrong, but because each operationalizes a different transmission mechanism between hazard and financial variable. This divergence, referred to throughout the cookbook as model risk, creates a genuine danger for practitioners: concluding there is no risk when one exists, or the reverse.

This Cookbook sets out to bring clarity to the modelling landscape. It reviews and classifies existing models, examines their underlying assumptions and limitations, and draws practical lessons from their differences. Its purpose is not to identify a single best model, but to provide a structured guide to choosing and using models appropriately. The aim is to turn a complex and expanding body of modelling approaches into actionable guidance for their intended users, from researchers to, above all, practitioners in financial institutions.

This cookbook is scoped primarily around a risk-management use case, helping an institution price, provision or capitalize for a physical hazard it is already exposed to. A related but distinct use case, identifying which assets most need adaptation or resilience investment, asks a different question of the same models and can point to different answers, since a model built to price risk accurately is not automatically built to rank adaptation opportunities. A bridge between the two exists through resilience. The models reviewed here will also support the resilience strategies at the core of the adaptation plans of financial institutions, provided they make the drivers of vulnerability and exposure actionable. **Key finding 10 sets** out what such a bridge requires, and Warning 3 its limits.

CLIMATE SIGNAL	PHYSICAL HAZARD	PHYSICAL DAMAGE	OPERATIONAL DISRUPTION	ECONOMIC CONSEQUENCE	FINANCIAL IMPACT
Rising temperatures, Shifting precipitation, Sea-level rise	Flood, Wildfire, Drought, heatwave, storm	Asset destruction, Crop loss, infrastructure failure	Production halt, supply chain break, labour loss	Revenue decline, productivity loss, fiscal stress	Default, impairment, claims, repricing, systemic risk

FIGURE 1. From climate signal to financial impact: the transmission chain of physical climate risk

Scope and method

This cookbook anchors its analysis in the hazard classification set out in Annex A of the EU Taxonomy's Do No Significant Harm criteria, distinguishing chronic hazards, such as changing temperature patterns, sea-level rise and shifting precipitation, from acute hazards, such as heatwaves, wildfires, storms, floods and droughts. This choice ensures direct relevance for European financial institutions operating under evolving prudential and disclosure requirements, without claiming to cover every hazard or modelling approach on the market. The panorama presented in this cookbook combines a systematic literature review conducted through Scopus and Web of Science, over thirty interviews with climate analytics providers and executives, ten interviews with end-users inside financial institutions, and recurring exchanges with a dedicated scientific working group and the wider actuarial and financial-risk research community. Exploratory clustering techniques, including Kohonen self-organizing maps and variational autoencoder projections, were then used to organize the resulting body of methodologies into the analytical structure presented here.

A parallel effort was conducted on the Pladifes platform, an Institut Louis Bachelier initiative that centralizes information on ESG and sustainable-finance data providers, including a dedicated physical risk section covering exposure, hazard and vulnerability datasets (*see the Climate Risk Models Hub*).

It complements the market panorama built for this cookbook and gives readers a useful entry point for exploring data providers beyond the models and providers reviewed in Parts 2 and 3, and it presents both the clusters and an overview of every model encountered while building this cookbook.

KEY OBJECTIVES

The Cookbook pursues four objectives, each responding to a practical difficulty raised by the fragmented modelling landscape described above: understanding how models work, using them, choosing and adapting them, and navigating the market as a whole.

Objective 1. Increasing transparency. We clearly specify, step by step, how physical hazards are translated into financial impacts. This cookbook does not treat each model as an opaque estimate. It makes explicit the hazard, exposure and vulnerability assumptions behind it, the proxies used at each stage of the transmission chain, and the financial output ultimately produced.

Objective 2. Providing an operational toolbox. We offer a practical toolbox for the banking sector, the insurance industry, and asset owners and asset managers, organized around the financial object each institution actually manages: a loan, an insurance policy, a bond, an equity holding, or a diversified portfolio.

Objective 3. Supporting model selection and adaptation. We equip practitioners with a structured map of the design choices, hazard representation, exposure and vulnerability modelling, financial proxy selection, and aggregation, that differentiate one model from another (Part 3), so that a model can be adapted to the hazard, the asset class, and the regulatory context at hand.

Objective 4. Organizing the ecosystem. We provide classification tools (Part 4), to explore the landscape and a transparent account of the underlying research protocol (Part 5) capable of organizing a rapidly expanding and heterogeneous market into coherent, comparable analytical families, instead of leaving practitioners to navigate an undifferentiated list of providers and methodologies.

How to talk about Physical Climate Risks Financial impacts models

First, in order to clarify how different models are used in different contexts, we will present the main ones in Part 2:

SECTION	MODEL FAMILY	OUTPUT	USEFUL FOR
2.1.1	Damage functions (cat model)	Asset-level loss	All downstream financial models, via AAL and PML
2.1.2	Actuarial models	Compound loss distribution	EP curves; PD models; capital requirements
2.1.4	Loss and frequency curves (OEP/AEP)	PML at return period T; TVaR	Stress tests; reinsurance pricing; Solvency II SCR
2.2	Probability of default models	PD uplift; LGD uplift; credit VaR	IFRS 9 provisioning; Basel capital; bond pricing

SECTION	MODEL FAMILY	OUTPUT	USEFUL FOR
2.3	Accounting and valuation-impact models	EBITDA impact; enterprise value change; ratio deterioration	Credit analysis; equity valuation; regulatory disclosure
2.4	Network, supply-chain and systemic models	Network loss multiplier; systemic VaR	Macro stress tests; supervisory risk assessment
2.5	Scenario, stress-test and catastrophe metrics	Scenario range; stress VaR; OEP/AEP/PML	All applications requiring forward-looking risk quantification

TABLE 1. Model families reviewed in Part 2, with their outputs and uses

Each of these model families is described step-by-step:

- By explaining what they are designed for
- By laying down their mathematical representation
- By offering examples of their application in the industry
- By putting forward providers present today in the market that use and sell these kinds of models.

These provider-facing sections are built with this question in mind, what the models actually sold on the market do, not only the methodologies published in the academic literature. Each major vendor model reviewed gets the same card-based treatment described below, so a reader can compare a commercial tool against a published damage function on the same terms.

Each card also answers a second question, about scale. At which scale is the result directly estimated, and from which scale onward is it aggregated, extrapolated or propagated? An excellent asset-damage model is not automatically a sound portfolio model, and a macroeconomic model is more than the sum of asset-level models. Each card therefore carries a scale suitability field, which records the levels at which a model is robust and the levels at which its use becomes an extrapolation.

For readers less familiar with actuarial and credit terminology, the table above (Table 1) uses AAL (Average Annual Loss) and PML (Probable Maximum Loss) for expected and worst-case asset losses, EP (Exceedance Probability) and TVaR (Tail Value at Risk) for loss-distribution metrics, OEP and AEP (Occurrence and Aggregate Exceedance Probability) for return-period losses, PD, LGD and credit VaR (Probability of Default, Loss Given Default and credit Value at Risk) for credit-risk outputs, and IFRS 9 and Solvency II SCR for the accounting and insurance capital rules those outputs feed into.

Second, in order to clarify how these individual models from the industry or from the literature work and what they produce, we offer a simpler representation of them.

Models compress a physical observation into a financial number, anchored first in a geographical location, a data source and a scenario, then carried through a physical side, the resolution at which the hazard is observed, how many hazards are considered, how the data is accessed and how often it refreshes, through a single method function, a damage function, a regression, a structural model, anything that the paper actually uses, into a financial side, the granularity of the output, its accuracy, and the sector it applies to.

Control variables sit underneath both sides, the elements that condition the transformation, land use, prices, construction cost, alongside their sources. Each card offers a simplified reading of the underlying mathematical model, keeping what determines whether it is usable, comparable, and trustworthy for a given financial application

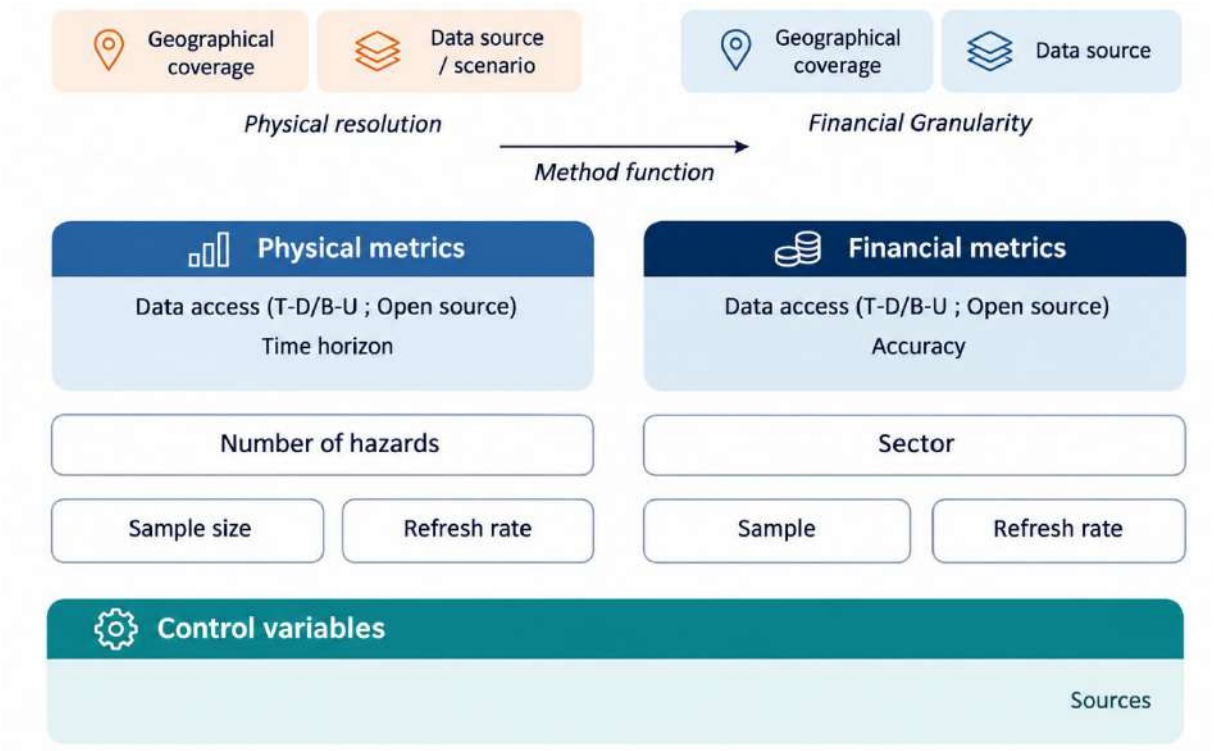


FIGURE 2. Structure of the model card, from the physical signal to the financial output

The left and right sides of the card mirror the two halves of the same transformation.

The physical signal is on the left, the financial output is on the right, with the method function as the hinge between them.

This placement lets a reader compare models by scanning straight across, resolution against granularity, time horizon against time horizon, access type against access type. When the physical and financial blocks report different geographies, resolutions, or scenarios, that asymmetry becomes visible at a glance instead of staying buried inside a paper's methodology section, and this kind of mismatch is identified in Part 1 as a structural feature of the market. It drives much of the divergence between models built on the same hazard.

KEY FINDINGS

Ten key findings emerge from the review of over 200 models and the interviews conducted for this Cookbook. Findings 1 to 3 concern the nature of physical risk and the uncertainty attached to its measurement. Findings 4 to 8 address how models differ and how to match them to the question at hand. Findings 9 and 10 look at what institutions are now expected to do with them, from regulatory compliance to adaptation. The analysis is anchored in the hazard classification of the European Union (EU) Taxonomy, which ensures direct relevance for European institutions; the framework applies beyond Europe, but emerging economies, where observed hazard, exposure and loss data are scarcer, will meet more of the data limitations discussed below, and the uncertainty disclosure card (see Warning 2) matters all the more there.

Key finding 1. Physical risk is measurable, and extreme events belong in reference stress scenarios. Physical risk is a quantifiable risk factor that transmits to existing risk categories (credit, market, liquidity, underwriting), captured through the combination of hazard, exposure and vulnerability. Approaches calibrated on historical data and designed for stationary environments systematically underweight the low-probability, high-severity events that matter most under a changing climate. These events should be treated as reference stress scenarios, not discounted as outliers.

Key finding 2. Turning physical damage into financial loss is the hardest step to model. The transformation of physical damage into financial loss is a key modelling challenge. Between the physical hazard and the financial metric lies a chain of intermediate mechanisms, assumptions and proxies. The further a model moves along this chain, the greater the uncertainty generally becomes. The closer a variable sits to the physical process, the more its representation should be constrained by physics wherever possible; the closer it sits to financial variables, the greater the weight of behavioral and statistical assumptions, and the wider the uncertainty that should be disclosed with the result.

Key finding 3. Uncertainty compounds along the chain and mostly goes undisclosed. Uncertainty compounds across hazards and stages, and today it goes largely undisclosed. Modelling several hazards together raises the risk of double-counting, while modelling too few misses their interactions, and tipping points remain inherently unpredictable. Most models stay deterministic, delivering a single point estimate, and available evidence points to a systematic underestimation of physical financial impact, consistent with warnings from the Network for Greening the Financial System (NGFS) and the Financial Stability Board (FSB). The NGFS itself acknowledges that its long-term scenarios tend to underestimate acute physical risk (NGFS, 2024). Physical risk modelling carries a structural uncertainty that calls for documented and validated expert judgement, and its outputs should be presented with corresponding transparency (see Warning 2).

Key finding 4. The financial object determines the choice of model more than the hazard does. Physical climate risk analysis is inherently multi-scale. The same hazard may simultaneously affect a single physical object, an income-generating asset, a loan, a company, a portfolio, a national economy, or the financial system as a whole. The financial object under analysis, more than the hazard itself, determines which type of model, and which approach, is most relevant.

ANALYSIS LEVEL	FINANCIAL OBJECT AT RISK	EXAMPLES	FINANCIAL OUTPUTS	SPATIAL SCALE
Physical object or good	Single physical item	Inventory, machinery	Replacement cost	Site

ANALYSIS LEVEL	FINANCIAL OBJECT AT RISK	EXAMPLES	FINANCIAL OUTPUTS	SPATIAL SCALE
Asset	Income-generating asset	Building, plant	Asset loss, downtime	Asset
Infrastructure	Network asset	Road, power grid	Repair cost, service loss	Network
Loan	Credit exposure	Mortgage, corporate loan	PD, LGD, LTV	Borrower
Company	Operating firm	Manufacturer, utility	Revenue, EBITDA, cash flow	Firm
Portfolio	Aggregated holdings	Loan book, fund	VaR, loss, capital	Multi-region
National economy	National economy	Sovereign debt, public finances	GDP, fiscal balance, spread	National
Regional / international network	Cross-border economic network	Supply chain, commodity market	Trade loss, price shock	Regional / international
Financial system	System-wide finance	Banking system, capital markets	Systemic risk, liquidity, solvency	International / global

TABLE 2. Financial objects at risk, from the most granular to the broadest scale of analysis

Table 3 turns this multi-scale diagnosis into an evaluation criterion. A model is direct at a level when it estimates its output there from observed or simulated data, aggregated when finer-level results are combined under correlation and diversification assumptions, extrapolated when its output is carried to a level it was never calibrated for, and propagated when the shock is transmitted through a network, a balance sheet or the financial system (see Part 3.4).

MODEL FAMILY	ASSET	LOAN / COMPANY	PORTFOLIO	NATIONAL ECONOMY	FINANCIAL SYSTEM
Damage functions and cat models (2.1)	Direct	Aggregated	Aggregated	Extrapolated	Out of scope
Actuarial models and EP curves (2.1)	Aggregated	Aggregated	Direct	Extrapolated	Out of scope

MODEL FAMILY	ASSET	LOAN / COMPANY	PORTFOLIO	NATIONAL ECONOMY	FINANCIAL SYSTEM
PD and credit models (2.2)	Out of scope	Direct	Aggregated	Extrapolated	Propagated
Accounting and valuation models (2.3)	Aggregated	Direct	Aggregated	Out of scope	Out of scope
Network, supply-chain and systemic models (2.4)	Out of scope	Extrapolated	Extrapolated	Direct	Propagated
Scenario and stress-test models (2.5)	Extrapolated	Extrapolated	Direct	Direct	Direct

TABLE 3. Scale suitability of the model families

Key finding 5. Diverging outputs and opaque methods create model risk. Model outputs diverge, and opacity turns that divergence into model risk. Outputs vary along four axes: the scenario, the hazard, the method or damage function, and the financial impact targeted. Two models sharing the same climate scenario can produce very different results without either being wrong. As much of the market relies on partially proprietary methodologies, this divergence creates a real risk of wrong conclusions, assuming a portfolio is safe when it is not, or the reverse, and correlated models should not be read as independent evidence (Part 4).

Key finding 6. Models improve when they are updated with new local information. Models should be continuously refined as new information becomes available. Local claims data, new satellite observations, updated engineering surveys or new scenarios can all sharpen an existing estimate. Bayesian updating is particularly valuable, allowing local signals to refine national or sectoral priors where the historical record is thin.

Key finding 7. There is no single best model, so the question should decide which model to use. There are only models suited to a specific financial object, hazard, data environment and regulatory use. Insurers, banks and asset managers ask different questions of the same hazard, and the choice of model should follow from the question, not the reverse. The five model families reviewed in Part 2 are complementary rather than competing: damage functions feed actuarial and credit models, which in turn feed portfolio and macro-level assessments, while scenario and stress-test methods supply the forward-looking inputs to all of them.

Key finding 8. Hazard data and financial decisions rarely work at the same resolution. Spatial and temporal resolution rarely match between hazard data and financial decision-making. Climate models operate at tens to hundreds of kilometers, while underwriting and credit decisions require asset- or address-level precision. Downscaling, interpolation and aggregation bridge the gap but add uncertainty, making high-resolution geospatial data a key competitive differentiator among providers.

Key finding 9. Physical risk now drives credit quality and portfolio performance, beyond compliance. Physical climate risk is no longer solely a compliance matter. Prudential and disclosure requirements, from the European Central Bank (ECB) and the French Prudential Supervision and Resolution Authority (ACPR) to the Basel Committee on Banking Supervision (BCBS), the European Banking Authority (EBA) Guidelines on the management of environmental, social and governance (ESG) risks under the sixth Capital Requirements Directive (CRD VI), the climate scenario analysis requirement introduced in the Own Risk and Solvency Assessment (ORSA) by the Solvency II review, the supervisory texts of the International Association of Insurance Supervisors (IAIS), the IFRS S2 climate-related disclosure standard and the French Energy-Climate Law (LEC) Article 29, reviewed in Part 1, require financial

institutions to embed physical risk into governance, existing risk categories, stress testing and disclosure. Physical risk is increasingly a direct driver of credit quality, collateral value and portfolio performance.

Key finding 10. Pricing risk and guiding adaptation call for different model capabilities. Risk quantification capability and adaptation decision-support capability are two distinct qualities of a provider. A model built to price risk accurately is not automatically built to rank adaptation options. For that, it must let users modify the drivers of vulnerability and exposure, such as flood protection, the location of critical equipment or business continuity arrangements, and observe the effect on financial loss (see Part 3.2).

The paradox of climate-risk precision

The findings above describe real progress in the measurement of physical climate risk. That progress carries risks of its own: precision can be mistaken for certainty, uncertainty can be hidden rather than disclosed, and better measurement, acted upon by each institution in isolation, can weaken the very resilience it is meant to support. The three warnings below address each of these risks in turn.

Warning 1. Better models must not breed false confidence.

A decade of work since the launch of the Task Force on Climate-related Financial Disclosures (TCFD) has produced an impressive body of research and significant advances in our understanding of the financial consequences of climate change. Financial institutions now have access to an expanding range of models, datasets and analytical tools that would have been difficult to imagine ten years ago. Yet it would be a serious mistake to confuse this accumulation of knowledge with a corresponding degree of control over climate-related financial risks. If anything, the growing sophistication of these tools should make us more aware of the limits of what we know.

The greatest danger may therefore be overconfidence: the belief that increasingly sophisticated models necessarily provide an increasingly reliable picture of the risks ahead. The parallel with the 2008 financial crisis is difficult to ignore. Then, too, complex pricing models created an illusion of precision and control, encouraging financial institutions to underestimate risks that ultimately proved systemic. Climate risk modelling must not repeat that mistake. The lesson is not to distrust models, but to resist the temptation to mistake their mathematical sophistication for certainty.

Warning 2. Uncertainty is not a modelling flaw to be eliminated. It is a fundamental feature of climate risk.

Unlike many conventional financial risks, physical climate risks cannot be reliably inferred from historical observations alone. Climate change is altering the frequency, intensity and interaction of hazards, while the economic and financial consequences of these changes depend on complex, often non-linear transmission mechanisms.

Uncertainty accumulates at every stage, from climate projections and hazard intensity to asset vulnerability, physical damage, economic disruption and financial loss. Each step introduces assumptions, proxies and modelling choices. Spatial resolution, aggregation effects, interactions between hazards and the possibility of tipping points add further layers of uncertainty. Yet most models still produce single-point estimates, often without disclosing how uncertain those estimates are.

This is more than a technical limitation. It creates the risk of a sudden and potentially destabilizing reassessment of asset values, echoing the climate-related “Minsky moment” warned of by Mark Carney when he chaired the FSB: a sharp repricing of risks that financial markets had previously underestimated. The response cannot simply be to build more sophisticated models. It must also be to make uncertainty visible and actionable. Providers should systematically disclose the sources of uncertainty in their models, the assumptions tested, the range of scenarios considered, the treatment of extreme events and the confidence that can reasonably be placed in their outputs.

To make this disclosure systematic, this Cookbook proposes an uncertainty disclosure card for each provider, which documents the information supplied without scoring the precision of the output: the sources of uncertainty acknowledged, the parameters subject to sensitivity analysis, the availability of an interval or a full distribution, the range of scenarios covered, the uncertainty introduced by aggregation and the treatment of distribution tails. The card is designed to be read alongside existing model risk management frameworks, such as the Federal Reserve's SR 11-7 guidance in the United States and the ECB supervisory expectations on climate-related and environmental risks, so that it supports model validation rather than replacing it. The key message for institutions is that a model can be outsourced, but the responsibility for understanding it cannot.

This has a direct implication for policy. Regulators generally have no direct hold on climate data and model providers, so the lever runs through the supervised institutions. Regulators and standard setters should expect banks, insurers and asset managers to require providers to disclose the confidence margins, limitations and scope of their models alongside the results, to request a scale suitability statement and an uncertainty disclosure card with each model they use, and to document this within their model risk management and governance frameworks, so that no result is used at a scale, or with a confidence, that the model does not support.

Crucially, quantifiable risk must be distinguished from deep uncertainty, where neither probabilities nor the underlying model structure can be established with confidence. Approaches combining physical and statistical models, alongside frameworks such as CLIMAFIN that explicitly address non-linearity, endogeneity and compound risks, offer promising ways forward. But no modelling architecture can eliminate uncertainty altogether.

A useful principle follows: the closer a variable is to the physical process, the more its representation should, wherever possible, be constrained by physical science. The further the analysis moves towards financial outcomes, the greater the role of behavioral and statistical assumptions, and the greater the need to disclose uncertainty.

Warning 3. Better risk pricing does not necessarily produce better outcomes.

There is a final paradox. Even if financial institutions become substantially better at measuring physical climate risks, this does not guarantee that the financial system, or society, will become more resilient.

For an individual institution, particularly an insurer, more accurate risk assessment can justify higher premiums, reduced coverage or withdrawal from the most exposed areas. Banks may respond by restricting lending, while investors may redirect capital away from vulnerable assets and territories. Each decision may be rational from the perspective of the institution concerned. Collectively, however, they can amplify the very risks they seek to manage.

The danger is twofold. As insurers withdraw, the pool over which risks are mutualized shrinks. At the same time, the volume of uninsured risk grows and becomes increasingly concentrated among households, businesses, banks holding exposed collateral and, ultimately, public finances. These two dynamics reinforce one another, potentially creating a feedback loop between declining insurability, falling asset values, reduced access to finance and growing financial instability.

This is another route through which the climate-related Minsky moment could materialize: not only through the sudden recognition of previously underestimated risks, but also through the collective consequences of institutions acting on more precise risk information.

The central lesson is that better measurement is necessary, but not sufficient. Greater precision in climate-risk assessment must be matched by stronger adaptation and resilience strategies, as well as collective mechanisms for sharing risks that individual institutions can no longer bear alone. This is where the next stage of work must begin, and where our ongoing IST research on risk-sharing and mutualization will seek to contribute. Together, these warnings and the findings of this Cookbook provide essential building blocks for financial institutions to design effective adaptation and resilience strategies, while making no claim to address the full range of challenges involved in developing and implementing them.

The Climate Risk Models Hub

The three warnings share one idea. Knowledge about physical climate risk is produced by many providers, and the financial system's ability to use it well depends on how openly it is shared, compared and questioned. Models evolve, new providers enter the market and methodologies change faster than any publication cycle can follow. The Cookbook is therefore be extended by the Climate Risk Models Hub, an open platform that gathers the models identified in this review, in the academic literature and on the market.

The Hub works in two directions. **Providers will be able to add and update their own models** using the descriptive framework developed in the Cookbook, covering the hazards modelled, spatial resolution, damage and financial transmission channels, data sources and validation practices. **Users such as banks, insurers, asset managers and supervisors are able to search and compare models** according to their exposures, the hazards they face, their time horizon and the regulatory use they have in mind. A conversational assistant will guide users who are less familiar with the field toward the models that fit their needs.

The Hub is designed as a direct response to the three warnings. Each model entry displays its assumptions and limits next to its outputs, so that sophistication is never read as certainty. Providers report the sources of uncertainty, the scenarios tested and the treatment of extreme events, in line with Warning 2, and placing models side by side shows where they diverge, which is itself a measure of the uncertainty that single-point estimates conceal. By treating knowledge about climate risk as a shared resource, the Hub gives smaller institutions the same reference base as larger ones and offers providers, users and supervisors a common vocabulary. This is a first step toward the collective response called for in Warning 3, since sharing knowledge comes before sharing risk. The Hub makes no claim to replace expert judgement. Its purpose is to make that judgement better informed, more transparent and easier to share.

What this means for you

The key findings do not carry the same message for every reader. The points below set out, for each type of user, what to do first.

Banks. Treat physical risk as a factor that feeds credit, market and liquidity risk rather than as a separate risk category, and test it through reference stress scenarios. Choose the model from the financial object (loan, collateral, counterparty, portfolio), and check its scale suitability (Table 3) before an asset-level result is used at portfolio level. A provider model used in credit processes should sit inside the existing model risk management framework, with the uncertainty disclosure card as part of its documentation.

Insurers. Underwriting, pricing, reinsurance and capital ask different questions of the same hazard, and call for different models. Check how each model treats distribution tails, correlated hazards and adaptation assumptions, and reflect the climate scenario analysis requirement of the ORSA in the choice of scenarios. More precise risk pricing can accelerate the withdrawal of cover described in Warning 3, which is a reason to read precision together with disclosed uncertainty.

Asset managers. A model built to price a single asset is not automatically a sound portfolio model. Ask how uncertainty is aggregated from asset to portfolio, and compare at least two model families when outputs diverge, rather than relying on one point estimate. Where models are highly correlated, agreement between them is not independent evidence.

Supervisors. Because supervisors have little direct hold on climate data and model providers, the lever runs through supervised institutions. Ask them to show how provider outputs are validated, documented and challenged within model risk management, use the scale suitability statement and the uncertainty disclosure card as a common basis for supervisory dialogue, and expect low-probability, high-severity events to appear among the reference stress scenarios.

For any user, the path from question to model follows four steps. Start from the financial object at risk (Table 2), state the question being asked of it, identify the model family suited to that question and the levels at which it is direct, aggregated, extrapolated or propagated (Table 3), then check in the uncertainty disclosure card whether the result can be used at that scale and with what confidence (see Appendix for more questions to ask providers).

READER'S GUIDE

PART / SECTION	QUESTIONS ANSWERED BY THE PART	PAGE
PART 1 — THE MENU: UNDERSTANDING PHYSICAL CLIMATE RISK AND FINANCIAL IMPACT		
1.1 WHAT IS A PHYSICAL CLIMATE RISK?	What is a physical climate risk, how does it differ from transition risk, and how is it measured?	19
1.2 FROM PHYSICAL DAMAGE TO FINANCIAL IMPACT	How does physical damage translate into financial impact, and how do the FSB and NGFS frameworks bridge the two?	28
1.3 FINANCIAL OBJECTS AT RISK	Which financial objects, from a single asset to the financial system, are exposed to physical risk?	31
1.4 PHYSICAL RISK AND FINANCIAL IMPACTS: MODELLING ONTOLOGY	What is a model, how can it be improved, and what are its limits?	35
1.5 THE PURPOSE OF THIS COOKBOOK AND ITS INTENDED USERS	Who is this cookbook for, insurers, banks, asset owners and asset managers, and how should each use it?	38
1.6 RESEARCH QUESTIONS	How can financial impacts be predicted from physical risks, including new or rare hazards, using a Bayesian approach?	40
PART 2 — THE RECIPES: FAMILIES OF PHYSICAL RISK FINANCIAL IMPACT MODELS		
2.1 DAMAGE-BASED MODELS	How is asset-level physical damage estimated, and how do actuarial models and loss and frequency curves turn it into a loss distribution?	49
2.2 PROBABILISTIC AND DEFAULT-BASED MODELS	How does a physical hazard affect a borrower's probability of default, system resilience, and liquidity risk?	101
2.3 ACCOUNTING AND VALUATION IMPACT MODELS	How does a physical hazard affect price, turnover, enterprise value and financial ratios?	136
2.4 NETWORK AND SUPPLY-CHAIN IMPACT ASSESSMENT MODELS	How do physical shocks propagate through supply chains, and what systemic and lagged effects follow?	169
2.5 SCENARIO-BASED AND PROJECTIVE MODELS	How are climate scenarios, SSP, RCP, NGFS, and catastrophe metrics built, and what are their predictive limits?	215
PART 3 — THE INGREDIENTS: DESIGN CHOICES		
3.1 THE HAZARD CHOICES AND THEIR REPRESENTATION	How are the hazards chosen and represented: single or multiple hazards, spatial resolution, historical or projected?	236

PART / SECTION	QUESTIONS ANSWERED BY THE PART	PAGE
3.2 EXPOSURE AND VULNERABILITY MODELLING CHOICES	How are asset location, socio-economic variables, and adaptation and resilience hypotheses modelled?	282
3.3 FINANCIAL PROXY CHOICES	Which financial, market or actuarial proxy should express the impact, and over what time horizon?	296
3.4 AGGREGATION CHOICES	How should individual results be aggregated, and what compound-risk, non-linearity, or Type I, II and III errors can this introduce?	306
PART 4 — MODEL CLUSTERING		
4.1 CLOSED CLUSTERS	How can models be classified using a fixed, predefined structure such as Kohonen self-organising maps?	345
4.2 OPEN CLUSTERS	How can models be classified using open, exploratory techniques such as PCA, t-SNE and variational autoencoders?	350
4.3 SUPERVISED CLUSTERISATION AND FILTERING	How can clustering of this kind be re-used and steered toward a specific analytical objective?	356
PART 5 — THE CLIMATE FINANCIAL RISK HUB, REFERENCES AND APPENDIX		
5.1 THE CLIMATE FINANCIAL RISK HUB	How can we build an organized and clustered model bank in an online platform?	360
5.2 A A PANORAMIC VIEW OF THE FINANCIAL RISK MODELLING ECOSYSTEM	How is the physical risk modelling market structured, what is model risk, and how was this panorama built?	365
5.3 REFERENCES	The full bibliography used to build this cookbook	370

TABLE 4. Reader's guide: questions answered by each part and section of the cookbook

1 THE MENU – UNDERSTANDING PHYSICAL CLIMATE RISK AND FINANCIAL IMPACT

The objectives of this first part are to:

1. Define what is physical climate risk and how it can cause significant financial impacts

2. Suggest a conceptual framework to compare the models used to assess these financial impacts

3. Explain the research protocol followed in order to better understand the market that is being built to assess financial impacts from physical risks.

1.1 WHAT IS A PHYSICAL CLIMATE RISK?

A risk is the possibility that an undesirable event may occur. In the context of global warming, two types of risks are rising.

Climate risks are gaining momentum in the financial sector. More specifically, physical climate risks (such as natural catastrophes) are increasingly important for every major financial institution and activities: asset management and asset owners (AO/AM) (MSCI, 2026), banks (Banque de France, 2025) and insurance companies (WWF, 2026) are at the heart of this turmoil.

Due to the combined increase of climate change and human activity, physical risks are following the trend. The compliance policies are evolving rapidly, especially in Europe (ACPR, 2021).

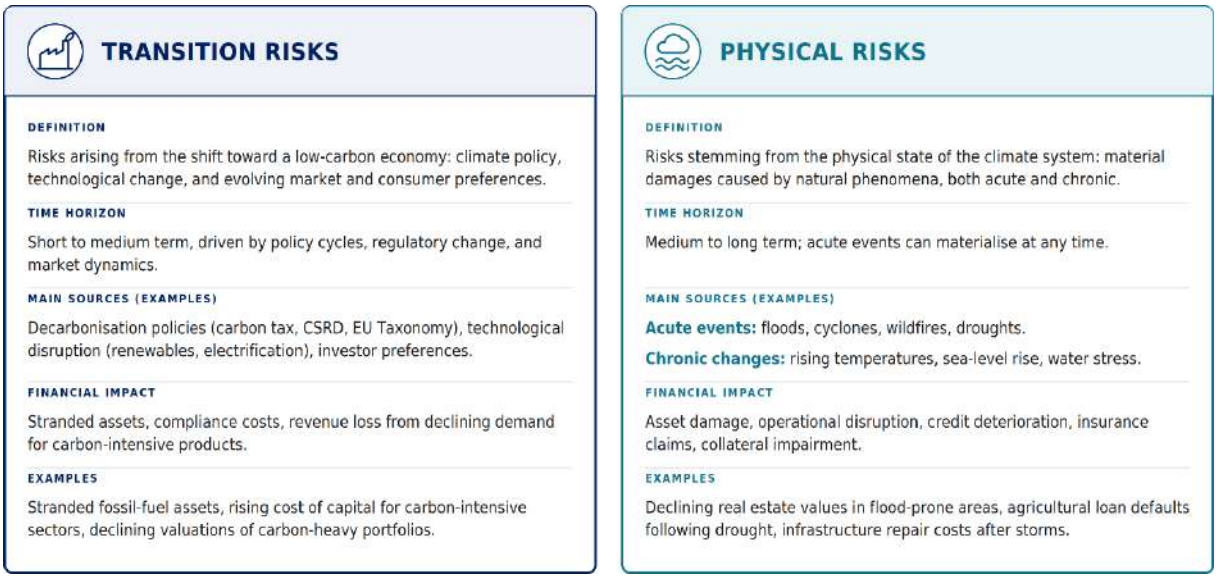


FIGURE 3. Physical Risks and Transition Risks

Physical risks (PR) are distinct from transition risks (TR). TR relates to the adjustment toward a lower-carbon economy. PR arises from the state of the climate system itself and from the material damage it inflicts on the world that financial institutions depend upon.

PR is the possibility that physical forces and natural processes will cause harm to people, built infrastructure, ecosystems, or economic activity. The importance of these events on people cannot (and perhaps should not) be quantified. However, the intensity they have and the threat they represent on regular activity is quantifiable in the context of financial institutions. The concept is sharpened into a specific analytical category: the financial losses or impairments to asset values, creditworthiness, and economic productivity that arise when physical hazards materialize.

Physical Climate Risks – How can one identify them?

Physical risks are distinct from transition risks. While transition risks are generated by policy, technology, and market responses to decarbonization, physical risks originate in the biophysical state of the planet.

Physical Risks are a function of the climate system and its interaction with modern economies. They depend on what we have built and where we have built it, infrastructures, natural systems, or economic activity arising from the physical dynamics of natural phenomena. They are to be considered as part of the life of our financial system, long before transition risks arise.

Underwriting

In a way, physical risks have been one of the driving forces behind the rise of the insurance industry. The very word "underwriting" traces back to the late 17th century, when merchants and ship owners would gather at Edward Lloyd's Coffee House on Tower Street, recorded in the London Gazette in 1688 (Raynes, 1950). Those willing to share in the risk of a maritime voyage would literally write their names beneath the description of the voyage and cargo, indicating how much of the loss they were prepared to absorb if the ship failed to return. The coffee house effectively became the first organized marketplace for the pricing and distribution of physical risk. It was essentially a wager on whether a vessel would survive storms, winds, and other natural elements on the way to the East Indies or the Americas. Then it was formalized into the contractual structures that underpin modern insurance and reinsurance. Lloyd's of London actually grew directly from this practice and remains today one of the world's leading insurance markets. Its activity has fundamentally not changed:

A) a physical hazard exists,
B) exposure (i.e. what's at risk financially) can be quantified
C) and the financial consequences can be distributed across multiple risk-bearers.

Non-linearity

However climate change has changed the game. As a mechanism through which physical risk enters financial analysis, by raising average temperatures and in some areas increasing the frequency and severity of extreme events, it shifts the entire distribution of hazard occurrence. It makes events much more difficult to predict.

For financial actors, it breaks a foundational assumption of traditional risk modelling: that the future resembles the past. Historical loss data, calibrated against a stationary climate, become progressively less reliable as the climate state drifts. Some of these difficulties can take the form of an unsolvable problem, such as tipping points (Condon, 2023).

The physical risks are usually divided into two categories, depending on the time they take to make a "visible damage": chronic (long-term) and acute (fast).

Both chronic (e.g. gradual sea-level rise, persistent heat stress, shifting growing seasons) and acute events (e.g. floods, cyclones, wildfires, and droughts) translate into physical damage with direct financial consequences.

The critical insight of climate risk analysis is that climate change is a directional, foreseeable shift in the probability landscape of physical loss, as well as the frequency, intensity, and geographic reach of physical hazards. Driven primarily by anthropogenic greenhouse gas emissions, it alters long-run climatic averages as well as the distribution of extreme weather events. For financial actors, this matters because it introduces non-stationarity into the risk environment: historical loss data, calibrated on a stable climate, become progressively less reliable as the climate state drifts. Both chronic changes (rising mean temperatures, shifting precipitation regimes, sea-level rise) and acute events (storms, floods, wildfires, heatwaves) translate into physical damage with direct financial consequences on financial assets, liabilities, and the broader economy.

Measuring physical risk requires combining three analytically distinct but interdependent components. Extensive work has already been conducted on this topic, making this approach a standard in the financial sector (TCFD, 2017; CISL and ClimateWise, 2019; UNEP FI, 2020, 2022, 2023, 2024; NGFS, 2020; NGFS & FSB, 2022; WRI, 2020).

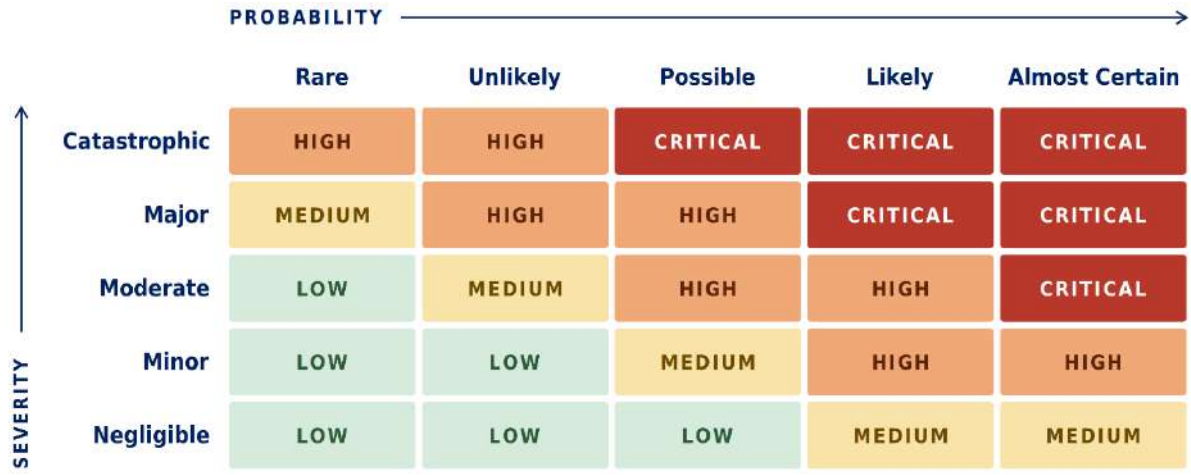


FIGURE 4. Standard Risk Matrix (adapted from ISO 31000 and ISO/IEC 31010)

In the standard ISO 31000 framework, risk = probability × severity. Colour bands: LOW (green), MEDIUM (yellow), HIGH (orange), CRITICAL (red).

“Traditional” risk approaches are not the best fit to capture the financial impacts of physical climate risks, whether in terms of probability or severity.

Calibrated on historical data and designed for stationary risk environments, they systematically underweight tail events, precisely the category of risk that matters most in a changing climate. This is why, even where the probability of occurrence is low, a high-severity event renders a portfolio or any asset exposure genuinely at risk.

A 1-in-100-year flood that destroys a logistics hub, or a prolonged drought that triggers cascading agricultural defaults, may each carry a low annual probability while generating losses that are financially catastrophic and, in some cases, unrecoverable.

Physical climate risk assessment therefore requires frameworks that explicitly account for the asymmetry between probability and impact treating low-frequency, high-consequence events not as outliers to be discounted, but as central scenarios to be stress-tested.



FIGURE 5. Reversed Risk Matrix: a Physical Climate Risk Perspective

In the physical climate risk framework, a low-probability but highly severe event (e.g. a 1-in-200-year flood) still constitutes a HIGH RISK because its potential financial impact is extreme. Traditional risk matrices underweight tail events; this reversed perspective aligns with prudential and actuarial practice for catastrophic hazards.



FIGURE 6. Hazard, exposure and vulnerability, the three components of physical risk

- **The Hazard (H)** component determines the likelihood of an event. Hazard refers to the physical event itself: its type, intensity, frequency, and spatial extent, as characterised by meteorological, hydrological, and geophysical observations.
- **The Exposure (E)** component measures the people, assets, or economic activities located within the reach of a given hazard. A flood becomes a financial event only when it reaches insured or productive assets. Exposure therefore represents the maximum potential loss in a given area. The relevant exposure differs by actor, for example loans for a bank, insured assets for an insurer, or physical assets for a company.
- **The Vulnerability (V)** component describes the degree of damage to the exposed assets and socioeconomic elements that can be expected at different hazard intensities. Vulnerability quantifies how much damage a given unit of exposure sustains per unit of hazard intensity. It is typically formalized through damage functions that translate physical impact into monetary loss.

The combination of hazard, exposure, and vulnerability constitutes the standard risk assessment framework adopted across the natural hazard sciences. It is already embedded in financial supervisory guidance (ECB, 2020; European Parliament, 2020, 2022; European Commission, 2021; BCBS, 2025; FSB, 2025). A recurring challenge is that the spatial and temporal resolution at which hazard data are available often does not match the resolution at which financial decisions are made. Global climate models produce outputs at scales of tens to hundreds of kilometres; financial underwriting and credit assessment require address-level precision (Condon, 2023).

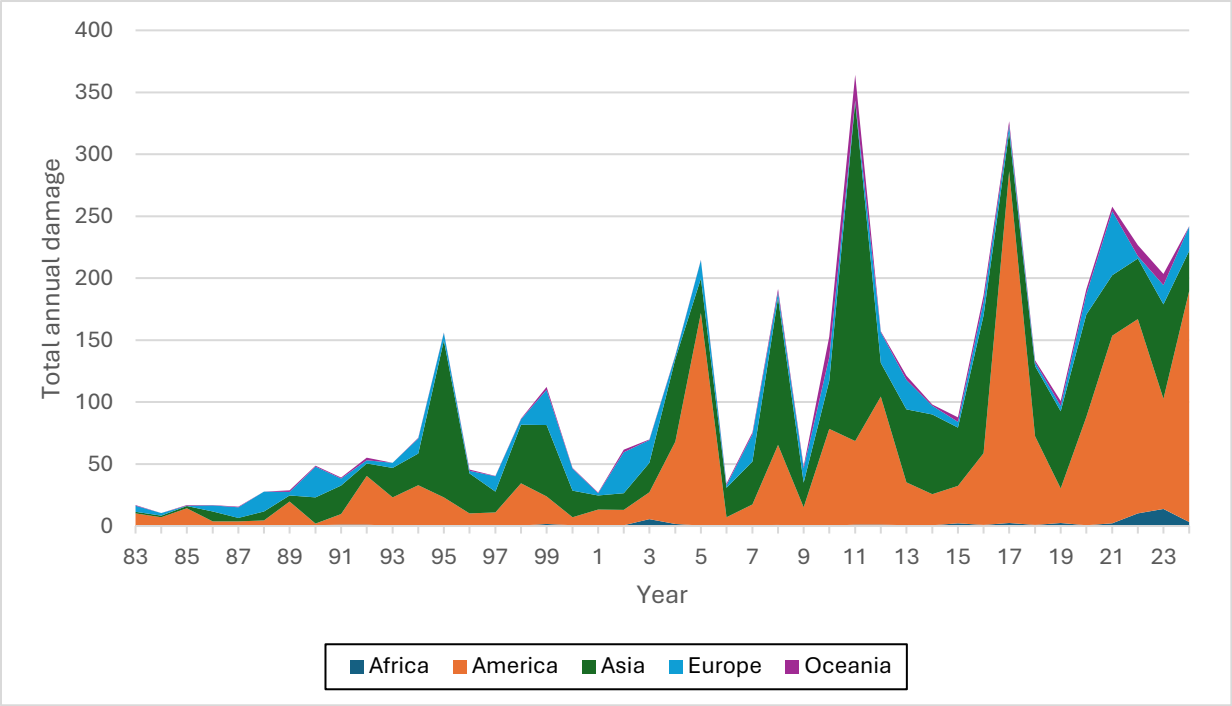


FIGURE 7. Total damage related to natural catastrophes from 1983 to 2024 (in $Bn).
Source: EM-DAT and author calculations

Hazard measurement relies on meteorological, hydrological, and geophysical observations and/or models, often at high spatial resolution.

Exposure measurement maps the location and characteristics of assets, populations, and economic activities within the reach of the relevant hazards. Vulnerability measurement quantifies the sensitivity of exposed elements to a given hazard intensity, typically expressed through damage functions that relate physical impact to asset loss.

The European Union's Sustainable Finance Taxonomy provides the operational hazard classification that frames this cookbook. Annex A of the Do No Significant Harm (DNSH) criteria lists the climate-related physical hazards against which economic activities must be assessed. The list distinguishes chronic hazards such as changing temperature patterns, sea-level rise, increasing variability of precipitation, ocean acidification, and shifting wind patterns, from acute hazards such as heat waves, wildfires, cold spells and frosts, storms, floods, droughts, and glacial lake outbursts.

	Temperature-related	Wind-related	Water-related	Solid mass-related
Chronic	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	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	

TABLE 5. DNSH's Annex A table listing the most widespread climate hazards

As specified in the document "The list of climate-related hazards in this table is non-exhaustive, and constitutes only an indicative list of most widespread hazards that are to be taken into account as a minimum in the climate risk and vulnerability assessment."

This classification reflects a consensus on the categories of hazard. It defines the scope of disclosure obligations for firms subject to EU Taxonomy regulation.

We anchor our analysis in DNSH Annex A to ensure that its methodological guidance is directly applicable to practitioners operating within European prudential and disclosure frameworks of physical hazards relevant to financial and economic activities. Annex A of the Do No Significant Harm (DNSH) criteria enumerates the climate-related physical hazards that economic actors must assess when determining whether their activities are aligned with EU Taxonomy requirements. These hazards are grouped into chronic categories (such as changing temperature patterns, sea-level rise, and shifting precipitation) and acute categories (including cyclones, floods, wildfires, and storms). This cookbook adopts the European Taxonomy's hazard classification as its primary reference frame, ensuring that the methodologies discussed are directly applicable to regulatory compliance and disclosure contexts relevant to European financial institutions.

Physical Climate Risks: Is the issue a compliance one, a financial one, both or neither?

Physical climate risks affect financial institutions through two channels, regulatory compliance and financial exposure. The boundary between them has narrowed over the past decade. Insurers are subject to Solvency II, whose 2025 revision requires climate change scenario analysis within the Own Risk and Solvency Assessment (ORSA). Banks fall under the Basel III standards, transposed in the European Union through the CRR and CRD. Asset managers are governed by the UCITS and AIFMD directives, with disclosure obligations set by SFDR and, in France, by Article 29 of the Energy-Climate Law. Bank-owned asset managers also enter the Basel perimeter through the prudential consolidation of their parent group, although the funds they manage for clients remain outside it.

Supervisors have progressively embedded climate considerations into each of these frameworks, and physical risk assessment has moved from a voluntary exercise to a supervisory expectation and, in several cases, a binding requirement. These binding requirements remain concentrated in the European Union, through Pillar 3 disclosures, the EBA guidelines, the revised Solvency II ORSA and Article 29. The picture is uneven. The Basel Committee kept its 2025 disclosure framework voluntary, climate risk carries no dedicated Pillar 1 capital charge, and several supervisory stress tests rely on voluntary participation.

Physical climate risk remains a financial issue everywhere, since floods, storms and droughts affect collateral values, insurance claims, default probabilities and asset prices whatever the regulatory regime. The compliance channel sets a floor, and the financial channel determines the actual exposure. Both place the same demand on models, which financial institutions expect to be auditable, interpretable and consistent with their existing risk management frameworks (ClimInvest, 2019). The table below summarises the main supervisory expectations that now apply to banks, insurers and asset managers.

COMPLIANCE AREA	REGULATORY OR PRUDENTIAL REQUIREMENT	SOURCE
Climate Risk Governance	Institutions must integrate climate-related risks into their governance, business strategy and risk appetite framework.	ECB, 2020
Integration into Prudential Risk Categories	Climate risks are treated as drivers of existing risk categories (credit, market, operational, liquidity) and must be reflected in internal risk management frameworks.	ECB, 2020; BCBS, 2021a; BCBS, 2021b
Risk Management and Internal Controls	Banks must identify, measure and manage physical and transition risk exposures through internal procedures, adequate data infrastructure and associated controls.	ECB, 2020; ECB, 2022b
ESG Risk Management (EU Banks)	Banks must identify, measure, manage and monitor ESG risks over short, medium and long-term horizons, including a horizon of at least ten years, and prepare plans addressing these risks.	EBA/GL/2025/01; CRD VI Art. 87a
Climate Scenario Analysis	Institutions must use climate scenarios to assess financial impacts, including effects on probability of default (PD), loss given default (LGD), solvency ratios and portfolio values.	ECB, 2022a; ACPR, 2021
Supervisory Climate Stress Tests	Selected institutions take part in supervisor-led exercises and estimate losses under NGFS-based climate scenarios, largely on a static balance sheet. Participation in the ACPR insurance exercises was voluntary.	ECB, 2022a; ACPR, 2021; ACPR, 2023; ACPR, 2024
Climate Exposure Data	Institutions must collect geographic and sectoral data to identify physical and transition risk exposures, geolocate their assets and classify exposures by sector and asset type.	ECB, 2022b; BCBS, 2022

COMPLIANCE AREA	REGULATORY OR PRUDENTIAL REQUIREMENT	SOURCE
Integration into Stress Tests and ICAAP	Banks must incorporate physical and transition climate risks into their stress testing programmes and Internal Capital Adequacy Assessment Process (ICAAP) where those risks are material over relevant time horizons.	BCBS, 2022
Prudential Disclosure of Physical Climate Risks (Pillar 3)	Institutions must publish standardized quantitative disclosures on exposures to acute and chronic physical climate events, using templates set by the EBA ITS. The requirement first applied to large listed institutions and CRR3 extends it proportionately to other institutions.	CRR Art. 449a; Implementing Regulation (EU) 2022/2453
International Climate Disclosure Framework for Banks	The Basel Committee provides qualitative tables and quantitative templates, including exposures subject to physical risk. The framework is voluntary and each jurisdiction decides whether to implement it.	BCBS, 2025
Climate Risk in Insurers' ORSA (EU)	Insurers must assess material climate change risks within their ORSA, including climate change scenario analysis over long-term horizons.	Directive (EU) 2025/2; EIOPA, 2021
International Supervision of Insurers	Supervisors expect insurers to identify climate-related risks within existing risk categories and to extend ORSA scenario analysis beyond typical business planning horizons.	IAIS ICP 16 guidance (revised December 2024)
ESG Ratings and Data Providers	Regulators are recommended to monitor ESG rating and data providers and to require transparency on their methodologies.	IOSCO, 2021a (FR09/2021)
Sustainability Practices in Asset Management	Regulators are recommended to set requirements on sustainability-related disclosures by asset managers and to prevent greenwashing in sustainability-linked financial products.	IOSCO, 2021b (FR08/2021)
Sustainability Risk Disclosure for Investors	Financial market participants, including asset managers and institutional investors, must disclose how they integrate sustainability risks, physical climate risks included, in their investment decisions and in pre-contractual and periodic reporting.	SFDR, Regulation (EU) 2019/2088
Climate Risk Reporting for French Financial Entities	French institutional investors and asset management companies with more than €500 million in assets must publish annually their exposure to physical risks, defined as exposure to the direct physical consequences of climate change, and to transition risks.	LEC Art. 29; Décret n° 2021-663

COMPLIANCE AREA	REGULATORY OR PRUDENTIAL REQUIREMENT	SOURCE
Corporate Disclosure of Physical Climate Risks	Entities must disclose their acute (storms, floods, droughts) and chronic (sea-level rise, reduced water availability) physical climate risks where these could reasonably be expected to affect their prospects. IFRS S2 applies where jurisdictions adopt it, and ESRS E1 is its counterpart for companies within the CSRD scope.	IFRS S2 (ISSB, 2023); ESRS E1 (CSRD)

TABLE 6. Main supervisory requirements on climate risks for banks and insurers

Other insights:

The EU Taxonomy classifies climate-related hazards in Appendix A of the Climate Delegated Act (Regulation (EU) 2021/2139), which crosses the acute and chronic distinction with four hazard families (temperature, wind, water and solid mass). This classification is one node in a wider ecosystem of frameworks. The Task Force on Climate-related Financial Disclosures (TCFD), created by the Financial Stability Board in 2015, brought the acute and chronic distinction into mainstream corporate disclosure. Its recommendations now underpin mandatory reporting in the United Kingdom, in the European Union through the ESRS, and in jurisdictions adopting IFRS S2, the ISSB having taken over TCFD monitoring in 2024. The NGFS scenarios provide the standardized climate pathways used in supervisory stress tests and a common reference for cross-institutional comparison.

Other frameworks serve different purposes. The UNDRR and the International Science Council maintain Hazard Information Profiles, a scientific classification designed for multi-hazard risk governance and aligned with the Sendai Framework for Disaster Risk Reduction (UNDRR, 2025). Catastrophe modelling firms apply proprietary peril classifications calibrated to actuarial and underwriting practice, while the Insurance Development Forum promotes open modelling standards for the insurance sector. Practitioners working across these frameworks therefore face varying terminologies, geographic scopes and severity metrics for nominally identical hazards. A flood may carry a different definition, return period and spatial resolution in the EU Taxonomy, in a catastrophe model and in the Hazard Information Profiles.

This fragmentation is now recognized at the international policy level. The T7 task force on insurance solutions, convened for the 2026 French G7 Presidency, calls for a shared international understanding of physical risks, including a free-access database of natural catastrophe risk indicators and an open-source hazard exposure platform based on CLIMADA and EM-DAT (Think 7, 2026). Until such tools exist, practitioners need to understand the intended use of each framework and the assumptions embedded in its hazard definitions before translating methodologies across regulatory contexts. The stakes extend beyond disclosure, since the widening insurance protection gap documented by the T7 transfers uninsured losses to households, firms, lenders and public budgets (T7, 2026).

1.2 FROM PHYSICAL DAMAGE TO FINANCIAL IMPACT

Physical climate risks only become financially relevant once physical disruptions are translated into economic and monetary consequences. A flood is a hydrological event. A wildfire is a combustion process. A drought is a climatic anomaly. None of these phenomena are financial by themselves, in theory. Financial institutions therefore face a difficult analytical problem: transforming physical observations into estimates of losses, impairments, defaults, repricing effects, or broader macroeconomic consequences.

	BANK	INSURANCE COMPANY	INVESTOR PERSPECTIVE (AO/AM)
TIME HORIZON	Long-term	Short-term	Depends on the strategy
SPACE CONSIDERED	Specific and broad	Specific and broad	Unspecific, very dynamic
DAMAGE TYPE	Loan AND collateral	Pay-off equivalent	Asset value

TABLE 7. Time horizon, space considered and damage type by type of financial actor

This translation process is the central challenge of physical climate risk modelling. Financial actors do not directly observe "financial climate risk". Instead, they observe or estimate chains of causality linking environmental disturbances to financial outcomes. Between the initial hazard and the final financial metric lies a succession of intermediate mechanisms, assumptions, proxies, and modelling choices.

At the broadest level, the transmission chain can be represented as follows:

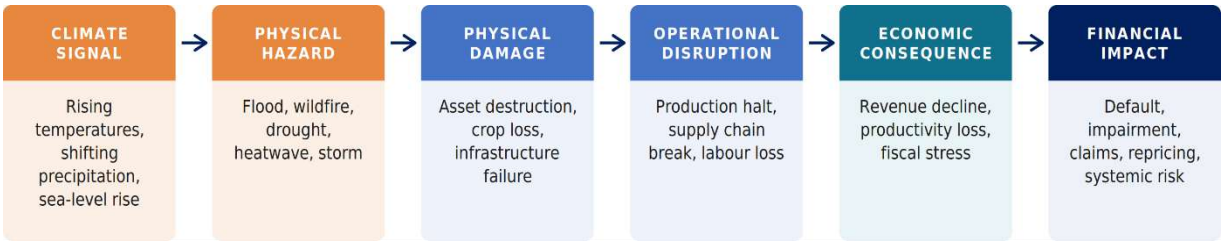


FIGURE 8. Transmission chain from climate signal to financial impact

Each layer introduces uncertainty, simplification, and methodological trade-offs. A cyclone may damage industrial infrastructure, interrupt production, reduce revenues, weaken repayment capacity, increase default probabilities, trigger insurance claims, and eventually affect asset valuations or sovereign spreads.

However, each of these effects depends on assumptions regarding geography, adaptation capacity, leverage, insurance penetration, reconstruction speed, and market reactions.

The same physical event may therefore generate radically different financial consequences depending on the context in which it occurs. A flood affecting an industrial district with high insurance coverage and strong public infrastructure may lead to temporary operational losses followed by rapid reconstruction.

The same flood affecting uninsured agricultural regions or highly leveraged municipalities may generate persistent economic stress and long-term financial deterioration. Physical intensity alone is therefore insufficient to estimate financial materiality.

For this reason, most physical climate risk models rely on layered architectures combining hazard, exposure, vulnerability, and financial translation modules.

Hazard models estimate the probability and intensity of climatic events, for certain events like floods. Hazard models are typically based on open-source methodologies published in academic papers (Wing et al., 2024). Exposure models identify which assets, infrastructures, firms, or populations are located within the affected area (Ehrlich et al., 2018). Vulnerability models estimate the degree of damage generated by a given hazard intensity (Huizinga et al., 2017). Financial modules then attempt to convert these damages into monetary indicators relevant for investors, banks, insurers, corporations, or public authorities.

The financial outputs produced by these models vary considerably depending on the intended use case. Insurance-oriented models frequently estimate insured losses, probable maximum losses, or claims distributions. Banking models often focus on credit deterioration, collateral impairment, probability of default, or loan delinquency. Asset managers may seek to estimate valuation impacts, using a number of different accounting variables, but also abnormal returns, volatility shifts, or Climate Value-at-Risk metrics. Macroeconomic institutions use models aggregating different impact variables at a national or world-level such as productivity losses, GDP contractions, fiscal stress, or sovereign vulnerability.

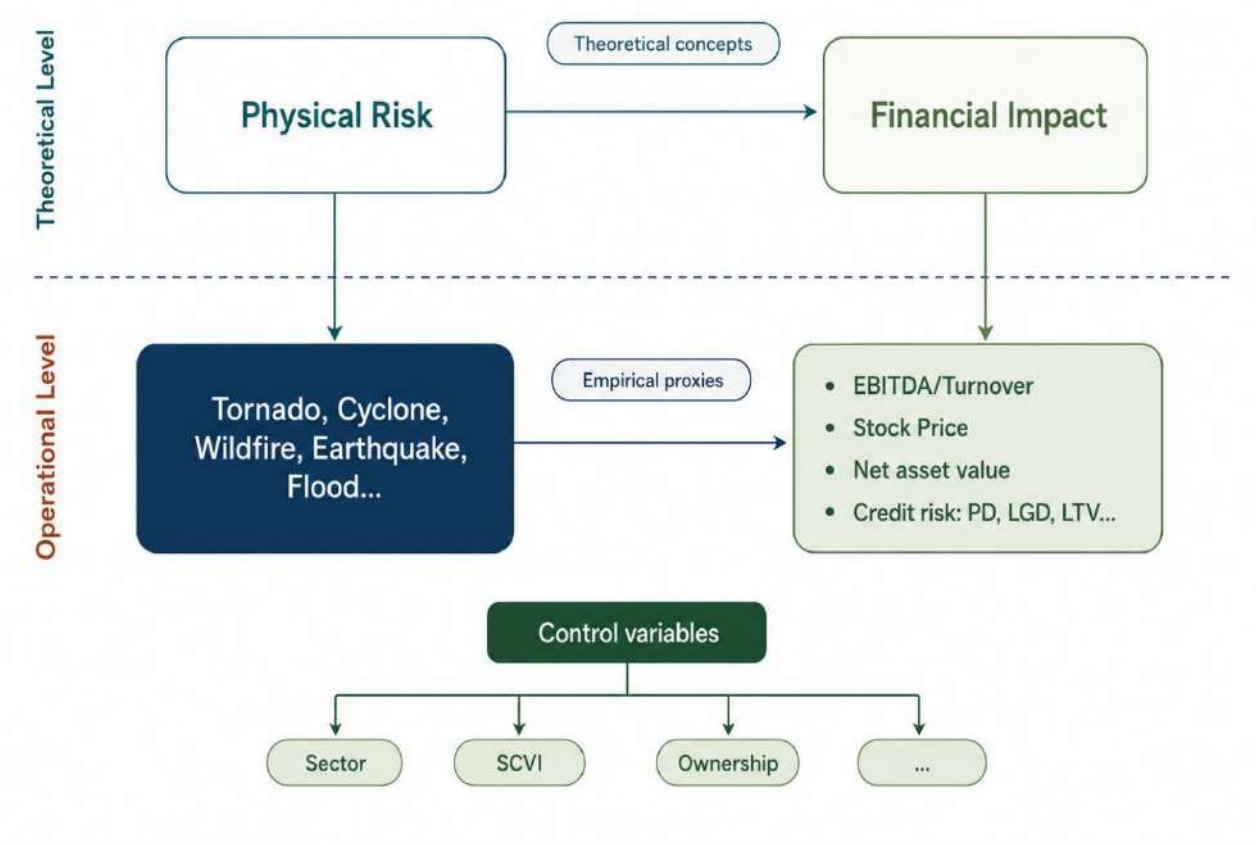


FIGURE 9. The Libby boxes of physical risk financial impact models

Therefore, two models may rely on identical climate scenarios while producing radically different financial outputs because they operationalize different transmission mechanisms between physical hazards and financial variables (UNEP FI, 2023; GARP,2025). The divergence between models is a matter of data quality, sophistication and of course, transparency. This divergence reflects different assumptions concerning causality, temporal dynamics,

geographical coverage, vulnerability considerations, financial propagation mechanisms and it's telling us that this flourishing field will continue to produce more and more insights in the near future as it "boomed" since 2015 (Lin et al., 2023).

This difficulty is reinforced by the **extensive use of proxies throughout the modelling chain**. Physical climate risk cannot usually be measured directly through a single observable variable. Instead, models approximate reality through physics-based indicators. Physical proxies may for example include flood depth, burned area and vegetation or plant density, wet-bulb temperature, precipitation anomalies, drought indices, wind speed or sea-level rise and so on and so forth.

Financial proxies may include insurance claims, EBITDA losses, municipal bond spreads, abnormal stock returns, expected shortfall, sovereign spreads, infrastructure repair costs, or default rates. The use of these proxies is creating an extremely heterogeneous modelling landscape.

A wildfire model based on satellite burned-area observations and insurance losses (Axa, 2024) differs fundamentally from a drought model relying on Palmer Drought Severity Index (PDSI) (Palmer, 1965) agricultural yield reductions and sovereign credit deterioration. The diversity of possible combinations partly explains the rapid proliferation of methodologies currently observed in the market.

An additional difficulty lies in the problem of spatial and temporal resolution. Climate models often operate at scales ranging from tens to hundreds of kilometres and over long temporal horizons extending toward 2050 or 2100. Financial decision-making, however, frequently requires highly granular and short-term information. Mortgage underwriting depends on individual buildings. Corporate lending depends on facility-level exposure. Infrastructure insurance may require parcel-scale precision.

Bridging these incompatible resolutions requires multiple layers of interpolation, downscaling, aggregation, and approximation (Chegwidden et al., 2022; Frank and Chegwidden, 2022). Each transformation introduces uncertainty. A climate scenario produced at continental scale must eventually be translated into assumptions concerning a specific industrial site, logistics hub, agricultural region, or municipal territory. The further models move from physical observation toward financial estimation, the greater the uncertainty generally becomes.

This resolution problem also contributes to the growing strategic importance of geospatial datasets (Colman, 2021). Satellite imagery, remote sensing technologies, cadastral databases, geolocated corporate disclosures, infrastructure maps, and proprietary asset databases increasingly form the backbone of modern physical risk analytics. Control over high-resolution exposure data has therefore become a major competitive advantage for specialized climate analytics firms and financial data providers.

Physical climate risks should not be understood solely as isolated sources of loss. They may also generate broader systemic effects through financial feedback loops. Repeated climatic events may alter insurance affordability, reduce property values, weaken municipal finances, increase sovereign fiscal burdens, or accelerate capital reallocation across regions and sectors. In certain cases, the financial system itself may amplify vulnerability through the withdrawal of insurance coverage, reductions in credit availability, or declining investment in exposed territories (Herweijer & al., 2009; T7, 2026).

Physical climate risk modelling is essentially a translation problem. Models convert evolving biophysical systems into economic and financial representations under conditions of deep uncertainty. Understanding how this translation works is essential to interpret model results, compare methodologies, and assess their relevance for financial decisions, such as whether to lend, insure, adapt or divest.

1.3 FINANCIAL OBJECTS AT RISK

Physical climate risks propagate through financial systems across multiple scales of economic organization, from individual physical objects to the global financial system itself. The nature of the object under analysis strongly influences the modelling framework, the spatial resolution required, the relevant hazards, and the financial metrics ultimately produced.

ANALYSIS LEVEL	FINANCIAL OBJECT AT RISK	TYPICAL EXAMPLES	FINANCIAL OUTPUTS	SPATIAL SCALE
PHYSICAL OBJECT OR GOOD	Single physical item	Inventory, machinery	Replacement cost	Site
ASSET	Income-generating asset	Building, plant	Asset loss, downtime	Asset
INFRASTRUCTURE	Network asset	Road, power grid	Repair cost, service loss	Network
LOAN	Credit exposure	Mortgage, corporate loan	PD, LGD, LTV	Borrower
COMPANY	Operating firm	Manufacturer, utility	Revenue, EBITDA, cash flow	Firm
PORTFOLIO	Aggregated holdings	Loan book, fund	VaR, loss, capital	Multi-region
NATIONAL ECONOMY	National economy	Sovereign debt, public finances	GDP, fiscal balance, spread	National
REGIONAL OR INTERNATIONAL ECONOMIC NETWORK	Cross-border economic network	Supply chain, commodity market	Trade loss, price shock	Regional or International
FINANCIAL SYSTEM	System-wide finance	Banking system, capital markets	Systemic risk, liquidity, solvency	International or Global

TABLE 8. Financial Objects at Risk

At the most granular level lie physical and economic objects. Here, we can consider and include individual goods, inventories, raw materials, industrial components, containers, agricultural products, machinery, or commodities stored within logistical systems. Taken individually, these objects may possess limited financial significance.

However, the value of the smallest economic units becomes significant and more easily accounted for once we aggregate them. They become inventories, collateral structures, assets; they acquire economic and financial value. A flood damaging a single product that only incurs a cost of replacement has little systemic importance when compared with a flood damaging millions of euros of inventory within a logistics platform that will incur an important cost of replacement along with a business interruption which might also provoke a loss of market share.

These objects are therefore progressively transformed into assets. In financial and accounting terms, assets correspond to economic resources capable of generating future value, revenues, productivity, or cash flows. Real estate assets, industrial assets, infrastructure assets, agricultural assets, inventories, and energy assets all depend on physical conditions for their functioning and preservation. Physical climate risks may damage these assets directly through destruction or indirectly through operational disruption, accelerated depreciation, rising maintenance costs, or declining economic productivity.

Infrastructure assets occupy a particularly important position because they support broader economic systems. Roads, railways, ports, airports, energy grids, telecommunications systems, pipelines, and water infrastructure function as connective structures linking production, logistics, labour, and trade. Damage affecting these assets may therefore propagate across entire economic networks. The financial significance of infrastructure risk often exceeds the direct replacement cost of the damaged asset itself because indirect economic losses may become considerably larger than physical repair costs.

Loans constitute another level of financial aggregation. Physical climate risks affect loans indirectly through borrower vulnerability and collateral deterioration. Mortgage loans depend on the long-term value and insurability of underlying properties. Commercial and industrial loans depend on the operational continuity of firms exposed to climatic hazards. Agricultural lending depends heavily on climatic stability and water availability. At this stage, physical climate risks become credit risks because environmental disruptions weaken repayment capacity and increase default probabilities.

At company level, climate risks become embedded within broader organizational and financial structures. Firms may experience direct damage to facilities and inventories, interruptions in production, supply chain disruptions, labour productivity losses, rising insurance costs, or declining revenues linked to climatic disturbances. Public traded firms may additionally face valuation adjustments as investors reassess long-term exposure to vulnerable regions or sectors.

Portfolios introduce another scale of aggregation. Banks, insurers, pension funds, sovereign funds, and asset managers rarely hold isolated assets. Instead, they manage diversified portfolios combining thousands of loans, securities, infrastructures, firms, or real estate exposures distributed across multiple regions and sectors. At this level, diversification effects, concentration risks, geographic clustering, and cross-asset correlations become central analytical challenges. A climatic event affecting one region may simultaneously impact multiple supposedly independent exposures within the same portfolio.

Physical climate risks may also be analysed at national scale. Entire countries may exhibit structural vulnerability depending on geography, infrastructure quality, economic specialization, adaptation capacity, fiscal resilience, and demographic pressures. Climatic disruptions may weaken tax revenues, increase reconstruction expenditures, reduce productivity, destabilize agricultural systems, or deteriorate sovereign creditworthiness. Sovereign-level modelling therefore attempts to capture the macroeconomic consequences of physical climate risks.



FIGURE 10. Geographical and financial scales of analysis

The two analytical scales of physical climate risk assessment: geographical (top) and financial (bottom). Granularity increases from left to right.

International, national, local, different scales can be used for different financial impacts precision. Risk models operate at subregional, regional, continental, international, or global scales. The ambitions of such models at these levels differ. Models with increasing complexity focus on interconnected systems such as global trade networks, commodity flows, energy systems, food systems, and international financial linkages. Localized climatic disruptions may therefore propagate through supply chains and financial markets far beyond the original hazard zone.

At the broadest level lies the financial and economic system itself. Physical climate risks may generate systemic effects when repeated climatic shocks simultaneously affect multiple sectors, regions, institutions, or infrastructures. Insurance withdrawal, declining collateral values, sovereign fragility, infrastructure degradation, migration pressures, and large-scale capital reallocation may collectively contribute to financial instability and structural economic transformation.

This effect can be particularly observable in labour-intensive economies, especially in Asia. However, capital-intensive economies can suffer from repercussions from these events. Physical climate risk analysis is therefore inherently multi-scale. The same climatic event may simultaneously affect objects, assets, loans, companies, portfolios, sovereign entities, and financial systems. Understanding the scale at which a model operates is essential for interpreting its outputs and comparing methodologies across the rapidly expanding landscape of physical climate risk analytics.

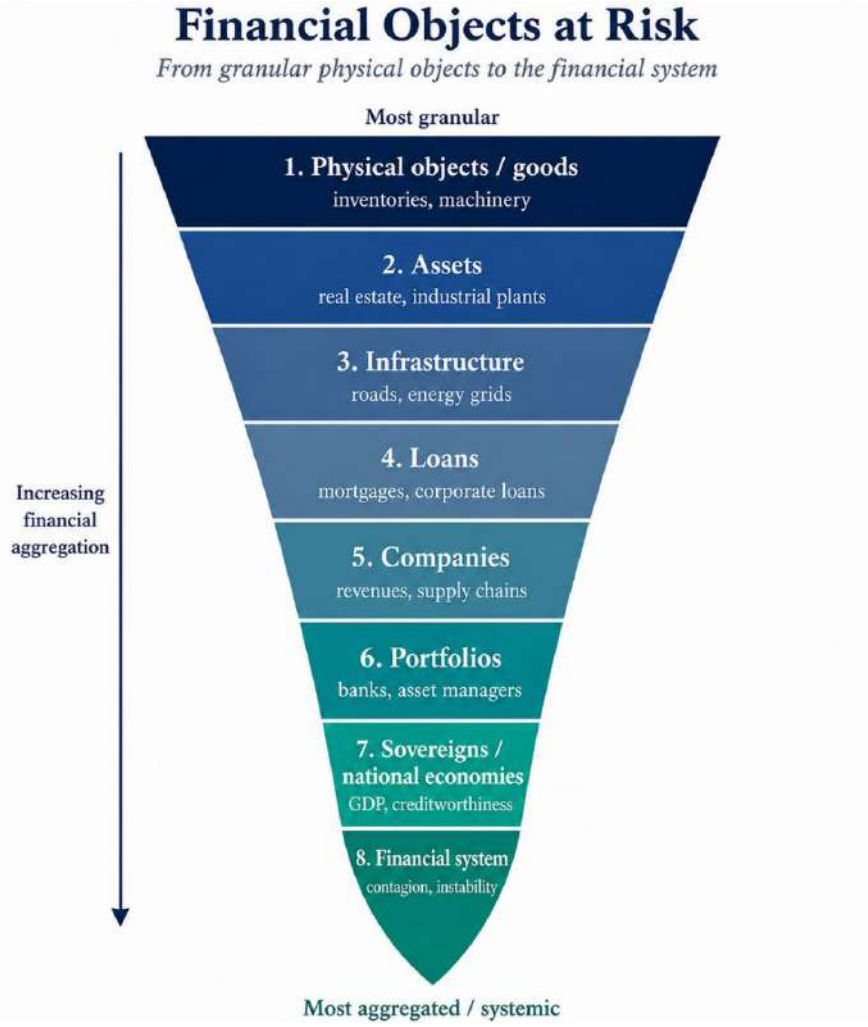


FIGURE 11. Financial objects at risk, from granular physical objects to the financial system

The proliferation of approaches raises a basic question about what these models actually contain. Three points summarize why this question matters.

Opacity. Physical climate risk modelling has become a competitive race (Keenan, 2019). Most commercial models operate as black boxes, and only a few providers disclose how they work.

Cost. Commercial models are expensive, yet many of their components are open source (Condon, 2023). Climate scenarios are published by the IPCC, the NGFS and the IEA. Many damage functions are also freely available from the NGFS and the JRC (see Scientific Climate Ratings).

Risk of model error. Choosing the wrong model has direct consequences. A model may signal no risk when a risk exists (type II error), or signal a risk when none exists (type I error) (Fiedler et al., 2021).

1.4 PHYSICAL RISK AND FINANCIAL IMPACTS: MODELLING ONTOLOGY

1.4.1 WHAT IS A MODEL?

A model means different things in different contexts. For a biologist it is a laboratory animal. For an artist it is a person posing to make a portrait. For AI scientists, it is a set of weights. The SR 26-2 supervisory guidance of the Federal Reserve defines the model as a complex quantitative method, system, or approach that applies statistical, economic, or financial theories to process data into quantitative estimates. For Diaconescu et al. (2025) it is a formal representation of a system (and its environment) that is used to study the system and issue predictions about its behaviour. Regardless of the specific definition, three properties are important. First, a model is necessarily a selective and simplified representation of reality. It retains some features of the object of interest while neglecting others. Second, model results depend on the inputs, including assumptions, data and expert judgments. Third, every model is built for a given purpose and must be validated with this purpose in mind.

1.4.2 THE SPECIFICITIES OF PHYSICAL CLIMATE RISK MODELS

As we discussed above, a physical risk model typically describes a chain of mechanisms from the climate signal to the financial impact. A physical risk model therefore combines several distinct model components, relying on different scientific methods, data, and assumptions.

Physical risks are particularly difficult to model for several reasons (Bolton et al., 2020). First, physical risk models have very long-time horizons. Over these horizons both the climate system is non-stationary, evolving due to natural climate variability and the anthropogenic climate change, and our economic system also evolves in a multitude of unpredictable ways. This non-stationarity implies that physical risk models cannot be based purely on historical data. They need to incorporate future long-term projections of the Earth climate and the economic system. However, the future evolution of these systems is subject to deep uncertainty, which cannot always be quantified in probabilistic terms. Fundamental parameters of the evolution of the climate system such as the climate sensitivity are not known precisely. Moreover, the dynamics of the climate system is subject to structural changes known as tipping points. The ability of climate science to predict the occurrence of tipping points and their impact is limited. Economic dynamics on the other hand is determined by human decisions and geopolitical circumstances, which may be even harder to predict. Interconnectedness and systemic features can further amplify the impact of physical climate risks making the task of quantifying them even harder. However, not all aspects of physical climate risk are equally unpredictable. Many relationships between warming and physical hazards can be studied rigorously and quantified with meaningful uncertainty ranges. Deeper non-quantifiable uncertainties often emerge when assessing mitigation and adaptation scenarios and systemic financial consequences.

1.4.3 MODEL RISK IN PHYSICAL CLIMATE RISK ASSESSMENT

Financial supervisory guidance defines model risk as the potential for adverse financial consequences which may result from decisions made based on model output. Model risk is influenced by model's inherent risk (assumptions, complexity, quality of inputs); its exposure (size of affected portfolios); its purpose (what the model is used for); and its use (for example, understanding of model limitations). Model risk is not the same as uncertainty in model outputs.

However, model risk may arise if this uncertainty is not reported correctly. Model risk is managed through robust model development, model validation and sound model governance; it may be mitigated through limits on model use, monitoring model performance, revising models over time and supplementing models with other analysis.

This sound advice can be applied almost unchanged in any quantitative modelling setting. However, the non-stationary nature of physical climate risks makes it difficult if not impossible to use the back-testing for monitoring model performance. Indeed, since the climate and economic dynamics are non-stationary and can exhibit structural changes, good past performance of a model is not a guarantee of its future reliability.

If back-testing is impossible, how to recognize a good physical risk model?

First of all, while it may be difficult to back-test the full modelling chain, it is usually possible to back-test the individual components, such as, for example, the damage functions, or the company-level financial impacts. Second, models can be validated by benchmarking them against appropriate scientific literature. Assumptions and data sources used by the model should also be analysed and compared to state of the art.

Following this model risk guidance principles ensures that the models are sound and used appropriately. However, even sound and properly used physical risk models cannot fully describe the possible futures that matter for the decision maker.

Diaconescu et al. (2025) discuss how quantitative models can be used in the context of so-called hyper-complex systems, which are not fully predictable. These are precarious complex systems undergoing constant evolution under the effect of novel and unpredictable phenomena, such as the Earth climate system. They identify three modelling pitfalls which can invalidate model predictions for such systems.

The first one is causal rest, which refers to excluding any phenomenon relevant for the system in question from the model. In physical climate-risk modelling, these might include changing social behaviours, infrastructure failures, political responses, migration, financial network effects, or interactions between environmental systems. Some omissions are deliberate and justified. Others reflect incomplete scientific knowledge or mechanisms that cannot currently be quantified.

The second one is quantophrenia or undue reliance on quantitative representations where such representations are not justified or cannot be obtained. A particularly interesting example is the tendency to reduce complex climate-related financial risks to a single monetary loss estimate or risk score. Such indicators may be useful for communication and comparison, but their precision may conceal substantial uncertainty about data, assumptions and omitted mechanisms. Numerical results should therefore not be confused with comprehensive knowledge of the risks under consideration.

The third one is misleading objectivity when formal modelling approaches are applied to analyse subjective phenomena which may be perceived through different viewpoints.

1.4.4 WHAT TO EXPECT FROM PHYSICAL RISK MODELS AND HOW TO USE THEM RESPONSIBLY?

In the face of deep uncertainty and inability to formulate precise, long-term predictions it is important to understand what physical risk models can, and cannot do:

- Such models can help us identify exposures and vulnerabilities. They cannot, however, identify all possible future risk sources.
- Models can help us explore the consequences of chosen scenarios to which they cannot, however, assign specific probabilities.

- Models allow us to understand important risk transmission mechanisms, but it is hopeless to try to represent all interactions and feedback of an economic system in a model.
- We can use models to compare different mitigation and adaptation strategies, but it is impossible to determine a uniquely optimal strategy among an infinity of possible decisions.
- Models can be used to assess resilience of a system to specific adverse events, but cannot guarantee resilience to future risks, whose nature is not yet known.
- We can obtain conditional estimates of financial loss in given scenarios, but cannot reliably forecast the full loss distribution.

Drawing on both complex systems and climate risk literature, one can formulate the following best practices for physical risk models:

1. Use multiple scenarios and modelling approaches. Different scenarios can reveal distinct vulnerabilities. Comparing results from different models can help identify conclusions that are robust to modelling assumptions and others that are highly sensitive.
2. Complement quantitative models with qualitative assessments. Expert judgement, case studies, historical experience, and knowledge of institutional and behavioural responses can help identify risks insufficiently represented by quantitative models.
3. Focus on robustness and resilience rather than relying exclusively on optimality. When future probabilities cannot be established reliably, a strategy that performs reasonably well across several plausible futures may be preferable to one that is optimal under a particular assumed model.
4. Make risk assessment an iterative process. Models should be updated as scientific knowledge, observed climate impacts, economic exposures, and adaptation strategies evolve, perhaps using Bayesian approaches (see below). Decisions should themselves retain some flexibility to respond to new information.
5. Quantify uncertainty but be transparent about the assumptions used in uncertainty quantification and about the fact that this uncertainty is not probabilistic. For example, if ten different risk models have been used to obtain ten different estimates of a future loss value, the distribution of the ten numbers gives an idea of uncertainty but cannot be interpreted as the probabilistic distribution of future loss values.

1.5 THE PURPOSE OF THIS COOKBOOK AND ITS INTENDED USERS

The rapid development of physical climate risk analytics has generated a fragmented ecosystem of models. Financial institutions are now confronted with a rapidly expanding market composed of catastrophe models, climate scenarios, geospatial datasets, vulnerability functions, financial translation tools, and proprietary analytics platforms. While this expansion is positive, as it may reflect an increasingly important awareness of climate-related financial risks, it also creates substantial difficulties for practitioners attempting to compare methodologies, interpret outputs and to select appropriate modelling approaches.

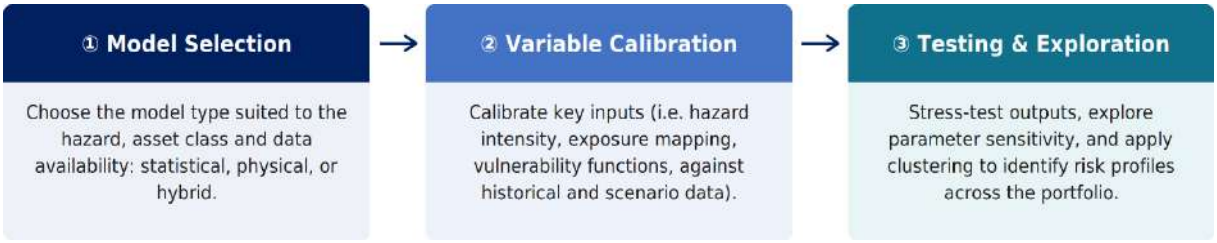


FIGURE 12. Three-step approach to modelling: selection, calibration, testing and exploration

The purpose of this cookbook is therefore not to propose a single "best" model for assessing physical climate risks. Such an objective would be neither realistic nor analytically sound given the diversity of financial objects and use cases involved. Instead, we provide a structured framework for understanding how physical climate risk models operate, how they translate physical hazards into financial consequences, and how their methodological assumptions shape their outputs.

More specifically, the objective is to clarify the architecture of the emerging market for physical climate financial impact modelling. The cookbook aims to identify the main modelling families currently used across the financial sector, examine the assumptions embedded within these models, and compare the different ways through which physical hazards are transformed into economic and financial metrics. Particular attention is devoted to proxies, transmission channels, spatial resolution, financial translation mechanisms, uncertainty propagation, and methodological trade-offs.

This document also aims to reduce the opacity surrounding climate analytics markets. Many methodologies currently used by financial actors rely on partially proprietary datasets, black-box assumptions, or complex modelling chains that remain difficult to interpret from outside specialized technical communities. One objective of this cookbook is therefore to provide financial practitioners with conceptual tools enabling them to better understand what models measure, what they approximate, and what their outputs truly represent.

The cookbook is intended for decision-makers and modellers in insurance companies as well as for anyone interested in climate risk financial impacts modelling. Insurers and reinsurers are historically among the most advanced actors in physical risk modelling because their business models are directly exposed to catastrophic losses. Acute hazards such as floods, cyclones, storms, hail, and wildfires already generate major underwriting and claims impacts. Climate change increases the strategic importance of understanding how hazard distributions, exposure accumulation, and vulnerability dynamics evolve over time. Insurance actors therefore require increasingly sophisticated approaches capable of combining geospatial precision, catastrophe modelling, and financial loss estimation.

The banking industry constitutes another central audience for this work. Physical climate risks increasingly affect credit risk, collateral valuation, operational resilience, and regulatory stress testing exercises. Mortgage portfolios, commercial real estate lending, agricultural finance, infrastructure finance, and corporate lending may all be affected by climatic disruptions. Banks must therefore understand how physical hazards may influence borrower

solvency, asset values, and regional economic activity across different temporal horizons. The growing integration of climate risks into prudential supervision further reinforces the need for robust analytical frameworks.

This cookbook is also intended for asset owners and asset managers. Pension funds, sovereign wealth funds, investment funds, and institutional investors are increasingly expected to evaluate the long-term resilience of their portfolios to climate-related physical risks. These actors face particularly complex challenges because they must aggregate exposures across multiple sectors, geographies, asset classes, and temporal horizons. Physical climate risks may affect portfolio performance through infrastructure damage, declining productivity, supply chain disruptions, repricing effects, sovereign deterioration, or broader macroeconomic instability (see section 2.4.1 for more information on climate physical risks propagation).

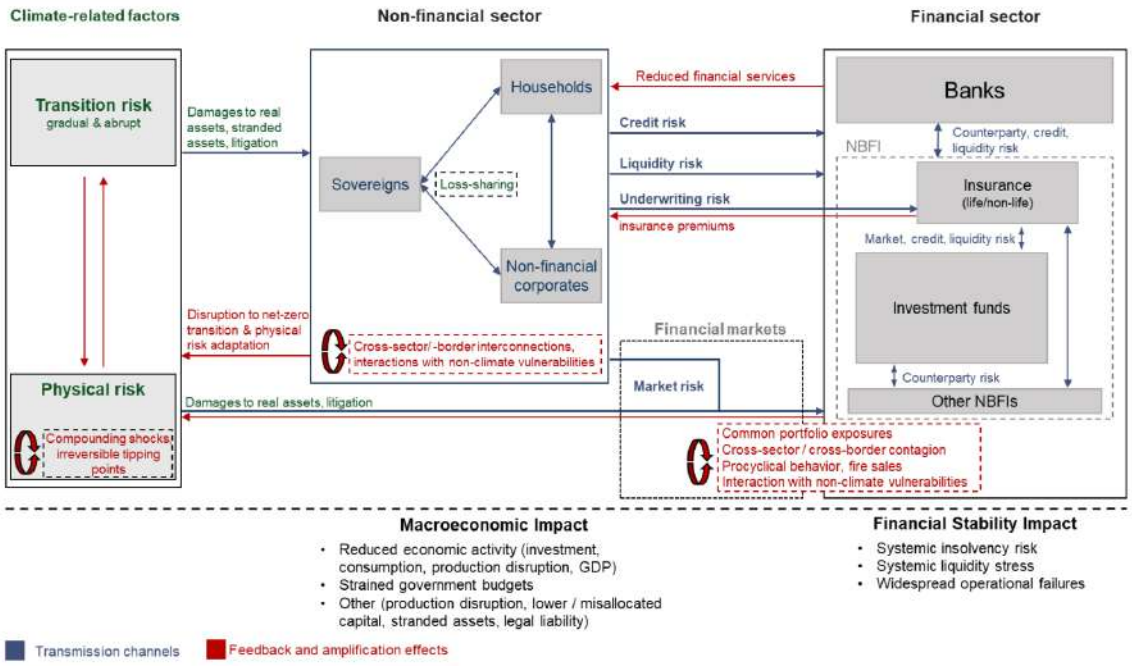


FIGURE 13. Framework for the assessment of climate-related vulnerabilities

Source: FSB, 2025

More broadly, the document is for financial institutions seeking to better understand the rapidly evolving landscape of physical climate risk modelling. This includes regulators, supervisors, public financial institutions, development banks, consultants, risk officers, sustainability teams, quantitative analysts, and strategic decision-makers. Different users in the financial sector may require different levels of granularity, precision, explainability, or computational complexity depending on their operational needs and regulatory obligations.

This cookbook does not assume that physical climate risks can be reduced to a single universal metric or solved through a purely technical exercise. Physical climate risk modelling involves uncertainty, incomplete data, competing assumptions, spatial mismatches, and evolving scientific knowledge. Models are simplifications of reality. In the next sections, we will provide several ideas and clues on how to address those gaps.

As the financial institutions are the first line of defence against chaos and global collapse of the economy, our objective is to provide an operational analytical framework capable of helping practitioners navigate the growing complexity of physical climate financial impact modelling. Understanding how models are built, what assumptions they rely upon, which financial objects they target, and where their limitations lie will be a core competency of tomorrow's financial systems.

1.6 RESEARCH QUESTIONS

1.6.1 HOW TO PREDICT FINANCIAL IMPACTS FROM PHYSICAL RISKS?

Financial impacts from physical risks cannot be predicted perfectly. Any model of financial loss produces an estimate that differs from the realized value. The prediction error is written

$$e = l - \hat{l}$$

Where e is the prediction error, l the realised loss, $\hat{l}$ is the loss estimated by the model.

This error has two sources. The first is the irreducible randomness of the physical system, which no model can remove. The second is reducible uncertainty, which comes from incomplete data, errors in parameter estimation, or an inadequate model structure.

Under standard statistical assumptions, the expected squared error of a predictor decomposes into three terms.

$$\mathbb{E}\left[(l - \hat{l})^2\right] = Bias(\hat{l})^2 + Var(\hat{l}) + \sigma^2$$

- $\mathbb{E}\left[(l - \hat{l})^2\right]$ is the mean squared error, the average squared gap between realised and estimated losses
- $Bias(\hat{l})^2$ is the systematic gap between the average estimate and the true expected loss
- $Var(\hat{l})$ is the variance of the estimate, its sensitivity to the particular sample used to fit the model
- σ^2 is the variance of the irreducible noise, the part of the loss that no model can predict

This decomposition assumes that the true process generating losses is known, at least in principle. In practice, this process is unknown. A further source of error, model uncertainty, comes from the choice of the model itself, from omitted physical or economic mechanisms, and from unknowns that cannot be quantified. This component lies outside the formula above and is usually assessed by comparing several models or through expert judgement.

A model can be biased (it systematically underestimates or overestimates losses), imprecise (its variance is large), or both. Bias and variance are linked. Flexible models with few structural assumptions tend to reduce bias and increase variance. Strong structural assumptions or regularization tend to reduce variance at the cost of higher bias. More data reduces variance and allows the use of more flexible models. Selecting a modelling strategy therefore means finding an acceptable balance between bias and variance, given the data available.

Financial impacts can nonetheless be approximated and measured. The general prediction problem consists in estimating a function f such that

$$l = f(H, E, V) + \varepsilon$$

- l is the financial loss
- f is the function linking the three risk components to the loss
- H is a hazard descriptor (e.g. flood depth, wind speed, temperature anomaly)
- E is an exposure variable (e.g. asset value, insured sum, notional)
- V is a vulnerability function that converts physical intensity into a fractional loss
- ε is the residual term, the part of the loss the model does not explain

Different modelling choices correspond to different assumptions about the form of f and the distribution of ε.

Parametric regression, damage functions calibrated from engineering surveys, and machine-learning methods are all particular implementations of this general structure. They share the same aim, which is to make the relationship between physical hazard and financial loss tractable, even if never exact.

Models can be linear or nonlinear. In a linear specification, the expected loss is a weighted sum of predictors.

$$\hat{l} = \mathbb{E}[l|x] = \beta_0 + \beta_1 x_1 + \cdots + \beta_p x_p$$

- With $\hat{l}$ the estimated loss
- $\mathbb{E}[l|x]$ is the expected loss given the observed predictors
- $x_1, \ldots, x_p$ are observable covariates (hazard intensity, asset characteristics, geographical features), β_0 is the intercept and $\beta_1, \ldots, \beta_p$ are coefficients estimated from historical data
- p is the number of predictors

Linear models are interpretable, computationally simple, and well suited to datasets of moderate size. The relationship between physical damage and financial loss, however, is often strongly nonlinear. Damage functions from engineering studies can show threshold effects (damage that starts only above a given intensity), saturation at high intensities, and asymmetric tails. Nonlinear models, such as generalized additive models, tree-based ensembles and deep neural networks, aim to capture these features at the cost of greater complexity and lower interpretability.

The appropriate degree of nonlinearity therefore depends on the quality and volume of the data, the intended use of the model, and the interpretability requirements of the regulatory or institutional context.

1.6.2 HOW TO PROVIDE AN ACCURATE FINANCIAL ASSESSMENT IN CASE OF A NEW PHYSICAL RISK? THE BAYESIAN APPROACH TO PHYSICAL RISKS

One of the central difficulties of physical climate risk modelling lies in the treatment of rare, emerging, or previously unobserved events. Traditional statistical approaches rely heavily on historical observations in order to estimate probabilities, calibrate loss distributions, and infer future financial impacts. This creates a major challenge for climate-related physical risks because the most damaging future events may have no true historical equivalent within available datasets.

Extreme floods, unprecedented heatwaves, compound climatic events, or rapidly evolving wildfire regimes may therefore exceed the statistical boundaries of historical observations. In such contexts, purely empirical approaches rapidly encounter limitations. Historical data alone may become insufficient to estimate the probability or severity of future losses with acceptable confidence.

This problem is particularly important for prudential and actuarial applications because rare events often generate disproportionate financial consequences. Financial institutions are not only concerned with average losses, but also with tail risks, catastrophic scenarios, and low-frequency high-severity events capable of destabilizing portfolios or institutions.

Bayesian approaches provide an important conceptual framework for addressing these difficulties. Unlike purely frequentist approaches, Bayesian inference does not rely exclusively on direct historical observations. Instead, it combines existing knowledge, expert judgement, collective information, and observed data in order to progressively update probability distributions as new information becomes available.

In practical terms, Bayesian reasoning allows institutions to start from an initial probabilistic belief, commonly referred to as a prior distribution, and then revise this belief using additional evidence. This approach is particularly useful for physical climate risks because fully localized historical datasets are often incomplete, fragmented, or statistically insufficient for robust inference.

In actuarial sciences, large-scale datasets collected at national or sectoral levels may be used to construct baseline priors representing expected loss distributions, hazard frequencies, or vulnerability relationships. Local observations are then used to estimate deviations relative to these broader statistical baselines.

For example, a national flood loss distribution may provide an initial estimate of expected damage across a country. Yet local geographical characteristics such as river topology, urban density, soil permeability, infrastructure quality, or adaptation measures may significantly alter the actual exposure profile of a specific territory. Bayesian updating allows these local signals to progressively refine the broader prior assumptions.

This approach is particularly valuable when dealing with new or poorly documented risks. A region experiencing unprecedented wildfire activity may possess limited historical claims data. However, expert knowledge, engineering assessments, geospatial analysis, climate projections, and observations from comparable regions may still provide valuable information capable of informing probabilistic assessments.

Bayesian approaches therefore expand the notion of usable data beyond strictly observed historical events. Information may originate from multiple heterogeneous sources including scientific expertise, engineering studies, catastrophe models, collective insurance datasets, climate simulations, remote sensing observations, or expert elicitation processes. The objective is not to eliminate uncertainty, but to formalize it within coherent probabilistic structures.

This characteristic is particularly important for climate change because the underlying system itself is evolving. Historical frequencies may progressively lose predictive power as climatic conditions drift away from past norms. Bayesian reasoning provides a flexible framework capable of integrating new information continuously as hazard patterns evolve over time.

In physical climate risk modelling, Bayesian approaches may therefore serve several functions simultaneously. They may improve estimation under data scarcity, formalize expert judgement, integrate heterogeneous information sources, quantify uncertainty more explicitly, and facilitate adaptive updating as new climatic observations emerge.

However, Bayesian approaches also introduce important methodological challenges. Priors may strongly influence results when observational data remains limited. Expert judgement may contain implicit biases. Aggregated datasets may fail to capture local specificities. Different institutions may adopt different prior assumptions, leading to divergent outputs even under similar climatic conditions.

The challenge is therefore not simply technical. It also concerns governance, transparency, explainability, and model interpretability. Financial institutions and regulators increasingly require methodologies capable of justifying how assumptions are constructed, how priors are calibrated, and how uncertainty propagates throughout the modelling chain.

In practice, Bayesian approaches illustrate a key lesson of physical climate risk modelling: accurate financial assessment cannot rely exclusively on historical observations. It requires combining empirical evidence, scientific understanding, collective knowledge, and evolving expertise within probabilistic frameworks capable of adapting to a changing climate system.

Bayesian inference rests on a single identity. Let θ denote the parameter of interest, for instance the expected annual loss rate or the tail probability of a catastrophic event and let x denote the observed data. Bayes' theorem states:

$$p(\theta \mid x) = \frac{p(x \mid \theta)\, p(\theta)}{p(x)}, \qquad p(x) = \int p\,(x \mid \theta)\, p(\theta)\, d\theta$$

where $p(\theta)$ is the prior distribution encoding existing knowledge before the data are observed, $p(x \mid \theta)$ is the likelihood, i.e. the probability of observing x given θ,

$p(x)$ is the marginal likelihood acting as a normalizing constant,

$p(\theta \mid x)$ is the posterior distribution, which combines prior knowledge and new evidence into an updated distribution over θ.

Since $p(x)$ does not depend on θ, the posterior is often written $p(\theta \mid x) \propto p(x \mid \theta)\, p(\theta)$. The posterior is proportional to the likelihood times the prior.

Once the posterior is obtained, prediction for a future loss $\tilde{L}$ is obtained by integrating over the full posterior uncertainty rather than conditioning on a single point estimate:

$$p(\tilde{L} \mid x) = \int p(\tilde{L} \mid \theta) \cdot p(\theta \mid x)\, d\theta$$

This predictive distribution $p(\tilde{L} \mid x)$ is the quantity of operational interest for financial institutions. It yields a central estimate of future losses but also an important characterisation of residual uncertainty. In other words, it will be essential for computing Value-at-Risk, Expected Shortfall, or probability of ruin measures under data-scarce conditions.

In the context of a relatively known physical climate risk, the prior $p(\theta)$ may be constructed from national or sectoral loss databases, catastrophe model outputs, or expert elicitation.

The likelihood $p(x \mid \theta)$ is calibrated on locally available observations, which may be sparse. As local data accumulate over time, successive updates progressively shift the posterior toward the data, while retaining the stabilising influence of the prior in regions where evidence remains thin.

This "updating" feature makes Bayesian inference well suited to evolving hazard regimes, where new climatic observations can be incorporated without discarding the full historical modelling structure. But it is one way among many to approach risks.

2 THE RECIPES – FAMILIES OF PHYSICAL RISK FINANCIAL IMPACT MODELS

In this part, we'll dive into the recipes for building financial impact models. We will see how financial impacts can be assessed in a very direct manner, on the field via damage-based models, how the impact can threaten the banking industrial via probability-of-default-based models, and how the impact can threaten an investment portfolio via accounting-based models. We will also explore a field relatively new that aims to assess the physical risks financial impacts on supply chains and the "indirect effects" a climate hazard can have throughout the planet.

Each part is composed of an explanation about how the model works and an example use case.

AN OVERVIEW

Part 2 of this cookbook describes 5 groups of physical climate risk models, each with its own mathematical foundations, empirical grounding, and financial application.

The progression is sequential: each section builds on the outputs of the preceding ones. In order to better understand how a model can be chosen, you need to ask yourself, what is the model aiming at?

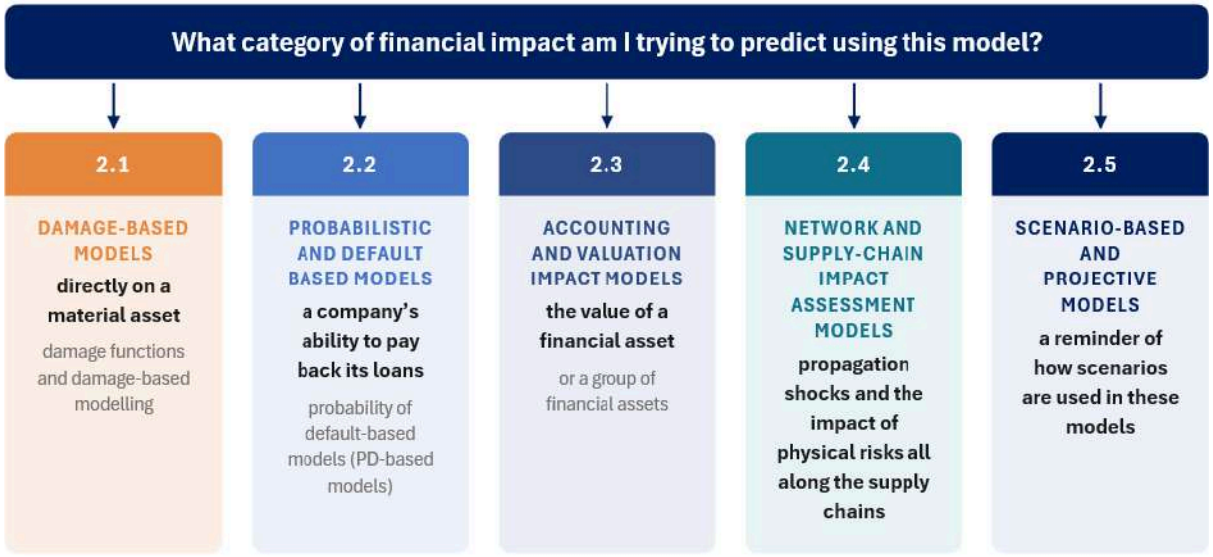


FIGURE 14. The Families of Models

If the model aims to predict the impact of any physical risk (based on historical data or physics-based data, simulations…) **directly on a material asset** such as a building, in the most direct way, you will use damage functions and damage-based modelling.

If the model aims to compute and understand the impact of a physical risk on **a company's ability to pay back its loans** and to survive financially to the current climate trends, you will need to consider probability of default-based models (PD-based models).

If the model aims to compute the impact of a physical risk on **the value of a financial asset**, or a group of financial assets, the most useful ones will be in 2.3.

Finally, if the model aims to take into account propagation shocks and the impact of physical risks all along the supply chains, then the section you will need is 2.4

Section 2.5 will only serve as a reminder of how scenarios are used in these models.

SECTION	MODEL FAMILY	OUTPUT	USEFUL FOR
2.1 DAMAGE-BASED MODELS			
2.1.1	Damage functions (cat model)	Asset-level loss	All downstream financial models via AAL and PML
2.1.2	Actuarial models	Compound loss distribution; Return levels; tail index; TVaR	EP curves; capital requirements; All models needing tail characterization
2.1.3	Loss and frequency curves (OEP/AEP)	PML at return period T; TVaR	Stress tests; reinsurance pricing; Solvency II SCR
2.1.4	Real Effects	Mortality and morbidity; labour productivity loss; displacement; health and social costs	Double materiality (CSRD); adaptation planning; social risk assessment; public policy evaluation
2.2 PROBABILISTIC AND DEFAULT BASED MODELS			
2.2.1	Physical Risks and direct probability of default	PD; LGD; credit VaR	IFRS 9 provisioning; Basel capital; bond pricing
2.2.2	Probit, logit and survival models	Conditional PD given hazard covariates; hazard rate; time to default	PD calibration; IFRS 9 lifetime PD; early warning systems; mortgage and SME scoring
2.2.3	Structural Credit Models	Distance to default; asset value shock; credit spread change	Corporate bond and CDS pricing; equity-credit linkage; climate-adjusted ratings
2.2.4	Portfolio Credit Risk and Climate Credit VaR	Portfolio loss distribution; expected and unexpected loss; climate credit VaR; concentration	Economic capital; ICAAP and Pillar 2; portfolio allocation; supervisory stress tests
2.2.5	From Credit Risk to Liquidity Risk	Deposit outflows; collateral haircuts; LCR and NSFR deterioration; funding gap	Liquidity stress tests; ILAAP; ALM; contingency funding plans

SECTION	MODEL FAMILY	OUTPUT	USEFUL FOR
2.3 ACCOUNTING AND VALUATION IMPACT MODELS			
2.3.1	Price Loss, Turnover Loss, Capex and Opex	Composite hazard score S(x)	Portfolio screening; ESG disclosure; credit pre-screening
2.3.2 (accounting)	Enterprise value impact	EBITDA impact; ΔEV; ratio deterioration	Credit analysis; equity valuation; regulatory disclosure
2.3.3	Financial ratios	RoA; RoE; RoI; Debt/Equity; LTV; interest coverage	Covenant monitoring; credit rating; lending decisions; collateral valuation
2.3.4	Macro and Micro models	Firm and sector output loss; regional value-added loss; input-output multipliers	Sector exposure analysis; regional lending; bridge between firm and macro levels
2.4 NETWORK AND SUPPLY-CHAIN IMPACT ASSESSMENT MODELS			
2.4.1	Physical Risks and macroeconomic frameworks	Network loss multiplier; RAR; systemic VaR	Macro stress tests; supervisory risk assessment
2.4.2	Shock propagation	Indirect loss amplification factor	Network models; supply chain; systemic risk
2.4.3	Supply chain impact	ΔGDP = f(temperature, scenario)	Sovereign risk; macro stress tests; IAM integration
2.4.4	Systemic Risk	CoVaR; SRISK; contagion and fire-sale losses	Macroprudential policy; financial stability reports; G-SIB supervision
2.4.5	Lagged effects	Impulse response functions; recovery time; persistent GDP and productivity loss	Long-horizon stress tests; sovereign risk; reconstruction finance

SECTION	MODEL FAMILY	OUTPUT	USEFUL FOR
2.5 SCENARIO-BASED AND PROJECTIVE MODELS			
2.5.1	SSP-RCP projections	Projected hazard intensity and frequency by scenario and horizon; climate change multipliers	Defining the climate pathway that feeds every downstream model; long-horizon physical risk assessment
2.5.2	Stress-test analysis	Scenario range; stress VaR; OEP/AEP/PML	All applications requiring forward-looking risk quantification
2.5.3	Cat Models	Stochastic event loss table; AAL; OEP/AEP curves; PML at given return periods	Insurance and reinsurance pricing, capital and reserving; portfolio tail risk
2.5.4	ML-Based Models	Predicted hazard, damage or loss at asset level; risk classification scores	Large portfolios with rich observed data; downscaling; filling gaps in exposure or vulnerability data
2.5.5	Predictive limits	Uncertainty ranges; model spread; identification of non-stationarity and tipping-point risks	Interpreting and communicating the confidence of any model output; model risk management

TABLE 9. Model families, outputs and uses

The practitioner who has read Part 2 in full will possess the foundations to better understand how the models he's using work. For each part we will explore the general mechanism, the mathematical foundation on which it is built, the data it can use and the different techniques through which it can be processed, we will then finish the part with a field review and a model card showing what kind of model has been used on the field.

The standardized model representation cards are used as examples of models across families. We will also consider key limitations to avoid the most common misapplications.

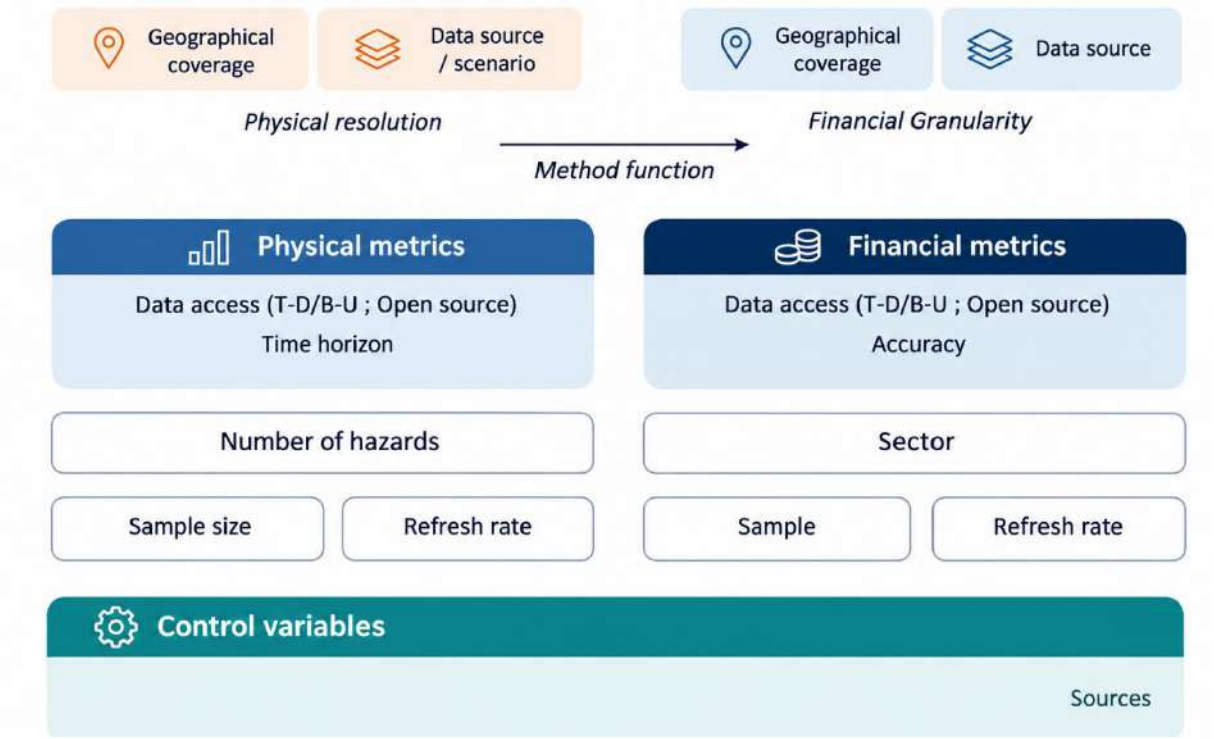


FIGURE 15. Model card template used to describe models

There's one more thing to consider. These models are comparable because they all study the same object, a financial impact, using the same kind of input, the physical risk. You probably already have the data for most of these proxies.

In empirical research, you ask a question and look at how real-world observations answer it. All these models ask the same question. What are the financial impacts of physical risks? We've already covered the context, meaning how varied and complex physical risks and financial impacts can be. That leaves three things that can change in a model. The data, the technique, or both (Bloomfield et al., 2016).

If the data changes but the technique stays the same, it's called generalization. It happens when the geographical coverage changes, or the data source, the physical or financial metrics, the number of hazards, the sector or the sample size.

If the data stays the same but the technique changes, it's called robustness. Say you first ran your model with Monte Carlo simulations. Later you realize the same data is enough to train a neural network that predicts damages years ahead, and you have a good reason to think it fits better.

When both the data and the technique change, it's called triangulation. The question stays the same. Take an OEM's assets in Germany. Over the years you gather new types of data, starting with financial information, then topography, precipitation and maybe even real-time sonar data (to obtain a 3D real-time digital twin of the assets). The data and the techniques both get challenged along the way until the predictions are right. Triangulation is a good way to assess the financial impacts of physical risks on a portfolio when you want to see how much the models differ.

If you'd prefer to pick the best model based on the data you already have instead of your objectives, parts 3 and 4 are for you. They focus on an old question Atkinson asked in 1970: "Under the assumption that one particular model is the true one, is there evidence that one of the other models fits the data better?"

2.1 DAMAGE-BASED MODELS

Damage-based models are the very basis on which are built modern models. They estimate the direct physical degradation of assets exposed to climate hazards. They represent the most granular and physically detailed family of physical climate risk models, and are widely used in the insurance and reinsurance industries to quantify losses affecting buildings, industrial facilities, infrastructure, and other material assets. Rather than inferring losses indirectly from macroeconomic observations, these models simulate the physical destruction process itself, following a four-stage chain:

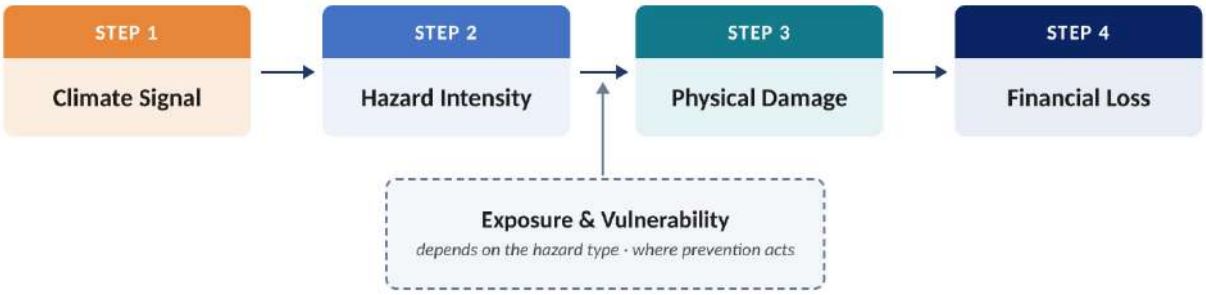


FIGURE 16. The modelling chain of Physical Risks Financial Impact models

Hazard intensity alone does not determine damage. The same flood depth produces very different losses depending on the exposure of assets and on their vulnerability, which covers construction type, flood defences, soil conditions and maintenance. Vulnerability is specific to each type of climate signal, since a building resistant to wind may be fragile under water. It is also the step where prevention acts. Adaptation measures reduce damage without changing the hazard, which is why the damage function sits at the center of damage-based models (see section 3.2.4)

The mechanism at the core of this process is the damage function, which links hazard intensity to the degree of degradation suffered by an asset. In its most compact form, the logic is: probability and intensity of the hazard × estimated level of physical damage × economic value of the exposed asset = expected loss.

A flood reaching 1.5 metres in a logistics warehouse might yield a 40% destruction rate for electrical systems and stored equipment, which then translates into repair costs, replacement costs, operational disruption, and insurance claims.

An asset has lost value. In such an event, there is always damage. But if the hazard risk has been well-anticipated, there might not be a significant effect on the company's ability to be solvent nor to be profitable. In other words, its Return on Investments (RoI), Return on Equity (RoE) and Return on Assets (RoA), may not be impacted.

The objective of these models is therefore to transform a physical climate signal into a quantified direct financial consequence, step by step.

2.1.1 DAMAGE FUNCTIONS AND PRECISE SIMULATION MODELLING

What is a damage function?

A damage function (also called a vulnerability function or loss function) is a relationship representing a measure of hazard intensity at a given location to a predicted level of physical degradation for an asset of a given type. It is the core translation mechanism of every damage-based model. Without it, there is no path from the physical world to the financial one.

Fragility functions are a related but distinct tool: they express the probability of reaching or exceeding a defined damage threshold, whereas damage functions express the expected damage ratio itself (Ramanathan, 2017)

The damage function outputs a Mean Damage Ratio (MDR): the expected fraction of an asset's replacement value that is destroyed or rendered non-functional by a hazard event of a given intensity. In other words, more specifically in catastrophe models, it is considered as the expected repair cost of a structure divided by its replacement value, as a function of hazard intensity h (wind speed, flood depth, ground acceleration).

$$MDR(h) = E\left[\frac{Cost\ of\ repair}{Cost\ of\ replacement} \mid H = h\right]$$

A flood depth of 1.2 metres applied to a reinforced concrete light-industrial building might yield an MDR of 0.18, meaning that on average 18% of the building's replacement cost value must be spent on repair or replacement. Multiplied across the insurer's or investor's full exposure, the MDR is the mechanism by which a physical signal becomes a financial number.

One of the most famous tools using damage function that is open source is CLIMADA (Aznar-Siguan & Bresch, 2019).

Verisk's anatomy of a damage function

What can we say about damage functions, in simple terms? Verisk (Ramanathan, 2017) has written a great blog post that makes several good points about damage functions in general. It's a good read and its main idea is simple. Please consider the following:

- Damage functions are not always independent of other model components. Model calibration against real loss experience introduces an implicit dependence between hazard and vulnerability, even when each component is developed separately.
- Damage functions for a given construction type and peril are not similar across countries and regions. Local building code enforcement, claims settlement practices, replacement cost versus actual cash value policy terms, and socioeconomic factors (maintenance quality, risk tolerance) all shift the curve for nominally identical structures.
- Engineering analysis is not sufficient to develop damage functions. Monetary loss also depends on claims adjustment behaviour and preparedness or response actions, none of which are captured by structural response alone.
- Damage functions are not deterministic in nature; once the wind speed is known, it is impossible to calculate damage severity with certainty. At fixed hazard intensity there is a full distribution of possible damage ratios, not a single value. Part of this comes from measurement error itself: wind speed instruments carry roughly 5 to 10 percent error, which can propagate into 50 to 100 percent error in damage estimation.
- All structures within a model hazard footprint do not suffer damage at damaging wind speeds or ground shaking. Construction quality, age, and maintenance vary within a portfolio even at identical hazard exposure.
- Changes in building codes will not have an immediate effect on vulnerability. There is a lag between code adoption and actual field enforcement, so the effect propagates gradually through the building stock rather than at the moment the code changes.

THE GENERAL FORM AND THE FOUR CANONICAL FUNCTIONAL FORMS OF A DAMAGE FUNCTION

A damage function can be logistic or score-based, for example:

$$MDR = f(proximity, fuel\ load, construction\ class)$$

Consider MDR, a Mean Damage Ratio where $MDR \in [0, 1]$. It is the proportion of the object or the asset that is

$\boldsymbol{I}$ is the hazard intensity metric (e.g. flood depth in metres, wind speed in m/s, ground motion in PGA g if you wish to consider earthquakes). Keep in mind, it's usually a physics-based metric, but it can also be an index (PDSI, actuarial weather index...)

$\boldsymbol{v}$ is the vulnerability class of the exposed asset (construction type, age, occupancy, design standard)

$\boldsymbol{\theta 1}, \ldots, \boldsymbol{\theta k}$ are the other fitted parameters that will be useful to determine (calibrated from engineering data, claims records, laboratory experiments, post-event surveys)

1. Logistic (sigmoidal) curve

$$MDR_I = \frac{1}{1 + e^{\left(-\frac{I-\mu}{s}\right)}}$$

where $\boldsymbol{\mu}$ is the median damage intensity (50% MDR point) and s is the scale parameter controlling the steepness of the curve.

It is most widely used for wind and earthquake because the damage process has a natural S-shape. This way, the damage becomes negligible below a threshold, but there is a rapid increase through what is called the "vulnerability range" before it attains a saturation at total loss.

More generally, any sigmoid curve can be used. For tropical cyclones, Emanuel (2011) proposes

$$f = \frac{{v_n}^3}{1 + {v_n}^3}$$

with $v_n = \frac{max(V - V_{thresh},\ 0)}{V_{half} - V_{thresh}}$

- $\boldsymbol{f}$ is the fraction of the property value destroyed, equivalent to the MDR.
- $\boldsymbol{V}$ is the maximum sustained wind speed at the asset location, which serves as the intensity I.
- $\boldsymbol{V_{thresh}}$ is the wind speed below which no damage occurs.
- $\boldsymbol{V_{half}}$ is the wind speed at which half of the property value is destroyed, in other terms where f = 0.5.
- $\boldsymbol{v_n}$ is the normalised excess wind speed. It equals 0 at V_{thresh} and 1 at V_{half}.

The exponent 3 reflects the assumption that damage scales with the cube of wind speed, since wind power grows as V^3.

Writing $v_n^3 = e^{3 \ln v_n}$ shows that this is form 1 applied to a transformed intensity:

$$x = \ln(V - V_{thresh})$$

With:

- $\boldsymbol{\mu} = \boldsymbol{ln}(\boldsymbol{V_{half}} - \boldsymbol{V_{thresh}})$ and $s = \frac{1}{3}$ is the scale parameter, the inverse of the cubic exponent,
- $\boldsymbol{x}$ is the logarithm of the wind speed in excess of the threshold. It replaces I in form 1.
- $\boldsymbol{\mu}$ is the logarithm of the excess wind speed at half damage.

The Emanuel function is therefore a log-logistic curve. It is the logistic counterpart of form 4, which is a probit in $\ln(I)$

Unlike form 1, it has a strict threshold, with zero damage for $V \leq V_{\text{thresh}}$

Forms 1 and 4 are both binomial generalized linear models (GLMs) when fitted on the share of damaged assets or on claim counts. The link function, which maps the probability of damage onto a linear predictor, is the logit in I for form 1 and the probit in ln I for form 4. When fitted on a continuous MDR, the two forms become non-linear regressions or fractional logit regressions.

2. Power law

$$\boldsymbol{MDR_I = a \times I^{\alpha}}$$

Where I is still the hazard intensity (e.g. wind speed),

$\boldsymbol{\alpha}$ governs the speed at which the damage increases with the intensity,

and a is a proportionality coefficient, fixing the damages global level.

Power law models are often used for flood depth–damage relationships, with α below 1. Damage rises quickly in the first tens of centimetres of water and then flattens (Huizinga et al., 2017). Claims data is a good reality check. On more than 2 million US flood insurance claims, standard curves overestimated damage at shallow depths (Wing et al., 2020). The power law also breaks down at high depths, where losses approach total, so it needs capping at 1.

In practice, the function should therefore be capped at 1:

$$\boldsymbol{MDR(I) = min(1, a \times I^{\alpha})}$$

The exponent can also reflect the energy released by the hazard, which drives the damage:

wind power $P = ½ \rho$

A v^3 scales with the cube of wind speed, and hail impact energy scales with the fourth power of hailstone diameter. A power-law MDR can thus mimic these relationships in a simplified way, with a acting as a vulnerability factor.

3. Piecewise linear

$$\boldsymbol{MDR_I = \min(1, \max(0, \alpha_0 + \alpha_1 \times I))}$$

Here I is the hazard intensity, α_0 the intercept (baseline damage at zero intensity) and α_1 the slope (damage increase per unit of intensity). The result is clipped to the [0, 1] range.

This is used in rapid-assessment contexts and in early cat model implementations. It's useful because it helps to calibrate the model when using sparse data. In a way, it lacks the physical realism of the sigmoidal form but it is analytically tractable and transparent. Other explanatory variables can be added (bounded linear

regression), but the [0, 1] clipping prevents a generalization as a GLM; the logistic form (form 1) handles additional covariates more naturally within a GLM framework.

4. Lognormal fragility curve

$$MDR_I = \Phi\left(\frac{\ln I - \lambda}{\zeta}\right)$$

Where Φ is the standard normal Cumulative Distribution Function (CDF)

$$\Phi(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{-t^2/2}\,dt$$

λ is the log-median intensity

ζ is the log-standard deviation.

This is standard in earthquake engineering (FEMA P-58, OpenSees) and increasingly adopted for wind and wildfire (Mitchell-Wallace et al., 2017) but also for earthquakes (Dowrick & Rhoades, 1995). The lognormal form is theoretically motivated by the multiplicative nature of structural failure mechanisms.

This form can also be embedded in a generalized linear model (GLM) by using a probit link (instead of a logit link), which allows additional covariates such as building codes or building age to be included.

How to turn MDR into an expected financial loss

$$EL = MDR(I, v) \times TIV$$

where:

EL is the expected loss for a single event at a given intensity level

TIV is the Total Insured Value (replacement cost) of the exposed asset. This is a simplification here as TIV should be split between buildings and contents, and it excludes business interruption (BI) (see section 2.3). The MDR applies to building values, while contents (including machinery) follow different curves reflecting on-site installations, with threshold effects. BI depends on the peril and is triggered above one or more severity (MDR) thresholds; the indemnity then depends on the time needed to resume activity, usually subject to a deductible and a limit (see the section on secondary uncertainty)

For a portfolio of *n* assets with exposures TIV_i and damage ratios $MDR_i(I, v)$:

$$Portfolio\ EL = \sum_i MDR_i(I, v) \times TIV_i$$

Aggregated across all events in the stochastic event catalogue, weighted by their annual rates λe, this gives the Average Annual Loss (AAL), which is the entry point to the actuarial metrics (see section 2.1.2).

How Damage Functions Are Built

Damage functions are "calibrated" against observed loss data. Four main sources are used, each with distinct strengths and limitations:

- **Engineering studies and laboratory experiments**: These are controlled physical tests on building components, structures, stored goods or material samples at known hazard intensities. They provide

precise, repeatable damage ratios for specific construction types, but at high cost and only for narrow asset classes. They are widely used in wind engineering (wind tunnel tests on roofing systems and cladding panels) and earthquake engineering (shake table tests).

- **Post-event field surveys**: These involve the systematic collection of damage data after actual events, with observed damage levels classified against measured or modelled hazard intensities. They are the most directly relevant calibration source, but they only cover perils and geographies where major events have occurred and where damage was systematically recorded. Several standard frameworks support this work. The ATC-20 protocols structure rapid post-earthquake building inspections in the field, the Copernicus Emergency Management Service (EU) produces damage assessments from satellite imagery, and the HAZUS methodology (FEMA, US) incorporates such observations into standardized loss estimates.
- **Insurer and reinsurer claim records**: They form by far the largest dataset for calibrating vulnerability functions, particularly for flood, wind and hail. Moody's RMS and Verisk (formerly AIR Worldwide) have calibrated their vulnerability models against tens of millions of individual claims across several decades and geographies. These databases are commercially sensitive and not publicly available, which gives established catastrophe model vendors a structural advantage.
- **Simulation-based and physics-based methods**: These include computational fluid dynamics (CFD) for wind and water, finite element analysis (FEA) for structural response, and fire behaviour models (FlamMap, PHOENIX) for wildfire. They generate synthetic damage data where empirical observations are sparse, particularly for rare events beyond historical experience or for new asset types that existing calibration datasets do not cover.

HAZARD	INTENSITY METRIC	FUNCTIONAL FORM	KEY VARIABLES	CALIBRATION SOURCE
River flood	Depth (m) is *d*, velocity (m/s), duration (hours)	$MDR(d) = a \cdot d^2 + b \cdot d + c$ So, this is our basic damage function. *d* is defined but *a, b* and *c* are coefficients built on historical data and they depend on the construction type.	Construction type (Reinforced concrete (RC), masonry, timber), basement presence, elevation, floor layout.	JBA, FLEMOps, USGS, EU HANZE data base
Tropical cyclone / wind	Wind speed (m/s or mph), gusts, pressure	$MDR(v) = \frac{1}{1 + exp\left(-\frac{v - \mu}{s}\right)}$	Roof type, connection standard, glazing ratio, age, wind pressure design code by geographic zone; see also the Emanuel (2011) form	RMS, AIR, claims from Katrina, Harvey, Irma, Ian
Earthquake	Peak Ground Acceleration (PGA, g) or Modified	$MDR(PGA) = \Phi\left(\frac{\ln PGA - \lambda}{\zeta}\right)$	Seismic design code vintage, structural system	FEMA P-58, HAZUS, SYNER-G, claims from Northridge,

	Mercalli Intensity		(RC frame, URM, wood), soil class	Tohoku, Christchurch
Wildfire	Flame zone proximity (m), ember density, fire behaviour index	Logistic or score-based $$MDR = f\left(\begin{matrix}proximity, fuel\ load, \\ construction\ class\end{matrix}\right)$$	Proximity to fire perimeter, roof material, vents, defensible space, vegetation type	Insurance Council of Australia, USFS, claims from Camp Fire, Black Saturday
Heat stress (chronic)	Cooling degree days, WBGT (Wet Bulb Globe Temperature) (°C), exceedance hours above threshold	$$\Delta Productivity = \gamma \times (T - T_{ref}) + \delta \times \max(0, T - T_{threshold})^2$$	Outdoor vs. indoor work, sector (construction, agriculture, logistics), cooling infrastructure	Burke et al. 2015, ILO heat stress studies, Kjellstrom 2016
Drought	PDSI (Palmer Drought Severity Index), SPI (Standardized Precipitation Index), soil moisture anomaly (σ), crop-specific indices	$$\frac{\Delta Y}{Y} = f(PDSI, growth\ stage, variety)$$	Irrigation coverage, crop variety, soil type, growth stage at onset of stress	FAO GAEZ, IIASA EPIC, historical crop insurance claims

TABLE 10. Examples of Hazard-Specific Damage Functions

Secondary Uncertainty: from a Mean Damage value to a Loss Distribution

In its basic form, a damage function returns the expected damage ratio for a given hazard intensity and vulnerability class. It does not capture the scatter around that expectation.

In practice, two assets of the same construction type exposed to the same flood depth will sustain different levels of damage. This is due to idiosyncratic differences in state of maintenance, contents placement, drainage effectiveness and micro-topography. This scatter is known as secondary uncertainty, and it materially affects portfolio loss distributions, particularly in the tail.

Secondary uncertainty is parameterized by the coefficient of variation of loss given event, βe.

In the ELT (Event Loss Table) framework used by Moody's RMS, AIR Worldwide (Verisk), and Oasis LMF, each event is characterized not only by its mean loss $\underline{L}e$ but also by βe. For a given event, the loss is then drawn from a distribution with mean $\underline{L}e$ and standard deviation $\beta e \times \underline{L}e$.

Secondary uncertainty also has consequences at the portfolio level. It widens the loss distribution of each event, which thickens the tails of both the OEP (Occurrence Exceedance Probability) and AEP (Aggregate Exceedance Probability) curves. The size of this effect depends on how strongly losses are correlated across assets. For independent assets, the scatter largely averages out across a large portfolio. When damage is correlated, for instance among residential properties in the same floodplain that share drainage conditions and building practices, the scatter persists. The tail of the portfolio loss distribution can then be significantly heavier than a mean-damage view suggests, even when no single property dominates the total exposure.

Advanced Simulation Techniques

The simplest damage functions are deterministic: a given hazard intensity maps to a fixed damage ratio. More sophisticated implementations rely on simulation-based approaches to propagate uncertainty through the damage estimation process. Three techniques dominate in the insurance and engineering industries: Monte Carlo, Finite Element Analysis and Dynamic Hazard Modelling

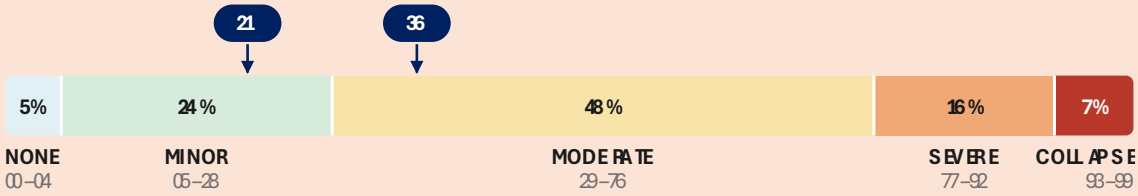
Monte Carlo simulations are repeated stochastic sampling across the joint distribution of hazard intensity, asset characteristics, and secondary uncertainty parameters to generate a full empirical distribution of possible loss outcomes for a given event or scenario. For a large portfolio, a Monte Carlo run may generate 100,000 or more simulated loss realizations for a single event, providing precise estimates of the full loss distribution including the tail. This is the workhorse method for cat model loss calculations (Grossi & Kunreuther, 2005).

Example Box: A single-family residential structure and a hurricane

Suppose a single-family residential structure is exposed to hurricane wind speeds of 100 (≈ 44.7 m/s, ≈ 161 km/h), and the model defines five discrete damage states: none, minor, moderate, severe, and collapse.

A damage function first assigns a probability to each state at that wind speed, for example 5 percent none, 24 percent minor, 48 percent moderate, 16 percent severe, and 7 percent collapse.

These probabilities are converted into cumulative ranges, and a random number generator draws a value within that range for each simulated trial.



A draw of 36 lands in the moderate band;

A draw of 21 lands in minor.

Repeating this draw many times, typically thousands of iterations for a single event, produces an empirical distribution of damage outcomes that approximates the underlying probability distribution implied by the damage function.

Hence the often continuous, hazard and vulnerability distributions are approximated through repeated random sampling).

DAMAGE STATE	DAMAGE STATE PROBABILITY	CUMULATIVE PROBABILITY	RANDOM NUMBER LOWER BOUND	RANDOM NUMBER UPPER BOUND
None	0.05	0.05	00	04
Minor	0.24	0.29	05	28
Moderate	0.48	0.77	29	76
Severe	0.16	0.93	77	92
Collapse	0.07	1.00	93	99

TABLE 11. Damage state probabilities and repeated random sampling
Source: Adapted from Grossi & Kunreuther (2005, p. 77)

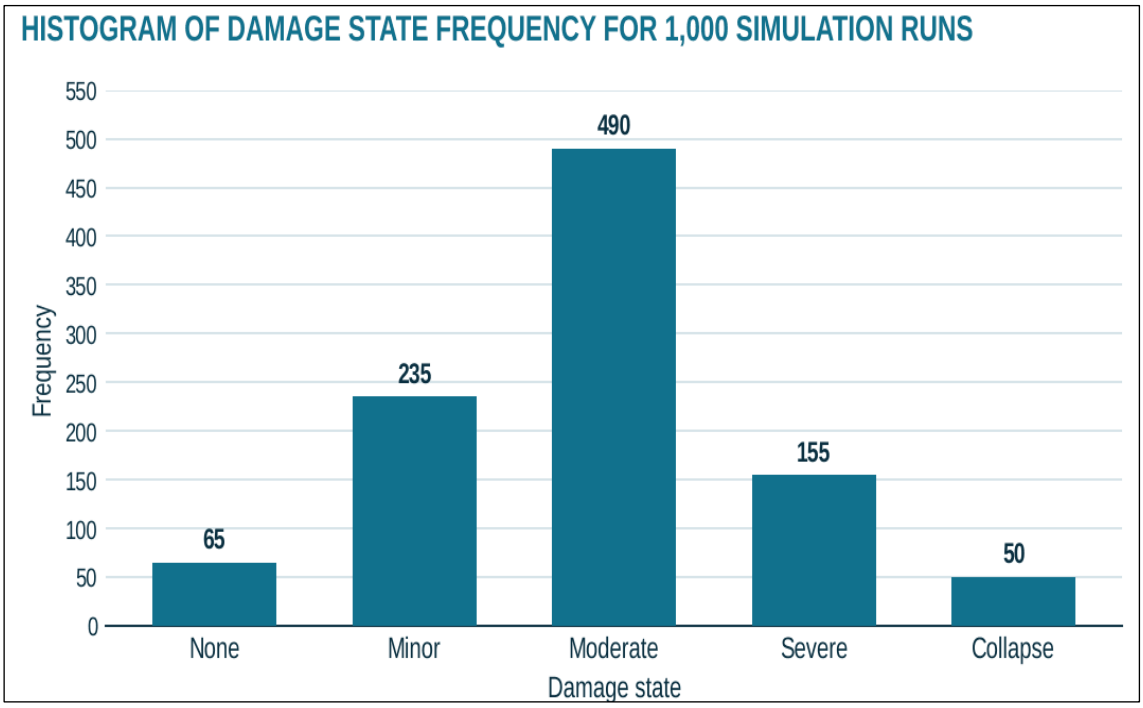


FIGURE 17. Histogram of damage state frequency for 1000 simulation runs

Other advanced simulation techniques can include:

- **Finite element analysis (FEA):** structural engineering simulation that models the response of a building or infrastructure component to applied loads (wind pressure, seismic ground motion, hydrostatic pressure) at the level of individual structural members. FEA produces damage estimates at the sub-component level (column failure, wall cracking, roof uplift), which can be aggregated to asset-level MDRs. Used principally to calibrate the parameters of the analytical fragility curves (λ, ζ in our lognormal form) and to extend damage estimates to construction types or design standards with limited empirical claims data.
- **Hydrodynamic modelling:** computational simulation of water flow dynamics using the shallow water equations or full Navier-Stokes equations (Hervouet, 2007), producing spatially explicit flood depth, velocity, and inundation duration maps at metre-scale resolution. Tools such as HEC-RAS (USACE), MIKE FLOOD (DHI), and Delft3D are standard in engineering consultancy and are increasingly used as the hazard input layer for insurance-grade vulnerability assessments. The computational cost can be high as hydrodynamic simulation of a river catchment at 1m resolution may require days of run time on a high-performance computing cluster.
- **Fire behaviour models:** physically-based simulation of combustion dynamics, fuel consumption, and thermal radiation to estimate ignition probability and structural exposure for wildfire. Tools such as FlamMap, FARSITE, and PHOENIX RCM generate spatially explicit estimates of fire intensity metrics (flame length, radiant heat flux) under given weather and fuel conditions, which are then passed to structural vulnerability functions to estimate damage ratios for the built environment.

Where can I find data?

The European joint research center (JRC) worked on flood risk and drought impact functions (Huizinga et al., 2017). In 2017, they published damage functions for each continent using gathered data from various authors.

DAMAGE CLASS	FLOOD DEPTH [M]	DAMAGE FUNCTION					
		EUROPE	NORTH AMERICA	CENTRAL & SOUTH AMERICA	ASIA	AFRICA	OCEANIA
Residential buildings	0	0,00	0,20	0,00	0,00	0,00	0,00
	0,5	0,25	0,44	0,49	0,33	0,22	0,48
	1	0,40	0,58	0,71	0,49	0,38	0,64
	1,5	0,50	0,68	0,84	0,62	0,53	0,71
	2	0,60	0,78	0,95	0,72	0,64	0,79
	3	0,75	0,85	0,98	0,87	0,82	0,93
	4	0,85	0,92	1,00	0,93	0,90	0,97
	5	0,95	0,96	1,00	0,98	0,96	0,98
	6	1,00	1,00	1,00	1,00	1,00	1,00

TABLE 12. JRC damage functions linking flood depth to damage by region

It is an example of an open-source solution for you, the reader, who would want to better understand how damage-based models operate, and how to use them.

But there are other resources that render the making of a basic climate risk financial impact model costless. However, such an approach could make your assessment very flawed and uncertain.

TOOL	DEVELOPER	PERILS	WHAT IT PROVIDES	FINANCIAL OUTPUT
Damage curves				
JRC global flood depth-damage functions	European Commission JRC	River and coastal flood	Depth-damage curves by continent and occupancy type (residential, commercial, industrial, agriculture, infrastructure), with maximum damage values per country (Huizinga et al., 2017)	Damage ratio, damage in €/m^2
Hazus	FEMA (United States)	Flood, hurricane wind, earthquake, tsunami	Free GIS-based software with a large library of depth-damage and fragility curves by building type, calibrated on US data	Direct building and contents losses, business interruption
Impact engines				
CLIMADA	ETH Zurich	Tropical cyclone, flood, storm surge, wildfire, heat, drought	Python platform combining hazard, exposure and impact functions in a probabilistic setting, with a module for the cost-benefit of adaptation measures (Aznar-Siguan & Bresch, 2019)	Asset-level damages, average annual loss, exceedance probability curves
Delft-FIAT	Deltares	Flood mainly	Python tool computing object-level damages from inundation maps and user-defined damage curves. FloodAdapt uses it for adaptation planning	Damage per building, expected annual damage
physrisk	OS-Climate	Flood, heat, wind, drought, wildfire	Python library linking open hazard indicators to vulnerability models for assets held in financial portfolios	Damage and disruption distributions per asset, portfolio aggregates
CAPRA	World Bank, IDB, UNDRR	Earthquake, hurricane, flood, tsunami, landslide	Probabilistic risk platform built for public authorities in developing countries	Average annual loss, probable maximum loss
Catastrophe model frameworks				

TOOL	DEVELOPER	PERILS	WHAT IT PROVIDES	FINANCIAL OUTPUT
Oasis LMF	Oasis Loss Modelling Framework	Any peril, depending on the model plugged in	Open platform for running catastrophe models. Hazard, vulnerability and financial modules are separate, so policy terms and reinsurance structures can be applied	Ground-up and insured losses, event and year loss tables, OEP and AEP curves
OpenQuake	GEM Foundation	Earthquake (non-climatic)	Open risk engine, often used as a reference for probabilistic loss calculation	Average annual loss, loss exceedance curves
Hazard inputs				
Aqueduct Floods	World Resources Institute	River and coastal flood	Global flood hazard maps under climate scenarios, with an analyzer estimating damages and the benefits of flood protection	Expected annual urban damage, affected GDP
Explore2	INRAE, Météo-France and partners	River flows, floods, low flows (France)	Hydrological projections to 2100 for French catchments	None. Explore2 is a hazard input only

TABLE 13. Examples of open-source tools for damage-based modelling

A step-by-step application: How to go from Hazard Intensity to Asset-Level Financial Loss

Step 1 – Identify the exposed asset and its vulnerability class.

Use the asset register: construction type (reinforced concrete, unreinforced masonry, wood frame, steel frame), occupancy (residential, commercial, industrial, critical infrastructure), age and design standard (pre-code, code-compliant, post-code upgrade), and location (geocoordinates with precision ≤100m). Consider any other information that might be important to add here.

Step 2 – Select and apply the hazard intensity metric.

For the event being modelled (or the stochastic event being sampled from the cat model catalogue), identify the hazard intensity at the asset location. This comes from the hazard module of the cat model (for insured losses) or from a deterministic hazard footprint (for scenario analysis). Express as the appropriate metric: flood depth in metres, wind speed in m/s at 3-second gust, Peak-Ground Acceleration (PGA) in g, proximity to fire perimeter in metres.

Step 3 – Apply the damage function to obtain MDR.

Using the functional form (logistic, power law, lognormal) calibrated for the asset's vulnerability class, compute

$MDR(I, v)$. For example, a reinforced concrete commercial building in a 1.2m flood depth with a power-law curve

$$MDR = 0.12 \times d^2 \text{ gives } MDR = 0.12 \times 1.44 = 0.17$$

Step 4 – Apply secondary uncertainty.

If secondary uncertainty is modelled (as in full ELT-based calculations), sample the realized loss ratio from the distribution f$(MDR, \beta e)$.

For $\beta e = 0.5$ and $MDR = 0.17$, the distribution has mean 0.17 and standard deviation 0.085. The realised loss for this event and this asset is drawn stochastically.

Step 5 – Compute the gross financial loss.

$$Gross\ loss = realised\ MDR \times TIV$$ **(Replacement Cost Value of the asset).**

If you consider a warehouse with $TIV = €8m$ and $MDR = 0.17$ the $gross\ physical\ loss = €1.36m$.

- This represents the total repair and restoration cost before insurance conditions are applied.

Step 6 – Apply financial conditions.

Now the next step depends on whether you invested, if you lent or if you insured

A) For insurance applications, you'll need to review your contract. You'll need to apply a deductible, co-insurance, and policy limit to convert ground-up loss to insured loss.

B) For banking applications, you'll need to convert gross physical loss to revenue loss (via the P&L bridge in section 2.3.1) and to PD uplift (via the credit model in section 2.2).

C) For corporate disclosure, you'll need to convert to any type of value of your choice, such as RoI or RoE or RoA impact and balance sheet impairment per the accounting model in section 2.3.

Step 7 – Aggregate across all assets and all events.

Summing $MDR_i \times TIV_i$ across all assets i exposed to a given event gives the event-level portfolio loss Le. Weighting Le by the annual occurrence rate λe and summing across all events in the stochastic catalogue gives the Average Annual Loss (AAL). Ordering the event losses and constructing exceedance probabilities gives the OEP curve. These are the entry points to all the metrics described in sections about actuarial models (see 2.1.2 and 2.1.3).

An Example Scenario: Flood Damage Function Applied to an Industrial Portfolio

A reinsurer holds aggregate stop-loss protection for an insurer's portfolio of 340 light-industrial properties across the Rhine valley corridor in Germany. A 1-in-30-year Rhine flood event produces a spatially variable inundation footprint. Consider $\beta e = 0.5$

Asset-level flood depths range from 0.2m to 2.8m depending on elevation and proximity to the river. The insurer uses the German FLEMOps residential vulnerability curve adapted for light industry:

$MDR = min(1, 0.12 \times d^2)$ for reinforced concrete

$MDR = min(1, 0.24 \times d^2)$ for unreinforced masonry.

ASSET TYPE	COUNT	AVG DEPTH (M)	MDR	AVG TIV (€M)	EL (€M)
RC (Reinforced Concrete) light industrial ≥ 0.5m depth	142	1.05	0.13	4.2	77
Unreinforced masonry ≥ 0.5m depth	89	0.85	0.17	3.8	57
RC, depth < 0.5m (negligible damage)	71	0.25	0.008	4.0	2
Not inundated	38	0	0	3.9	0
Portfolio total	340	–	0.13 (weighted)	4.1	136

TABLE 14. Flood damage function applied to an industrial portfolio

The total event-level ground-up loss is estimated this way.

Here we can see a €136m on a total portfolio Total Insured Value (TIV) of €1,394m

Hence a portfolio-level MDR of 9.8% for this event.

So now, we consider the insurance recovery contract. Consider we have an 80% primary recovery and a stop-loss attach at €120m with a €50m limit.

In other words, the maximum that the insurance is okay to pay is €170m. We do not attain that limit. In our estimation the insured loss reaches the stop-loss attachment with €16m into the stop-loss layer.

A further question follows: how severe must the event be before the stop-loss layer is exhausted? Secondary uncertainty provides one way to answer it.

Applying $\beta e = 0.45$, the coefficient of variation appropriate to industrial flood, yields a loss distribution around the €136m mean.

At the 90th percentile of this distribution, the realized loss reaches €171m, of which €51m falls within the stop-loss layer, more than three times the €16m implied by the mean loss alone.

This is a 3.2 times amplification. It will arise entirely from the dispersion of possible outcomes around the mean; it reflects no change in the expected loss itself, only in the range of loss the layer must be prepared to absorb.

Provider Box – Climada and Climada Tech

CLIMADA (CLIMate ADAptation) is an open-source Python platform developed by the Weather and Climate Risks group at ETH Zurich (Aznar-Siguan & Bresch, 2019). It computes impacts from three separate objects, a hazard set, an exposure set and a set of impact functions, and returns an expected annual impact, exceedance frequency curves and return-period losses. Its distinctive module appraises adaptation measures. Each measure is described by its cost and by the way it modifies the hazard, the exposure or the impact function,

and its benefit is the damage it averts over time (Bresch & Aznar-Siguan, 2021). In simplified form, the benefit and the benefit-cost ratio are written as follows.

For each event and each asset, CLIMADA returns the impact. It then aggregates these impacts into an expected annual impact, the equivalent of the AAL, and into exceedance frequency curves and return-period losses. Two modules go further than a standard catastrophe model. The first one appraises adaptation measures. Each measure is described by its cost and by the way it modifies the hazard, the exposure or the impact function, and its benefit is the damage it averts over time (Bresch & Aznar-Siguan, 2021). In simplified form, the benefit and the benefit-cost ratio are written as follows.

$$B = \sum_{t=1}^{T} \frac{\left(EAI_t^{base} - EAI_t^{m}\right)}{(1+r)^t} \text{ and } BCR = \frac{B}{C_m}$$

- $\boldsymbol{B}$ is the present value of the damage averted by measure m.
- $\boldsymbol{EAI_t^{base}}$ is the expected annual impact in year t without the measure, and EAI_t^m is the expected annual impact in year t with it. r is the discount rate and T is the appraisal horizon in years. C_m is the cost of the measure, and BCR is the benefit-cost ratio, with a value above 1 meaning the measure pays for itself.

The second module, *unsequa*, propagates uncertainty in the hazard, exposure, impact function and discount rate inputs through the whole chain and ranks them by their contribution to the variance of the output (Kropf et al., 2022).

Finance researchers use the platform as an open alternative to commercial catastrophe models. One study used CLIMADA to model tropical cyclone damage to the geolocated mines and power plants of listed Mexican companies. It then passed these asset-level shocks to a climate-adjusted dividend discount model. Investor losses were underestimated by up to 70% when assets were placed at headquarters locations, and by up to 82% when acute tail risk was left out (Bressan et al., 2024). In 2022, Sebastian Glink, Simone Thompson and David Bresch founded CLIMADA Technologies, an ETH spin-off that sells hazard ratings and financial metrics built on the open-source core. Because the code is fully open, every assumption in the chain can be read and changed. In return, the user must take on the calibration and the choice of hazard resolution that a commercial vendor would otherwise handle.

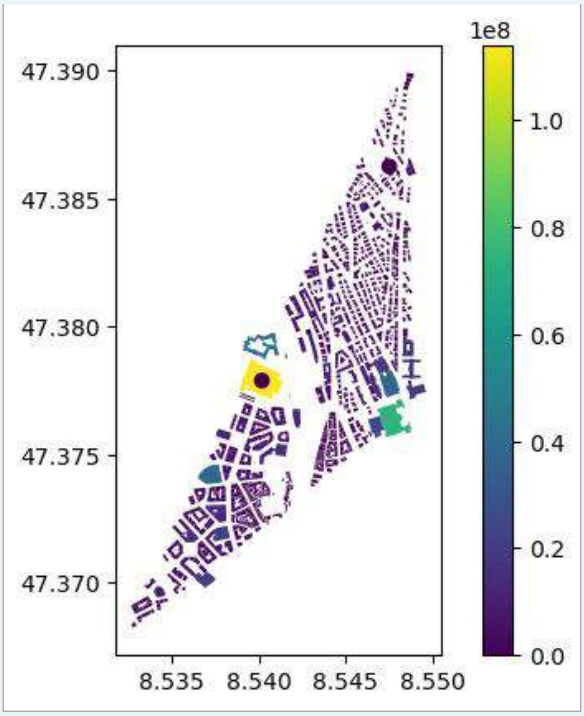


FIGURE 18. An example of output using Climada open-source damage functions.
Source: Climada documentation

The open-source version of CLIMADA draws its flood damage functions from the JRC (Huizinga et al., 2017), among other sources depending on the peril. The JRC curves come from a review of regional and national depth-damage studies, harmonized into one concave curve per continent and rescaled per country with maximum damage values in euros per square metre.

CLIMADA then multiplies this curve by the exposure value of each asset to produce a loss. The hazard layers usually paired with the open-source core are coarser than purpose-built models such as Fathom, which resolves flood extent at 30 m (Wing et al., 2024). The core is therefore a reference tool for understanding asset-level "raw" damage, and it leaves calibration and resolution to the user.

CLIMADA Technologies closes most of this gap in its commercial platform, Delta Climate. It covers about thirty hazards, including hail and avalanches, with global coverage for all of them except hail. Flood hazards reach 30 m resolution, temperature-related hazards 90 m and wind 1 km. Floods, wind and hail rely on forward-looking stochastic event sets under every SSP scenario up to 2100.

The generic JRC-type curves are replaced by impact functions specific to the region, the hazard and the building characteristics. About half of the calibration data comes from public sources and half from insurers' claims and clients' records of downtime and production losses. These curves are perturbed to quantify vulnerability uncertainty alongside hazard uncertainty. The platform reports the AAL and return-period losses, and the company recommends return-period losses because the AAL hides the tail. Model outputs are backtested once a year against insurers' claims.

Model Representation Card – Damage Function / Vulnerability Model - Salt Farming in Indonesia

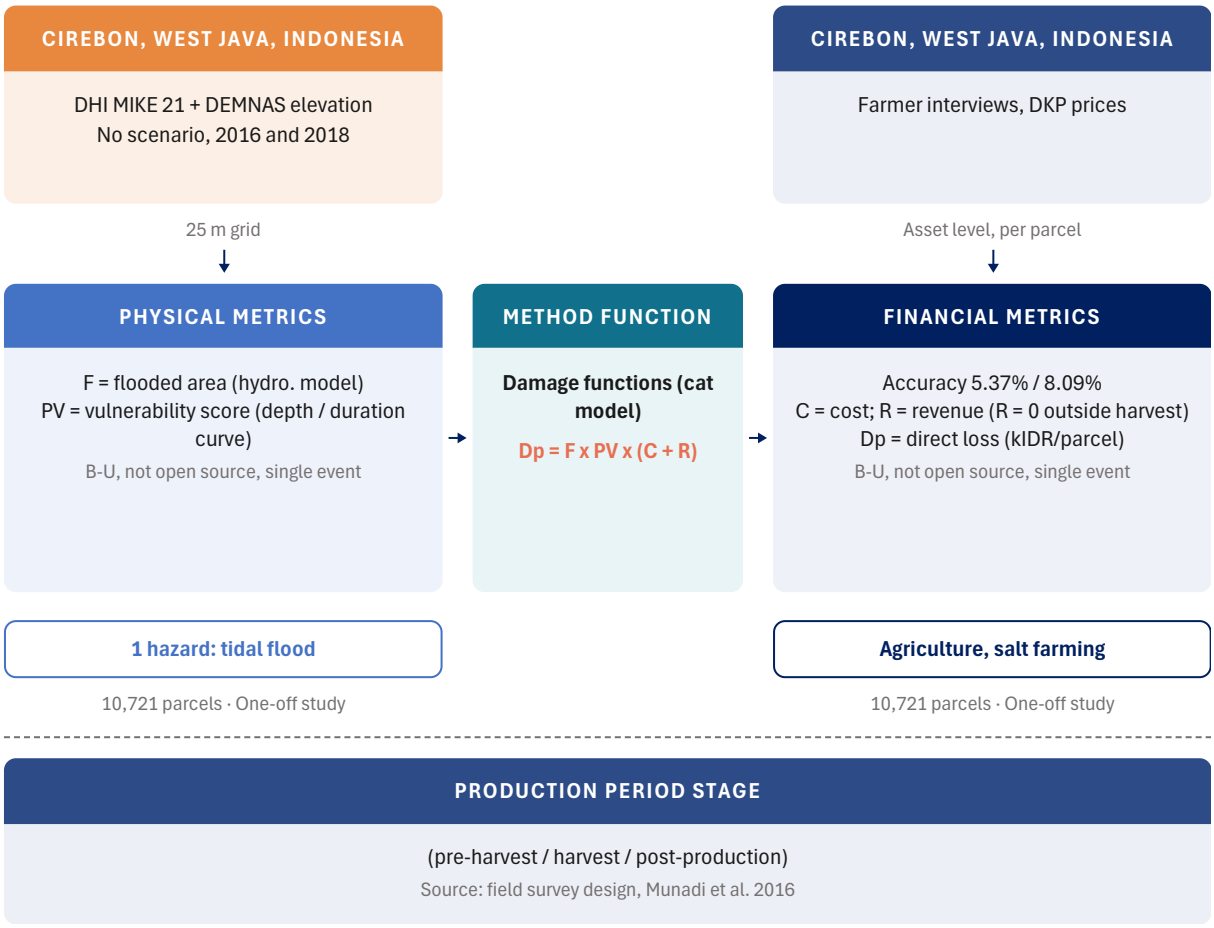


FIGURE 19. Model representation card – Damage function / vulnerability model (salt farming in Indonesia)

Why is this model useful?

This model can be used if you are working in a coastal, low-data region where a specific sector needs a defensible loss estimate but no formal damage records exist.

It only requires a handful of expert or local knowledge interviews and one or two reconstructed historical events, which makes it fit for small-scale, data-scarce production contexts rather than large institutional portfolios.

It produces individual, parcel-level loss figures, which is useful when you need to understand who specifically is affected and by how much.

What does this model teach us?

This model teaches us that the timing of a hazard within a production cycle can matter more than the hazard's raw physical severity, since the same flood depth produced very different losses depending on whether it hit before or during harvest. It also shows that a workable depth duration damage function does not require a rich historical loss dataset, fourteen interviews and two reconstructed events were enough to calibrate one across more than ten thousand parcels.

What it does not teach us anything about is the future, since there is no scenario and no time horizon, it only tells you what a specific past event cost.

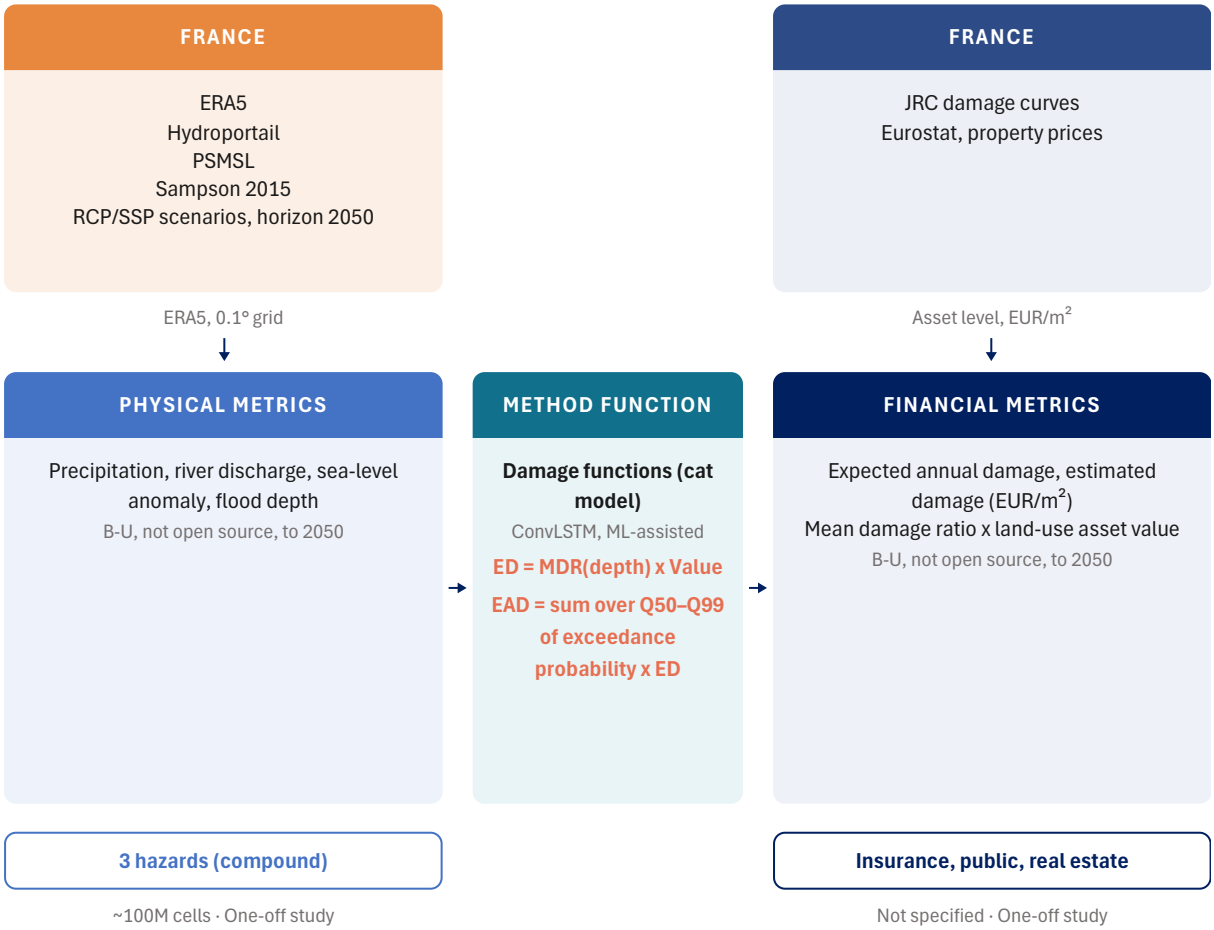


FIGURE 20. Model representation card – Damage functions (cat model) applied to flood risk in France

Why is this model useful?

This model can be used if you need a forward-looking, asset-level view of flood damage across a large territory, and if the intended audience is institutional, insurers, public sector bodies, or real estate holders who need a euro-per-square-metre figure they can plan around. It is a fit when you want to lean on an established, standardized damage curve methodology (JRC), while still capturing evolving hazard probability through a machine learning forecast layer.

What does this model teach us?

This model teaches us that combining a probabilistic forecasting layer with a fixed, well-established damage curve is a workable way to project asset-level damage forward in time without having to build a new vulnerability function for every hazard. It also shows that land use class is not a minor detail, it is the axis the entire damage curve is built around, since the same flood depth causes very different damage depending on what sits on that land. And because the hazard here is explicitly compound, pluvial, fluvial, and coastal combined, it teaches us that treating these hazards separately would likely understate the real risk.

Provider Box – refinq

refinq is primarily a damage-based model. Its core is a damage function defined for each combination of hazard, economic activity, location and time horizon. Activities are identified by their NACE or EU Taxonomy code, and the model covers 24 physical climate hazards for return periods of 1 to 100 years, alongside 11 nature and biodiversity metrics. The damage functions are empirical and statistical.

They are calibrated against observed loss records and benchmarked against the Munich Re NatCat database, and modelled loss ratios are compared with observed sector-level losses for each hazard and NACE sector. refinq also produces a Nature Value at Risk (or NVaR), which combines physical climate risk with nature dependencies and impacts, and which the provider considers to go beyond a standard damage-based approach.

refinq uses machine learning. But it is not used in the core risk modelling. Expected annual loss and physical risk scores therefore keep the properties of a damage-based model. They can be traced back to a damage function, and they are compared with observed losses where clients share their loss data.

Instead of using ML-based model for projective analysis, the solution uses a learning component: a large language model in the GaiaGuide module. This component intervenes once the risk metrics are computed, interprets them, and generates costed adaptation recommendations for each site.

Clients can then record the measures they have implemented or plan to implement. These measures reduce the vulnerability component and produce a net risk score after adaptation. The language model adds no information about the future climate or about losses. Its role is to translate existing metrics into decisions. It is an original way to use AI, especially compared to other providers.

Keep in mind when AI is put forward that the key question is where the algorithm sits in the modelling chain, in the generation of hazards, in the damage function, or in the interpretation of results (see Section 2.5.4 on ML-based models).

Key Limitations

Data availability constrains calibration quality. Vulnerability functions for well-studied perils (North Atlantic hurricane, California earthquake, European flood) are calibrated against millions of historical claims across multiple decades. Functions for emerging perils (wildfire damage to industrial facilities, heat-stress productivity loss for specific occupational categories, permafrost thaw damage to Arctic infrastructure) will rest on far smaller empirical bases, sometimes fewer than a hundred observations.

The Mean Damage Ratio (MDR) uncertainty will be characterized by a 95% confidence interval spanning $MDR \pm 0.15$ at typical intensities.

Stationarity of vulnerability. Damage functions are calibrated on historical building stock. As building codes, materials, and construction practices evolve, the vulnerability of the current stock progressively diverges from the historical sample on which functions were calibrated.

If the vulnerability was calibrated based on masonry buildings from the 1960s–1990s, we will systematically overestimate damage to post-2010 RC-frame construction and underestimate it for pre-1950 unreinforced masonry. However, consider that updating vulnerability functions to track building stock evolution is a very demanding process that most CAT model vendors perform only every three to five years.

Compound and cascading hazards. Damage functions are almost universally calibrated for a single hazard acting in isolation. Physical climate events are rarely single-hazard: tropical cyclones bring wind, rainfall, storm surge, and

flooding simultaneously; drought weakens foundations before an earthquake; wildfire destroys vegetation that previously protected slopes from landslide and it also renders the soil less moist, which makes the terrain more prone to be flooded. The interactions between concurrent hazards can produce MDRs that exceed the sum of individual-hazard MDRs, because structural damage from one hazard (e.g. cracked walls from earthquake) dramatically increases vulnerability to the next (e.g. subsequent flood inundation). To our knowledge, no widely deployed vulnerability model currently captures multi-hazard interactions in a physically consistent way.

Computational cost scales with resolution. Asset-level damage modelling with hydrodynamic inputs at 1–10m resolution requires substantial computational resources. For a national-scale portfolio of 500,000 residential properties subject to a full stochastic event catalogue of 200,000 flood events, the full MDR computation may require several thousand CPU-hours. So, of course these kinds of models are not meant to be used by anyone. It's quite an obstacle for smaller institutions (regional banks, SME insurers, corporate treasury teams) to run granular damage-based assessments without access to cloud-scale computing infrastructure or vendor platforms. Score-based approaches (section 2.3) sacrifice precision for computational tractability in these contexts.

The exposure data problem. A damage function is only as useful as the exposure data it is applied to.

The weakest link in most real-world implementations is the data, the asset register: incomplete geo-location data (billing addresses which is less precise than production facility addresses), missing construction attributes (year built, construction type, elevation), and inaccurate replacement cost values. For corporate physical risk assessments, the International Sustainability Standards Board's IFRS S2 standard now requires disclosure of material physical risks, creating a regulatory incentive for firms to invest in asset register quality. The quality gap between insurance-grade exposure data and the data available for most corporate or banking-sector physical risk assessments remains large.

Provider Box – The case of Adaptation/s – Using inquiries about industrial sites to fit the model

Adaptation/s does not produce climate data. It buys its exposure layers from external providers, selected after an eight-month benchmark, and concentrates its own work on vulnerability.

Each site is represented as a digital twin of the infrastructure, broken down into systems and subsystems such as roofs, walls, machines, stocks, people, industrial processes and external dependencies. In total, more than 100 vulnerability parameters describe a site, from roof and material types to the elevation of stocks, the dependencies between machines and the governance of the site.

These parameters are collected through detailed questionnaires and engineering documents, and machine learning helps extract them from the client's documents.

Damage functions here are deterministic and engineering-based. They are built in house from the engineering literature and from prevention engineering expertise, and they do not rely on historical insurance claims. The model follows how a failure in one system spreads to others, for instance a flooded electrical room that stops a production line, and it computes the resulting business interruption. Outputs include the average annual loss, net loss estimates, replacement costs, business interruption costs and cascading operational losses. Adaptation is modelled within the same chain. Each measure changes the vulnerability parameters it acts on, so the model gives a gross risk, a net risk after adaptation and the return on investment of each measure.

Validation currently draws on site teams and operational experts, backtesting of the financial impact of past events on specific infrastructures is under way, and discussions with insurers aim to bring claims data into the comparison. The clients were historically industrial corporates, and the provider is now extending toward insurers and banks, with insurers seen as the natural users of such a granular vulnerability framework.

> Their approach is a good example to understand that all the methodologies on the market today do not necessarily control the whole modelling chain. Some approaches focus on one or two specific in-house expertise and the rest is externalized (usually the climate side of the model).

2.1.2 ACTUARIAL MODELS

Actuarial models are one of the historical foundations of catastrophe risk quantification. Unlike "simple" damage-based models attempting to simulate the physical destruction process asset by asset, they answer a whole new set of questions such as:

- How many events of a given severity will a portfolio experience over a year?
- What is the distribution of losses that a cedant should expect across ten thousand possible futures?
- What capital reserve does a portfolio need to survive the worst 0.5% of years?

Its main tools (i.e. compound stochastic processes, extreme value distributions, credibility theory, and copulas) remain indispensable in modern catastrophe modelling (Embrechts et al., 2013; Mitchell-Wallace et al., 2017; Hochrainer-Stigler et al., 2019; Wing et al., 2020).

Because we want to keep this cookbook short and avoid you reading 5000 pages, we are not going to cover in-depth the mathematical foundations behind the actuarial models. We will only describe of they work and the interpretation of their outputs today.

The engineering layer of a cat model (the damage function) and the actuarial layer (the loss distribution) are complementary, they are sequential stages of the assessment pipeline. This section will describe the actuarial layer in full.

The Two-Component Decomposition

Every actuarial cat model decomposes total portfolio loss into two stochastic components that are modelled and combined:

- **Event frequency:** How many events occur in a given year? If it's governed by a counting distribution, these are most commonly used through a Poisson process.
- **Loss severity:** Given that an event occurs, how large is the loss? This will be governed by a severity distribution chosen to capture the heavy tail of catastrophic losses.

The aggregate annual loss is then the random sum of N independent severity draws, where N itself is random. This is the compound distribution framework, and it is the engine of every actuarial catastrophe model in existence.

In home insurance, the same decomposition can be estimated directly on claims data, with a logistic model for the probability of a claim and a Gamma model for its cost, it has been done for subsidence in France at the building level (Chatelain, 2023).

MATHEMATICAL REPRESENTATION

1. The homogeneous Poisson process to assess the events frequency

$$N \sim Poisson(\lambda)$$

$$N = number\ of\ events\ in\ one\ year\ (random\ variable) = k$$

MATHEMATICAL REPRESENTATION

$\boldsymbol{\lambda}$ is the mean annual occurrence rate (events per year), it is estimated from a historical catalogue or climate-adjusted stochastic event set.

The probability of exactly k events in one year is:

$$P(N = k) = \frac{\lambda^k \cdot e^{-\lambda}}{k!}$$

Remember: the key property of the Poisson is

$$E[N] = Var[N] = \lambda$$

This memoryless property means the time between events is exponentially distributed. It's a good first approximation for independent catastrophe events. For overdispersed data,

$$Var[N] > E[N]$$

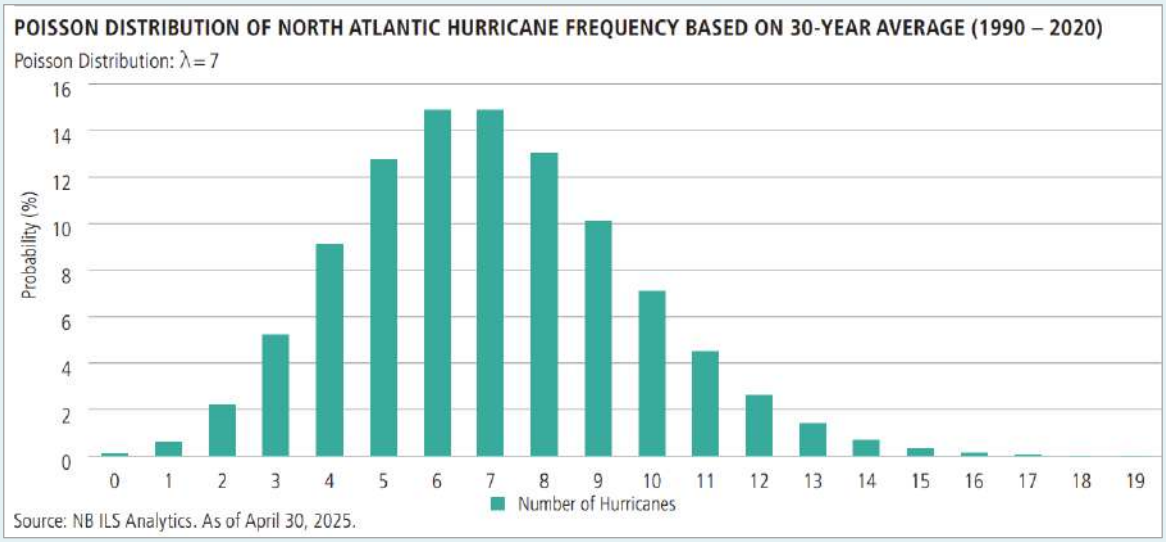


FIGURE 21. Poisson Distribution of North Atlantic Hurricane Frequency based on 30-Year Average (1990 – 2020)
Source: Catastrophe Modeling 101, O'Rourke and Burdick, 2026, Neuberger Berman.

However, the Poisson distribution assumes that events occur independently of one another, so that one storm makes the next neither more nor less likely. But in real life, events often cluster. Several occur within a short period because they share a common driver, such as a persistent atmospheric pattern or unusually warm sea surface temperatures. European winter windstorms are a typical case. Lothar and Martin struck France within two days of each other in December 1999, and the 2004 and 2005 Atlantic hurricane seasons both produced an exceptional number of landfalling storms. Clustering makes the annual number of events more variable than a Poisson process predicts. In that case the Negative Binomial distribution is preferred, because it allows the variance of the event count to exceed its mean (see below for more details on Negative Binomial processes).

2. Loss severity distribution: heavy-tailed families

For each event e, the loss X^e is drawn from a severity distribution $F(x)$

The choice of *F* must reflect the heavy-tailed nature of catastrophe losses. Four families are standard in actuarial practice:

Log-normal distribution

MATHEMATICAL REPRESENTATION

This is the "default" choice for everyday property damage and repair costs. It says: losses result from many small multiplying factors (e.g., cost overruns compounding on each other), which naturally produces a curve with a long-but-not-extreme tail. It's lighter-tailed than GPD, which means it's good for typical claims (Tuo and Guo, 2011) but it might be less suited to true catastrophe extremes.

This approach is widely used for property damage, repair costs, and health claims:

$$\ln(X) \sim \mathcal{N}(\mu, \sigma^2)$$

Here are a few reminders;

- $\boldsymbol{X}$ is the loss, the severity variable itself (for example, the dollar loss from a property damage event).
- $\boldsymbol{ln(X)}$ is the natural log of that loss.
- $\boldsymbol{\mu}$ is the mean of $ln(X)$, not the mean of X. It is the location parameter of the underlying normal distribution.
- $\boldsymbol{\sigma}$ is the standard deviation of $ln(X)$ hence σ^2 is its variance. This is also called the scale (spread) parameter.

This distribution has a sub-exponential (lighter) tail than GPD. Consider that it will be useful when the loss process is multiplicative (when many small factors are compounded). It's commonly used for assessing hail damage (Burnecki et al., 2000; Gao & Shi, 2025).

Pareto distribution

The "80/20 rule" distribution (i.e. we assume a small number of events will account for a large share of total losses).

It's the return of a power-law tail that is scale-free:

$$F(x) = 1 - \left(\frac{x_m}{x}\right)^{\alpha}, \quad x \geq x_m$$

Here, x_m is the minimum possible loss, the scale parameter. It sets the threshold below which no loss can occur, since the Pareto distribution only applies once losses exceed this floor.

The tail index α controls how heavy the tail is.

For $\alpha < 2$ the variance is infinite; for $\alpha < 1$ the mean is infinite. Extreme earthquake and mega-storm losses often imply $\alpha \in [1.5, 2.5]$.

The smaller alpha (α) is, the heavier the tail (i.e. there is more risk of catastrophic outliers).

If α drops below 2, the losses are so volatile that a mathematical "average loss" barely means anything (it is called an infinite variance); below 1, even the average itself breaks down (infinite mean). Big earthquakes and mega-storms can theoretically fall in this extreme range.

Generalized Pareto Distribution (GPD)

Instead of modelling every loss, GPD focuses only on losses above a high threshold (e.g., only losses over €10M). It's the mathematically "correct" shape for extreme, rare losses (i.e. what regulators and reinsurers will want to use for worst-case tail modelling).

The parameter ξ tells you how heavy the tail is: higher ξ means more risk of a truly extreme, damaging event. Earthquakes and hurricanes tend to have heavier tails ($higher\ \xi$) than floods.

MATHEMATICAL REPRESENTATION

This asymptotically correct distribution for exceedances above a high threshold u (if you want to sound cool in front of your colleagues, you can tell them it's called the "Pickands-Balkema-de Haan" theorem):

$$F(x) = 1 - \left[1 + \xi \frac{x-u}{\sigma}\right]^{-\frac{1}{\xi}},$$

where $\boldsymbol{\xi}$ is the shape (tail) parameter and σ is still the scale parameter.

For $\xi > 0$ the distribution has a power-law tail.

For most nat-cat perils, $\xi \in [0.1, 0.5]$. Note that there is a heavier tail for earthquake and hurricane than for flood (See Ling et al., 2021 for an application with floods for example. Also see Dey & Das, 2016 for an application to hurricanes).

Cauchy distribution

This distribution can be useful in a limiting case ($\alpha = 1$, infinite mean and variance). It's rarely used directly for loss modelling but it signals that the standard actuarial Law of Large Numbers breaks down entirely. In practice it is referenced as a stress test for the assumption that reserves are sufficient to cover expected losses.

Compound distribution: aggregate annual loss

So how do you consider compound events and aggregate them in actuarial models?

$$S = X_1 + X_2 + \cdots + X_N = \sum_{i=1}^{N} X_i$$

$$S = aggregate\ annual\ loss$$

- $\boldsymbol{N \sim Poisson(\lambda)}$ = random number of events

$$X_i \approx iid \text{ (independent and identically distributed)}$$

- $\boldsymbol{F = loss\ from\ event\ i}$, still iid (independent and identically distributed) from severity distribution F

Moments of the compound distribution are:

$$E[S] = \lambda \cdot E[X]$$

$$Var[S] = \lambda \cdot E[X^2] = \lambda \cdot (Var[X] + E[X]^2)$$

The moment generating function (sometimes called MGF) of S will be:

$$M_S(t) = exp(\lambda(M_X(t) - 1))$$

For heavy-tailed X (i.e. when using Pareto and GPD), $M_X(t)$ does not exist for $t > 0$

If $t > 0$ the MGF approach breaks down. In that case, the distribution of S is computed numerically via the Panjer recursion or via Monte Carlo simulation (Embrechts et al., 2013).

Panjer recursion: another way of assessing the possible total loss

The Panjer recursion builds up the probability of every possible total loss amount, without simulation (Panjer, 1981).

MATHEMATICAL REPRESENTATION

It starts from "zero events happened" and then works outward, adding one loss unit at a time, reusing each answer to get the next one.

It gives you the exact answer with no random sampling error, which is why insurance companies use it for pricing and solvency calculations as long as the losses can be rounded onto a discrete grid. Consider you have a Poisson distribution.

When N is Poisson and X is discretised to a lattice $h, 1, 2, \ldots$, the aggregate loss distribution can be computed exactly:

$$f_S(kh) = \frac{\lambda}{k} \cdot \sum_{j=1}^{k} j \cdot f_X(jh) \cdot f_S((k-j)h)$$

- $\boldsymbol{h}$ is the discretization step, the fixed lattice spacing you choose to round losses onto (for example 1000 euros), so that losses only take values $0, h, 2h, 3h$, and so on.
- $\boldsymbol{k}$ is the lattice index for the aggregate loss S, so kh is the actual aggregate loss value you are computing the probability for.
- $\boldsymbol{f_S(kh)}$ is the probability mass function of the aggregate annual loss S, evaluated at the point kh. This is the quantity you are solving for recursively.
- $\boldsymbol{\lambda}$ is the mean annual occurrence rate of the Poisson frequency distribution, as we have seen in $N \sim Poisson(\lambda)$ at the beginning of the section.
- $\boldsymbol{j}$ is the summation index, running from 1 to k, it steps through each possible discretized severity value below or equal to kh.
- $\boldsymbol{f_X(jh)}$ is the probability mass function (pmf) of the severity distribution X, evaluated at the discretized point jh. This is the discretized version of your severity distribution F, whichever family you chose, GPD, log-normal, or Pareto.
- $\boldsymbol{f_S((k-j)h)}$ is the aggregate loss probability mass function evaluated at the earlier lattice point $(k-j)h$. This is the recursive part; you are building up f_S at kh using previously computed values of f_S at smaller lattice points.

This recursion starts at;

$$f_S(0) = e^{-\lambda}$$

Where the chance a Poisson variable equals 0.

It then computes the full aggregate loss probability mass function (*pmf*) to any desired precision. It is $O(N^2)$ in computation but eliminates sampling error entirely, making it the method of choice for pricing and solvency when the loss amounts can be discretised.

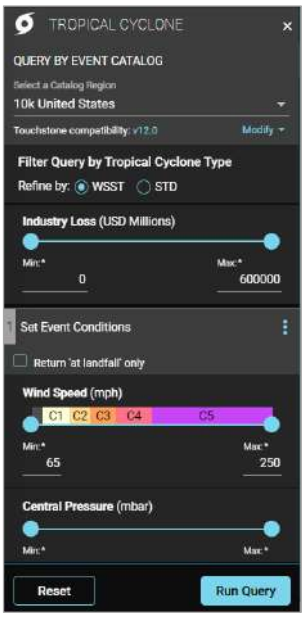
How many events are expected? Frequency Distribution in Detail: Poisson, NB, and Hawkes Processes

The choice of frequency distribution matters more in catastrophe modelling than in most actuarial applications because climate change is modifying the parameters of these distributions in non-stationary ways. Understanding what each distribution assumes is therefore essential for applying them under changing climate conditions.

Homogeneous Poisson process (HPP)

Homogeneous Poisson process will set the number of events per year in the frequency layer of a damage model that will feed directly into the compound loss.

Events occur at constant rate λ, independently of each other and of time. Its key limitation under climate change is that λ is constant. Empirical observation of tropical cyclone frequency (Knutson et al., 2020) and compound flood events (Moftakhari et al., 2017) increasingly suggests time-varying rates and inter-event dependence.



An illustration from a reinsurance provider's UI

Non-homogeneous Poisson process (NHPP)

Non-homogeneous Poisson process will replace the constant event rate with a time-varying one, so damage models can project rising annual expected losses under a different trend.

The rate here is a deterministic function of time: $\lambda(t)$

Used in climate-adjusted cat models where the event frequency is projected to change with global mean temperature. The expected number of events in $[0, T]$ is

$$\Lambda(T) = \int_0^T \lambda(t)\, dt$$

- t is time, the continuous variable running along the horizon (in years, typically).
- $\lambda(t)$ is the instantaneous event rate at time t. Unlike the homogeneous Poisson process where λ was a single constant, here it is a function that can change value at every point in time.
- T is the length of the time horizon you are evaluating, the upper bound of the integral (for example, 10 years, or 50 years for a long-term climate projection).

For a linear climate trend:

$$\lambda(t) = \lambda_0(1 + \gamma t)$$

$$\Lambda(T) = \lambda_0\left(T + \frac{\gamma T^2}{2}\right)$$

- λ_0 is the baseline rate, the value of $\lambda(t)$ at $t = 0$, before any climate trend is applied. This is your starting point, equivalent to the constant λ from the homogeneous Poisson process.
- γ is the climate trend parameter. It controls how fast the rate grows per unit of time.

A positive γ means the event rate increases over time, consistent with a warming driven rise in catastrophe frequency. If γ were zero, you would be back to the homogeneous Poisson process with constant rate λ_0.

RMS and AIR's climate change modules implement variants of this.

Negative Binomial process

Negative Binomial process swaps in for the Poisson frequency layer when a portfolio's historical event counts are more volatile than Poisson allows, giving damage models wider loss distributions and fatter tails on the frequency side (Meyers et al., 2003; Rijal et al., 2018).

When event counts are overdispersed (clusters of storms in active seasons followed by quiet years), $N \sim NegBin(r, p)$ with

$$E[N] = \frac{r(1-p)}{p}$$

Where p is the success probability parameter; together with r it controls the degree of overdispersion

and

$$Var[N] = \frac{r(1-p)}{p^2} > E[N]$$

The extra parameter r controls clustering. For Atlantic hurricane seasons, historical data suggest significant overdispersion consistent with climate oscillations operating on multidecadal (AMO) and interannual (ENSO, 2–7 years) timescales (Bove et al., 1998; Yang, 2013; Klotzbach et al., 2018). The same holds for European winter storms, whose overdispersed annual counts are largely explained by the phase of the North Atlantic Oscillation (NAO) (Economou et al., 2015).

Hawkes Processes

Hawkes processes let damage-based models raise the event rate right after a loss occurs. It captures clustering effects like aftershock damage or back-to-back hurricane and flood losses that a memoryless Poisson process would miss.

$$\lambda(t) = \mu + \sum_{\tau i < t} \alpha \cdot e^{-\beta(t-\tau_i)}$$

- where **μ** is the base rate,
- α is the branching ratio (how much each event raises the rate), and
- β is the decay rate.
- τ_i are the occurrence times of the previous events, each past event i has its own timestamp τ_i

The sum runs over all $\tau_i < t$ (i.e. only events that have already happened before the current time t contribute to raising the rate at time t). Hence the process is "self-exciting", every past event leaves a decaying trace on the current intensity, governed by the $e^{-\beta(t-\tau_i)}$ term.

Hawkes processes are increasingly used in earthquake sequences and in modelling compound climate events (back-to-back hurricanes, sequential flood-storm events).

Provider Box – The case of RMS Climate on Demand

RMS is, with AIR, one of the historical actors of catastrophe modelling, and both firms emerged in the late 1980s. Moody's, which acquired RMS in 2021, now offers Climate on Demand, a product that carries a catastrophe model into the future climate.

Climate on Demand quantifies physical risk for seven perils, namely tropical cyclones, inland floods, sea level rise, wildfires, heat stress, water stress and earthquakes. It covers the period 2020 to 2100 under four RCP scenarios and two NGFS scenarios, across several time horizons, and distinguishes commercial, industrial and residential buildings. Results are available at four levels of aggregation. These are the individual location

or real asset, the portfolio, the company (based on the location of its facilities), and the area, from sovereign to sub-sovereign units such as postcodes.

The platform rests on the RMS probabilistic catastrophe models, whose event catalogues have been calibrated over decades of insured losses. Global climate models describe the large-scale climate of 2050 or 2100, yet they do not resolve the track of a hurricane or the pattern of rainfall over a river basin. Climate on Demand combines the two. The climate models supply the change, and the catastrophe model supplies the events.

Each catalogue is conditioned to represent the events that could occur in a given future climate. Because the model simulates the full life cycle of each event, the physical variables that change during it can be adjusted directly while their interactions are preserved. Examples are the radius of maximum winds, the track of a hurricane or the spatial pattern of rainfall. For floods, the conditioning is applied to high-resolution rainfall, and water depths follow from the hydraulic simulation. Sea level rise is added to coastal flooding.

The catalogues are very large. The US inland flood model simulates 50,000 years of continuous precipitation, and the US wildfire model more than 18 million fires over 50,000 years. Each simulated event is then converted into damage and downtime from the location, replacement cost and building characteristics of each asset.

Outputs are provided for each peril and for all perils combined. They include hazard metrics, financial impact metrics and impact scores on a 0 to 100 scale. The financial metrics comprise the mean annualized damage ratio (ADR) and return-period tail risk metrics.

The ADR is the expected aggregate annual damage expressed as a fraction of total exposed value. It results from the combination of hazard, vulnerability and financial modelling, and therefore reflects both the frequency and the severity of events. It is reported per 1,000 dollars of asset value, so an ADR of 0.52 corresponds to about 520 dollars of average annual damage per million dollars of assets.

Moody's adds its uncertainty and reports a risk metric equal to the mean plus one standard deviation of the annual loss.

$$RM = ADR + \sigma_L$$

- where $\boldsymbol{RM}$ is the risk metric,
- $\boldsymbol{ADR}$ is the expected aggregate annual damage expressed as a fraction of total exposed value. It results from the combination of hazard, vulnerability and financial modelling, and therefore reflects both the frequency and the severity of events. It is reported per 1,000 dollars of asset value, so an ADR of 0.52 corresponds to about 520 dollars of average annual damage per million dollars of assets.
- and $\boldsymbol{\sigma_L}$ is the standard deviation of annual loss on the same scale. This standard deviation combines three sources.

Primary uncertainty is the variation of climate from one year to the next and whether an event occurs at all. Exposure uncertainty reflects how well the value of each location is known.

Secondary uncertainty is the size of the loss once an event occurs, which depends on construction quality and on the correlation between nearby locations. Somewhere in this simple approach lies the logic of actuarial cat modelling: a frequency, a severity and a dependence structure, each kept separate and each conditioned on the future climate.

Extreme Value Theory: if you want to go further

The mathematical justification for using GPD and Pareto distributions for catastrophe losses comes from Extreme Value Theory (EVT), specifically the Pickands-Balkema-de Haan theorem and the Fisher-Tippett-Gnedenko theorem.

Fisher-Tippett-Gnedenko theorem. The maximum of n i.i.d. random variables, when properly normalised, converges to one of three distributions as $n \to \infty$:

Ξ is a shape parameter.

Gumbel ($\Xi = 0$)

Fréchet ($\Xi > 0$)

Weibull ($\Xi < 0$)

Together these form the Generalized Extreme Value (GEV) distribution:

$$G(x) = exp\left(-\left(1 + \xi \cdot \frac{x - \mu}{\sigma}\right)^{-1/\xi}\right)$$

Empirically, many nat-cat perils, particularly hurricane and earthquake losses, are found to lie in the Fréchet domain ($\Xi > 0$), with a heavy-tailed behaviour (Dey & Das, 2016). The theorem only establishes that a limit exists among three possible types, but this behaviour is an empirical finding specific to the peril and dataset. Flood losses, for instance, have sometimes been found closer to the Weibull domain (Li et al., 2023).

Pickands-Balkema-de Haan theorem. For large thresholds u, the distribution of exceedances $X - u \mid X > u$ converges to the Generalized Pareto Distribution (GPD (ξ, σ)). This is the theoretical foundation for the Peaks Over Threshold (POT) method, which is the standard EVT approach for estimating return levels of extreme losses from a historical claims dataset.

If you are interested in a more detailed application, the Peaks Over Threshold protocol below follows standard practice in hydrology and meteorology (Dutfoy et al., 2014) and has been applied to the financial market risk of extreme sea-level rise (Garcia-Jorcano & Sanchis-Marco, 2024).

EVT ESTIMATION PROTOCOL - PEAKS OVER THRESHOLD: A QUICK GUIDE ON HOW TO CONSIDER EXTREME EVENTS

Step 1: Choose a threshold u. Then plot the Mean Excess Function

$$e(u) = E[X - u \mid X > u]$$

It should be approximately linear above the chosen u as linearity signals GPD fit. Typically, by default, u can set at the 90th–95th percentile of the loss distribution.

Step 2: Fit GPD to exceedances. Use Maximum Likelihood Estimation (MLE) or Probability-Weighted Moments to fit (ξ, σ) to the n_u exceedances above u. Profile likelihood confidence intervals are strongly preferred over normal approximation because the MLE of ξ has non-standard asymptotic behaviour for $\xi > 0.5$.

Step 3: Estimate return levels. The T-year return level (loss exceeded with probability 1/T per year) is:

EVT ESTIMATION PROTOCOL - PEAKS OVER THRESHOLD: A QUICK GUIDE ON HOW TO CONSIDER EXTREME EVENTS

$$x_T = u + \frac{\sigma}{\xi} \cdot ((T \cdot \lambda_u)^{\xi} - 1)$$

Where:

$$\lambda_u = P(X > u)$$

it is the exceedance rate above u.

Step 4: Uncertainty quantification. Bootstrap or profile likelihood confidence intervals on x_T. For insurance pricing, use the upper bound of the 90% CI for prudent reserving.

Step 5: Account for non-stationarity (optional). Under climate change, the stationarity assumption may not hold. The threshold can be made time-varying, $u(t)$, using quantile regression, and the GPD scale parameter $\sigma(t)$ can be modelled as a function of time or global mean temperature; the exceedance frequency $\lambda(t)$ then follows the NHPP described above.

Credibility Theory and Portfolio Heterogeneity

A challenge in actuarial catastrophe modelling is that individual geographical cells or lines of business often have sparse loss histories. Fitting a separate severity distribution to five years of observations from a Midwest US hail portfolio can be unreliable.

Credibility theory aims to solve this by combining the cell-specific observations with broader portfolio or industry data using a weighted average:

$$P_{credibility} = Z \cdot X_{own} + (1 - Z) \cdot \mu_{global}$$

- $P_{credibility}$ is the credibility-weighted estimate, the final blended loss or rate estimate used for pricing a specific cell or line of business.
- Z is the Bühlmann credibility factor, a weight between 0 and 1 that sets how much their own data versus the broader portfolio data counts.
- X_{own} is the observed average loss (or claim frequency) computed from the cell's own limited data, for example the five years of Midwest hail observations in your example.
- μ_{global} is the mean loss (or frequency) computed from the broader portfolio or industry data, the fallback value used when own data is too sparse to trust on its own.

The Bühlmann credibility factor is Z, comprised between 0 and 1:

$$Z = \frac{n}{n + k}$$

- n the number of own observations available for the cell, already defined in your prose.
- k is the credibility constant, the ratio of within-cell variance to between-cell variance. A smaller k means fewer own observations are needed before the cell's own data dominates the blend.

$$k = \frac{\sigma^2\ within}{\sigma^2\ between}$$

- σ^2_{within} is the variance of losses within a single cell over time, how much one geography's own year-to-year experience fluctuates.
- $\sigma^2_{between}$ is the variance of the mean loss level across different cells in the portfolio, how much the underlying risk level differs from one geography or line of business to another.
- **Credibility** is used for example in property pricing and is the mathematical basis for the "load" adjustments that cat modellers apply when a specific geography has an unusually thin claims history.

Another approach was implemented for instance in the US National Flood Insurance Program, which used the square root rule to set credibility standards that vary by rating territory based on the territory's own heterogeneity (FEMA, 2022).

Square root rule represents a classical limited fluctuation credibility):

$$Z = \sqrt{\left(\frac{n}{n_{full}}\right)} \text{ where } Z \text{ is capped at } 1$$

- n is the actual number of observations available for the territory (for example, the number of claims or exposure years).
- n_{full} is the full credibility standard, the number of observations required before the territory's own data is trusted on its own, with no blending needed.
- Z is the credibility factor, same role as in the Bühlmann formula, but derived differently here, geometrically through a square root rather than linearly through $n/(n+k)$.
- n_{full} itself comes from a target confidence and precision level:

$$n_{full} = \left(\frac{z_{\alpha/2}}{p}\right)^2$$

- $z_{\alpha/2}$ is the standard normal critical value for the desired confidence level (for example 1.96 for 95 percent confidence).
- p is the precision tolerance, how close the sample estimate must fall to the true mean, expressed as a proportion (for example 0.05 for within 5 percent).

How can Monte Carlo simulations be implemented?

For most real-world portfolios, the compound distribution S cannot be computed analytically because: (i) the severity distribution F has no closed-form; (ii) the portfolio contains multiple correlated perils; and (iii) the event catalogue has complex spatial structure. Monte Carlo simulation is the standard computational engine.

Algorithm

1. Sample $N \sim Poisson(\lambda)$ events from each peril.
2. For each sampled event, draw an intensity from the hazard module and compute asset-level losses via the damage function.
3. Aggregate across assets to obtain the event loss L^e.
4. Sum across all N events to obtain the annual aggregate loss $S = \Sigma L^e$.
5. Repeat for 100,000 to 500,000 simulated years.
6. Sort the simulated annual losses to construct the empirical loss distribution: OEP and AEP curves, AAL, VaR at any return period.

The standard error of the $T - year$ return level estimate from M simulated years is approximately proportional to $1/\sqrt{(\frac{M}{T})}$

For a 250-year return level, $M = 100{,}000$ simulated years gives a standard error of about 5% of the estimated loss.

This is sufficient for pricing although borderline for regulatory capital calculations requiring high precision at 200-year return periods (EU Parliament and Council, 2009).

Provider Box – Morningstar's Sustainalytics: a rapid evolution in Physical Risks Portfolio Management

Morningstar Sustainalytics measures physical risk from the investor's side. Its Physical Climate Risk Metrics (PCRM) start from more than 12 million geolocated physical assets, linked to more than 12,000 listed issuers across 135 sectors and 235 countries. The hazard and damage layer comes from XDI's Climate Risk Engines, which are engineering-based. They combine peril-specific hazard models and climate projections with the design and construction of each asset, such as building standards, material types, height above ground and upstream flood defences. The output for each asset is a direct damage estimate and a measure of operational disruption.

The financial outputs used Morningstar's standardized financial databases. Expected losses are accumulated year by year up to 2050 and divided by the company's operating cash flow:

$$DLR = \frac{\sum_{t=1}^{T} L_t^D}{\sum_{t=1}^{T} OCF_t}$$

where $\boldsymbol{DLR}$ is the direct loss ratio,

$\boldsymbol{L_t^D}$ is the expected direct loss in year t from damage to the company's assets and loss of productive capacity,

$\boldsymbol{OCF_t}$ is the company's operating cash flow in year t, and T is the horizon, 2050 by default.

The indirect loss ratio follows the same form, with the losses that come from damage to local critical infrastructure and to the surrounding region.

The total loss ratio is the sum of the two. A fifth signal gives the share of a company's assets whose average annual damage exceeds 1% of their replacement cost.

The results are expressed under RCP 2.6 and RCP 8.5, and companies can be benchmarked against peers, regions and horizons. They describe gross risk. Company-level adaptation is not included, although hazard maps already reflect public defences such as the Thames Barrier in London.

Portfolio Information	RCP2.6	RCP8.5	Coverage %
	Total Loss Ratio Discounted	Total Loss Ratio Discounted	
	Expected Total Loss Amount Discounted USD	Expected Total Loss Amount Discounted USD	
Portfolio	0.04	0.05	
PCRM Test Portfolio	92,197.61	103,305.55	79.00%
Benchmark	0.01	0.01	
Morningstar Global Markets	975,549,031,504.47	1,184,056,670,673.03	63.00%

Data correct as 16 July 2024, 09:25:58 UTC

FIGURE 22. An example output from PCRM1.0

PCRM2.0

Sustainalytics' Physical Climate Risk Metrics 2.0, resented in 2026, address three problems that Morningstar sees in current tools. Hazard models are hard to audit, company asset data are incomplete and lack materiality context, and public companies, private companies, real estate and infrastructure are rarely assessed consistently. The product rests on three partners, and each one owns one layer of the data.

Veridion supplies the geolocation layer. Its AI platform collects public filings and digital signals such as job postings, social media and map pins. From these it identifies the operational sites of each company, where people work and what they produce there. This matters for exposure, because a storm that hits an unoccupied registered office has little economic consequence.

XDI supplies the hazard and damage layer through its engineering-based Climate Risk Engines. These combine sub-asset design and construction data, hazard models and forward-looking climate scenarios.

Damage is estimated from engineering archetypes, which XDI describes as digital twins of each building type and its major components, adjusted to the building standards of each country. These are distinct from vulnerability functions calibrated on the average behaviour of a building class.

The construction period, height above ground and expected materials of an asset therefore shape its damage estimate. For each asset, the output is a projection of physical damage and operational disruption.

The product assesses 11 hazards: coastal inundation, tropical cyclone storm surge, riverine flooding, surface water flooding, extreme wind, extreme heat, freeze-thaw, forest fire, soil movement, tropical cyclone wind and landslide. The climate signal currently comes from CORDEX regional climate models. A transition to the CMIP6 multi-model ensemble is planned from mid-2027. Resolution depends on the hazard and on the geography.

HAZARD	RESOLUTION (CORDEX-BASED)
Coastal inundation	30 m to 0.5 m
Tropical cyclone storm surge	30 m to 0.5 m elevation, 450 m bathymetry
Riverine and surface water flooding	30 m to 5 m
Extreme wind, extreme heat, freeze-thaw	Climate grid cell (0.44° for CORDEX)
Forest fire	30 m
Soil movement	250 m
Tropical cyclone wind	0.25°
Landslide	1 km

TABLE 15. Hazards and resolutions covered by Morningstar's Sustainalytics

Provider Box – Swiss Re turns its natural catastrophe expertise into a set of climate risk scores

Swiss Re hazards fall into three families: a wet score for river and flash floods, a sea level rise score for storm surge and coastal flooding, and a dry score for heat, drought and wildfire.

Each score combines climate model outputs with Swiss Re's proprietary flood and storm surge zones at 30 m resolution. Non-linear transfer functions then translate changes in the mean and the extremes of temperature and precipitation into a dimensionless score, under SSP1-2.6, SSP2-4.5 and SSP5-8.5 and for near, medium and long-term horizons.

Because the scores have no unit, they rank sites and portfolios by relative exposure and point to hot spots and key perils without stating a loss. Monetary figures come in a second step, as a bespoke financial impact assessment. Swiss Re's natural catastrophe models produce loss costs for the perils the scores have identified, and the risk engineering team adds site-level mitigation recommendations and risk transfer solutions.

In a published case, a French insurer held close to 80% of its property portfolio in three regions, mostly Nouvelle-Aquitaine. The scores showed that river flood risk in the Gironde would stay stable over the coming decades, while coastal flood risk there would rise with sea level, which Swiss Re expects to become critical along the whole Atlantic coast by 2070. A single "flood" figure would have hidden this split between two hazards moving in opposite directions, and the insurer used the scores to choose where to diversify its portfolio. The purpose of this kind of approach is evidence that sometimes **simplifying the metric can sometimes mean simplifying the decision making**.

Provider Box – EQECAT / Corelogic / Cotality

EQECAT, later known as CoreLogic is a US property data and analytics company that rebranded as Cotality in March 2025. Its catastrophe models rest on a very large property database, with fields such as first floor height, roof age, construction materials and building codes, and on its own reconstruction cost and automated valuation models. The peril models cover flood, wildfire, severe convective storm, winter storm, earthquake and tsunami.

They run on the Risk Quantification and Engineering (RQE) platform, where clients can change exposures, damage functions and other inputs, and each peril relies on a stochastic event set of 300,000 simulated years. The Climate Risk Analytics product couples these catastrophe models with IPCC AR6 global climate models, which CoreLogic dynamically downscales under SSP1-2.6, SSP2-4.5, SSP3-7.0 and SSP5-8.5. It reports more than twenty financial measures for every structure in the continental United States, among them the AAL and several probable maximum losses (PML). These results can be aggregated to a company, a sector or a postcode.

The example of Texas, a Czajkowski model

This model sits at the frontier between industry and academia. The Wharton Risk Management and Decision Processes Center used CoreLogic's probabilistic flood model to price residential flood insurance in two Texas counties (Czajkowski et al., 2011). Galveston County, on the coast, was modelled for riverine flooding and storm surge. Travis County, inland around Austin, was modelled for riverine flooding only. The portfolio covered 226,407 single-family homes in Travis and 89,046 in Galveston. The model follows the usual chain of hazard,

exposure, vulnerability through a depth-damage curve, and financial loss. It produces an exceedance probability curve and an AAL for each home and each FEMA flood zone.

The pricing logic applies the compound distribution of this section directly. The pure premium of a home is its expected annual loss, $E[S] = \lambda \cdot E[X]$. The authors then add a loading for taxes, the cost of capital, shareholder returns and the correlation between wind and water losses.

$$P_{loaded} = (1 + \ell) \cdot AAL$$

In this model, P_{loaded} is the premium a private insurer would charge. AAL is the average annual loss of the home, which equals the pure premium. ℓ is the loading factor, tested at 50%, 100%, 200% and 300%.

Comparing these prices with National Flood Insurance Program (NFIP) rates showed that the public programme misprices risk in both directions. In Travis County the NFIP undercharged the high-risk A zones and overcharged the moderate and low-risk X500 and X zones. In Galveston it underpriced storm surge in the A, X500 and X zones and overpriced the V zones.

Where the NFIP overcharges, a private insurer could apply a loading close to 200% and still match its price. The tail explains much of the gap between the two counties.

A 10,000-year event produces about $2.3 billion of residential losses in Galveston from storm surge, against $890 million in Travis from riverine flooding. On average, the pure premium per $1,000 of exposure is more than twelve times lower in Travis. The same model was used to test home elevation. Raising homes by 8 feet cut the AAL by 92% in Travis and 82% in Galveston. For existing homes, though, the cost was high (about $58,000 to raise a 2,000 sq ft house by 2 feet), so elevation was rarely cost-effective on average. The study shows how a catastrophe model's AAL becomes a price, and how the loading carries the tail risk that the average leaves out.

Model Representation Card – Compound Poisson / Actuarial Catastrophe Model of Corelogic

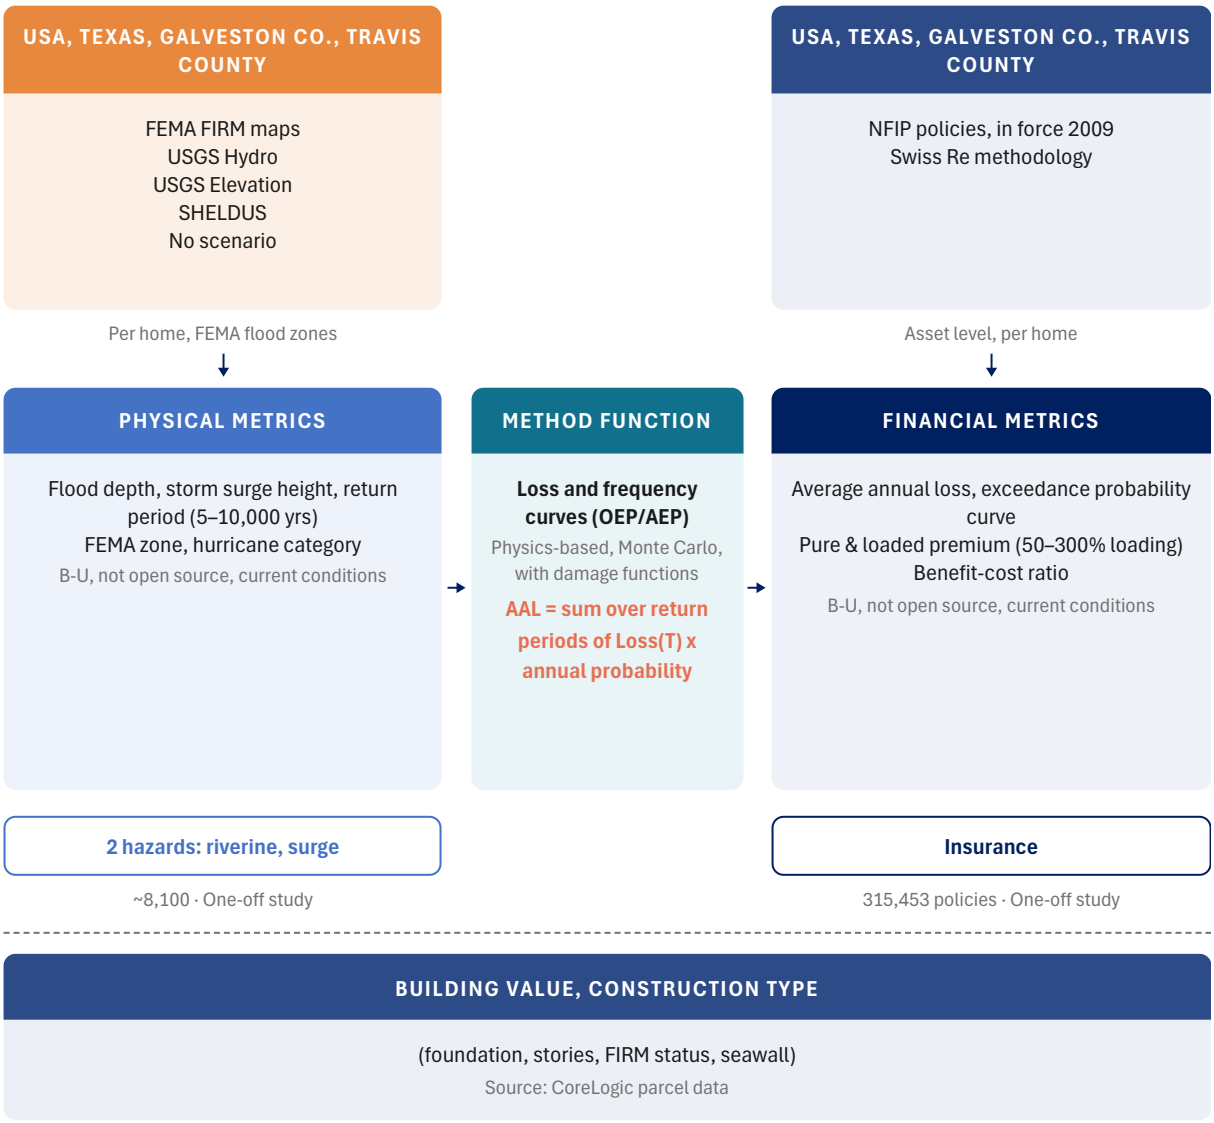


FIGURE 23. Model representation card – Compound Poisson / actuarial catastrophe model (Texas, Czajkowski et al., 2011)

Why is this model useful?

This model can be used if the goal is to price an actual insurance product or testing whether a mitigation measure is worth its cost. It fits situations where you need the full range of event severities, from a five-year flood up to a ten thousand-year one, because pricing an insurance premium requires knowing the whole distribution of possible losses, not just an average.

What does this model teach us?

This model teaches us that an average annual loss number on its own is not enough for insurance, what actually drives a premium is the full exceedance probability curve behind it. It also shows that the same underlying model can directly answer a policy question, is raising a house on stilts worth it, by expressing the mitigation's benefit as a ratio against the loss it avoids across that whole curve. And it is a reminder that even a model producing sophisticated actuarial output can be built on a fairly standard combination, physics-based hazard simulation, Monte Carlo sampling across return periods, and a conventional damage function underneath it all.

Flood Zone*	# of Single-Family Residences	Total Insured Value**	Total Premium	Average NFIP Premium	Average Premium per $1000 Insured Value	Minimum Premium per $1000 Insured Value	Maximum Premium per $1000 Insured Value
V	2,938	$ 704,947,600	$ 9,751,336	$ 3,319	$ 14.17	$ 2.62	$ 73.47
A	16,286	$ 3,297,771,000	$ 12,968,810	$ 796	$ 5.12	$ 0.21	$ 118.00
X500 / B	5,988	$ 1,237,473,000	$ 1,923,657	$ 321	$ 1.92	$ 0.93	$ 17.10
X / C	43,432	$11,335,174,200	$ 14,048,862	$ 323	$ 1.44	$ 0.86	$ 15.47
Galveston County Total	**68,644**	**$16,575,365,800**	**$ 38,692,665**	**$ 564**	**$ 2.90**	**$ 0.21**	**$ 118.00**

* The NFIP rate structure does not distinguish between X500/B and X/C flood hazard areas, and charges the same premium to structures in both of those areas.
** NFIP Insured Value = Total Building + Total Content insured values. Deductibles are not subtracted.

FIGURE 24. Galveston County 2009 NFIP Premium Data

Source: Czajkowski, J., et al., 2011

Key Limitations

- **Stationarity assumption.** Classical actuarial models calibrate λ and F from historical data assuming the underlying process is stationary. Climate change requires a non-stationary approach, for example Atlantic hurricane frequency is projected to change under warming, European flood recurrence intervals are shortening, and wildfire seasons are lengthening. Any actuarial model that does not explicitly adjust $\lambda(t)$ and F for climate change is implicitly assuming a world that no longer exists.
- **Tail extrapolation uncertainty.** The GPD tail index ξ is estimated from a handful of extreme observations. Small changes in ξ produce large changes in return levels at 200+ year periods. A 95% confidence interval on ξ often spans $[0.1, 0.7]$, which translates to a 500-year loss estimate ranging over an order of magnitude.
- **Independence of events.** Compound Poisson assumes event independence. Compound events (for example the apparition of a hurricane and then a storm surge and then an inland flood) and sequence dependencies (soil saturation from one flood increasing the severity of the next) make this approach obsolete. Copula-based approaches partially address it but add parametric uncertainty.
- **Model risk: distribution choice.** Log-normal and GPD can produce very different tail estimates from the same data. In practice, cat modellers often run multiple distributions and use expert judgment to select.

2.1.3 LOSS AND FREQUENCY CURVES

Loss and frequency curves are the standard language in which catastrophe risk is communicated between insurers, reinsurers, investors, and regulators.

They are not the canonical output format of the actuarial model described above: a complete, auditable representation of the loss distribution expressed as a function of either frequency (how often) or return period (how rare). A core competency for any practitioner working with physical climate risk in the insurance industry is to be able to read, construct, and critically evaluate these curves. To better understand what they are, we need to start with looking at four “canonical curves”.

The Four Canonical Curves

A) Occurrence Exceedance Probability (OEP)

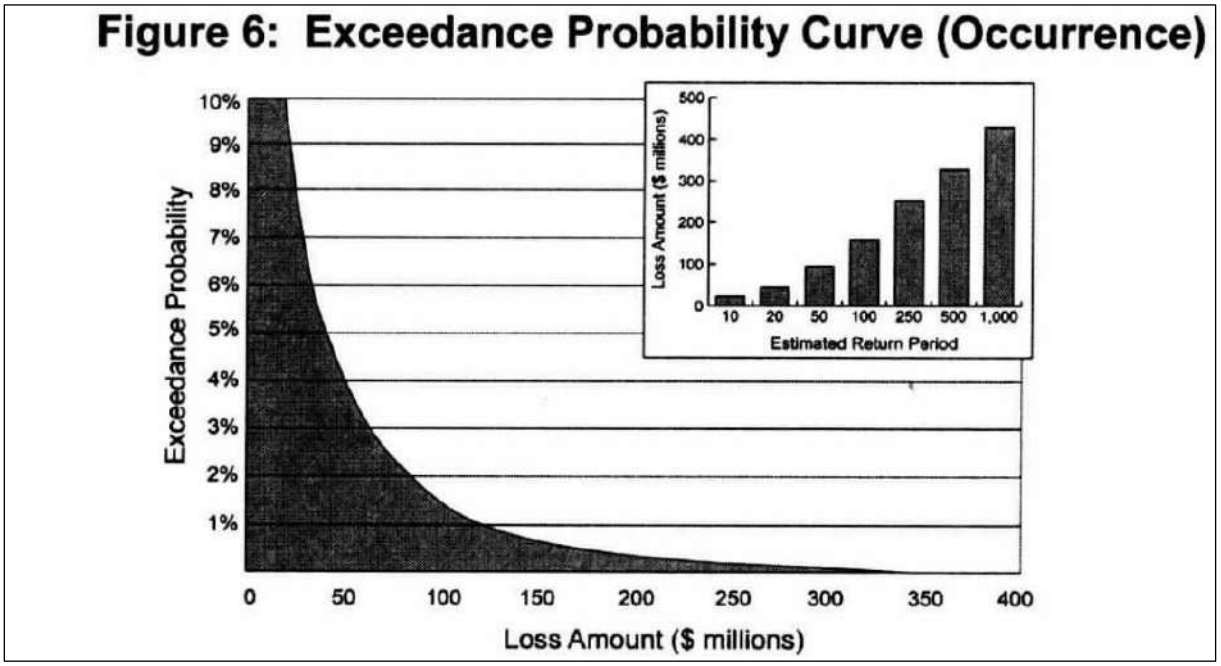


FIGURE 25. An OEP curve example

Source: "The Aggregation and Correlation of Reinsurance Exposure." Klinker, Frederick L, Glenn G Meyers, and David A Lalonde. CAS Forum 2003 (May).

The OEP curve gives, for each loss level L, the annual probability that at least one event in the year produces a loss exceeding L.

For a fixed set of events, each occurring once in the year, and independently of the others, the formula is as follows:

$$\boldsymbol{OEP(L) = 1 - \prod_{e} F_e(L)}$$

- $\boldsymbol{e}$ is the index running over each individual event in the stochastic event set for the year.
- $\boldsymbol{F_e(L)}$ is the probability that event e, if it occurs, produces a loss less than or equal to L, in other words the cumulative distribution function of event e's severity evaluated at L.

For a random number of events N,

$$\boldsymbol{OEP(L) = 1 - PGF_N\left(F_X(L)\right)}$$

Under the Poisson assumption, this becomes $OEP(L) = 1 - \exp\left(-\lambda\left(1 - F_X(L)\right)\right)$, i.e. for an ELT,

$$\boldsymbol{OEP(L) = 1 - \exp\left(-\Sigma_e \lambda_e \left(1 - F_e(L)\right)\right)}$$

- where $\boldsymbol{\lambda_e}$ is the annual occurrence rate of event e, with $\Sigma_e \lambda_e = \lambda$
- Here, F_X is the cumulative distribution function of the loss X of a single event drawn at random from the event set, and

$$\boldsymbol{PGF_N(z) = E[z^N]}$$

is the probability generating function of the annual event count N, with z ∈ [0,1] its argument; for a Poisson count with mean λ,

$$PGF_N(z) = \exp\left(-\lambda(1 - z)\right)$$

Note that the area under the OEP curve is the expected annual maximum loss, not to be confused with the AAL, which is the area under the AEP. The OEP curve is the relevant metric for per-occurrence reinsurance treaties and for regulatory stress tests that ask: "What is the single worst event we might face in a given year?"

Caution: because the OEP only retains the largest event of each year, it can be distorted by the clustering effects discussed earlier. For instance, the major hail episodes of 2022 included three extreme events: the strongest one masks the other two in the OEP, although these were as severe as (if not more severe than) the largest past events (Ela 2014, Felix 2009). Adding all 2022 events strongly shifts return periods at lower loss levels, whereas their removal has less impact on the OEP.

Under a negative binomial frequency,

$$OEP(L) = 1 - PGF_{NB}\left(F_X(L)\right)$$

Where:

$$PGF_{NB}(z) = \left(1 + \frac{\lambda}{r}(1 - z)\right)^{-r}$$

is the probability generating function of a negative binomial count with mean λ and dispersion parameter $r > 0$. The variance of N is $\lambda + \frac{\lambda^2}{r}$, so a small r means strong clustering, and the Poisson case is recovered as $r \to \infty$.

For the same mean frequency, the OEP is lower than under Poisson.

B) Aggregate Exceedance Probability (AEP)

The AEP curve gives the probability that the sum of all losses S in the year exceeds L:

$$AEP(L) = P(S > L)$$

The AEP curve lies on or above the OEP curve at every loss level, which places it on or to the right of the OEP curve at every probability level (annual aggregate losses are at least as large as the largest single event). The gap between the two curves quantifies the contribution of multiple events to total annual losses. AEP is the relevant metric for aggregate stop-loss treaties, annual results volatility, and portfolio-level capital requirements.

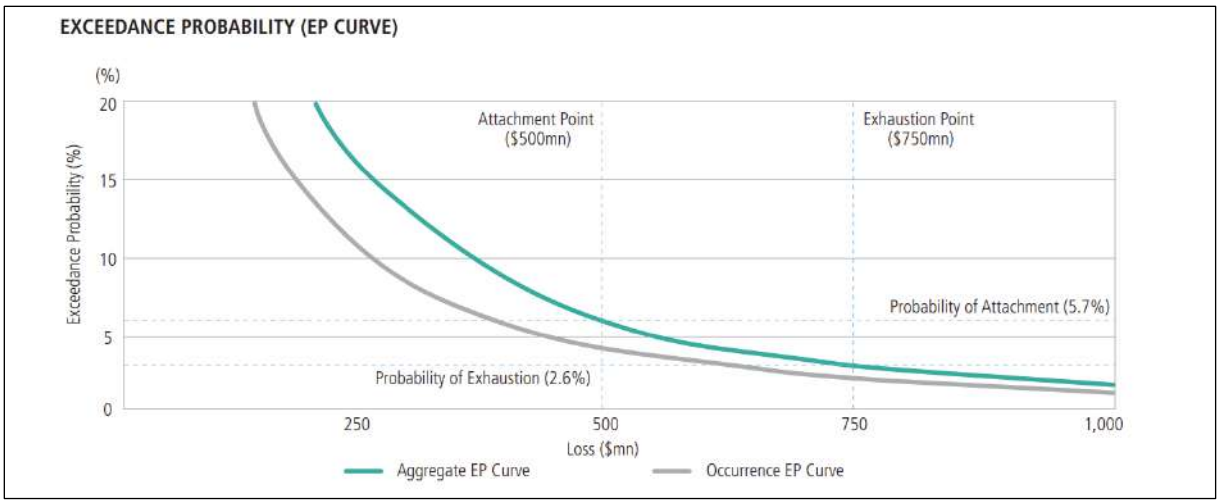


FIGURE 26. An EP Curve example

Source: Catastrophe Modeling 101, O'Rourke and Burdick, 2026, Neuberger Berman.

C) Average Annual Loss (AAL) curve

AAL is the first moment of the aggregate loss distribution:

$$\boldsymbol{AAL} = \boldsymbol{E[S]} = \boldsymbol{\lambda} \cdot \boldsymbol{E[X]} = \int_0^{\infty} \boldsymbol{\left(1 - F_S(x)\right) dx} = \int_0^{\infty} \boldsymbol{AEP(x)\, dx}$$

The last equality shows that AAL is the area under the AEP curve, which gives an intuitive geometric interpretation: the AAL is the "average height" of the loss exceedance function.

- $\boldsymbol{AAL}$ is the average annual loss, the single number insurers use as the pure premium, the expected cost of risk before any loading.
- $\boldsymbol{E[S]}$ is the expected value of the aggregate annual loss S, the same S as in the AEP formula.
- $\boldsymbol{\lambda}$ is the mean annual event rate, which is the same $\boldsymbol{\lambda}$ from the Poisson frequency process above.
- $\boldsymbol{E[X]}$ is the expected value of the severity distribution, the average loss size per event.
- $\boldsymbol{x}$ is the integration variable, running over every possible loss value from 0 to infinity.
- $\boldsymbol{F_S(x)}$ is the cumulative distribution function of the aggregate loss S evaluated at x, so $1 - F_S(x)$ is the survival function, i.e. the probability that S exceeds x. This is equal to the AEP.
- An AAL curve plots the AAL against a driver such as the global warming level and shows how the expected annual cost of risk evolves as the driver changes.

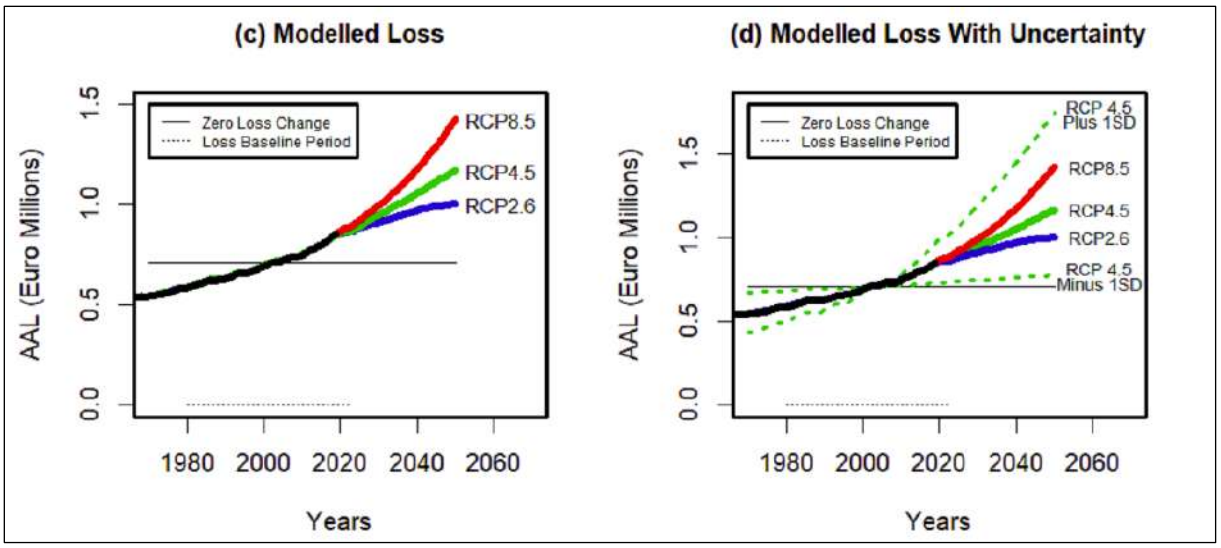


FIGURE 27. Example of an AAL curve

Source: Jewson et al., 2023, A service to help insurers understand the financial impacts of changing flood risk in Europe, based on PESETA IV, Climate Services

D) Return period (RP) curve

The return period T corresponding to a loss level L is simply:

$$T = \frac{1}{OEP(L)} \text{ (or } \frac{1}{AEP(L)} \text{ for aggregate).}$$

A '100-year loss' is the loss with a 1% annual exceedance probability. The Probable Maximum Loss (PML) at return period T is:

$$PML_T = OEP^{-1}\left(\frac{1}{T}\right) = VaR_\alpha(X_{max}) \quad where \quad \alpha = 1 - \frac{1}{T}$$

- Where $\boldsymbol{T}$ is the return period, in years, how rare the event of interest is (a "100-year loss" means T = 100).
- $\boldsymbol{PML_T}$ is the Probable Maximum Loss at return period T, the loss level you would expect to see exceeded with an annual probability of $\frac{1}{T}$ (annual convention, $T = \frac{1}{OEP}$).
- This differs from the event-based definition $T = \frac{1}{\Lambda(L)}$, the mean time between two events exceeding L; since $OEP = 1 - e^{-\Lambda}$, both converge for long return periods.
- $\boldsymbol{OEP^{-1}}$ is the inverse of the OEP function, instead of putting in a loss level and getting a probability out, you put in a probability (here $\frac{1}{T}$) and get the corresponding loss level out.
- $\boldsymbol{VaR_{alpha}}$ is the Value at Risk at confidence level alpha, the standard risk metric, the loss level that won't be exceeded with probability alpha.
- $\boldsymbol{X_{max}}$ is the per-occurrence maximum loss variable.
- $\boldsymbol{\alpha}$ is the confidence level, $1 - (\frac{1}{T})$so a 100-year return period corresponds to α = 0.99, a 99% confidence level.

In practice, consider the following key calibration points:

- PML_{100} (common market benchmark),
- PML_{200} (Solvency II reference, 99.5% VaR),

- PML_{250} (common rating agency benchmark),
- PML_{500} and above (extreme scenario stress).

As mentioned in the previous section, the Solvency II SCR is calibrated at the 200-year level (99.5th percentile) of the one-year loss distribution, which covers all events of the year. An internal model therefore reads the 200-year PML on the AEP, while the 200-year OEP remains the reference for per-occurrence reinsurance.

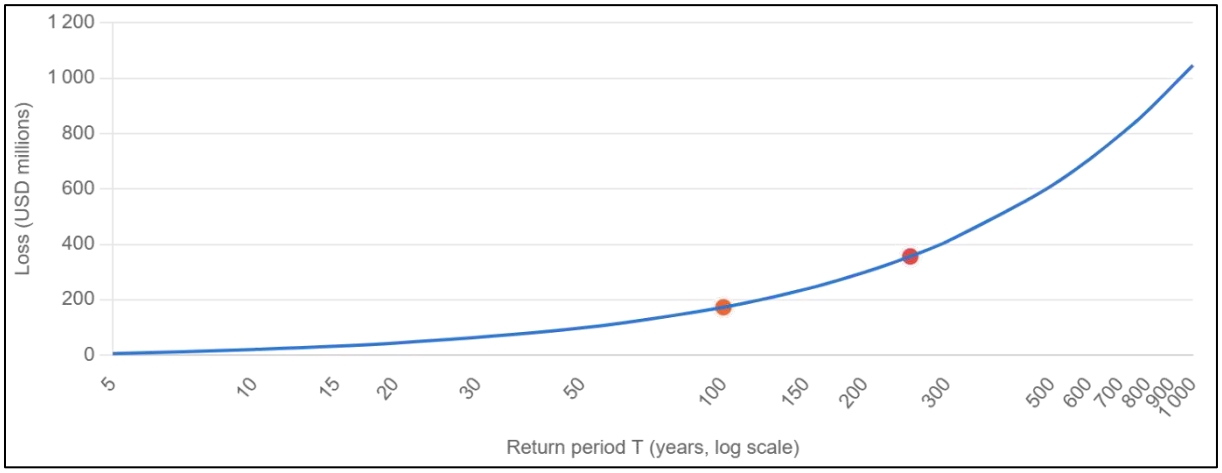


FIGURE 28. An example of an RP curve (log scale)

Mathematical Representation

OEP AND AEP CONSTRUCTION FROM A STOCHASTIC EVENT LOSS TABLE (ELT)

Given an Event Loss Table with M simulated events, each characterised by rate λ^e and loss X^e, the OEP and AEP curves are constructed as follows:

Step 1: OEP via order statistics.

$$OEP(L) = 1 - exp\left(-\sum_{e:\, X_e > L} \lambda_e\right)$$

$$\text{with } \Lambda(L) = \Sigma\, \lambda^e$$

- This is exact for a Poisson event set without secondary uncertainty, i.e. when each event loss is deterministic. The probability that no event exceeds L is $e^{-\Lambda(L)}$, where:

$\Lambda(L) = \Sigma_{e:\, X_e > L}\, \lambda_e$ is the total rate of events exceeding L.

Step 2: AEP via Monte Carlo.

Simulate Y years (typically $Y = 100{,}000$).

For each simulated year y: draw $N_Y \sim Poisson(\Lambda_{total})$ events; where $\boldsymbol{\Lambda(L) = \Sigma_{e:\, X^e > L}\, \lambda^e}$ is the total rate of events exceeding L.

For each event, draw its identity from the event frequency weights and its loss $X_e \sim F_e$ (secondary uncertainty), then sum losses to get S_Y.

Sort S_Y and read off the observed quantiles.

OEP AND AEP CONSTRUCTION FROM A STOCHASTIC EVENT LOSS TABLE (ELT)

Step 3: Confidence intervals.

The PML estimates of Step 2 come from a finite number of simulated years, so they carry sampling error. Following the annual convention of the RP curve section, the PML at return period T is the loss whose annual exceedance probability equals $\frac{1}{T}$. Its 95% confidence interval is roughly approximated by:

$$\mathrm{PML}_T \pm 1.96 \times \frac{\mathrm{PML}_T}{\sqrt{Y\,p_T}} = \mathrm{PML}_T\left(1 \pm \frac{1.96}{\sqrt{k}}\right), \qquad p_T = \frac{1}{T}, \qquad k = \frac{Y}{T}$$

- $\boldsymbol{PML_T}$ is the PML estimated at return period T.
- $\boldsymbol{Y}$ is the number of simulated years.
- $\boldsymbol{p_T}$ is the annual exceedance probability of PML_T
- $\boldsymbol{k}$ is the expected number of simulated years with a loss beyond PML_T
- **1.96** is the 97.5% quantile of the standard normal distribution.

This approximation rests on two assumptions.

- First, it assumes:

$$\left|\frac{d \ln L}{d \ln p}\right| = 1$$

  $\boldsymbol{L}$ is the loss level.

  $\boldsymbol{p}$ is the annual exceedance probability of L.

  In other words, a 1% change in p moves L by 1%. This rarely holds.

- Second, it relies on the asymptotic normality of empirical quantiles.

  That holds only if the loss density is positive at the target quantile and k is large (around 200). At $k = 200$, the half-width of the interval is about $1.96/\sqrt{200} \approx 14\%$ of the PML.

  That holds only if the loss density is positive at the target quantile k is large (around 200). A rank-based (binomial) interval or a bootstrap over simulated years is therefore preferred. With Y = 100,000, results beyond $T = 500$ (k = 200) should not be relied upon. $T = 200$ is the usual regulatory requirement.

Return period conversion

$$\mathrm{TVaR}_\alpha = \mathbb{E}[S \mid S > \mathrm{VaR}_\alpha(S)] = \mathrm{VaR}_\alpha(S) + \mathbb{E}[S - \mathrm{VaR}_\alpha(S) \mid S > \mathrm{VaR}_\alpha(S)]$$

S is the aggregate annual loss,

$VaR_\alpha\,(S)$ is the loss level that S does not exceed with probability α, and the second term is the mean excess of S over that level, i.e. the average amount by which losses exceed VaR_α (S) when they do.

$Tail\ Value - at - Risk$ (TVaR), also called Conditional Tail Expectation (CTE) or Expected Shortfall (ES), is a coherent risk measure, unlike VaR.

OEP AND AEP CONSTRUCTION FROM A STOCHASTIC EVENT LOSS TABLE (ELT)

It is the reference measure of the Swiss Solvency Test (99% Expected Shortfall), whereas Solvency II relies on the one-year 99.5% VaR. $TVaR$ captures the full severity of tail losses beyond the threshold, which makes it better suited to capital adequacy assessment.

Reading and Constructing EP Curves: Step-by-Step

Step 1 – Obtain or construct the Event Loss Table (ELT).

The ELT is a list of $\{event_id, rate\ \lambda^e, mean_loss\ X^e, loss_cv\ \beta^e\}$ tuples representing the stochastic event set.

Sources include vendor models (Moody's RMS, Verisk, formerly AIR), the open-source Oasis LMF, or a table built manually from historical nat-cat data with Monte Carlo augmentation.

Step 2 – Sort events by loss and compute OEP. Sort events in descending order by X^e. Compute the cumulative rate $\Lambda(L) = \Sigma\ \lambda^e$ for all events with $X^e > L$.

$$OEP(L) = 1 - e^{-\Lambda(L)}$$

Plot OEP(L) on the y-axis against L on the x-axis: this is your OEP curve.

Step 3 – Read off PMLs. The PML at return period T is the loss L_T where:

$$OEP(L_T) = \frac{1}{T}$$

For T = 200 (Solvency II reference): find L where OEP = 0.005, or read the same probability on the AEP for an aggregate capital view. With fewer than 100,000 simulated years, the T = 500 estimate is not robust.

Similarly:

- For T = 100: find L where OEP = 0.01.
- For T = 250: OEP = 0.004.
- For T = 500: OEP = 0.002.

Step 4 – Compute the AEP via simulation. Simulate $Y = 100{,}000$ years as described. Sort S_y and read quantiles. The AEP curve lies on or to the right of OEP (higher or equal loss at the same probability level) because aggregate losses include multiple-event years.

Step 5 – Compute the AAL as area under the AEP.

Numerically integrate the AEP curve:

$$AAL = \int_0^\infty AEP(L)\, dL \approx \sum_i AEP(L_i) \cdot \Delta L_i$$

Cross-check:

$$AAL = \sum_e \lambda_e \cdot E[X_e]$$

from the ELT.

Step 6 – Apply secondary uncertainty.

Secondary uncertainty (β^e in the ELT) is the dispersion of the loss of a given event around its mean X^e Each event loss then follows a distribution with mean X^e and coefficient of variation β^e , often a beta distribution in vendor models. The OEP of Step 2 becomes

$$\mathbf{OEP}(L) = 1 - \exp\left(-\sum_e \lambda^e \left(1 - F_e(L)\right)\right)$$

Where:

- $\boldsymbol{\lambda^e}$ is the annual rate of event e
- $\boldsymbol{F_e(L)}$ is the cumulative distribution function of the loss of event e evaluated at L.
- $\boldsymbol{X^e}$ is the mean loss of event e.
- $\boldsymbol{\beta^e}$ is the coefficient of variation of the loss of event e, its standard deviation divided by X^e

Without secondary uncertainty, $1 - F_{e(L)}$ equals 1 if $X^e > L$ and 0 otherwise, which gives back the formula of Step 2. In the simulation of Step 4, each event loss is drawn from F_e instead of being set to X^e.

Secondary uncertainty leaves the AAL unchanged, since each event keeps the same mean loss. It redistributes probability toward the tails, so high return period losses on both the OEP and AEP curves exceed those of a mean-loss-only computation.

Portfolios with high β^e (sparse exposure data, uncertain vulnerability functions) show significantly fatter tails than the mean-loss ELT suggests. This matters most for pricing excess of loss layers attached high in the tail, where a mean-loss ELT understates the probability of the layer being hit.

Climate Non-Stationarity: Adjusting EP Curves

The most important practical challenge for EP curves under climate change is that the event rates λ^e are no longer constant. Two adjustment approaches are standard:

Frequency-only adjustment.

The first thing you can do is scale each event rate by a climate multiplier:

$$\lambda^e(t) = \lambda_0^e \times \mathbf{m_{freq}}(\mathbf{hazard}, \mathbf{region}, t, \mathbf{scenario})$$

λ_0^e is the adjusted annual rate of event e at horizon t, $\lambda_{e,0}$ is its baseline rate, and m_{freq} the frequency multiplier, which depends on the hazard, the region, the horizon and the climate scenario.

This shifts the OEP curve upward, with a higher exceedance probability at every loss level, so every PML moves to a higher loss. The shape of individual event losses stays unchanged. Moody's RMS publishes frequency multipliers for its climate change module; CLIMADA provides open-source equivalents calibrated to CMIP6.

Frequency + intensity adjustment.

Both rates and loss amounts change:

$$\lambda^e(t) \textbf{ and } X^e(t) = X^e{}_0 \times m_{int}(hazard, region, t, scenario)$$

$X^e(t)$ is the adjusted mean loss of event e at horizon t, $X^e{}_0$ its baseline mean loss, and m_{int} the intensity multiplier.

We inflate both the frequency and the severity of events, producing a more dramatic shift of the OEP/AEP curves. The two multipliers act differently along the curve. On the AAL they compound fully, so a 10% frequency increase and a 5% intensity increase raise the $AAL\ by\ 1.10 \times 1.05 - 1 \approx 15.5\%$.

In the tail, the intensity multiplier still scales losses one-for-one, but the frequency effect is damped. For a generalized Pareto tail and long return periods,

$$\frac{\mathrm{PML}_T^{\mathrm{adj}}}{\mathrm{PML}_T^{\mathrm{base}}} \approx m_{\mathrm{int}} \cdot m_{\mathrm{freq}}^{\xi}$$

where $\mathrm{PML}_T^{\mathrm{adj}}$ is the climate-adjusted PML at return period T, PML_T^{base} the PML under the baseline catalogue, and $\xi > 0$ the shape parameter of the generalised Pareto tail.

With $\xi = 0.3$, the same 10% and 5% multipliers raise PML_{250} by about: $1.05 \times 1.10^{0.3} - 1 \approx 8\%$, roughly half the AAL uplift.

The interaction is therefore sub-additive in the tail. For a lighter, exponential tail ($\xi = 0$), the frequency multiplier adds a constant amount to every return level, so its relative effect fades as T grows. On insured losses, deductibles and limits alter both effects further.

Provider Box – MunichRe loss frequency approach to Risk

Munich Re's loss frequency work rests on NatCatSERVICE, its natural catastrophe database maintained since 1974. It has formed the basis of Munich Re's risk management and research since that year, and sits alongside EM-DAT and Sigma as one of the three global reference loss databases, with over 30,000 recorded events.

This historical base has a known limit: traditional actuarial methods built on historical experience do not work for low frequency, high severity catastrophes. The company addresses this by pairing NatCatSERVICE with probabilistic catastrophe models, which output exceedance probability curves linking annual frequency to loss severity, built from large-scale Monte Carlo simulation.

The two feed each other: historical losses calibrate and validate the simulated curves. The provider also treats these curves as non-stationary, since climate change shifts the hazard side while exposure growth and changes in building vulnerability shift the exposure and vulnerability sides, and the two need to be separated before any trend is drawn.

Recent tools like Location Risk Intelligence push this down to portfolio and site level.

In 2025, global insured losses from natural catastrophes reached about 108 billion dollars without any hurricane making landfall in the United States. Wildfires, floods and severe convective storms accounted for about 98 billion dollars of this total, led by the Los Angeles wildfires. That is the case for frequency-severity curves broken out by peril and zone, not just aggregated regional totals.

The reinsurer sells this expertise through Location Risk Intelligence, a cloud platform that can be reached through a web application or an API. The platform is split into editions (Natural Hazards, Climate Change, Wildfire HD, Reporting and Climate Financial Impact). The Climate Financial Impact Edition is the one that turns hazard data into money. It delivers three metrics.

Climate Expected Loss (CEL). CEL is the expected annual loss caused by physical damage to a building and its contents. It is expressed as a rate of the total rebuilding cost. The calculation runs in three steps:

- Munich Re's hazard models give the intensity of the hazard in each grid cell for each return period, such as flood depth or peak gust.

- Vulnerability curves convert each intensity into a damage ratio. These curves vary by region to reflect building type, materials and local building codes.
- The result is an exceedance probability curve specific to the location, and CEL is its expected value, the area under that curve.

CEL is therefore the average annual loss described in Section 2.1. The simulated distribution includes extreme events that have never been observed, so the rate is not limited by the historical record. In the river flood example given in the brochure, CEL rises from 4.30‰ in the current climate to 5.11‰ in a future projection. CEL currently exists for four perils: tropical cyclone, river flood, storm surge and extratropical storm.

1-in-100 Year Damage. This metric gives the share of an asset's value destroyed by the local 100-year event, from 0% to 100%. It complements CEL by describing the tail of the distribution. It is currently available for standard residential buildings only.

1-in-100 Year Hazard Intensities. This metric gives the raw intensity of the 100-year event: flood depth for river flood and storm surge, and peak 3-second gust for cyclones and winter storms.

Each peril behind CEL relies on its own hazard model and its own approach to climate projection. For more information on perils, see section 3.1.1.

Model Representation Card – EP Framework

An interesting model that we used to illustrate AAL curves is based on aggregation by Jewson et al., 2023.

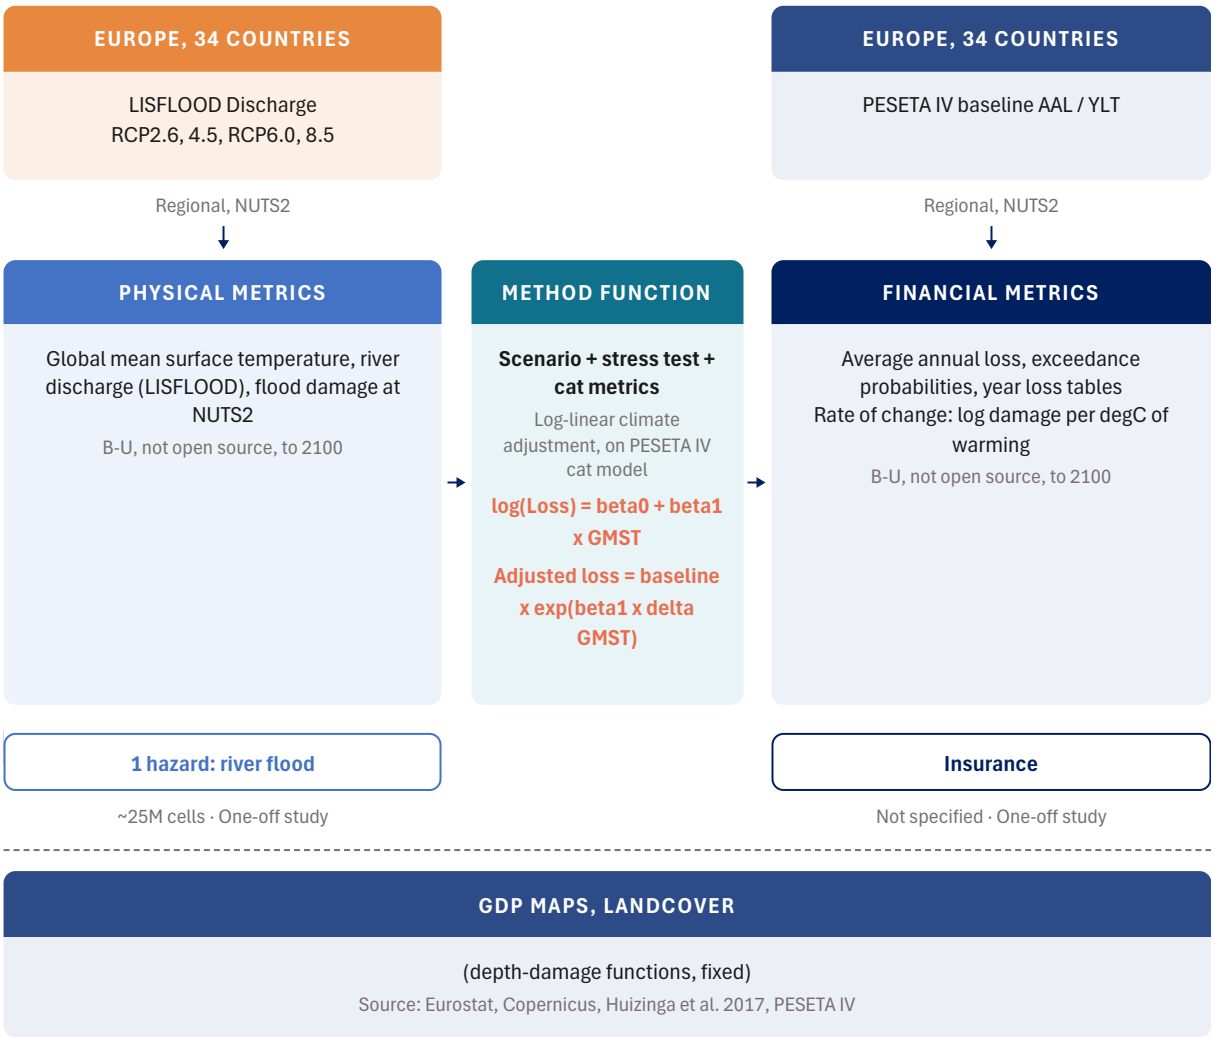


FIGURE 29. Model representation card – EP framework (log-linear warming scaling, Jewson et al., 2023)

Why is this model useful?

This model can be used if you need to know not just how much flood damage costs today, but how that cost shifts as warming increases, expressed in a form insurers already work with, average annual loss, exceedance probabilities, and year loss tables. It is a fit when you already have or trust an existing catastrophe model's baseline output and want a lightweight way to adjust it for climate change, which is less time consuming than building a new hazard or damage model from scratch. Because it works at NUTS2 regional resolution across 34 countries, it suits portfolio-level or regulatory questions more than a single-site assessment.

What does this model teach us?

This model teaches us that you do not need to rebuild an entire catastrophe model to make it climate aware, a single regression relating log damage to global mean surface temperature is enough to turn a static exceedance probability curve into one that shifts under different warming levels. It also shows that this kind of adjustment holds up best only over a limited warming range, the paper's R-squared figures degrade beyond about 1.2 degrees, which is a useful reminder that a simple climate multiplier is a reasonable near-term tool but that even sophisticated models lose power when they aim to extrapolate blindly into a much hotter world.

Key Limitations

- **Tail uncertainty dominates at high return periods.** The difference between a 200-year and a 500-year PML estimate is driven almost entirely by the assumed shape of the tail distribution above the last observed historical event. For most peril-region cells, there are zero or one observations above the 100-year level, making the entire 200–500-year range a pure extrapolation.
- **OEP vs. AEP divergence.** For portfolios concentrated in a single geography and peril, OEP ≈ AEP (most annual losses come from a single dominant event).
- **For diversified global books.** AEP markedly exceeds OEP at short and moderate return periods because many moderate events accumulate. The two curves converge at long return periods, where one large event dominates the annual total. Treaty structures designed around OEP provide inadequate protection for multi-event years: the 2011 year (New Zealand, Japan, Thailand, US tornadoes, Australia floods) was an AEP-driven loss year, not an OEP-driven one.
- **Model-specific curves.** OEP/AEP curves from RMS, AIR, and KCC for the same portfolio can diverge by 30–50% at the 250-year level due to differences in hazard model, vulnerability functions, and event catalogue construction. Using a single vendor's curve without cross-checking creates significant model risk.

2.1.4 REAL EFFECTS: AN APPROACH TO PHYSICAL RISKS TO GO BEYOND THE FINANCIAL IMPACTS

Most physical climate risk models focus primarily on monetary losses: damage ratios, expected annual losses, insured claims, balance-sheet write-downs. This is a necessary simplification, as finance can be seen as the “first defence line” against these events (i.e. if the financial impact is low, chances are that the human impact will be low as well), but it is incomplete.

Physical climate hazards also generate a wide range of non-financial consequences that affect individuals, communities, labour systems, and public health infrastructure. These consequences are collectively described in the literature as "real effects". They frequently persist long after the triggering event has passed and cannot be captured by accounting losses alone.

Floods, heatwaves, storms, droughts, and wildfires produce direct effects on human health, psychological wellbeing, labour productivity, displacement patterns, and mortality.

Heatwaves raise mortality rates among vulnerable populations and reduce outdoor labour productivity across entire regions. Floods generate long-term psychological trauma, housing insecurity, and chronic stress disorders. Wildfires cause respiratory damage through prolonged smoke exposure.

Extreme weather events disrupt healthcare systems, transportation networks, education access, and essential services.

This dimension connects with the "real effects" literature developed in empirical accounting and finance research from the mid-2010s onwards. That literature demonstrated that corporate disclosure requirements, regulatory constraints, and supervisory oversight produce measurable effects on real economic behaviour beyond financial reporting (for a good example looking at the link between firm disclosures, So2, life expectancy and firm profitability, see Chen et al., 2018).

Applied to physical climate risk, the implication is that climate disclosure frameworks and regulatory stress-testing exercises may influence not only what firms report, but how they organize their operations. Disclosure pressure has been associated with firms taking concrete adaptation measures (Li, 2025), including:

- Relocating facilities out of high-exposure flood or wildfire zones
- Strengthening adaptation infrastructure (flood barriers, cooling systems, backup power)
- Improving emergency planning and business continuity frameworks
- Reducing employee exposure to hazardous heat or air-quality conditions
- Modifying supply chain structures to reduce concentration in climate-vulnerable regions

Life and health insurers are directly exposed to the human consequences of climate hazards through mortality, morbidity, mental health, and long-term healthcare costs.

Heatwaves, air pollution, infectious disease dynamics, and natural disasters progressively alter mortality tables and health expenditure structures.

Extreme climate events generate claims across a wide range of lines: hospitalization, respiratory and cardiovascular conditions, psychological trauma, disability, occupational injuries, population displacement, and chronic mental health disorders.

Physical climate risks therefore generate financial effects and human effects simultaneously, and in most real-world events the two are deeply interconnected.

A severe heatwave may at the same time raise mortality rates among the elderly, reduce labour productivity in outdoor industries, increase electricity demand and grid stress, trigger insurance claims, impair loan collateral values, and slow macroeconomic activity. Purely financial risk models that capture only the asset-damage component of such an event will systematically underestimate the full social cost.

A growing body of research therefore argues for integrating broader human, health, and social dimensions into physical climate risk analysis to ensure that the models deployed by insurers, banks, and regulators reflect the full spectrum of climate-related vulnerabilities (Seklecka et al., 2017; Boumezoued et al., 2022; Bhattachaya et al., 2024; Aprea et al., 2025).

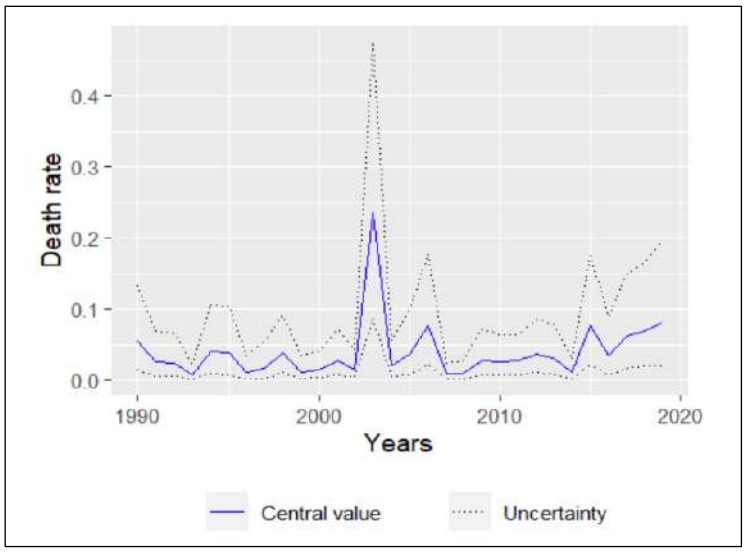


FIGURE 30. French Death Rate attributable to high temperature exposure (both sexes, France, for 100,000)
Source: Boumezoued et al., 2022 (Milliman); the data is from the Global Burden of Disease (GBD) database

In many papers, the Lee-Carter's model is being used. We will not dive into too many details of these kinds of models, as we will focus mainly on financial impacts. But this literature is worth mentioning as it is becoming increasingly important to understand that there will be a shock beyond the financial impact of climate change and climate-related events.

Beyond literature, the model is also actually used on the field.

Provider Box – Milliman's approach to mortality

A Lee-Carter mortality model can be combined with a climate index (Boumezoued et al., 2022). For France, the index is built from the number of summer days above 35°C and above 40°C, and calibrated on mortality data from 1990 to 2018. The climate effect becomes significant above age 65. Mortality shocks computed with the EIOPA method come out on average 6.12% larger than those from a standard Lee-Carter model.

Milliman extended the standard mortality models used by life insurers so that they carry a climate signal. The starting point is the Lee-Carter model, which writes the log death rate at each age as an age profile plus a common time trend weighted by age (Lee & Carter, 1992):

$$\ln(m_{x,t}) = a_x + b_x k_t + \varepsilon_{x,t}$$

- where $m_{x,t}$ is the death rate at age x in year t,
- a_x is the average log death rate at age x,
- k_t is the period index that drives the general fall in mortality over time,
- b_x measures how strongly age x responds to that index, and $\varepsilon_{x,t}$ is the error term.

Milliman adds a heat-wave component to this structure, so that years and regions with more intense heat waves produce mortality shocks on top of the long-term trend. The method was first applied to France and the United States, alongside an alternative approach for vector-borne diseases.

A later study derives mortality scenarios directly from national climate projections. For the Netherlands, it uses the KNMI scenarios of the Dutch meteorological institute. Its variables cover average, maximum and night-time temperatures, the number of tropical days above 30°C, and air pollution.

Under the high-emission scenario with dry summers, tropical days rise from fewer than 10 a year today to about 50 by 2100. Life expectancy at age 65 is then 0.7 years lower in 2050 than in the baseline. A life insurer or a pension fund can feed these mortality shifts into its reserves and stress tests, in the same way that a property insurer feeds a damage curve into its loss estimates.

Milliman adds a heat-wave component to this structure, so that years and regions with more intense heat waves produce mortality shocks on top of the long-term trend. The method was first applied to France and the United States, alongside an alternative approach for vector-borne diseases.

A later study derives mortality scenarios directly from national climate projections. For the Netherlands, it uses the KNMI scenarios of the Dutch meteorological institute. Its variables cover average, maximum and night-time temperatures, the number of tropical days above 30°C, and air pollution.

Under the high-emission scenario with dry summers, tropical days rise from fewer than 10 a year today to about 50 by 2100. Life expectancy at age 65 is then 0.7 years lower in 2050 than in the baseline. A life insurer or a pension fund can feed these mortality shifts into its reserves and stress tests, in the same way that a property insurer feeds a damage curve into its loss estimates.

2.2 PROBABILISTIC AND DEFAULT BASED MODELS

The preceding sections described how physical risk is quantified as an insurance or engineering loss: a damage ratio, an annual expected loss, a PML. The probabilistic and default-based family of models asks a very different question for a different sector:

When physical climate hazards affect the cash flows, asset values, or operating conditions of a firm or household, how does that translate into their probability of defaulting on a financial obligation?

This is the translation layer between the physical risk world and the credit risk world. It connects the output of cat models and actuarial frameworks to the balance sheets, rating methodologies, and capital adequacy frameworks of banks, bond investors, and prudential supervisors. It is also, as of 2026, one of the most actively developed areas in climate finance, with central banks (ECB, DNB, Banque de France), the NGFS, the Basel Committee, and dozens of academic groups publishing new methodologies each year.

2.2.1 PHYSICAL RISKS AND PROBABILITY OF DEFAULT

The probabilistic and default-based family of models asks a fundamentally different question: when physical climate hazards affect the cash flows, asset values, or operating conditions of a firm or household, how does that translate into their probability of defaulting on a financial obligation?

The translation layer between the physical risk world and the credit risk world is not easy to make, as many models are “internal”. It connects the output of cat models and actuarial frameworks to the balance sheets, rating methodologies, and capital adequacy frameworks of banks, bond investors, and prudential supervisors. It is also, as of 2026, one of the most actively developed areas in climate finance, with central banks (ECB, DNB, Banque de France), the NGFS, the Basel Committee, and dozens of academic groups publishing new methodologies each year.

Consider for example an interesting take on climate-related probability of default in the agricultural sector (Blanc Blanc-Blocquel et al., 2024) through an experiment that an increase from an “initial PD” of 4% to a 6% PD due to physical climate risk. Their first model is fairly simple:

$$Stock = \beta 0 + \beta 1 \times SPX + \beta 2 \times TA + \varepsilon$$

Where:

- $Stock$ is the firm’s monthly stock value
- SPX is the S&P 500 index value for that same month
- TA is the temperature anomaly for that same month, in Celsius
- $\beta 0$ is the intercept, the stock value when SPX and TA are both zero
- $\beta 1$ is how much the stock moves for each one-unit move in SPX
- $\beta 2$ is how much the stock moves for each one-degree move in TA
- ε is the part of the stock’s movement not explained by SPX nor TA.

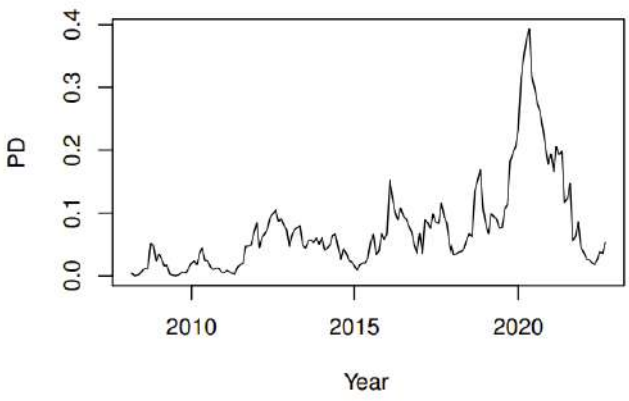


FIGURE 31. The link between default probabilities and temperature anomalies
Source: Blanc-Blocquel et al., 2024 (fig. 7)

This figure shows the resulting default probability over time, computed from the firm's real stock and debt data using this approach: a baseline PD around 4%, climbing to 6% once forecasted temperature anomalies are factored in.

From Physical Loss to Financial Default

The logical chain connecting a physical hazard to a credit event runs through several intermediate steps, each of which can be modelled explicitly or assumed away:

TRANSMISSION CHAIN: A 5-STEP TRAGEDY
1. Physical hazard event → asset damage / operational disruption (cat model output)
2. Asset damage → revenue loss, repair capex, insurance gap (financial impact model)
3. Revenue loss / repair costs → deterioration of debt service coverage, EBITDA compression, LTV deterioration (accounting model)
4. Financial deterioration → credit rating downgrade, covenant breach, refinancing risk (credit model)
5. Covenant breach / cash exhaustion → default, restructuring, resolution (PD model)

Each event in this chain can be represented through a model. Most existing approaches collapse several steps, either implicitly or explicitly.

The simplest approaches (score-based PD adjustments) jump from step 1 to step 5 using a statistical correlation observed in historical data. The most rigorous approaches, but harder to compute (structural credit models with physical risk integration) will model each step separately and propagate uncertainty through the full chain.

A “simpler” approach to physical climate risks using LGD (Loss Given Default)

When physical climate hazards occur, one of the most straightforward ways to reflect their impact on credit risk parameters is through climate-related collateral value adjustments.

Climate Physical Risk (CPR) can be quantified as an additional haircut applied to collateral, or as a revision of Loan-to-Value ratios. The effect must be used in the Loss Given Default (LGD) by reducing the expected recovery value in the event of default. Doing this will capture the exposure of collateral assets to potential physical damage.

IMPACT OF A HAZARD ON LOSS GIVEN DEFAULT (LGD) AND LOAN-TO-VALUE (LTV): "JUST ADD THE CLIMATE"

A. The collateral channel: from hazard damage to LGD

CPR consequences can be applied directly as a haircut on collateral value, which mechanically raises the Loss Given Default (Pozdyshev et al., 2025).

A.1 Expected damage over a time horizon

Before collateral value can be adjusted, the expected physical damage itself has to be estimated. A hazard model gives a set of possible damage outcomes and their probabilities under scenario j.

The expected annual impact is the probability-weighted average of those outcomes:

$$EAI_j = \mathbb{E}[X] = \Sigma_i x_{ij} \times Pr(x_{ij})$$

x_{ij} is one possible damage outcome under scenario j.

$Pr(x_{ij})$ is its probability of occurring.

Everything that follows uses this expected damage, or a related loss ratio, as an input.

A.2 From collateral value to recovery rate

Once the expected damage is known, the recovery rate, the share of the exposure a bank can expect to recoup at default, is adjusted down for physical damage, and up for any mitigation already in place, such as insurance, adaptation measures or guarantees:

$$RR_j = CV_0 \cdot (1 - h - \delta_j + m_j) / EAD$$

- CV_0 is the initial value of the collateral, before any damage.
- h is the haircut already applied to the collateral for reasons unrelated to climate, the standard Basel haircut.
- δ_j is the loss in collateral value from physical risk under scenario j, usually taken as a % of collateral value derived from EAI_j.
- m_j is the mitigation effect under scenario j : insurance proceeds, adaptation measures, guarantees, or other protection that offsets part of the damage.
- EAD is the Exposure at Default, the amount outstanding at the time of default.

The $Loss\ Given\ Default$ (LGD) follows immediately as the complement of the recovery rate:

$$LGD_j = 1 - RR_j$$

B. Downstream effect on Loan-to-Value (LTV) and risk weights

A drop in collateral value also raises the Loan-to-Value (LTV) ratio, and under the Basel standardized CRE20 approach, LTV directly sets the regulatory risk weight applied to the exposure. A shock to collateral value therefore affects both the bank's internal LGD model and its capital reporting.

IMPACT OF A HAZARD ON LOSS GIVEN DEFAULT (LGD) AND LOAN-TO-VALUE (LTV): "JUST ADD THE CLIMATE"

B.1 Basel CRE20 risk-weight bands

Two schedules apply, depending on whether the loan relies on the property's own cash flow generation. LTV is defined as:

$$LTV = {EAD}/{Collateral\ value}$$

The risk weight then follows from the LTV band. With cash-flow generation:

Loan-To-Value (LTV)	Risk weight
≤ 50%	30%
50% to 60%	35%
60% to 80%	45%
80% to 90%	60%
90% to 100%	75%
> 100%	105%

TABLE 16. Basel III risk weights for residential real estate exposures that depend on the property's cash flows, by LTV band (CRE20)

Without cash-flow generation:

LTV	Risk weight
≤ 50%	20%
50% to 60%	25%
60% to 80%	30%
80% to 90%	40%
90% to 100%	50%
> 100%	70%

TABLE 17. Basel III risk weights for residential real estate exposures that do not depend on the property's cash flows, by LTV band (CRE20)

IMPACT OF A HAZARD ON LOSS GIVEN DEFAULT (LGD) AND LOAN-TO-VALUE (LTV): "JUST ADD THE CLIMATE"

B.2 An example

All values are in euros. Consider a loan of €240,000. It is secured by collateral worth 300,000. A flood scenario cuts collateral value by 20,000, a 6.7% haircut on top of any standard haircut:

Item	Original estimate	Flood-adjusted estimate	Change
Exposure at Default	240,000	240,000	0
Collateral value	300,000	280,000	−20,000
LTV	80.0%	85.7%	+5.7 pts
Basel risk weight	30%	40%	+10 pts
RWA	72,000	96,000	+24,000

TABLE 18. Effect of a €20,000 flood shock on collateral value, LTV, risk weight and RWA for a €240,000 loan

RWA rises from 72,000 to 96,000. This 33.3% increase is because the LTV band shifted from 80% to 85.7% with no change in exposure size. A modest collateral shock can therefore produce a large, discontinuous jump in capital requirements once it crosses some LTV threshold.

C. The structural channel: from asset-value shock to PD

The collateral channel above treats physical risk as a write-down on collateral, affecting recovery only. The structural approach goes further: it lets physical risk hit the borrower's asset value directly, and asks how that changes the probability of default itself. The three formulas below build up to that result step by step.

C.1 Starting point: the plain Merton probability of default

In the Merton model, a firm defaults when the value of its assets falls below its liabilities at the chosen horizon:

$$Default \Leftrightarrow A_i < L_i$$

Where:

$$A_i = K_i + L_i$$

- A_i is the value of firm i's assets, and L_i its liabilities, the default barrier.
- K_i is the firm's capital, the cushion of assets above the default barrier.

Under the standard log-normal assumption for asset values, this default condition is a closed-form probability:

IMPACT OF A HAZARD ON LOSS GIVEN DEFAULT (LGD) AND LOAN-TO-VALUE (LTV): "JUST ADD THE CLIMATE"

$$PD_i = \Phi\left(\frac{\sigma_i}{2} - \left(\frac{1}{\sigma_i}\right) \cdot \left(\mu_i + ln\left(\frac{A_i}{L_i}\right)\right)\right)$$

- μ_i is the expected return of the firm's assets over the chosen horizon.
- σ_i is the volatility of the firm's asset value, over the same horizon.
- Φ is once again the cumulative distribution function of the standard normal distribution.

Physical risk is then introduced in the simplest way possible: the initial asset value is replaced by its post-shock value, everything else held fixed:

$$PD_i, climate = \Phi\left(\sigma_i/2 - (1/\sigma_i) \cdot [\mu_i + ln(A_i, climate/\ Li)]\right)$$

- $A_i, climate$ is the asset value after the physical climate shock.

In practice, $A_i, climate = A_i \cdot (1 - \delta_i)$.

Where δ_i is the physical climate damage rate applied to the firm's assets, the same role as δ_j in the collateral channel, but at the borrower level.

This gives a climate-adjusted probability of default: it is basically the same formula as before, just evaluated on a smaller asset value.

C.2 Making the shock dynamic: adding a jump term

The static shift above can also be built into a continuous-time asset process, so that ordinary market risk and the climate shock are modelled together:

$$\hat{A}_i(t) = A_i(0) \cdot exp[\ (r_i - \sigma_i^2/2 - \omega_i^2/2)t + \sigma_i W_i(t) + \omega_i Z(t) + J_i(t)\ ]$$

- $\hat{A}_i(t)$ is the value of firm i's assets at time t, after both financial and climate shocks.
- r_i is the expected return on firm i's assets.
- $W_i(t)$ is an idiosyncratic random shock specific to firm i, with sensitivity σ_i.
- $Z(t)$ is a common factor shared by all firms, for instance a macroeconomic shock, with sensitivity ω_i.
- $J_i(t)$ is the jump term representing physical climate risk specifically.

So, once again, this is the same asset process as in C.1, but with three sources of movement instead of one: firm-specific risk, a factor common to all firms, and a jump for physical climate risk. The jump pushes the whole distribution of asset values down, which shortens the distance to default and raises the PD.

The resulting probability of default takes the same Merton form as before, evaluated on the shocked asset path:

$$PD_i, jump = \Phi\left(\sigma_i/2 - (\mu_i + ln(\hat{A}_i(t)/L_i))/\sigma_i\right)$$

This is C.1's formula again, with the shocked asset value $\hat{A}_i(t)$ standing in for the plain A_i

C.3 The portfolio-level result: conditional PD with a climate factor

The two steps above describe a single borrower but as the real contribution is at the portfolio level the same climate jump is embedded inside the one-factor Vasicek model that underlies the Basel Internal Ratings-Based (IRB) approach, so the result plugs into existing capital formulas with little disruption.

IMPACT OF A HAZARD ON LOSS GIVEN DEFAULT (LGD) AND LOAN-TO-VALUE (LTV): "JUST ADD THE CLIMATE"

Physical risk is modelled as a coin flip: with probability q a climate shock occurs and shifts assets down by a normalised amount $\hat{\alpha}$ with probability $(1 - q)$ it does not.

The PD conditional on the state of the economy y is then a probability-weighted mixture of two ordinary Vasicek expressions, one for each state:

$$PD(default = 1 \mid S = -y)$$

$$= (1 - q) \cdot \Phi\left(\frac{\left[\sqrt{\rho} \cdot y + \Phi^{-1}(PD_0)\right]}{\sqrt{(1-\rho)}}\right)$$

$$+ q \cdot \Phi\left(\frac{\sqrt{\rho} \cdot y + \Phi^{-1}(PD_0) + \sqrt{1-\rho} + \hat{\alpha}}{\sqrt{(1-\rho)}}\right)$$

- y is the level of economic stress, a normalised systematic factor.
- PD_0 is the PD without climate impact, the bank's ordinary, pre-climate PD estimate.
- ρ is the asset correlation with the economy, set by Basel for the asset class.
- q is the probability that the climate event occurs, supplied externally by a climate model, not estimated firm by firm.
- $\hat{\alpha}$ is the size of the climate shock on assets: it is not calibrated directly, but backed out from the bank's own pre-climate and climate-adjusted PD estimates together with q.

The conditional PD is an average of two ordinary Vasicek PDs: one as if the climate event never happens, weighted by $(1 - q)$, and one as if it happens and raises the default threshold by $\hat{\alpha}$ weighted by q.

Setting $q = 0$ or $\hat{\alpha} = 0$ gives back the standard Basel formula exactly: the climate piece adds onto the existing framework.

If you managed to make sense of this equation congratulations! You now understand how the two channels here compound, physical risk simultaneously erodes collateral value (raising LGD) and asset value (raising PD), so a correct expected loss needs both moving together! Many thanks to Finalyze for explanations related to this model

Provider Box – Straight out of academia to the industry: Climacred / Climafin and a unified approach to PD-Based Modelling

I. Monasterolo, A. Mandel and S. Battiston built the CLIMACRED-PHYS methodology (Battiston et al., 2023; Mandel et al., 2025) on the basis of CLIMACRED-TECH, their transition-risk model. It starts from a common criticism of the standard toolkit. Integrated assessment models such as DICE derive aggregate financial losses top-down from country-level GDP losses, which leaves little room for the distribution of losses across firms and locations (Battiston et al., 2017).

CLIMACRED-PHYS is built bottom-up instead. It geolocates about 130,000 production facilities from company annual reports, and when a report gives no breakdown, the firm's revenue is split equally across its sites.

Hazard realizations are simulated at each hazard's own resolution, from flood depth maps at about 1 km to synthetic storm tracks. The published application uses the ISIMIP2b protocol, and the climate forcing now comes from the ISIMIP3b protocol with an ensemble of CMIP6 models. Country- and sector-specific vulnerability functions then convert hazard intensity into capital destroyed and output lost at each site. When several hazards hit the same facility in a year, the compound impact is the strongest hazard's damage scaled by the facility's tangible capital share.

$$D_{i,t}^{S} = g_i \cdot \max_{h \in H} f_{h,c_i,s_i}\left(x_{h,i,t}^{S}\right)$$

$D_{i,t}^{S}$ is the share of facility i's capital destroyed in year t under scenario S.

g_i is the facility's tangible capital share.

H is the set of hazards.

f_{h,c_i,s_i} is the vulnerability function of hazard h for the facility's country c_i and sector s_i. $x_{h,i,t}^{S}$ is the intensity of hazard h at the facility.

Where it departs from a standard cat-model pipeline is that rather than stopping at an expected annual loss or a probable maximum loss, their physical damage stream feeds directly into a structural credit model (see 2.2.3 for more information).

Each firm's climate-adjusted cash flow is accumulated against its debt service over the horizon, and default occurs when cumulative discounted cash flow falls short of cumulative discounted debt. The probability of default is read off empirically from the distribution of outcomes across simulated climate realizations, 1,000 per scenario in the published applications. The output is a portfolio-level, scenario-conditioned PD and VaR on bonds and equities, comparable directly against a bank's or insurer's existing credit risk metrics.

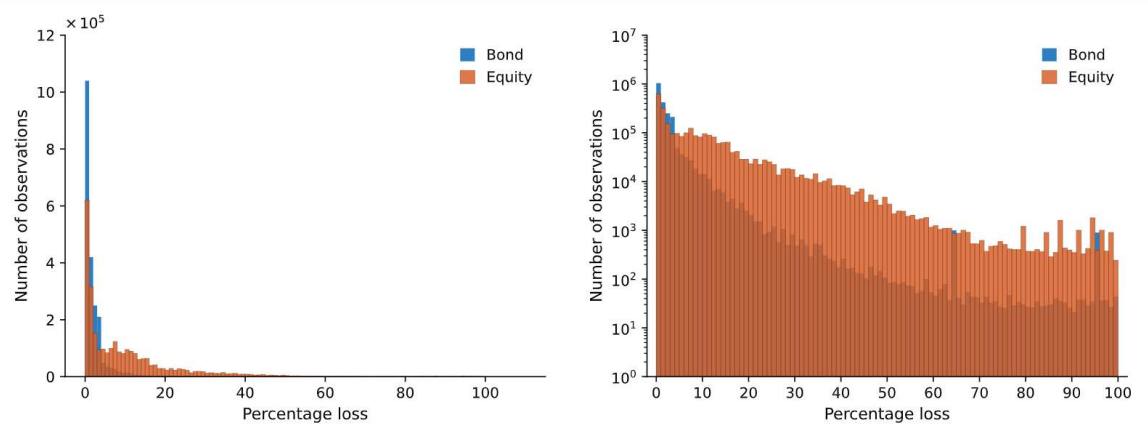


FIGURE 32. Distribution of losses on a global universe of bonds and equities across 1,000 climate realizations
Source: Adapted from Mandel et al., 2025 (fig. 1)

The method has already moved into supervisory and industry use:

- SEQUANTIS used it to compute VaR on French insurers' investment portfolios.
- Ceres applied it to about $2.2 trillion of syndicated loans held by major US banks, decomposed by peril and by scenario.
- The Monetary Authority of Singapore ran it as a physical-risk stress test on Singaporean banks.

The method also underlies the physical climate storylines (compound heatwave-drought-wildfire and storm-flood events) in the NGFS short-term scenarios.

In each application, the breakdown carries the signal. In the Ceres study, coastal flooding was the largest source of direct risk, while about two-thirds of the total risk came from indirect effects such as supply-chain disruption and lower productivity, a distinction any single blended PD adjustment would erase. (More details in section 3.1)

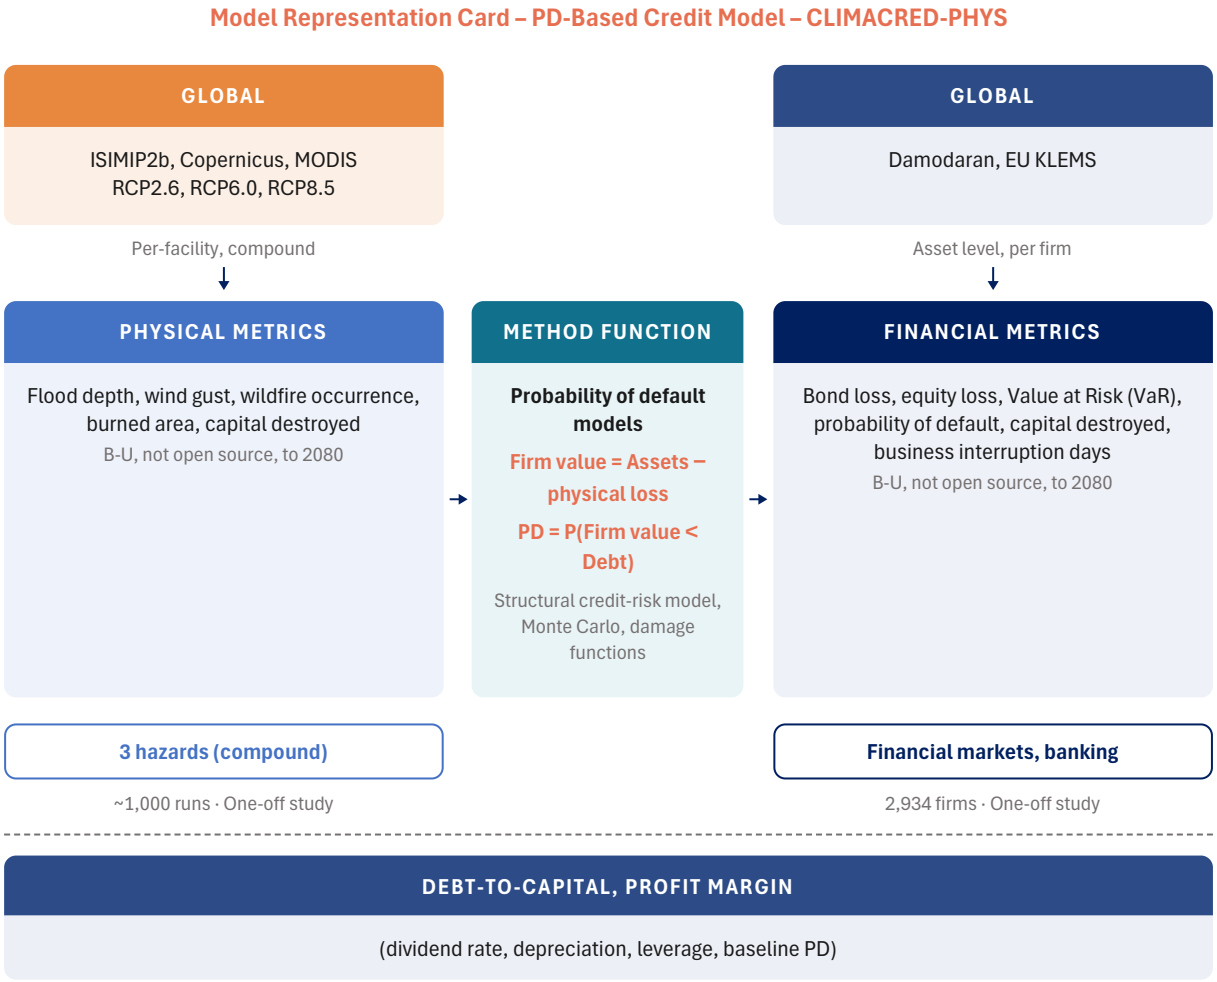


FIGURE 33. Model representation card – PD-based credit model (CLIMACRED-PHYS)

Why is this model useful?

This model can be used if the question is "can a company (i.e. a large group of assets) still service its debt in case of one or several occurring hazards?". It fits situations where the end user is a bank doing credit risk assessment, a bond investor or anyone needing physical risk translated into the language of default probability and credit VaR.

It goes beyond what we have seen so far, beyond a damage assessment or a loss figure. Because it runs on Monte Carlo simulation across thousands of firms and three RCP scenarios, it is built for portfolio-scale credit risk work.

What does this model teach us?

This model teaches us that physical damage only becomes a credit problem once it is run through a firm's balance sheet, the same climate hazard hitting two companies with different leverage or profit margins produces two very different default probabilities.

It also shows why the damage function step, converting flood depth or wind gust into a share of destroyed capital, is just an input here. But the real work happens afterward, when that capital loss gets weighed against debt to decide whether the firm actually defaults.

Other methods

Two other methods will be explored in this section.

A second approach operates through internal credit models, whether applied to wholesale or retail portfolios. The conventional risk drivers used in these models, such as revenue or asset values, can be directly impaired by climate-induced damages. Physical risk is thus introduced as an exogenous shock to the model's explanatory variables, propagating its effect onto the probability of default (PD) in a manner consistent with the existing internal ratings framework.

Finally, structural approaches using the Merton/Vasicek (see section 2.2.3) approach provide a more rigorous framework for formalising this transmission. These structural credit risk models can be used in combination with a shock applied to the borrower's asset value, or through the introduction of a direct increment on the PD or LGD, in order to capture the impact of physical risk. Grounded in robust theoretical foundations, this family of models represents one of the most promising avenues for modelling the transmission channel from physical climate hazard to default risk in a transparent and internally consistent way.

2.2.2 PROBIT, LOGIT AND SURVIVAL MODELS

Probit and Logit Models: Statistical PD from Physical Risk Scores

What are those?

Probit and logit models are the standard econometric tools for modelling binary outcomes (default / no default) as a function of explanatory variables. With the rise of physical climate risk, the explanatory variables include both traditional credit determinants (leverage, debt service coverage, sector, size) and physical risk metrics (hazard exposure score, AAL as a fraction of asset value, climate scenario VaR). The model estimates the probability that a borrower defaults within a given time horizon (typically 1 year for pricing, up to 30 years for mortgage risk) conditional on these covariates.

Survival models (also called duration models or hazard rate models) take a different approach: they model the time until default as a continuous random variable and estimate how physical risk exposure affects the instantaneous rate of default (the hazard rate) at each point in time (a very interesting application for the banking sector has been published by Calabrese et al., 2024).

This makes those models more informative than a fixed-horizon PD when the physical risk is cumulative (a property in a flood zone faces incrementally higher flood risk each year as sea levels rise) or when the timing of the default matters for economic loss calculation (early default vs. late default has different present value implications).

Mathematical Representation

LOGISTIC REGRESSION (LOGIT) FOR PD

The logit model expresses the probability of default as a logistic transformation of a linear combination of credit, physical risk and macroeconomic variables. The logistic function keeps the predicted PD between 0 and 1, and each coefficient reads as a change in the log-odds of default.

$$PD_{x,s} = \sigma(\eta_{x,s})$$

$$\eta_{x,s} = \alpha + \beta_1 CR(x) + \beta_2 RISK_s(x) + \gamma^\top Z_s$$

$$\sigma(\eta) = \frac{1}{1 + e^{-\eta}}$$

- $PD_{x,s}$ is the probability that obligor x defaults within the horizon T (for example 1 year) under climate scenario s
- $\sigma(\cdot)$ is the logistic function, which maps the log-odds into a probability between 0 and 1
- x indexes the obligor (a firm, a mortgage, a loan)
- s indexes the climate scenario (for example SSP2-4.5, SSP5-8.5)
- $\eta_{x,s}$ is the linear index, equal to the log-odds of default.

$$\eta_{x,s} = \ln\left(\frac{PD_{x,s}}{1 - PD_{x,s}}\right)$$

- α is the intercept, the log-odds of default when all covariates equal zero
- $CR(x)$ is the baseline credit risk score of obligor x, built from leverage, coverage ratio, sector, size and rating
- β_1 is the coefficient on $CR(x)$, the change in log-odds per one-unit increase in the credit score. It measures the marginal effect of ordinary credit risk (leverage, coverage ratio, sector, size, rating) on the log-odds of default.
- $RISK_s(x)$ is the physical risk metric of obligor x under scenario s (a hazard score, an $\frac{AAL}{TIV}$ ratio, a climate *VaR*, or a composite index), expressed in a stated unit or standardized
- β_2 is our key parameter of interest: the marginal effect of physical risk exposure on log-odds of default; estimated from historical default data matched to physical risk observations, and e^{β_2} is the corresponding odds ratio
- Z_s is the vector of macroeconomic controls (GDP growth, interest rate, unemployment, real estate cycle), taken from history during estimation and from the scenario path during projection
- γ is the vector of coefficients on Z_s. Each component γ_k measures the change in the log-odds of default per one-unit increase in the k-th macroeconomicth variable, holding the other covariates fixed.

Another way to consider it is:

$$\frac{\partial PD_{x,s}}{\partial Z_{k,s}} = \gamma_k p_0 (1 - p_0)$$

Where:

$Z_{k,s}$ is the k-th macroeconomic variable under scenario s

γ_k is its coefficient, the change in log-odds per one-unit increase in $Z_{k,s}$

LOGISTIC REGRESSION (LOGIT) FOR PD

p_0 is the fitted PD of the obligor at the point where the effect is evaluated

α, β_1, β_2 and γ are estimated on historical default data matched to observed physical risk exposure. Scenario values $RISK_s(x)$ and Z_s are then substituted to project $PD_{x,s}$.

The marginal effect of physical risk on PD

The marginal effect on PD of a one-unit increase in $RISK_s(x)$ at the baseline PD value p_0 is:

$$\frac{\partial PD_{x,s}}{\partial RISK_s(x)} = \beta_2 p_0(1 - p_0)$$

Where:

$\frac{\partial PD_{x,s}}{\partial RISK_s(x)}$ is the change in default probability for a one-unit increase in physical risk, holding the other covariates fixed

$p_0 = \sigma(\eta_{x,s})$ is the fitted PD of the obligor at the point where the effect is evaluated

β_2 is the physical risk coefficient defined above

This effect depends on p_0. It is largest when $p_0 = 0.5$ and close to zero for very safe or very risky obligors. So the same β_2 raises the PD of a mid-risk borrower much more than that of a prime borrower. This is why it can be recommended to report the average marginal effect across the loan book over the effect at a single point (Norton et al., 2019).

$$AME = \frac{1}{N}\sum_{i=1}^{N} \beta_2 PD_i(1 - PD_i)$$

where:

AME is the average marginal effect of physical risk on PD across the portfolio,

N is the number of obligors in the loan book,

i indexes each obligor,

- PD_i is the fitted PD of obligor i.

PROBIT MODEL VARIANT

$$PD(x,s) = \Phi(\alpha + \beta_1 \cdot CR(x) + \beta_2 \cdot RISK_s(x) + \gamma \cdot Z)$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function (CDF).

Probit and logit produce nearly identical fitted values for most empirical datasets; the choice is rarely material. Probit has a natural connection to the Merton structural credit model (see section 2.2.3), making it theoretically preferred in a credit risk context.

PD UPLIFT: CLIMATE-ADJUSTED VS. BASELINE

The PD uplift compares the default probability of an obligor under a climate scenario with its default probability under current climate conditions. Every other covariate stays the same, so the difference isolates the share of PD attributable to the physical hazard.

$$\Delta PD_s(x) = PD_{x,s} - PD_{x,\text{base}}$$

$$\Delta PD_s(x) = \sigma(\alpha + \beta_1 CR(x) + \beta_2\, \text{RISK}_s(x) + \gamma^\top Z_s) - \sigma(\alpha + \beta_1 CR(x) + \beta_2\, \text{RISK}_{\text{base}}(x) + \gamma^\top Z_s)$$

Where:

- $PD_{x,s}$ is the PD of obligor x under scenario s,
- $\Delta PD_s(x)$ is the climate-driven change in default probability for obligor x under scenario s,
- $PD_{x,\text{base}}$ is the PD of obligor x under the baseline,
- $\text{RISK}_s(x)$ is the physical risk metric of obligor x under scenario s,
- $\text{RISK}_{\text{base}}(x)$ is the physical risk metric of obligor x under the baseline,
- $\sigma, \alpha, \beta_1, \beta_2, CR(x), \gamma$ and Z are defined as in the logit model above.

The macroeconomic term $\gamma^\top Z_s$ stays in both expressions. Since σ is non-linear, it does not cancel out, and the size of the uplift depends on the obligor's starting PD. In the probit variant, Φ replaces σ.

ΔPD_s is the value banks feed into IFRS 9 staging models (Stage 1 to Stage 2 migration triggers), supervisory climate stress tests such as the ECB 2022 climate risk stress test, and loan-level credit pricing adjustments.

Empirical evidence supports the existence of this effect. In the United States, mortgages in high-flood-risk areas are up to 1.8 percentage points more likely to default than comparable mortgages in low-risk areas under current conditions (Dice, Hossain & Rodziewicz, 2024). Using a survival model, Calabrese et al. (2024) project a systematic increase in mortgage default risk by 2050 under RCP 4.5.

COX PROPORTIONAL HAZARD MODEL

A Cox model fits here because a logit or probit estimated once per loan answers a single question. Will this loan default by a fixed date, say year 1 or year 30? Such a model treats risk as constant over the whole period, yet physical risk is dynamic.

A mortgage in a flood zone becomes riskier every year as sea levels rise, so the risk of default in year 25 differs from the risk in year 1. A survival model therefore asks how risky the loan is at each point in time, and lets that risk change year by year as flood exposure grows. The same logic applies to a logit estimated on loan-year data, which becomes a discrete-time hazard model.

The Cox model has a second advantage. It does not require the shape of default risk over the life of a loan as an input. Physical risk multiplies whatever the baseline risk is, and the effect of one unit of physical risk is the same at every point in time. The modeller does not need to know how mortgage default risk naturally evolves over 30 years. The model estimates how much climate exposure increases that risk, whatever the shape of the baseline.

COX PROPORTIONAL HAZARD MODEL

$$h(t \mid x,s) = h_0(t) \cdot exp(\beta_1 CR(x) + \beta_2 RISK_s(x,t))$$

- $h(t \mid x,s)$ = instantaneous default hazard rate at time t for obligor x under scenario s.
- $h_0(t)$ is the baseline hazard function, left unspecified in the Cox model and recovered non-parametrically (Breslow estimator). Specifying it as Weibull (see below), Gompertz or log-logistic gives a parametric proportional hazards model instead.
- $RISK_s(x,\ t)$ is a time-varying physical risk covariate: for example, the flood depth at the asset location in year t under climate scenario s, increasing with sea-level rise; or the temperature anomaly exceeding a critical production threshold
- β_1 and β_2 are the coefficients estimated by partial likelihood.

The proportional hazards assumption means that each unit of physical risk multiplies the baseline hazard by the same factor $exp(\beta_2)$ at every point in time. Since $RISK_s(x,t)$ rises over time, the total multiplier $\exp(\beta_2 RISK_s(x,t))$ rises with it.

$h(t \mid x,s)\ dt$ is the probability of defaulting in $[t, t+dt]$ given survival to t.

If $\beta_2 = 0.3$ and $RISK_s$ increases by 1 unit, the default hazard at every time point is multiplied by

$$exp(0.3) \approx 1.35, i.e. 35\%.$$

Survival function and probability of default by horizon T

$$S(T \mid x,s) = P(\tau > T \mid x,s) = exp\left(-\int_0^T h(t \mid x,s)\ dt\right)$$

$$PD(T \mid x,s) = 1 - S(T \mid x,s)$$

Where: τ is the default time of obligor x

- $S(T \mid x,s)$ is the probability that obligor x survives (does not default) until horizon T under scenario s
- $PD(T \mid x,s)$ is the probability that obligor x defaults before horizon T under scenario s

For time-varying $RISK_s(x,t)$ varies over time, for instance with a flood probability that rises linearly, the integral has no closed form. With yearly values of the covariate, it reduces to a sum over years.

This time dependence makes survival models well suited to mortgage books with 25 to 30-year tenors, where physical risk exposure in year 25 can be much higher than in year 1. The Breslow estimator, however, only covers the loan ages observed in the data. Projections beyond them require a parametric baseline or an explicit assumption on its shape.

A discrete-time survival model has been estimated on European SME loans exposed to coastal, flash and river floods, and default probability trajectories have been projected under several flood severity scenarios. The projected increases in default probability reach up to 24 basis points in Spain and 18 basis points in France (Chen et al., 2026).

THE PARAMETRIC WEIBULL VARIANT

The Weibull variant uses a simple rule to predict how a loan's ordinary risk changes as it gets older, and climate risk then raises that risk each year according to how exposed the property is to floods or other hazards.

It replaces the free-form baseline risk curve with a simple formula, so that a single number tells us whether a loan's default risk tends to rise, fall or stay flat as the loan ages. Because the baseline follows a formula, it can be extended beyond the loan ages observed in the data. This removes the limit of the Breslow estimator noted above and makes 25 to 30-year mortgage projections possible.

This model variation is used for "closed-form" tractability. The Weibull hazard is just a flexible closed-form stand-in for $h_0(t)$ where a single parameter α controls whether baseline mortgage risk rises, falls, or stays flat over the loan's life, instead of having to estimate that shape non-parametrically point by point:

$$h_0(t) = \left(\frac{\kappa}{\theta}\right)\left(\frac{t}{\theta}\right)^{\kappa-1}$$

Where κ is the shape parameter, it describes how default risk evolves with loan age in the absence of climate effects:

If $\kappa > 1$ signals a hazard increasing with the loan age.

If $\kappa < 1$: it signals a hazard decreasing with the loan age.

If $\kappa = 1$: the hazard is constant and the model reduces to the exponential case

For climate-exposed mortgages, increasing $\kappa > 1$ is theoretically motivated: default risk increases over the loan life as the physical risk intensifies.

- θ is the scale parameter here, controlling the time-scale over which the hazard evolves.
- t is the time since loan origination
- $h_0(t)$ is the baseline hazard at time t

What Data Do You Need?

Estimating β_2, the marginal effect of physical risk exposure on log-odds of default, requires linking historical default records to physical risk observations. This is the primary methodological challenge, because:

- Historical loan-level default data with geographic identifiers is proprietary and rarely shared outside the institution.
- Physical risk observations (hazard events, AAL scores) must be matched to each obligor's primary business location or collateral location, which requires precise geocoding.
- Climate-related defaults are rare events: most standard credit datasets contain very few defaults attributable to physical climate hazards as a primary cause, requiring either long time series or large cross-sections.

This type of model has been estimated on real default data. About one million securitized loans to small and medium-sized firms in Belgium, Italy and Spain were matched to the flood risk indicator of the JRC Risk Data Hub at NUTS3 level and to historical flood events. A Cox proportional hazard model was fitted on loan default, controlling for interest rate, loan balance, residual loan term, collateralization and regional GDP and employment growth. Loans to firms hit by a flood in the previous 24 months were found about 1.5 times more likely to default, and banks were

found to charge an average risk premium of about 6 basis points to firms in high-flood-risk areas (Barbaglia, Fatica & Rho, 2023).

Most of the frequently cited supervisory exercises do not do this. The ECB/ESRB (2021) report matches AnaCredit loan exposures to JRC and Four Twenty Seven flood data, but the output is descriptive, the share of bank credit exposure located in low, medium or high flood-probability areas. They do not consider a fitted PD equation with an estimated physical-risk coefficient.

Caloia and Jansen's (2021) DNB stress test for the Netherlands estimates a PD regression on interest rate, LTV and regional GDP growth, then feeds flood-induced macro shocks through it; flood risk does not enter as a covariate in the PD equation itself, it enters directly only through the LGD module, via the effect of property damage on loan-to-value.

The ACPR pilot exercise (2021) followed a similar logic. It was a bottom-up, scenario-driven stress test organized by the ACPR and run by participating banks and insurers.

Direct estimates of a β_2 coefficient from matched default and hazard data are therefore rarely made public. Most institutional exercises propagate flood damage through existing PD and LGD models via macro or collateral channels instead of estimating physical risk as its own covariate.

Provider Box – Insights from the field – ODACC

ODACC (Outil de Diagnostic pour l'Adaptation au Changement Climatique) is a free service developed by the Banque de France to help companies understand their exposure to physical climate risks. It is accessible through the Banque de France's online platform for company managers. The tool uses climate data produced by Météo-France (DRIAS), which the Banque de France processes to make them directly usable by companies.

ODACC covers six hazards, namely heat, cold, precipitation, drought, wind and wildfires. **Projections rely on 17 climate models and are provided on the 8 km × 8 km SAFRAN grid**. They follow the French reference warming trajectory for adaptation (TRACC), which corresponds approximately to +2°C in 2030, +2.7°C in 2050 and +4°C in 2100.

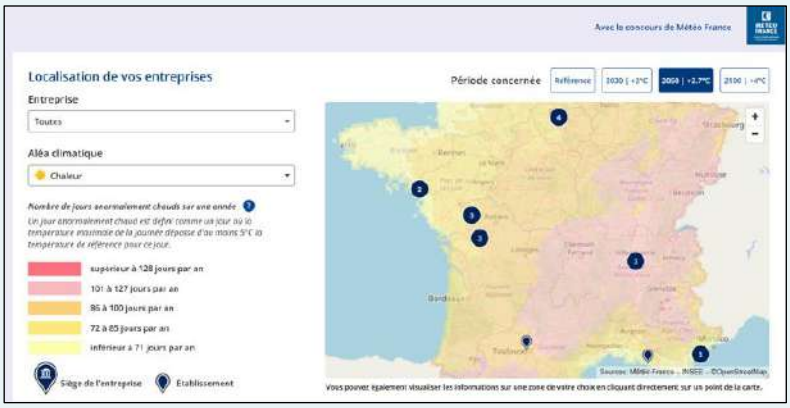


FIGURE 34. ODACC interface

For each establishment, the **tool provides hazard indicators** and a **criticality index on five levels, from low to critical**, which helps companies identify their most exposed sites. ODACC stops at hazard and exposure. It applies no damage function and produces no financial loss estimate or default modelling. Results are displayed on an interactive map of projected changes across the territory, together with a summary table for comparing a company's sites.

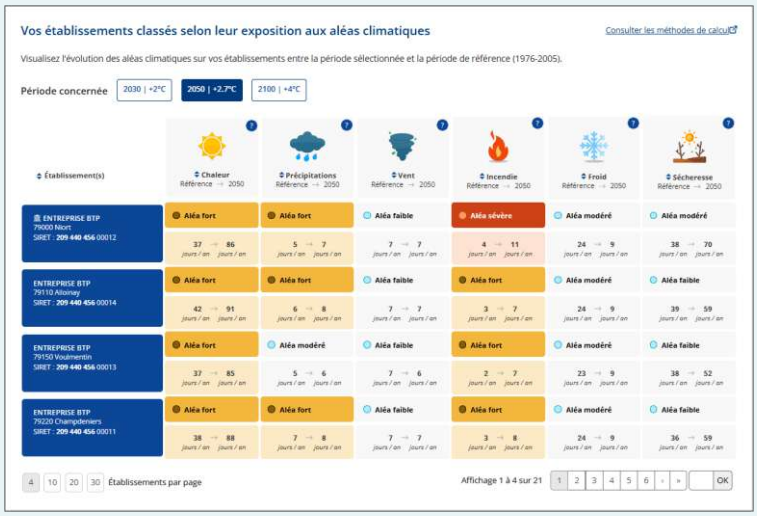


FIGURE 35. ODACC exposed assets interface

The tool is open to about five million French companies of all sizes and sectors. It is designed as an entry-level tool to raise awareness of physical climate risks and support adaptation decisions. ODACC is part of the French National Climate Change Adaptation Plan (PNACC) and will evolve as new information becomes available. Flood hazards are expected to be added by the end of 2027.

In parallel, the Banque de France is conducting exploratory work on a Physical Climate Risk Indicator (ICP). This synthetic indicator of gross physical risk would combine the geographic exposure of establishments to climate hazards with the sectoral sensitivity of their economic activities.

Provider Box – The long road to Credit Assessment – S&P Global Sustainable1

Sustainable1 cannot really fit in a box, because it is composed of several models. It is an ensemble of models covering multiple hazards, built through platform integration.

S&P Global Sustainable1 Origins and architecture. Sustainable1 is S&P Global's sustainability data business, incorporating the expertise and capabilities of Trucost. Sustainable1 brought physical-risk analytics to market in 2019 with an asset-based approach. In January 2022 S&P Global acquired The Climate Service, the Durham, North Carolina firm behind the Climanomics platform, and a new physical-risk framework built on that firm's work followed the same year.

Today Sustainable1 is a family of products that share one hazard engine and one library of impact functions:

- the Climanomics platform for client-uploaded real assets;
- the Physical Risk Exposure Scores and Financial Impact dataset for listed and private companies;
- a Country and Subnational dataset for sovereign and sub-sovereign issuers;
- Climate Credit Analytics (CCA) and Climate Risk Gauge, carrying physical-risk outputs into credit metrics for corporates, real estate, mortgages and sovereigns.

Sustainable1 staff describe the offer as a scientific team of 20 to 30 climate scientists, hydrologists and economists, together with an optional consulting service.

Hazard modelling. All products draw on CMIP6 output from 35 models, downscaled by NASA NEX-GDDP to a 0.25° grid (about 25 km). Hazards are computed annually and presented for most purposes as decadal means from the 2020s to the 2090s under SSP1-2.6, SSP2-4.5, SSP3-7.0 and SSP5-8.5. The physical-risk dataset covers ten hazards: coastal flood, fluvial flood, pluvial flood, extreme heat, extreme cold, tropical cyclone, wildfire, water stress, drought and landslide.

Climanomics adds wind changes and heat stress. It drops extreme cold from the loss calculation, since no financial impact pathway exists for that hazard.

Resolution follows how much each hazard varies over short distances:

- Coastal flood runs at 30 m in the United States and 90 m elsewhere. It combines GTSR storm-tide levels with Kopp et al. sea-level rise, and that sea-level component still rests on CMIP5 (Kopp et al., 2014).
- Fluvial and Pluvial flood are modelled at 30x30m resolution, leveraging a partnership with specialist flood data provider, JBA Risk.
- Landslide runs at about 8 km, using Kirschbaum et al. susceptibility maps (Stanley and Kirschbaum, 2017).
- Water stress runs at hydrologic-basin level, using WRI Aqueduct 4.0.
- Other hazards stay at 25 km.

Tropical cyclone projections are modelled out to the 2040s and are held constant after that decade to be consistent with the limited range of sea-surface temperature rise thus far observed. Historically S&P Global has quantified hazard metrics as a change against a historical baseline, for example the projected frequency of the baseline 100-year flood, however from early 2027 S&P Global will offer absolute magnitude hazard modelling capturing both historical and incremental climate change induced hazard exposure (e.g. change in flood depth, extent and frequency).

Three families of output. Sustainable1 produces three kinds of output.

1. Asset-level impacts based on asset locations, hazard levels at each location, and impact functions relating each hazard to loss of asset value via multiple impact pathways.
2. Exposure scores rank how exposed a location is to each hazard relative to global conditions.
3. Financial impact metrics express the impact to revenue, capex and opex at the asset and company level, driven by a library of over 2,000 asset type and hazard specific damage functions.

The *exposure score* places each asset location on a 1 to 100 scale relative to global conditions across all scenarios and periods:

$$x_h^* = \min\left(\max\left(x_h,\, x_h^{\min}\right),\, x_h^{\max}\right)$$

$$S_h = 1 + 99 \times \frac{x_h^* - x_h^{\min}}{x_h^{\max} - x_h^{\min}}$$

where S_h is the exposure score for hazard h, from 1 to 100, x_h is the hazard metric at the asset location, x_h^* is that metric capped between the two thresholds, $x_h^{\min}$ is the lower threshold set by S&P, below which the score equals 1, and $x_h^{\max}$ is the upper threshold, above which the score equals 100.

The *composite score* adds up the hazard scores with equal weights and maps the total through a concave curve. With this curve, a company that is highly exposed to a single hazard, or moderately exposed to multiple hazards receives a moderate to high score:

$$C = g\left(\sum_{h=1}^{N} S_h\right)$$

where C is the composite score, from 1 to 100, S_h is the exposure score for hazard h, N is the number of hazards, and g is a concave scoring curve that rescales the sum to 1 to 100.

A *sensitivity-adjusted composite* first multiplies each hazard score by a company sensitivity score δ and divides by 100. Each sensitivity score is itself normalised to 1 to 100 within the Core Plus universe of 20,000+ companies, after trimming at the 1st and 99th percentiles. The proxy depends on the hazard:

- water intensity (direct and indirect water use / revenue) for water stress and drought;
- labour intensity (headcount / revenue) for heat and cold;
- tangible-asset intensity (property, plant and equipment plus inventories / total asset value) for wildfire, floods, tropical cyclone and landslide.

$$\sigma_h = 1 + 99 \times \frac{I_h - I_h^{P1}}{I_h^{P99} - I_h^{P1}}$$

$$I^{\text{water}} = \frac{W}{R} \qquad I^{\text{labor}} = \frac{L}{R} \qquad I^{\text{tangible}} = \frac{PPE + INV}{TA}$$

where σ_h is the company sensitivity score for hazard h, from 1 to 100, I_h is the company intensity proxy for hazard h, trimmed at the 1st and 99th percentiles, and I_h^{P1} and I_h^{P99} are the 1st and 99th percentiles of this proxy across the Core Plus universe.

- I^{water} is water intensity, used for water stress and drought, with W the direct and purchased water use in m^3 and R the revenue in USD million.
- I^{labour} is labor intensity, used for extreme heat and cold, with L the number of employees. I^{tangible} is tangible-asset intensity, used for wildfire, floods, tropical cyclone and landslide, with PPE the property, plant and equipment, INV the inventories and TA the total assets, all in USD million.

It is then aggregated:

$$\tilde{S}_h = \frac{S_h \times \sigma_h}{100}$$

$$\tilde{C} = g\left(\sum_{h=1}^{N} \tilde{S}_h\right)$$

where $\tilde{S}_h$ is the sensitivity-adjusted score for hazard h, S_h is the exposure score, σ_h is the sensitivity score, $\tilde{C}$ is the sensitivity-adjusted composite score, and g is the same scoring curve as in the composite score.

For an owner-occupied office exposed to extreme heat there are three pathways:

- cooling costs, rising 5% per °C and weighted at 1.19% of asset value;
- HVAC degradation, at 6.76% per °C and weighted at 13.29%;

- employee productivity, at 1.86% per °C and weighted at 7.84%.

A rise of 2 °C in median maximum temperature yields a loss of 2.16% of replacement value per year.

Consider that the replacement value per year would be 2.21% if the formula used was exactly linear.

$$FI = \Delta T \times \sum_{k=1}^{3} r_k \, w_k$$

where FI is the annual financial impact as a share of replacement value, ΔT is the change in median daily maximum temperature in °C, k indexes the three pathways (cooling, HVAC degradation, employee productivity), r_k is the impact rate of pathway k per °C, and w_k is the weight of pathway k as a share of asset value.

Several modelling conventions shape the figures:

- There are two categories of Risk Type:
  - Relative hazard-based: Losses are attributable to climate change only. The baseline assumes full adaptation to late 20th-century conditions, so the Modelled Average Annualised Loss (MAAL) is zero by construction in the baseline period.
  - Absolute hazard-based: Losses are attributable to hazardous climate conditions rather than changes relative to a location-specific baseline. Unlike relative metrics, baseline-period impacts are not zero by construction, and baseline-period MAAL may therefore be non-zero.
- MAAL is undiscounted.
- Gains are capped at −5% in the company dataset.
- Impact functions come in four ownership perspectives (investor-owner, owner-occupier, tenant, manufacturer-supplier).
- Climanomics can switch on adaptation functions for 23 asset types, such as flood walls, elevation, wet or dry flood-proofing, cool and green roofs, impact-resistant glass and greywater recycling.

For companies with A+ or A data (described below), a 2025 supplement splits losses into revenue, OPEX and CAPEX impacts. Coastal, fluvial and pluvial flood and tropical cyclone are also available in beta across nine return periods, from 2 to 1,000 years.

From assets to companies. The company dataset rests on 11 million asset locations (7.2 million unique assets) in the Core Universe and 13.7 million in the Expanded Universe. Each asset is linked to an owner and an ultimate parent through Capital IQ identifiers. Sustainable1 estimated that about 90% of asset data is produced internally.

Core coverage includes over 100,000 companies (over 98% of global market capitalisation) and an expanded universe of over 870,000 public and private entities. Company metrics are weighted averages over assets, with assumed values per asset type as weights, for example USD 25 m for an office and USD 1.2 bn for a hydropower plant.

$$S_i = \frac{\sum_{a=1}^{n} V_a \, S_a}{\sum_{a=1}^{n} V_a}$$

where S_i is the score or financial impact of company i, S_a is the score or financial impact of asset a, V_a is the assumed value of asset a by asset type, and n is the number of assets linked to the company.

Data quality falls into four tiers:

- **A+** is analyst-researched complete coverage (about 20,000 companies).
- **A** requires at least five assets, one of which is neither an office nor a headquarters.
- **B** and **C** apply below that threshold. They blend the few known assets with country-level values:

$$S_i^B = \lambda\, \bar{S}^{\text{assets}} + (1-\lambda) \sum_j \rho_j\, S_j^{\text{country}}$$

$$S_i^C = \lambda\, \bar{S}^{\text{assets}} + (1-\lambda)\, S_0^{\text{country}}$$

where S_i^B and S_i^C are the company scores under tiers B and C,

- $\bar{S}^{\text{assets}}$ is the value-weighted average over the known assets,
- ρ_j is the share of company revenue generated in country j,
- S_j^{country} is the GDP-weighted score of country j,
- S_0^{country} is the score of the headquarters country,

and λ is the blending weight between known assets and country values.

Tier C replaces the revenue mix with the country's headquarters. The same blend applies to financial impact.

Country and subnational issuers. A separate dataset covers 201 countries and 2,117 subnational regions on nine hazards. Hazard values are averaged within each boundary using 1 km GDP grids (Kummu et al., 2018). The dataset also reports the share of GDP and of population above expert-set materiality thresholds, for example three months of extreme-heat days or any Cat3+ cyclone exposure.

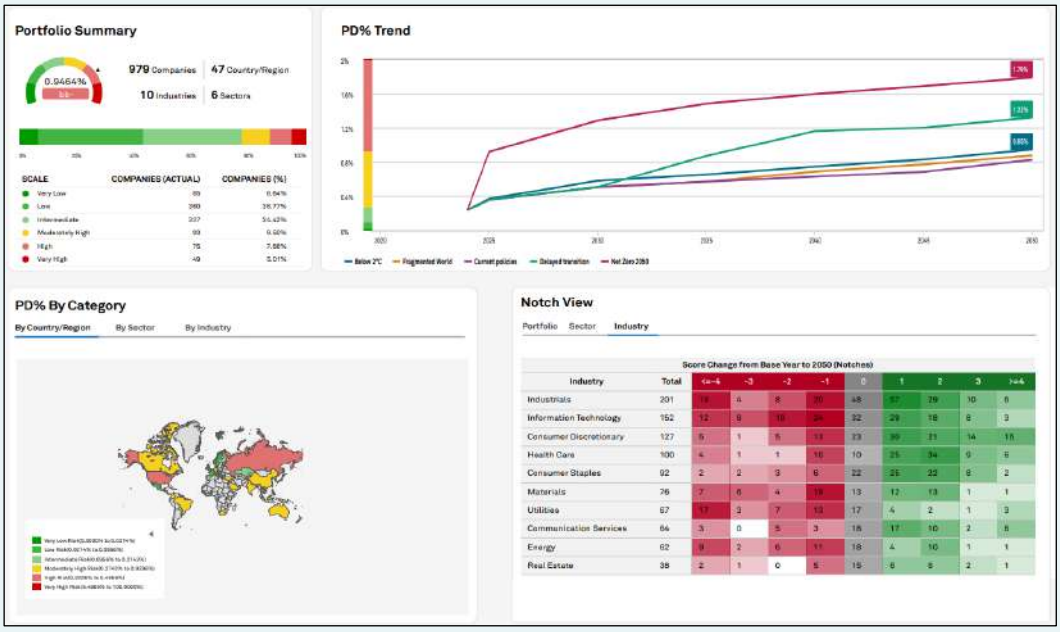


FIGURE 36. Portfolio summary

Source: S&P Global Sustainable1 Climate Credit Analytics | S&P Global;

Credit integration. This company-level physical-risk impact is what enters Climate Credit Analytics (CCA), a climate scenario analytics product that was developed by S&P Global with Oliver Wyman. Climanomics impacts are disaggregated into Revenue, Operating Expenses and Capital Expenditure at company level, then

flow into CCA through its financial drivers to generate climate-adjusted financial metrics. CCA then incorporates physical risk into its financial modelling framework alongside transition and other scenario drivers. CCA translates its financial performance drivers such as Unit Cost, Volumes, Price, Capital Expenditure and Asset Value into changes in company financial statements, including the income statement, balance sheet and cash-flow statement. CCA uses the resulting financial forecasts as inputs to its credit model to estimate the impact on forward PD. The credit model is calibrated to S&P ratings and default database. The scenario-adjusted financials can also be used independently in institutions' own financial/credit models.

Limitations reported by S&P. Hazards are modelled independently, so compound events are missing. Climate inputs are ensemble means, with no spread reported across models. GDP and population maps stay fixed at their 2010 to 2015 values. Company asset coverage is unknown for most sectors. Some impact functions are extrapolated from related asset types. The product is a climate model with no weather or real-time component. Sustainable1 said it is working on enhancing hazard and damage function resolution, integrating correlation and uncertainty quantification, and completing the workflow from hazard modelling to impacts to company financials, credit risk and valuation metrics.

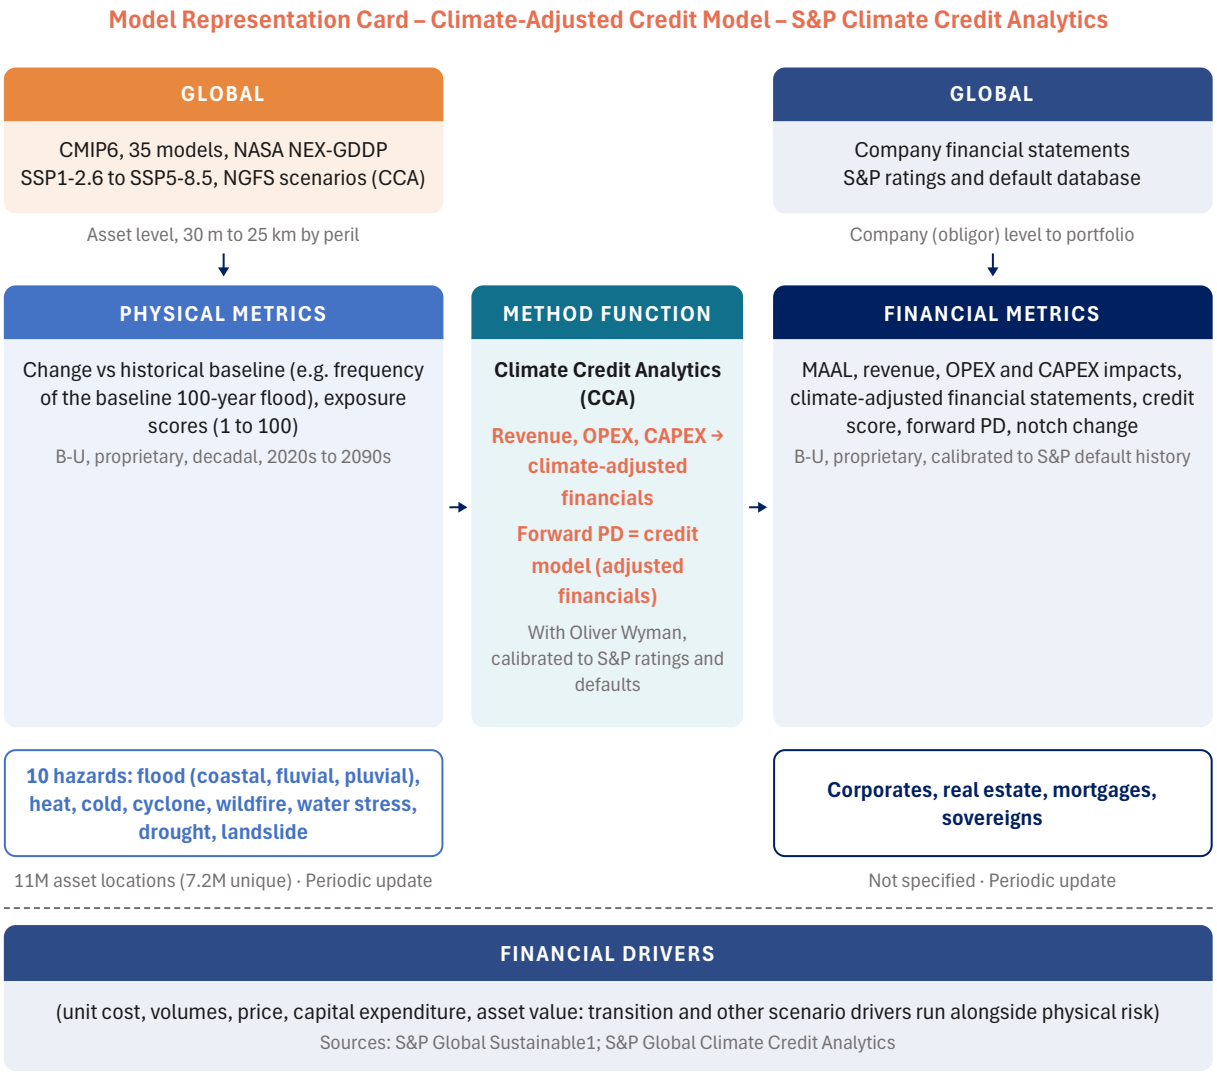


FIGURE 37. Model representation card – Climate-adjusted credit model (S&P Climate Credit Analytics)

Why is this model useful?

This model can be used by a bank needing TCFD-aligned disclosure numbers, a private equity firm needing to know what climate risk sits inside a target company's asset base, or an insurer needing a climate-adjusted view of a borrower's default risk. It fits when coverage breadth matters more than methodological novelty, with one platform spanning ten hazards, four scenarios, eight decades and over 2,000 damage functions.

It also fits when the end use needs to go straight into CAPEX and OPEX planning, or straight into a credit risk adjustment, without switching tools.

What does this model teach us?

This model shows what it actually takes to turn a hazard into a usable figure at scale.

It relies on hundreds of vulnerability functions calibrated per asset type, because a data center and an office respond to the same heatwave in different ways, through different pathways and weights.

S&P is upfront that the models build largely on public research. The commercial value lies in the proprietary asset database, with millions of locations already mapped, so a client asking about a single listed company gets an instant answer instead of a fresh analysis.

Model Representation Card – Reduced-Form Hazard Model – Flood Risk and Loan Pricing



FIGURE 38. Model representation card – Reduced-form hazard model (flood risk and loan pricing, Barbaglia et al., 2023)

Why is this model useful?

This model can be used to show that flood risk already affects lending outcomes today, without relying on a future scenario. It fits supervisory or cross-border analyses, and it adapts easily to the banking sector since it runs at NUTS3 resolution across three countries on about one million loans. It is the right choice when the deliverable needs to be a real, estimated effect on interest rates and default hazard.

What does this model teach us?

This model teaches us that physical risk already shows up in credit markets. Banks charge firms in high-flood-risk areas an average premium of about 6 basis points, and loans to firms hit by a flood in the previous 24 months are about 1.5 times more likely to default. No climate scenario is needed to make the case.

It also teaches us something about method choice. A reduced-form hazard model estimated directly on real default and pricing data addresses the same question as a structural credit model, namely how flood exposure changes default risk, without assuming a mechanism for how that happens.

What you gain is an empirically grounded number instead of an assumed one. What you lose is any ability to say something about a future you have not observed, since the model carries no projection and stops at current conditions.

2.2.3 STRUCTURAL CREDIT MODELS WITH PHYSICAL RISK: THE MERTON-CLIMATE FRAMEWORK

What is it?

Assessing a firm's ability to survive can also be done using firm asset valuation tools.

Structural credit models start from the firm's balance sheet and model default as the event that the firm's asset value V_t falls below its debt obligations D at the maturity date T (Merton 1974).

Physical climate risk enters by reducing V_t (through asset damage, revenue losses, or stranded assets) or by increasing the volatility of V_t (through greater uncertainty about future operating conditions).

This approach is consistent with standard credit pricing (CDS spreads, bond yields, loan pricing) because the default probability in the Merton model is directly linked to the value of equity as a call option on the firm's assets. It is the approach most naturally connected to market-based climate risk pricing.

Mathematical Representation

STANDARD MERTON MODEL (BLACK-SCHOLES-MERTON)

Firm asset value follows a Geometric Brownian Motion (GBM). We start from this model because it is the standard way to link a firm's asset value to its probability of default, and it gives us the baseline we then adjust for climate risk:

$$dV_t = \mu_V V_t\, dt + \sigma_V V_t\, dW_t$$

- V_t is the market value of the firm at time t
- V_0 is the current asset value

STANDARD MERTON MODEL (BLACK-SCHOLES-MERTON)

- μ_V is the instantaneous expected drift of the asset value. Under the risk-neutral measure, used for pricing, the risk-free rate r replaces μ_V.
- σ_V is the instantaneous volatility of asset returns
- W_t is the standard Brownian motion driving asset uncertainty

Default occurs at maturity T if V_T < D (face value of debt).

Real-world PD under a Geometric Brownian Motion (GBM):

$$PD = \Phi(-d_2)$$

$$d_2 = \frac{\ln\left(\frac{V_0}{D}\right) + \left(\mu_V - \frac{\sigma^2{}_V}{2}\right) T}{\sigma_V \sqrt{T}}$$

where Φ is the standard normal cumulative distribution function (CDF),

- D is the face value of debt (i.e. the default threshold at maturity),
- T is the debt maturity / risk horizon.
- d_2 is the 'distance to default': the number of standard deviations separating current asset value from the default threshold.

KMV, acquired by Moody's in 2002 and now part of Moody's Analytics, commercialised a calibrated version of this MODEL called the Expected Default Frequency (EDF™). It computes the distance to default from the Merton framework and maps it to an empirical default probability using a large historical default database (Crosbie & Bohn, 2003).

CLIMATE-ADJUSTED MERTON MODEL

Physical climate risk modifies the firm's asset value process in two ways:

1. Jump process for acute physical events (cat model integration)

We add a jump because climate disasters such as floods or storms strike suddenly and cause a sharp drop in asset value. This differs from the smooth day-to-day fluctuations the model already captures.

$$dV_t = \mu_V V_{t^-}\, dt + \sigma_V V_{t^-}\, dW_t - J_t V_{(t^-)}\, dN_t$$

Where N_t a Poisson process with intensity λ_s, counting climate events up to time t

- λ_s is the expected number of events per year under scenario s calibrated from sources such as CLIMADA, Munich Re or Swiss Re sigma
- $J_t \in [0,1]$ is the fractional asset loss from a single event, drawn from a distribution F_J derived from the cat model damage function
- $V_{(t^-)}$ is the asset value just before a possible jump

CLIMATE-ADJUSTED MERTON MODEL

The compound Poisson structure is the same as in the actuarial model of Section 2.1.2. However, here, it is embedded directly in the firm's asset dynamics.

2. Drift reduction for chronic physical risk (temperature, drought, sea level):

We lower the drift because chronic climate change works differently from a disaster. It is a slow, steady drag that reduces a firm's expected growth over time.

$$dV_t = \big(\mu_V - \delta_s(t)\big)V_t\,dt + \sigma_V V_t\,dW_t$$

Where $\delta_s(t)$ is the drag on asset value growth at time t under scenario s from chronic climate deterioration, such as rising temperatures reducing productivity, water stress reducing agricultural output, or sea-level rise eroding coastal property values.

- $\delta_s(t)$ is calibrated from the reduced-form damage functions in section 2.1.1.

Climate-adjusted PD (combined acute and chronic events)

We can use a scenario index because climate change can follow several pathways. Depending on the pathway, the chronic drag, the frequency of extreme events and the volatility all take different values, so the distance to default has to be computed separately for each scenario.

We also consider a climate scenario index (e.g. an RCP/SSP pathway) noted s and $\bar{\delta}_s$ is the chronic drift drag under scenario s

$$PD_{\text{climate}}(T,s) \approx \Phi\big(-d_2(T,s)\big)$$

$$d_2(T,s) = \frac{\ln\left(\frac{V_0}{D}\right) + \left(\mu_V - \bar{\delta}_s - \lambda_s \mathbb{E}[J] - \frac{1}{2}\sigma_{V,s}^2\right)T}{\sigma_{V,s}\sqrt{T}}$$

There are some key effects of climate adjustment on the distance to default:

- $\bar{\delta}_s$ as mentioned, is the chronic drag. It is the average chronic drag over the horizon,

$$\bar{\delta}_s = \frac{1}{T}\int_0^T \delta_s(t)\,dt$$

- λ_s is the event intensity under a scenario s and $\lambda_s E[J]$ is the expected annual loss rate from events. It also lowers the effective drift.

A higher climate volatility $\sigma_{V,s}$ such as $\sigma_{V,s} > \sigma_V$ expands the denominator, lowers d_2 even if the drift is unchanged

Downward jumps in firm value were introduced early in the climate credit literature. We haven't mentioned bonds many times. However, a defaultable bond model was developed in which corporate earnings uncertainty is combined with climate-related events that trigger sudden drops in firm value, so that bond prices and credit spreads reflect the exposure of the issuer to acute climate shocks (Agliardi & Agliardi, 2021). The jump term in the climate-adjusted Merton model above follows the same logic, with the frequency and size of the jumps taken from catastrophe models.

CLIMATE-ADJUSTED MERTON MODEL

The calibration of these jumps can draw on asset-level physical risk assessments. Acute and chronic shocks have been quantified on localized productive assets across IPCC scenarios, and neglecting asset-level information or tail acute risks was shown to underestimate investor losses substantially (Bressan et al., 2024). Such estimates provide natural inputs for λ_s and F_J. On the transition side, structural credit models have also been extended to scenario uncertainty, with a carbon tax trajectory progressively revealing the scenario and affecting default timing and credit spreads (Le Guenedal & Tankov, 2025). A comparable treatment of physical risk, where the scenario s itself is uncertain, remains to our knowledge, largely unexplored.

Connecting to Credit Spreads and LGD

The Merton model produces not just PD but also a credit spread. The yield spread over the risk-free rate for a zero-coupon bond is:

$$spread = -\frac{1}{T} \cdot ln(1 - PD \times LGD)$$

For climate-exposed firms, LGD itself is affected by physical risk through collateral value deterioration. A flood-damaged property used as collateral may realise only 40% of its pre-event value (LGD = 0.60 instead of the typical 0.35 for prime residential mortgage).

The joint effect on the credit spread is multiplicative: higher PD and higher LGD both widen spreads. Empirically, Bernstein et al. (2019) found that US coastal properties with high sea-level rise exposure sold at a 7% discount relative to inland comparables with identical characteristics, consistent with a climate-adjusted LGD effect being priced by buyers.

Merton-Vasicek in practice: a recent extension (Azzone et al., 2025)

N=10,430 listed companies globally

Physical climate risk shocks (aka negative jumps in firm value) significantly increase financial risk metrics, with expected losses reaching up to 27.17% in high-risk clusters and portfolio Value-at-Risk increasing by 3.63% over 5 years and up to 11.66% over 20 years.

Physical climate risks can alter asset allocation and require additional capital buffers.

What it does

Azzone, Ghesini, Stocco and Viola (2025) address a specific gap in how firm-value credit models have handled climate risk: existing structural approaches typically treat physical climate impact as a smooth drift adjustment, chronic damage gradually eroding growth, without a mechanism for the sudden, discrete asset destruction that an acute event like a flood or wildfire actually produces. Their approach extends the Vasicek credit risk model into a multivariate firm value model with jumps: a company's asset value follows the usual geometric Brownian motion, but with an added downward jump term triggered by climate-related hazard realizations, consistent with a compound Poisson jump structure.

The calibration is what makes the model practical.

Jump frequency and severity are tied to two firm-level characteristics: asset intensity (how much of the firm's value sits in physical, hazard-exposed capital as opposed to intangible or financial assets) and geographical exposure (where those physical assets sit, matched to hazard data by location). This differentiates, for instance, a real estate

firm concentrated in a floodplain from a company with the same credit rating but negligible physical exposure, a distinction a single climate-adjusted PD multiplier applied uniformly across a sector cannot capture.

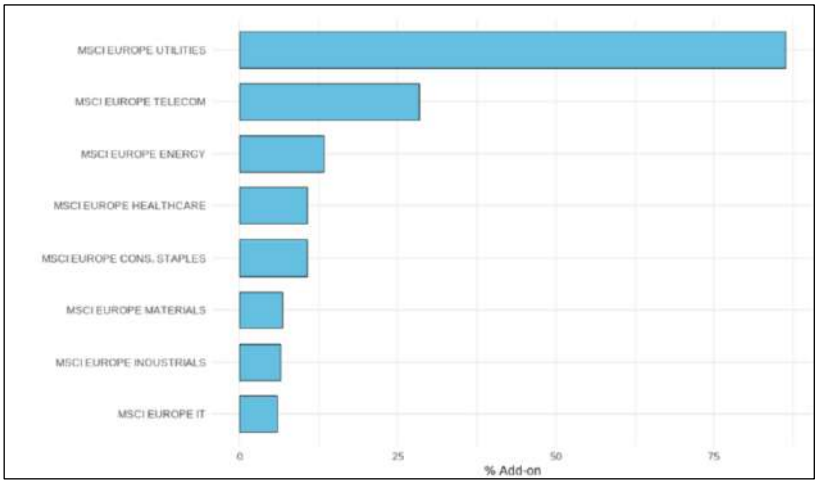


FIGURE 39. Comparison of % add-on on 5 year VaR 95%

Source: Azzone et al., 2025

Applied to sectoral portfolios, the model produces both expected and unexpected loss estimates under jump risk. Losses grow over time, consistent with intensifying hazard frequency under warming, and vary sharply by sector, concentrated in the highest asset-intensity industries such as real estate, utilities, and industrials. The authors' policy conclusion is direct: current capital buffer frameworks likely understate the capital needed for climate-exposed sectors, since standard credit risk capital charges are not built to capture jump risk of this kind, and in favour of a targeted sector-differentiated capital add-on.

Model Representation Card – Structural Credit Model with Climate Jumps – Azzone et al. (2025)

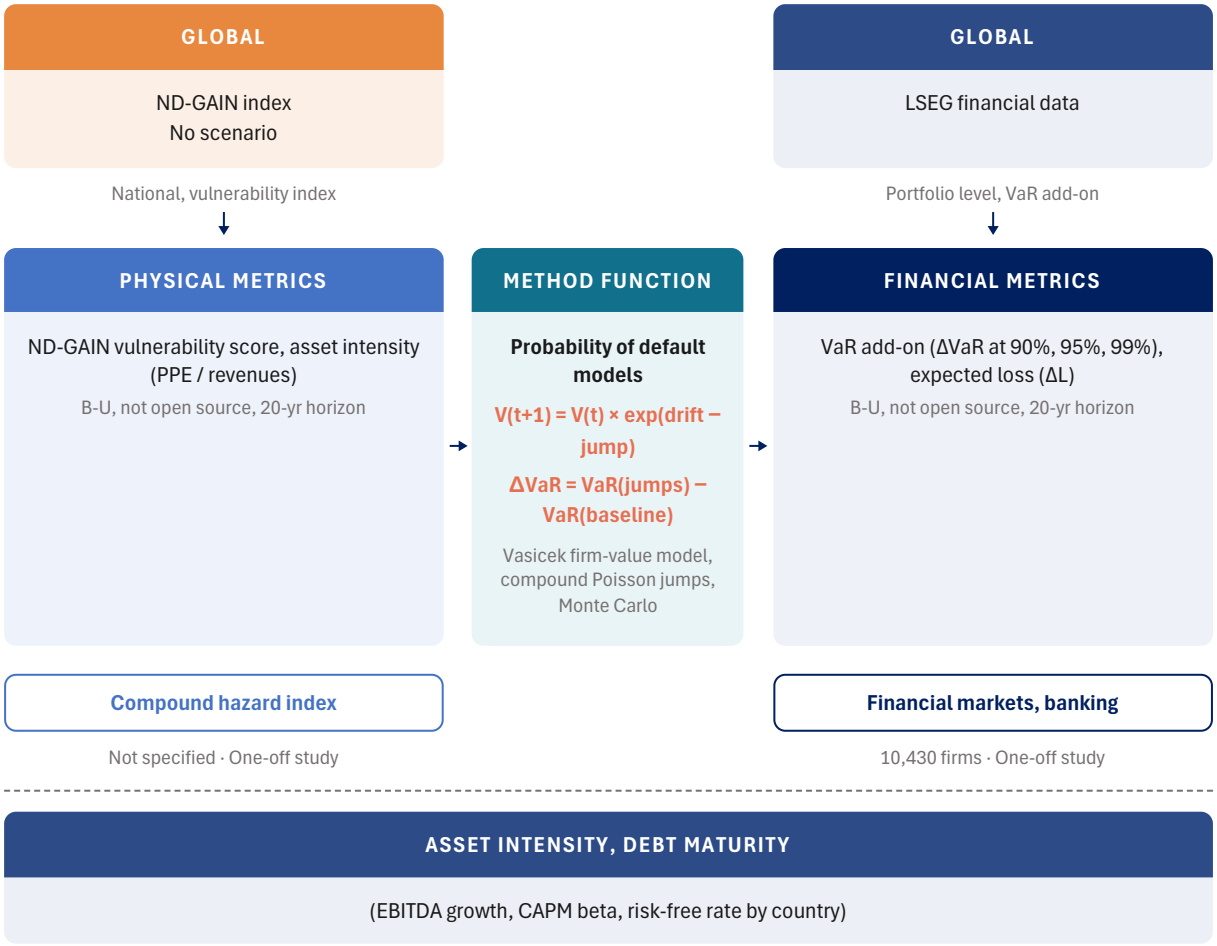


FIGURE 40. Model representation card – Structural credit model with climate jumps (Azzone et al., 2025)

Why is this model useful?

This model can be used if you manage a broad portfolio and need a fast, top-down way to gauge climate exposure without running physical hazard simulations for every holding.

It fits when the question is portfolio-level, how much extra Value at Risk does climate change add, as opposed to firm-by-firm loss estimation, and when you want something that plugs into an existing credit risk framework instead of building a new physical model from scratch. Because it leans on a public vulnerability index (i.e. not granular hazard data), it is especially useful for screening large, geographically spread portfolios quickly.

What does this model teach us?

This model teaches us that you do not need a full physical simulation to get a usable signal on climate risk, treating climate shocks as downward jumps in a standard credit risk model is enough to show a measurable VaR add-on.

It also teaches us something about what drives that exposure, since the model calibrates off asset intensity, how much of a firm's value sits in physical, fixed assets, which makes intuitive sense as a channel for climate damage but is a simplification.

The tradeoff is as follows: what you gain in speed and portfolio-wide coverage, you lose in physical grounding, this is a reduced-form shortcut when precision actually matters.

2.2.4 PORTFOLIO CREDIT RISK AND PHYSICAL CLIMATE CREDIT VAR

What is it?

From individual PD to portfolio credit VaR

A bank or insurer holds not one loan but a portfolio of N loans, each with its own probability of Default (PD) and exposure at Default (EAD).

The portfolio loss is:

$$L = \sum_{i=1}^{N} EAD_i \times LGD_i \times \mathbb{1}_i$$

where $L_i \sim Bernoulli(PD_i)$ is the default indicator for loan i.

- L is the total portfolio loss,
- EAD_i is the exposure at default of loan i,
- LGD_i is the loss given default of loan i, the share of the exposure lost if the borrower defaults,
- $\mathbb{1}_i \sim$ Bernoulli(PD_i) is the default indicator of loan i equal to 1 if the loan defaults and 0 otherwise.

The $credit\ VaR$ at confidence level α is:

$$\text{Credit VaR}_\alpha = VaR_\alpha(L) - \mathbb{E}[L]$$

where:

- $VaR_\alpha(L)$ is the portfolio loss exceeded with probability 1-α (for example α= 99.9%),
- $\mathbb{E}[L]$ is the expected loss, so the *Credit VaR* measures the unexpected loss that capital must cover.

Under physical climate change, PD_i(scenario) and LGD_i(scenario) both shift with the scenario. Default events also become correlated across loans exposed to the same hazard, since a flood in Munich hits every mortgage in the floodplain at once. This correlation is the channel through which climate risk enters portfolio credit models.

Modelling default correlation: factor models

The standard approach in credit portfolio models (CreditMetrics; Vasicek, 1987, 2002) is to introduce a common factor $Y \sim N(0,1)$ representing the state of the macroeconomy, and to assume that default occurs when a firm's latent creditworthiness index

$$Z_i = \sqrt{\rho}\, Y + \sqrt{(1-\rho)}\, \varepsilon_i$$

falls below a threshold (where $\varepsilon_i \sim N(0,1)$ is idiosyncratic). Under this model, the asset correlation ρ drives portfolio loss concentration:

VASICEK SINGLE-FACTOR MODEL WITH CLIMATE FACTOR

Extend the standard Vasicek model by adding a climate systematic factor C_s representing the aggregate physical risk shock under scenario s:

$$Z_i = \sqrt{\rho_M} \cdot Y + \sqrt{\rho_C} \cdot C_s + \sqrt{(1 - \rho_M - \rho_C)} \cdot \varepsilon_i$$

ρ_M = macro credit correlation (standard factor: GDP cycle, interest rates, ~0.10–0.25 for retail mortgages)

VASICEK SINGLE-FACTOR MODEL WITH CLIMATE FACTOR

ρ_C = climate factor loading: fraction of default variance explained by the shared climate exposure (depends on portfolio concentration in a single region or peril; can be 0.05–0.20 for a geographically concentrated coastal mortgage book)

C_s = climate systematic factor under scenario s. In a 2°C warming scenario, the variable shifts the threshold downward for all climate-exposed loans simultaneously.

The portfolio loss quantile under climate scenario s, given a realization c of the climate factor, is:

$$\mathbb{E}[L \mid C_s = c] = \sum_{i=1}^{N} EAD_i \cdot LGD_i \cdot \Phi\left(\frac{\Phi^{-1}(PD_i) - \sqrt{\rho_C}\, c}{\sqrt{1-\rho_C}}\right)$$

Where:

- Z_i is the latent creditworthiness index of borrower i. Default occurs when $Z_i < \Phi^{-1}(PD_i)$
- $Y \sim N(0,1)$ is the macroeconomic systematic factor, common to all borrowers
- $C_s \sim N(0,1)$ is the climate systematic factor, common to all borrowers exposed to the same hazard
- $\varepsilon_i \sim N(0,1)$ is the idiosyncratic shock specific to borrower i

Y, C_s and the ε_i are mutually independent, and $\rho_M + \rho_C < 1$ so that Z_i remains standard normal

The key insight is that a strongly negative realisation of the climate factor C_s such as a severe regional flood, raises the default probability of every loan in the exposed area at the same time. The climate loading ρ_C therefore concentrates portfolio losses, since defaults that would otherwise be spread out over time occur together. This is how physical climate risk will turn individually manageable loan losses into a systemic portfolio tail event.

Multi-factor extensions of this kind have been proposed to separate economic, physical and transition risk factors within the Basel ASRF framework (Garnier et al., 2022), and physical risk has also been introduced into the one-factor Vasicek model as a binary shock triggering a jump in asset value (Pozdyshev et al., 2025).

Provider Box – S&P PhysicalRisk and Climate RiskGauge

S&P Global Sustainable1 offers a complementary piece of this picture. Its Physical Risk Exposure Scores and Financial Impact dataset does not model default correlation directly. It produces the exposure data that a climate factor model like the one above needs as an input. Company assets and facilities are geolocated onto climate hazard maps built on CMIP6 output, covering ten hazards: coastal, fluvial and pluvial flood, extreme heat, extreme cold, tropical cyclone, wildfire, water stress, drought and landslide. Each location receives an exposure score from 1 to 100 per hazard, and the scores are combined into a composite score and aggregated from the asset level up to the company and portfolio level.

These scores can populate $RISKs(x)$ in the logit and probit models of Section 2.2.2. In the Vasicek extension above, they can be used to calibrate the climate loading ρ_C since they indicate which borrowers share exposure to the same hazard and how strongly.

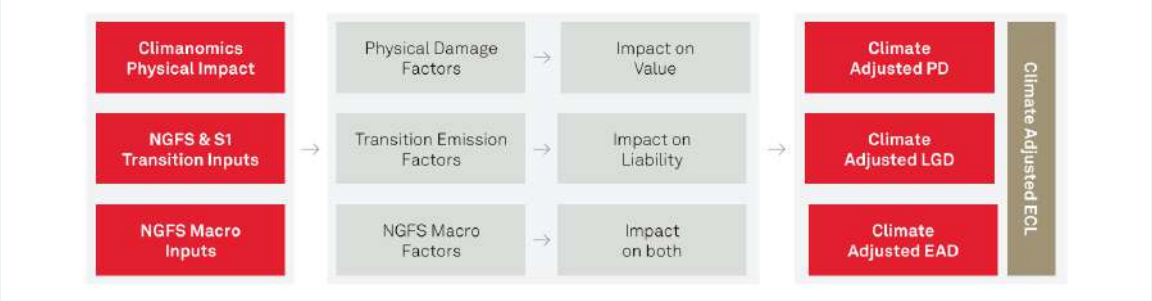


FIGURE 41. Analytical Dimensions and Steps Integrated in CRG RE (fig.1)
Source: S&P Global

These scores can populate $RISKs(x)$ in the logit and probit models of Section 2.2.2. In the Vasicek extension above, they can be used to calibrate the climate loading ρ_C since they indicate which borrowers share exposure to the same hazard and how strongly.

Climate Risk Gauge goes one step further and carries these exposures into credit metrics. In its real estate version, Climanomics physical impacts, transition inputs and NGFS macro scenarios are translated into impacts on asset value, liabilities or both. These impacts feed climate-adjusted PD, LGD and EAD, which combine into a climate-adjusted expected credit loss (ECL).

Model Representation Card – Survival Model (Cox) – Extreme Weather and Mortgage Default

FLORIDA (US)		FLORIDA (US)
Tropical cyclones, heavy rainfall First Street flood model, RCP 4.5, to 2050		Moody's Analytics mortgage data
Five-digit ZIP code		Loan level, monthly
PHYSICAL METRICS	**METHOD FUNCTION**	**FINANCIAL METRICS**
Tropical cyclone intensity (category), heavy rainfall, flood exposure × rainfall interaction Proprietary flood model, RCP 4.5 to 2050	**Survival model of default and prepayment** $h(t \mid x) = h_0(t) \times \exp(\Sigma_k f_k(x_k, t))$ Additive Cox proportional hazard, time-varying covariates, benchmarked against logit and XGBoost	Default and prepayment hazard, cyclone effect more than doubling from category 2 to category 3+, projected default risk in 2050 Proprietary loan data, scenario analysis
2 hazards: tropical cyclone, heavy rain / flood		**Residential mortgages**
Not specified · One-off study		69,046 loans (3.7M loan-months) · One-off study

LOAN AND BORROWER CHARACTERISTICS

(local economic conditions)

Source: Calabrese et al. (2024), European Journal of Operational Research, 314(1)

FIGURE 42. Model representation card – Survival model (extreme weather and mortgage default, Calabrese et al., 2024

What it does

The effect of extreme weather events on mortgage default and prepayment is estimated with an additive Cox proportional hazard model with time-varying covariates (Calabrese et al., 2024). Each month, the default hazard of a loan is modelled as a function of its usual credit characteristics and of the weather recorded in its five-digit ZIP code: tropical cyclone intensity, heavy rainfall, and their interaction with local flood exposure. The model is estimated on 69,046 Florida mortgages (about 3.7 million loan-month observations) from a Moody's Analytics dataset, and its predictive accuracy is compared with a flexible logistic model and an XGBoost algorithm.

A significant and non-linear effect of cyclone intensity on default is found. The impact more than doubles when moving from a category 2 hurricane to a category 3 or stronger. Heavy rain is shown to raise default risk in areas highly exposed to flooding, and to reduce the likelihood of early repayment. In a second step, the estimated flood effect is combined with First Street flood projections for 2050 under RCP 4.5, and a systematic increase in mortgage default risk under climate change is obtained.

Why is this model useful?

This model can be used by a mortgage lender or a supervisor who needs to bring event-level weather data into an existing credit scoring framework. It fits when the question is how a specific hazard, at a specific intensity, changes default risk month by month over the life of a loan. Since the structure of a standard survival model is kept, it can be added to existing IFRS 9 or stress testing tools with little change. Default and prepayment predictions are also shown to improve once weather variables are included, which gives a practical argument for collecting this data.

What does this model teach us?

This model teaches us that the effect of a hazard on credit risk depends on its intensity, and that the relationship is far from linear. A category 3 hurricane does much more damage to a loan book than a category 2 hurricane, so tail risk would be understated by a single "cyclone exposure" flag.

It also teaches us that hazards interact with local conditions. Heavy rain matters mainly where flood exposure is high, so the same rainfall has very different credit consequences from one ZIP code to the next.

The trade-off is scope. Only one US state, two hazards and residential mortgages are covered, and the forward-looking part relies on a single scenario (RCP 4.5 to 2050). What is gained in precision on the estimated effects is lost in generality across regions, hazards and asset classes.

Scale and the Question of Aggregation

One of the most important and underappreciated issues in PD-based physical risk models is the question of scale. The same physical risk can look very different depending on whether it is assessed at the level of an individual asset, an obligor, a sector, or the full economy, and the models appropriate for each level are structurally different.

- **Asset scale (micro).** The flood depth at a specific property location and the resulting damage to that property's value. This is the domain of damage functions and cat models. The PD implication is localised: one mortgage in the flood zone faces higher LGD; its default correlation with other mortgages outside the zone is low.
- **Obligor scale (firm/household).** The total financial impact across all of a firm's exposed assets, including business interruption, supply chain disruption, and insured vs. uninsured losses. The logit/probit and Merton-climate models operate at this scale.
- **Portfolio scale (bank/insurer).** The aggregate credit loss distribution across thousands of loans. The Vasicek factor model and portfolio CVaR operate here. The climate factor ρ_C is the critical new parameter that does not exist at the single-obligor level.
- **System scale (macro/sovereign).** The effect of widespread simultaneous defaults across a region or sector on the broader financial system: bank capital adequacy, credit supply, sovereign contingent liability,

central bank balance sheet. This is the domain of macroprudential stress tests (ECB, Bank of England, NGFS climate scenarios), network contagion models, and systemic risk frameworks, which are covered in section 2.4 of this cookbook.

2.2.5 FROM ONE RISK TO ANOTHER: FROM CREDIT RISK TO LIQUIDITY RISK

An important but often overlooked transmission channel from physical climate risk to financial instability runs not through credit default but through market liquidity. The mechanism works as follows:

When a major physical climate event materialises (or when forward-looking investors update their climate scenario probabilities), holders of climate-exposed assets may simultaneously attempt to sell or hedge. If the number of sellers greatly exceeds buyers, which is likely when the event affects a defined geography and asset class that is widely held, market liquidity evaporates and prices gap down sharply.

This creates a second-round financial amplification beyond the direct physical damage.

Empirically, homes projected to be underwater due to sea-level rise sell at a discount of about 7% in US counties where a larger share of residents believe in climate change, compared with similar homes in counties with more sceptics (Baldauf, Garlappi & Yannelis, 2020). This suggests that repricing can be triggered by shifts in beliefs and information, without a physical event.

The 2022 UK LDI crisis and the regional US bank failures of 2023 provide non-climate analogies for how rapid liquidity withdrawal can force fire sales and push otherwise solvent institutions to failure within days.

For climate physical risk, the relevant liquidity risk models are: (i) market microstructure models of asset fire-sale dynamics (Shleifer-Vishny 1992, Brunnermeier-Pedersen 2009); (ii) insurance run-off models for mass policy non-renewal in high-risk areas (already observed in California wildfire and Florida hurricane markets); and (iii) mortgage market models capturing appraisal-based credit tightening when collateral values fall below LTV covenants.

Provider Box – Quant Foundry – combining models to assess a company's survival capacity

Quant Foundry models each company as an autonomous agent that makes decisions under uncertainty. Each firm is a single object that holds its operating and variable costs, tax factors, supply chain vulnerabilities and capital structure constraints. It adjusts its capital deployment and strategy over time as the scenario unfolds.

RoE (Return on equity), FV (firm value), market cap, credit spread, PD and LGD all emerge from the simulation. None of them is calibrated from market-implied volatility or spreads.

Quant Foundry's argument is that markets cannot price climate risks they do not yet know about, so a structural model calibrated on market data would miss forward-looking climate exposure. Among the providers reviewed, Quant Foundry is the closest to the Merton structural logic detailed in this section, since default follows from a full model of the firm's balance sheet under stress.

Physical risk in the Quant Foundry's model comes through a partnership with XDI Systems. Its downscaled weather models and engineering structural models produce loss estimates for each property and asset, across the hazards of the EU Taxonomy. **Quant Foundry adds a firm-level layer, which turns these losses into days of business interruption and into damage and replacement costs, net of insurance where relevant and informed by insurance premium data. The firm sits inside a macro engine that combines a computable general equilibrium model with an integrated assessment model.**

The engine covers about 32 to 37 world regions and ten power generation technologies, and uses the MAGICC climate module. Projections run annually from 2020 to 2050 and can be extended to 2100. Results can be aggregated to the sector or the portfolio, which keeps the dependencies between holdings.

The provider tested the model on a real company with 36 global operations, using its public disclosures. Over ten years, the model's endogenous return on equity came out at about 12 to 12.5%, against a realized 12.4%. It also converged toward the company's own long-term internal plan of 8 to 9%. No formal statistical metric was reported, and the test rests on a single firm.

Key Limitations of PD-Based Physical Risk Models

- **Sparse default data.** Physical climate events severe enough to cause defaults are rare. Most existing datasets contain at most a few hundred defaults attributable primarily to a physical climate event. This makes β_2 estimates highly uncertain and sensitive to the choice of control variables.
- **Confounding with other default drivers.** Floods, storms, and droughts affect entire regions simultaneously. When firms in a flooded region default, it is often impossible to disentangle the physical damage channel from the economic recession channel (demand collapse, business disruption) and the credit supply channel (banks tightening lending in the aftermath). This confounding makes causal identification of β_2 extremely challenging.
- **Adaptation and insurance offsetting.** Firms and households that are aware of climate risk invest in adaptation (flood barriers, resilient construction, crop insurance).

This means that the β_2 estimated from historical data may overstate future PD uplift in scenarios where adaptation accelerates, or understate it where insurance becomes unavailable and unaffordable (the climate insurance protection gap, already widening in the US and Europe).

- **Horizon mismatch.** Standard credit models have 1-year PD horizons (Basel II/III, IFRS 9). Physical climate risk is most material at 10–30-year horizons (sea-level rise, chronic drought, gradual infrastructure degradation). Connecting the two requires either a lifetime PD framework (as in IFRS 9 for Stage 2 assets) or a separate scenario analysis that does not feed directly into Basel capital requirements.
- **Correlation with macroeconomic scenarios.** Physical climate scenarios (NGFS physical risk) and macroeconomic scenarios (GDP growth, interest rates) are not independent: a severe physical risk scenario also implies macro deterioration. Running them as independent inputs double-counts the macro channel and may significantly overstate PD uplifts, a problem explicitly noted in the ECB 2022 climate stress test methodology review.

2.3 ACCOUNTING AND VALUATION IMPACT MODELS

The model families described so far (such as damage functions, actuarial models, EP curves, econometric regressions, input-output frameworks, score-based tools, and probabilistic default models) all converge on a common problem: translating a physical hazard into a number that a financial decision-maker can act on. The accounting and valuation impact models described in this section occupy the final translation layer. They are the bridge between the physical and financial worlds, operating directly on the statements, ratios, and valuation multiples that investors, creditors, accountants, and regulators use every day.

Where a cat model produces an Average Annual Loss in dollars and a PD model produces a probability of default in basis points, the models in this section produce the quantities that appear in earnings releases, impairment tests, covenant calculations, credit rating rationale, and equity research reports: revenue loss, EBITDA compression, capex escalation, return on equity, loan-to-value ratio, and enterprise value change. These are the outputs that sit at the boundary of physical risk quantification and financial market pricing, and they are the outputs most directly relevant to corporate disclosure obligations under IFRS S2, TCFD, SEC climate rules, and CSRD.

The section is structured as follows. Section 2.3.1 addresses the P&L bridge: how physical events translate into revenue loss, cost escalation, and profitability impacts. Section 2.3.2 addresses enterprise value: how those P&L changes, combined with risk-adjusted discount rates, alter the fundamental value of a firm or asset. Section 2.3.3 addresses financial ratios: the specific indicators that trip covenants, trigger rating actions, and govern regulatory capital allocation. Section 2.3.4 addresses micro and meso models: the frameworks that operate at the level of individual firms, sectors, and local economic systems, connecting asset-level physical risk to broader economic outcomes without the abstraction of macroeconomic aggregation.

2.3.1 PRICE LOSS, TURNOVER LOSS, CAPEX AND OPEX

What is it?

The P&L impact model is the most direct accounting translation of physical climate risk.

It maps a physical event or chronic climate shift to the firm's revenues, cost of goods sold and operating expenses, and to its capital expenditure, which reaches the income statement through depreciation. It operates at the level of the business unit or firm, taking as inputs the physical damage or disruption estimates from upstream models (cat models, actuarial approaches, or hazard scores) and converting them into the financial quantities that accountants measure and auditors verify.

Four channels dominate the income statement impact of physical climate risk:

- **Revenue loss (ΔRev):** Physical damage reduces the capacity to produce and sell. A flooded factory cannot ship orders (Guzmán et al., 2022). A heat-stressed agricultural operation produces lower yields. A hurricane-damaged port delays cargo throughput. Revenue losses are the most immediately visible financial consequence of acute physical events and the most straightforward to model: volume lost multiplied by the average selling price and adjusted for partial recovery rates and alternative sourcing (Gould et al., 2018).
- **Price effects (ΔP):** Beyond volume, physical risk can alter the prices at which goods and services are transacted. Agricultural commodity prices spike after drought or flood events that reduce global supply. Construction materials surge in cost after major disasters due to a spike in repair demand against constrained local supply (demand surge). Energy prices spike during heatwaves as electricity demand outstrips grid capacity. These price effects can cut both ways: a commodity producer benefits from higher prices even as its own facilities are damaged; a commodity consumer faces margin compression from both higher input costs and lower production capacity simultaneously (Thompson and Lynne, 1994).

- **Operating cost escalation ($\Delta Opex$):** Physical climate risk increases operating costs through multiple channels (Aerts and Botzen, 2011; Ponticelli et al., 2023; Christiansen et al., 2024). Cleaning and minor repairs of damaged facilities. Increased insurance premiums post-event or post-risk reassessment. Higher energy costs from greater cooling demand or grid instability. Additional logistics costs from route disruption. Labor cost increases from health impacts, evacuation, and retraining. Regulatory compliance costs from new building codes or environmental requirements triggered by the event.
- **Capital expenditure escalation ($\Delta Capex$):** Physical climate risk forces capital expenditure that would not otherwise occur (Xiao, 2021). Remediation and reconstruction of damaged assets. Adaptation investment: flood barriers, elevated foundations, heat-resistant materials, fire-resistant landscaping, cooling systems. Relocation of exposed facilities. Premature retirement of assets rendered uneconomic by changed climate conditions (stranded assets). These capex items appear on the balance sheet but affect the income statement through depreciation, and affect cash flow directly in the year of expenditure.

THE P&L BRIDGE MATHEMATICAL REPRESENTATION

Step 1: from the physical event to revenue loss

For a single acute event affecting production capacity:

$$\Delta Rev_{acute} = -Rev_0 \times DR(I,v) \times \frac{DT}{365} \times (1 - \rho)$$

where:

- Rev_0 is the pre-event annual revenue (USD)
- $DR(I,v)$ is the damage ratio as a function of hazard intensity I and vulnerability class v (from the damage function in section 2.1.1); the fraction of productive capacity destroyed
- DT is the downtime in days: the period during which production is below normal
- ρ is the recovery rate during downtime: fraction of lost production recouped through inventory drawdown, alternative suppliers, overtime, or temporary facilities
- ΔRev is expressed as a change in revenue, so a loss is negative in both the acute and chronic formulas.

For chronic climate risk (e.g. gradual heat stress reducing labour productivity) at a rate δ:

$$\Delta Rev_{chronic}(t) = -Rev_0(t) \times \delta_{climate}(t,s)$$

$\delta_{climate(t,s)}$ is the productivity drag at time t under scenario s. It's calibrated from the reduced-form damage functions in section 2.1.1

Price effects (ΔP) introduced above are not modelled as a separate line item. They are folded into ΔRev by adjusting the average selling price assumption in both formulas above.

Step 2: from revenue loss to the gross profit impact

$$\Delta GP = \Delta Rev \times (1 - VC_{ratio})$$

$$VC_{ratio} = \frac{COGS}{Rev}$$

- ΔGP is the change in gross profit
- VC_{ratio} is the variable cost ratio
- $COGS$ is the cost of goods sold

THE P&L BRIDGE MATHEMATICAL REPRESENTATION

- Rev is revenue
- VC_{ratio} is the variable cost ratio computed as: $\frac{Cost\ of\ goods\ sold}{Revenue}$

Any firm with high operating leverage (high fixed cost share, low VC_{ratio}) suffers disproportionately large EBITDA losses for a given revenue decline, because the cost base does not contract with revenues.

This is why capital-intensive industries (airlines, steel, chemicals, utilities) are particularly vulnerable to revenue disruption from physical climate events.

Step 3: Direct cost escalation to operational expenses impact

$$\Delta Opex = C_{repair} + C_{insurance} + C_{energy} + C_{logistics} + C_{labour}$$

- $\Delta Opex$ is the change in operating expenses
- $C_{repair} + C_{insurance} + C_{energy} + C_{logistics} + C_{labour}$ are respectively the repair, insurance, energy, logistics and labour cost increases

Each component can be modelled and estimated separately.

Repair costs are derived from the damage function:

$$C_{repair} = DR(I, v) \times RCV$$

RCV stands for Replacement Cost Value.

Only the share of $DR(I, v) \times RCV$ that is expensed (cleaning and minor repairs) enters C_{repair}. The capitalised share (reconstruction) enters $Capex_{repair}$ in Step 5, so that the damage is counted once.

Insurance premium uplift is estimated from the post-event market hardening factor (consider for example a standard 15–25% premium increase in the policy year following a major loss event affecting the insurance class). Energy cost escalation under heatwave or grid instability scenarios uses the electricity demand elasticity to temperature.

Note that ΔFC_{opex} , used in the EBITDA bridge below, denotes this same $\Delta Opex$ quantity; the two are not additive and should be treated as a single term under one name.

Step 4: The EBITDA bridge

$$\Delta EBITDA = \Delta Rev - \Delta VC - \Delta FC_{opex}$$

$$= \Delta Rev \times (1 - VC_{ratio}) - \Delta FC_{opex}$$

EBITDA margin impact:

THE P&L BRIDGE MATHEMATICAL REPRESENTATION

$$\frac{\Delta EBITDA}{Rev_0}$$

Where $\Delta EBITDA$ is the change in EBITDA

and ΔVC is the change in variable costs, equal to $VC_{ratio} \times \Delta Rev$

For a firm with 30% EBITDA margin and a 10% revenue shock, a VC_{ratio} of 0.60 and

ΔFC_{opex} equal to 2% of Rev_0:

$$\Delta EBITDA = -10\% \times 0.40 - 2\% = -6\%,$$

$$m_1 = \frac{30\% - 6\%}{100\% - 10\%} \approx 26.7\%$$

Where m_1 is the post-shock EBITDA margin

The margin falls from 30% to 26.7%, as revenue also shrank by 10%. Expressed as a share of pre-shock revenue, EBITDA falls from 30% to 24%, which represents a 20% compression in absolute EBITDA.

Step 5: Capex Escalation model

$$\Delta Capex = Capex_{repair} + Capex_{adaptation} + Capex_{stranding}$$

- $Capex_{repair}$ is the capitalised share of reconstruction cost ($DR \times RCV$), net of insurance recovery
- $Capex_{adaptation}$ is the adaptation investment required to maintain the same level of physical protection under a changed climate: flood barriers, reinforced roofing, cooling systems, relocation costs
- $Capex_{stranding}$ is the replacement investment brought forward when exposed assets become economically non-viable before the end of their planned useful life. The value lost on the stranded asset itself is recorded as an impairment, a non-cash charge outside EBITDA and free cash flow.

If capex is rising, it will affect free cash flow directly in the year of expenditure and affects the income statement through higher depreciation in subsequent years.

The combined cash flow impact is:

$$\Delta FCF = \Delta EBITDA - \Delta Capex - \Delta Working\ Capital$$

(where working capital may increase if insurance receivables delay cash inflows).

White-Box: from Hazard to P&L, a step-by-step guide

Step 1 – Obtain the hazard input.

Source the physical intensity at the asset location from a cat model footprint, a hazard score overlay, or a historical event record. Express as the relevant metric: flood depth (m), wind speed (mph), temperature anomaly (°C), drought PDSI, wildfire perimeter distance (km).

Step 2 – Apply the damage function.

Convert hazard intensity to a damage ratio $DR(I, v)$ using the vulnerability curve for the asset type (construction class, sector, equipment type). The damage ratio drives both the physical repair cost and the production capacity loss.

Step 3 – Estimate downtime.

Downtime (DT) is a function of damage ratio. It considers the asset type, and its restoration capacity.

Results of longitudinal engineering studies usually formalize it as follows (Almufti and Willford, 2013; FEMA, 2018; De Paula Salgado et al., 2025):

$$\text{Downtime} = DT \approx f(DR, asset\ class)$$

For a commercial building at the damage ratio $DR = 0.40$:

$$DT \approx 60 - 90\ days.$$

For a power substation at $DR = 0.30$:

$$DT \approx 30 - 60\ days.$$

For an agricultural field at DR = 1.0 (total crop loss):

$$DT = remainder\ of\ the\ growing\ season.$$

These estimates are calibrated from post-event operational databases (EM-DAT, Swiss Re sigma, Zurich Flood Resilience Alliance data).

Step 4 – Compute the revenue loss.

Apply the formula:

$$\Delta Rev = -Rev_0 \times DR \times (DT/365) \times (1 - \rho)$$

In the example below, losses are reported as positive amounts for readability.

Where you can estimate ρ (the recovery rate during DT) from the firm's disclosed inventory levels, supply chain redundancy, and business continuity plan.

As an example, the ρ values can be 0.10 to 0.20 for capital-intensive manufacturing (as it's hard to substitute); 0.30 to 0.50 for logistics (can reroute); 0.60 to 0.80 for digital services (can redirect traffic).

Step 5 – Model the cost structure.

Decompose the firm's cost base into fixed and variable components using sector benchmarks (variable cost ratios by sector: manufacturing 55–65%, retail 70–80%, utilities 40–50%, agriculture 60–70%). Apply:

$$\Delta EBITDA = \Delta Rev \times (1 - VC_{ratio}) - \Delta FC_{opex}$$

Step 6 – Add direct cost escalation.

Estimate repair costs from the damage function, insurance premium uplift from market hardening data, energy cost changes from demand/supply models, and labour cost changes from disruption indices. Sum to obtain $\Delta Opex$ total.

Step 7 – Estimate capex requirements.

Quantify reconstruction capex (DR × RCV, net of insurance), adaptation capex (from engineering assessments of hazard-proofing requirements at the given location), and the replacement investment brought forward for assets that must be retired early. The net book value lost on those assets is an impairment, recorded outside capex.

Step 8 – Apply the IFRS accounting treatment.

Under IFRS, physical damage triggers:

(i) asset write-down to recoverable amount (IAS 36 impairment test);
(ii) recognition of business interruption insurance recoveries once the claim is virtually certain to be received (IAS 37);
(iii) reconstruction asset capitalisation;
(iv) provision for decommissioning if early retirement is likely (IAS 37).

Each of these affects the financial statements differently and in different time periods.

EXAMPLE SCENARIO - AN AUTOMOTIVE MANUFACTURING PLANT IMPACTED BY A FLOOD
Consider a Tier-1 automotive supplier operating an engine components plant in eastern Germany with annual revenue of €450m, Its EBITDA margin is 22% (€99m), and its replacement asset value is €320m. A 1-in-50-year Elbe River flood inundates the plant to a depth of 1.2 metres for 4 days, with receding waters and mud contamination extending effective downtime to 47 days. The plant holds 8 days of finished goods inventory. Insurance covers 80% of property damage and 60% of business interruption losses above a 5-day deductible.

A) Physical damage

Consider the damage ratio DR is at 0.18 from the German flood vulnerability curve for reinforced concrete light industrial at 1.2m flood depth.

$$\textit{Physical damage} = 0.18 \times €320m = €57.6m.$$

$$\textit{Insurance recovery} = 80\% \times €57.6m = €46.1m.$$

$$\textit{Net repair capex} = €11.5m.$$

B) Revenue loss

$$\textit{Effective downtime} = 47\ \textit{days}.$$

$\textit{Recovery rate}\ \rho = 8/47 \approx 0.17$ (as inventory covers 8 of 47 days).

$$\Delta Rev = €450m \times 0.18 \times (47/365) \times (1 - 0.17) \approx €450m \times 0.18 \times 0.129 \times 0.83 \approx €8.7m.$$

Business interruption (BI) insurance recovery is 60% of losses above 5-day deductible.

C) EBITDA impact

Let's consider a variable cost ratio (automotive manufacturing): 62%.

$$\Delta EBITDA_{volume} = €8.7m \times (1 - 0.62) \approx €3.3m \text{ lost contribution.}$$

$$BI\ recovery = 60\% \times €3.3m \times (47 - 5)/47 \approx €1.8m$$

Since the business interruption cover applies to lost profit after the 5-day deductible.

$$\Delta Opex: €2.1m$$

This is cleaning, overtime and temporary facilities.

$$Total\ EBITDA\ impact\ (net\ of\ BI\ recovery): €3.3m + €2.1m - €1.8m = €3.6m.$$

Measured on the pre-event revenue base of €450m, the EBITDA margin falls from 22.0% to 21.2% for the affected year.

D) Capex impact

Net repair capex (after insurance): €11.5m capitalised and depreciated over 7 years (≈ €1.6m additional depreciation p.a.).

The adaptation capex assessed is €3.2m for perimeter flood barriers rated to 1.5m depth (raising the plant's protection from a 1-in-50 to a 1-in-200 flood standard).

Total capex year of event: €14.7m vs. maintenance capex budget of €8.5m.

E) Cascading effects and scale question

The plant supplies 12 original equipment manufacturers (OEMs).

The 47-day production disruption propagates through the supply chain: OEMs drawing on alternative suppliers absorb €3.8m in premium freight costs. Three OEM customers trigger force majeure clauses.

The total supply chain economic loss is estimated at

$2.4 \times €8.7m \approx €20.8m$, using an illustrative supply chain multiplier of 2.4.

This helps to show how a firm-level revenue loss of €8.7m generates a system-level economic impact more than twice as large.

Model Representation Card – P&L Bridge Model – Drought and Water Utility Revenue in Brazil



FIGURE 43. Model representation card – P&L bridge model (the Brazilian case)

Why is this model useful?

This model can be used if you're dealing with a slow-building hazard, drought builds up over months, and you need to know what that gradual water shortage actually costs a utility in lost revenue and business interruption. It fits when the asset in question is a single, well-characterized system, in this case one water supply network serving a metro region and when the output needs to feed directly into a P&L, which goes beyond just a hazard map.

What does this model teach us?

This model teaches us that translating a physical shortage into a financial number requires more steps than the hazard model alone, streamflow and drought deficit have to pass through a threshold method, a frequency distribution, and finally a market price before they become a revenue loss figure. It also introduces something the other cards haven't shown yet, a resilience financial gap, the difference between what adaptation would cost and what inaction is already costing, which reframes the output from a passive damage estimate into something closer to an investment decision.

Provider Box – XDI (The Cross Dependency Initiative): from engineering damage to the P&L bridge

XDI is a physical climate risk analytics provider built around its Climate Risk Engines. They feed the first two steps of the P&L bridge described above, the damage ratio and the failure of an asset, and they also supply the physical layer behind the Quant Foundry model.

Hazard and climate modelling. For each hazard, three types of input are combined: historical climate observations at the asset location, local context data specific to the hazard (for example, soil composition maps for soil movement), and downscaled climate projections. The projections come from regional climate models of the CORDEX project, dynamically downscaled from CMIP5 and CMIP6 global simulations. These inputs are run through probabilistic models to estimate how often each hazard exceeds the damage and failure thresholds of the asset. XDI does not publish the specification of these models, such as the distributions used, the treatment of uncertainty or the confidence intervals. The method can therefore be described in its structure, but its calibration cannot be audited from public material (UNEP FI, 2024; XDI, 2024).

Only hazards that can damage physical assets and that climate change may influence are modelled. XDI's current methodology lists eleven of them: coastal inundation, riverine flooding, surface water flooding, extreme heat, extreme wind, forest fire, freeze-thaw, landslip, soil movement, tropical cyclone wind, and tropical cyclone storm surge. Earlier XDI material refers to ten hazards, with tropical cyclones counted as one. Each hazard is resolved at its own spatial granularity, for example between 5 and 50 metres.

XDI deliberately takes a stress-testing stance on scenarios. It favors high-emission pathways and, for each hazard, retains the regional climate model that produces the largest increase in that hazard.

Engineering vulnerability. XDI does not apply a single damage curve per asset class. Each asset is represented by an archetype, and the engines use engineering-based methods to assess exposure and vulnerability of asset archetypes to understand the likely damage and failure probability of assets caused by extreme weather and climate change hazards. The library covers over 140 individual asset types, and construction elements are analysed one by one, so an element can fail without being damaged, and its failure can cause other dependent elements to fail. Risk ratings are derived from US Federal Emergency Management Agency (FEMA) standards, which are used for pricing many insurance premiums in the USA, adjusted for the resilience of each property.

Financial metrics. The core output is the annual average loss (AAL) of an asset, from which three metrics follow:

$$DR = \frac{AAL}{RC}, \qquad TIP = AAL, \qquad VAR\% = \frac{TIP}{RC}$$

where:

- DR is the damage ratio
- AAL is the annual average loss from all hazards combined
- RC is the replacement cost of the property
- TIP is the Technical Insurance Premium, the AAL per representative property
- $VAR\%$is the value at risk, the TIP as a share of replacement cost over one year, without discounting

Since damage is expressed as a ratio, two properties can show the same damage ratio despite having different replacement values, which makes results comparable across countries. Reports usually show the maximum-to-date value at risk (MVAR), the highest annual VAR% reached up to a given year. This avoids misleading readings when a hazard temporarily decreases in some climate model runs.

Cross-dependency. The feature that gives XDI its name is the analysis of the infrastructure an asset depends on. Cross Dependency Testing identifies supply chain risks for assets. Services such as roads, rail, power, comms and water can all fail leading to losses at invested or owned sites. This is the channel through which a site that suffers no damage can still face business interruption.

Why it matters for the P&L bridge. XDI's damage ratio and failure probability are the inputs DR(I,v) and downtime in Steps 2 and 3 of the white-box guide. Its cross-dependency analysis extends business interruption beyond the site itself. The step from these physical metrics to revenue, opex and EBITDA is left to the user, or to a firm-level layer such as Quant Foundry's.

Key Limitations

- **Downtime estimation is the weakest link.**

DT is extremely sensitive to the facility type, supply chain structure, and restoration resource availability. Post-event studies consistently show that actual downtime exceeds pre-event engineering estimates by factors of 1.5–3.0 for complex industrial facilities. For financial modelling, this introduces a multiplicative uncertainty on the revenue loss estimate that dwarfs the uncertainty in the damage ratio itself.

- **Demand-side effects ignored.**

Most P&L models focus on supply disruption. But physical events also affect demand: a flooded retail area loses customers to alternative channels; a drought reduces consumer spending in rural economies; a prolonged heatwave shifts consumption patterns. Ignoring demand-side effects biases revenue loss estimates, sometimes upward (overestimate recovery) and sometimes downward (underestimate market share loss to competitors with undamaged capacity).

- **Insurance gap assumptions are critical.**

The net financial impact depends heavily on the gap between gross physical loss and insurance recovery. In emerging markets, insurance penetration for commercial property is often below 30%, meaning the firm bears the full gross loss.

Even in developed markets, business interruption sublimits, waiting period deductibles, and named-peril exclusions mean that insured losses are systematically below economic losses. The widening climate insurance protection gap (particularly for flood in Europe and wildfire in the US) makes the gross-to-net gap larger over time.

- **Chronic vs. acute effects require different models.**

The P&L bridge formula above is designed for acute events with a defined onset and recovery. Chronic effects (gradual heat stress, multi-year drought, slow-onset sea-level rise) require a different modelling structure: a continuous drag on productivity, a progressive increase in energy costs, and a gradual deterioration of asset values rather than a discrete damage event. Most firm-level models handle one or the other but not both in an integrated framework.

2.3.2 ENTERPRISE VALUE IMPACT

What is it?

Enterprise value (EV) is the present value of all future free cash flows generated by a firm, discounted at the firm's weighted average cost of capital. It is the canonical measure of a firm's total economic worth independent of its capital structure, and it is the starting point for equity valuation:

$$EV - Net\ Debt = Equity\ Value$$

Where $Net\ Debt = Total\ Debt - Cash\ and\ Cash\ equivalents$ (i.e. what is sufficiently "liquid").

This is the starting point to assess fair value pricing, credit analysis, the recovery rate, the infrastructure and asset valuation.

Physical climate risk affects enterprise value through multiple pathways simultaneously (Worthington and Valadkhani, 2004; Siddikee and Rahman, 2018; Mahalingam et al., 2018), which is what makes EV impact modelling both essential and technically demanding.

A single physical hazard event can:

- Reduce near-term free cash flows (via the P&L channels in 2.3.1);
- Increase the cost of capital (via higher climate risk premiums);
- Shorten the asset's economic life (through stranding);
- Reduce the terminal growth rate (via impaired competitive position);
- Require additional capital expenditure that dilutes returns.

Each of these effects shows up as a reduction in EV, but through different mechanisms and at different points in the Discounted Cash Flow (DCF) timeline.

The EV impact model is therefore able to integrate any framework for all the upstream physical risk models. It can consider the EP curve metrics from 2.1.3, the probability of default uplift from 2.2.2, the EBITDA bridge from 2.3.1, and the macro scenario impacts from 2.4.1. This impact model is flexible, and it assembles these variables into an auditable estimate of the change in firm value under each climate scenario. PlanetView is a good example of such models.

Provider Box – PlanetView

PlanetView is the interface of Planetrics, a climate modelling business acquired by McKinsey in 2021 and sold to SLR Consulting in April 2026. It follows the enterprise value logic of this subsection closely. Physical risk is assessed in a hazard, exposure and vulnerability framework. Direct damages from acute and chronic hazards are estimated bottom up at the asset level.

Chronic effects are also propagated top down through a computable general equilibrium model, which links sectors and countries through trade in intermediate goods. Losses in labour and land productivity, combined with supply and demand elasticities, shift each sector's output and costs in each country. These shifts are allocated to individual companies, which gives a path of future cash flows under each climate scenario. The impacts are summed at company level and discounted against a baseline scenario without climate change:

$$\Delta NPV = \sum_{t=1}^{T} \left(CF_t^S - CF_t^B\right)/(1+r)^t$$

where ΔNPV is the change in the company's net present value caused by climate risk, CF_t^S is its cash flow in year t under climate scenario S,

CF_t^B is its cash flow in the same year under the baseline, r is the discount rate, and T is the horizon.

The same cash flows are then priced differently for each asset class. For listed equities, the change in net present value gives the change in company value.

For corporate bonds, the model can translate the shock into a change in the Altman Z-score, then into a rating change and a probability of default, and discounts the expected bond payments accordingly.

Real estate, sovereign bonds and private equity each have their own valuation step. Scenarios include the NGFS and IEA pathways and the Bank of England's Climate Biennial Exploratory Scenario, among others, and can be customised. Several UK and EU institutions used the model in the climate stress tests of the Bank of England and the European Central Bank, and Aberdeen Investments has used it as its modelling partner since 2020. Two limits are documented. Supply chain disruptions caused by physical hazards are excluded from standard runs, and the models are updated every quarter, so results can move between releases.

Mathematical Representation

CONSIDER THE STANDARD DCF ENTERPRISE VALUE

$$EV_0 = \sum_{t=1}^{T} \frac{FCFF_t}{(1+WACC)^t} + \frac{TV_T}{(1+WACC)^T}$$

- EV_0 is the *baseline enterprise value* today, the value of the firm's operations to all capital providers (equity and debt)
- t is the *year index*, from 1 to T, the last year of the explicit forecast horizon (for example 10 to 30 years)
- $FCFF_t$ is the *free cash flow* to the firm in year t, the cash available to both shareholders and lenders after operating costs, taxes and investment

$$FCFF_t = EBITDA_t - Taxes_t - Capex_t - \Delta Working\ Capital_t$$

- $WACC$ is the discount rate applied to every cash flow. It stands for "weighted average cost of capital". It is commonly measured as:

$$w_E \cdot r_E + w_D \cdot r_D \cdot (1-\tau)$$

Where:

- w_E is the share of equity in total financing, $\frac{E}{(E+D)}$
- w_D is the share of debt in total financing, $\frac{D}{(E+D)}$, so that $w_E + w_D = 1$
- E is the market value of equity
- D is the market value of debt

CONSIDER THE STANDARD DCF ENTERPRISE VALUE

- r_E is the cost of equity, the return shareholders require (usually from the Capital Asset Pricing Model)
- r_D is the cost of debt, the pre-tax interest rate the firm pays on new borrowing
- τ is the corporate tax rate. The factor $(1-\tau)$reflects the tax deductibility of interest
- TV_T is the *terminal value* at year T, the value at T of all cash flows after the forecast horizon

$$TV_T = \frac{FCFF_{T+1}}{WACC - g}$$

- g is the long-run growth rate of $FCFF$ after year T, which must stay below $WACC$
- $FCFF_{T+1}$ is the first cash flow after the forecast horizon

Several authors find that physical climate risk raises the cost of capital. On the equity side, severe drought raises the cost of equity by about 92 basis points (Huynh et al., 2020), and a one standard deviation increase in heat stress exposure raises expected returns on S&P 500 stocks by about 45 basis points (Acharya et al., 2022). For a highly exposed firm, this suggests an equity premium of roughly 45 to 90 basis points, weighted by w_E in the WACC.

On the debt side, climate vulnerability has raised the cost of debt by about 0.68 percentage points across 15,265 firms in 71 countries, although the same study finds little effect on equity (Kling et al., 2021).

CLIMATE-ADJUSTED ENTERPRISE VALUE

$$EV_{climate}(s) = \sum_{t=1}^{T} \frac{FCFF_t + \Delta FCFF_t^s}{(1+WACC_s)^t} + \frac{TV_T^s}{(1+WACC_s)^T} + \Delta SA_s$$

Where $EV_{climate(s)}$ is enterprise value under scenario s

- $\Delta FCFF_t^s$ is the climate-driven change in FCFF in year t
- $WACC_s$ is the climate-adjusted discount rate
- TV_T^S is the scenario terminal value. It can be formalized also as: $\frac{FCFF_{T+1}^S}{(WACC_S - g_S)}$
- ΔSA_s is the stranded asset write-down

The climate adjustment has four additive components operating simultaneously (each is discounted and, where relevant, summed across t to obtain its enterprise-value-level counterpart used in the summary formula below $-\Delta EV_{CF}$ from $\Delta FCFF$, ΔEV_{TV} from ΔTV, while ΔEV_{WACC} and ΔEV_{SA} are already expressed at the enterprise-value level):

A) Cash flow adjustment: $\Delta FCFF$

Consider the sum of the P&L impacts from section 2.3.1:

$$\Delta FCFF_t^s = \Delta EBITDA_t^s - \Delta Taxes_t^s - \Delta Capex_t^s - \Delta NWC_t^s$$

- Each Δ term is the scenario value minus the baseline value in year t
- The ΔNWC_t^s is the change in investment in net working capital

CLIMATE-ADJUSTED ENTERPRISE VALUE

This component is quite simple to compute annually because it's anchored in the firm's disclosed financials. Each term above (EBITDA, Taxes, Capex, Working Capital) should be read as a climate-driven change versus the baseline path, i.e. $\Delta EBITDA_t$, $\Delta Taxes_t$, $\Delta Capex_t$, $\Delta(\Delta Working\ Capital)_t$ — not the baseline levels themselves.

B) WACC adjustment: $\Delta WACC$

Physical climate risk increases the required return on both equity (higher equity risk premium for climate-exposed firms) and debt (higher credit spread from PD uplift in section 2.2).

A higher WACC compresses enterprise value (EV) by increasing the discount rate applied to all future cash flows:

$$\Delta EV_{WACC} \approx -EV_0 \times \frac{\Delta WACC}{WACC - g}$$

- $\Delta WACC$ is equal to $WACC_s - WACC$
- $\frac{1}{WACC - g}$ is the duration of the cash flow stream

C) Terminal value adjustment: ΔTV

Chronic climate risk impairs the long-run growth rate g of the business, either by reducing the addressable market (sea-level rise stranding coastal real estate), reducing productivity (heat stress in high-temperature geographies), or increasing the cost of sustaining productive capacity (rising adaptation capex requirements). Hence:

$$\Delta TV = \frac{FCFF^s_{T+1}}{WACC_s - g_s} - \frac{FCFF_{T+1}}{WACC - g}$$

$$\Delta EV_{TV} \approx \frac{\Delta TV}{(1 + WACC)^T}$$

- $FCFF^s_{T+1}$ is the first post-horizon cash flow under scenario s
- g_s is the long-run growth rate under scenario s

D) Stranded asset write-down:

$\boldsymbol{\Delta SA}$ assets that become economically non-viable before their planned retirement must be written down to zero or salvage value.

$$\Delta SA_s = -NBV_{stranded} \times P(stranding|s) \times \frac{L_{lost}}{L_{rem}}$$

The formula above carries a remaining-life factor (remaining life lost / total remaining life), multiplying the $NBV \times P(stranding)$ term (as used in the worked example below, which applies a 15/25 factor rather than writing off the full NBV).

- $NBV_{stranded}$ is the net book value of the exposed asset
- $P(stranding \mid s)$ is the probability the asset becomes non-viable under scenario s
- L_{lost} is the remaining life lost (years)
- L_{rem} is the total remaining book life (years)

CLIMATE-ADJUSTED ENTERPRISE VALUE

For example, if you consider a coastal industrial facility with a 25-year remaining book life and a 40% probability of becoming uninsurable within 15 years under SSP5-8.5:

$$\Delta SA \approx -NBV \times 0.40 \times \left(\frac{10}{25}\right) = -16\% \text{ of the Net Book Value}$$

EV impact - Summary

$$\boldsymbol{\Delta EV(s) = \Delta EV_{CF}(s) + \Delta EV_{WACC}(s) + \Delta EV_{TV}(s) + \Delta EV_{SA}(s)}$$

Expressed as a percentage of baseline EV: $\frac{\Delta EV(s)}{EV_0}$.

This is the metric used by sell-side climate equity research, by the MSCI Climate VaR. It is also used in climate scenario analysis.

Consider a range for a highly exposed coastal industrial firm under SSP5-8.5 by 2050 as follows:

$$\boldsymbol{\Delta EV/EV_0 \in [-35\%, -15\%]}$$

However, for a low-exposure diversified technology firm:

$$\Delta EV/EV_0 \in [-2\%, -0.5\%].$$

THE WACC ADJUSTMENT IN DETAIL

$$\boldsymbol{WACC_s = w_E \cdot \left(r_f + \beta \cdot MRP + \pi_{climate,s}\right) + w_D \cdot \left(r_f + CS_s\right) \cdot (1 - \tau)}$$

Where:

- $\pi_{climate,s}$ is the physical climate risk premium on equity under scenario s
- $\beta_{climate,s}$ is the climate-adjusted equity beta: the standard market beta augmented by a physical climate risk loading. Empirically estimated from the cross-section of equity returns versus climate exposure scores.

For high exposure firms, this measure can be estimated most of the time $\beta_{climate} \approx 0.05 - 0.20$ above baseline (Laidlaw et al., 2023 – S&P).

The consequence of this is a raw uplift of $0.3 - 1.2\%$ at a 6% MRP before scaling by the equity weight; the actual WACC contribution is $w_E \times \beta_{climate} \times MRP$, which is smaller than the raw product unless the firm is 100% equity-financed.

- CS_s is the credit spread under scenario s. It is derived from the PD uplift model as follows:

$$\boldsymbol{CS_s = CS_0 + \Delta CS_s}$$

$$\boldsymbol{\Delta CS_s \approx \frac{\Delta PD \times LGD}{1 - PD_0} = \frac{0.8\% \times 45\%}{1 - 1\%} \approx 36\ bps}$$

- CS_0 is the baseline credit spread

THE WACC ADJUSTMENT IN DETAIL

- ΔPD is the climate-driven increase in probability of default
- LGD is loss given default
- PD_0 is the baseline probability of default

For a firm with PD of 1% that rises to 1.8% under SSP5-8.5, and LGD of 45%.

Note that this yields only the climate-driven uplift ΔCS_s, the total scenario credit spread used in the WACC formula above is the firm's baseline spread plus this uplift, $CS_{baseline} + \Delta CS_s$, not ΔCS_s alone).

White-Box approach to climate-Adjusted DCF: a step-by-step guide

Step 1 – Build the baseline DCF.

Construct a standard 10–30-year DCF using consensus revenue growth, sector EBITDA margin, maintenance capex as a percentage of revenue, and a normalized WACC. The baseline represents the firm's value absent physical climate risk beyond what is already priced by the market.

Step 2 – Overlay the FCFF adjustment path.

For each scenario (SSP1-2.6, SSP2-4.5, SSP5-8.5), compute the annual $\Delta FCFF_t$ using the P&L bridge from 2.3.1.

Acute events are modelled as point-in-time shocks in specific years (drawn from the OEP curve: a 1-in-20 event occurs in some year of the 30-year horizon with probability

$$P = 1 - (1 - 1/20)^{30} \approx 78\%$$

Chronic effects are modelled as smooth drags applied to the revenue and cost growth rates.

Step 3 – Adjust the WACC.

Apply the climate beta and credit spread adjustments. The WACC adjustment is scenario-specific: more warming implies higher physical risk, higher climate beta, wider credit spread, higher WACC.

The effect of a 50 bps WACC increase on a firm with a 10-year average DCF duration is approximately a 4–5% reduction in EV, independently of any cash flow impact.

Step 4 – Adjust the terminal value.

Reduce the long-run growth rate g by the chronic productivity drag $\delta_{climate}$ (scenario). Increase the terminal year capex intensity by the ongoing adaptation cost requirement.

The terminal value change is highly sensitive to g: a 0.5% reduction in g from 2.0% to 1.5% reduces TV by approximately 8% when

$$WACC = 8\% \left(\frac{TV_{new}}{TV_{old}} = \frac{8\% - 2.0\%}{8\% - 1.5\%} \approx 0.92 \right)$$

This amplifies even a small chronic drag into a large EV impact.

Step 5 – Add stranded asset adjustments.

For assets with material stranding probability under the scenario, compute the expected NBV write-down as:

$$NBV = P(stranding) \times NBV_{stranded}$$

Sources for the stranding probability can be considered as regulatory risk (new zoning restrictions, building codes), insurability risk (assets in zones where insurance is projected to become unavailable), and physical risk (assets below projected flood levels or in wildfire interface zones with no viable adaptation path).

Step 6 – Scenario weighting and expected EV.

Weight the scenario-specific EV impacts by scenario probability (or apply all scenarios as sensitivity analysis without weighting, as recommended by TCFD for scenario uncertainty).

$$\text{Expected } \Delta EV = \sum_s P(s) \times \Delta EV(s).$$

Present it as a range, given the scenario uncertainty.

Step 7 – Per-share and per-bond unit.

- $\Delta E = \Delta EV$ as debt value is senior and absorbs EV loss last in a going-concern firm.

Convert the EV impact to equity value impact to loan-level collateral value impact measured as the change in the loan-to-value (ΔLTV) ratio. These are the quantities used by equity investors and lenders respectively for portfolio climate stress testing.

$$\Delta LTV = -\frac{\Delta EV}{Debt}$$

Provider Box – Climate VaR: a fight for the right definitions

Schneider Electric's ClimVaR

ClimVaR is the financial valuation engine of Resource Advisor+ (RA+), the digital solution of Schneider Electric Advisory Services (SEAS), the climate consultancy of Schneider Electric's sustainability business. It sits on top of a physical exposure engine, called ECLR from 2017 and renamed RA+ Climate Risk in 2026, which covers 27 of the 28 hazards listed in the EU Taxonomy. The financial layer quantifies five hazard families: droughts, floods, wildfires, tropical and extra-tropical wind, and extreme heat. Its main users are the corporate finance departments of large companies, and its stated aim is to bring physical risk to the chief financial officer as figures in the income statement.

The model is a climate-adjusted discounted cash flow. Each hazard is translated into a change in one line of the income statement, and these changes are aggregated into a change in the value of the entity (Epelbaum, 2026).

$$\Delta V = \sum_{t=2025}^{T} \frac{\Delta R_t - \Delta OPEX_t - \Delta COGS_t}{(1 + WACC + \Delta WACC)^t} + \Delta TV$$

Where:

- ΔV is the change in the value of the entity caused by physical hazards, t is the year, from 2025 to the horizon T (often 2050),
- ΔR_t is the revenue lost in year t, mainly through business interruption,
- $\Delta OPEX_t$ is the increase in operating expenses, where building damage enters through an insurance premium proxy,
- $\Delta COGS_t$ is the increase in the cost of goods sold,
- $WACC$ is the weighted average cost of capital,
- $\Delta WACC$ is the risk premium investors add for climate exposure,

- ΔTV is the change in terminal value beyond T.

Direct damage to buildings and equipment enters as an operating cost, through the insurance premium that would cover it. Indirect damage enters as lost revenue. Business interruption functions are combined across hazards and spread along the supply chain with a multi-regional input-output model, so a flood at a supplier can reduce the revenue of its clients. A Shapley decomposition then splits the revenue loss across hazards, supply-chain ranks and risk factors.

Uncertainty is handled by Monte Carlo simulation over the climate and financial parameters, and results are reported as value at risk and expected shortfall. In the paper's worked example, better financial and spatial data reduce the VaR from 6.64 billion to 4.83 billion dollars, about 30%.

The unit of analysis is the site, described by more than ten systems such as buildings and equipment. Sites roll up to an entity defined by a country and a sector, then to the company, and finally to a portfolio. The minimum input is the revenue of each entity. When only revenue is provided, the other financial ratios come from country and sector benchmarks (Damodaran datasets, with at least 50 companies per country and sector). The precision of the result therefore grows with the share of the firm's own accounts fed into the model.

Adaptation measures are scored by comparing their cost, CAPEX plus OPEX, with the VaR they remove, which produces a ranking of interventions by marginal abatement cost. The same logic applies to site selection, by comparing the climate-adjusted value of locating a new plant in one country or another.

Two limits follow from this design. Balance sheet effects, namely asset impairment and loan loss provisions, fall outside the model's scope. And when client accounts are missing, results rest on sector and country averages, not on the firm's own cost structure.

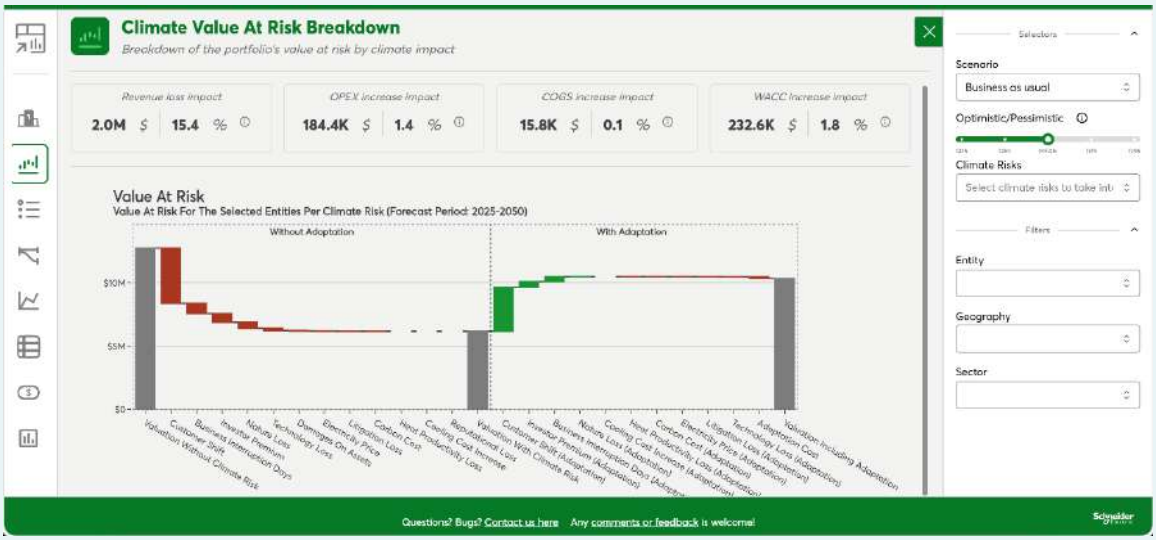


FIGURE 44. Schneider Electric Advisory Services' climate risk dashboard (formerly ECLR, now RA+ Climate Risk).

The MSCI case for a Climate VaR

Remember that traditionally, a value at risk is a quantile: the loss that is exceeded with a given probability over a given horizon. ClimVaR follows this definition. It simulates a distribution of losses and reports its quantile (VaR) and the mean loss beyond it (expected shortfall), and the paper explicitly reconciles its terms with these standard definitions. MSCI's Climate VaR is a scenario-conditional change in value: the percentage by which a company's discounted cash flows change under a given climate pathway. It carries no probability level, so it

is closer to a stress test result than to a VaR in the statistical sense. The same label therefore covers a risk measure in one case and a scenario impact in the other.

MSCI's Climate VaR (originally built by Carbon Delta, a Swiss startup MSCI acquired in 2019) is a discounted cash flow reframing: it takes a company's current valuation, understood as the present value of its expected future free cash flows, and subtracts the present value of new climate-related costs, or adds the present value of climate-related opportunities. The output for each company and scenario is a single percentage, how much the company's value changes once those costs or opportunities are priced in.

It has three components that sum together:

- Policy Risk Climate VaR (cost of decarbonization under a given carbon price path),
- Technology Opportunity Climate VaR (upside from low-carbon technology or patents),
- Physical Risk Climate VaR (asset-level hazard exposure translated into damage and business interruption costs, aggregated up to the company).

Physical risk modelling requires geolocated asset data, and the whole model needs company-level revenue, market cap, and WACC to run. It covers over 10,000 listed companies and 9,000-plus sovereign bonds, run under NGFS scenarios (Net-Zero 2050, Delayed Transition, Nationally Determined Contributions, and so on).

What is not disclosed however is the actual damage function library and the scenario-to-cash-flow translation model itself. In that sense it's a partial black box, transparent about structure, opaque about the internals that actually drive the number.

The fight for the rightful definition

Both descriptions rely on the providers' own material: a methodological paper for ClimVaR, public documentation for MSCI.

- **Reproducibility.** ClimVaR is documented in a paper that derives each equation, names each parameter and traces each sub-model to its source, for example heat-related productivity loss to Foster et al. (2021). MSCI publishes the structure, inputs and scenarios of its Climate VaR but not its damage functions or its translation from scenarios to cash flows. This protects its intellectual property but limits independent replication.
- **Physical risk scope.** ClimVaR's financial layer quantifies five hazard families (droughts, floods, wildfires, tropical and extra-tropical wind, extreme heat). Its exposure engine covers 27 of the 28 hazards in the EU Taxonomy. MSCI's physical component also covers a set of acute and chronic hazards, alongside a transition risk component and a technology opportunity component within the same metric. ClimVaR's documentation focuses on physical risk.
- **Supply chain.** ClimVaR propagates business interruption through a multi-regional input-output model, so a flood at an upstream supplier reduces the revenue of its clients. Shapley values attribute the loss to each hazard and each supply-chain rank. As far as its public documentation shows, MSCI's physical component works from the company's own geolocated assets, without a comparable treatment of upstream exposure.
- **Uncertainty.** ClimVaR runs Monte Carlo simulations over climate and financial parameters, reports results as VaR and expected shortfall, and shows how better data reduce the VaR. MSCI's public materials report results across NGFS scenarios. This captures scenario uncertainty but does not decompose it within a scenario.
- **Coverage and use.** MSCI covers a large universe of listed companies and bonds with pre-computed results, ready to query without in-house modelling. ClimVaR runs within a client-facing tool (RA+) fed

with each firm's own sites and accounts. Its precision depends on the data the client provides, and it falls back on sector and country benchmarks when accounts are missing.

The two tools therefore serve different needs. ClimVaR suits a firm or advisor that wants an auditable, site-level valuation of its own physical risk, including supply-chain effects. MSCI's Climate VaR suits an investor who needs comparable figures across a whole portfolio without building a model.

Model Representation Card – Enterprise Value Impact – Extreme Heat and Firm Performance

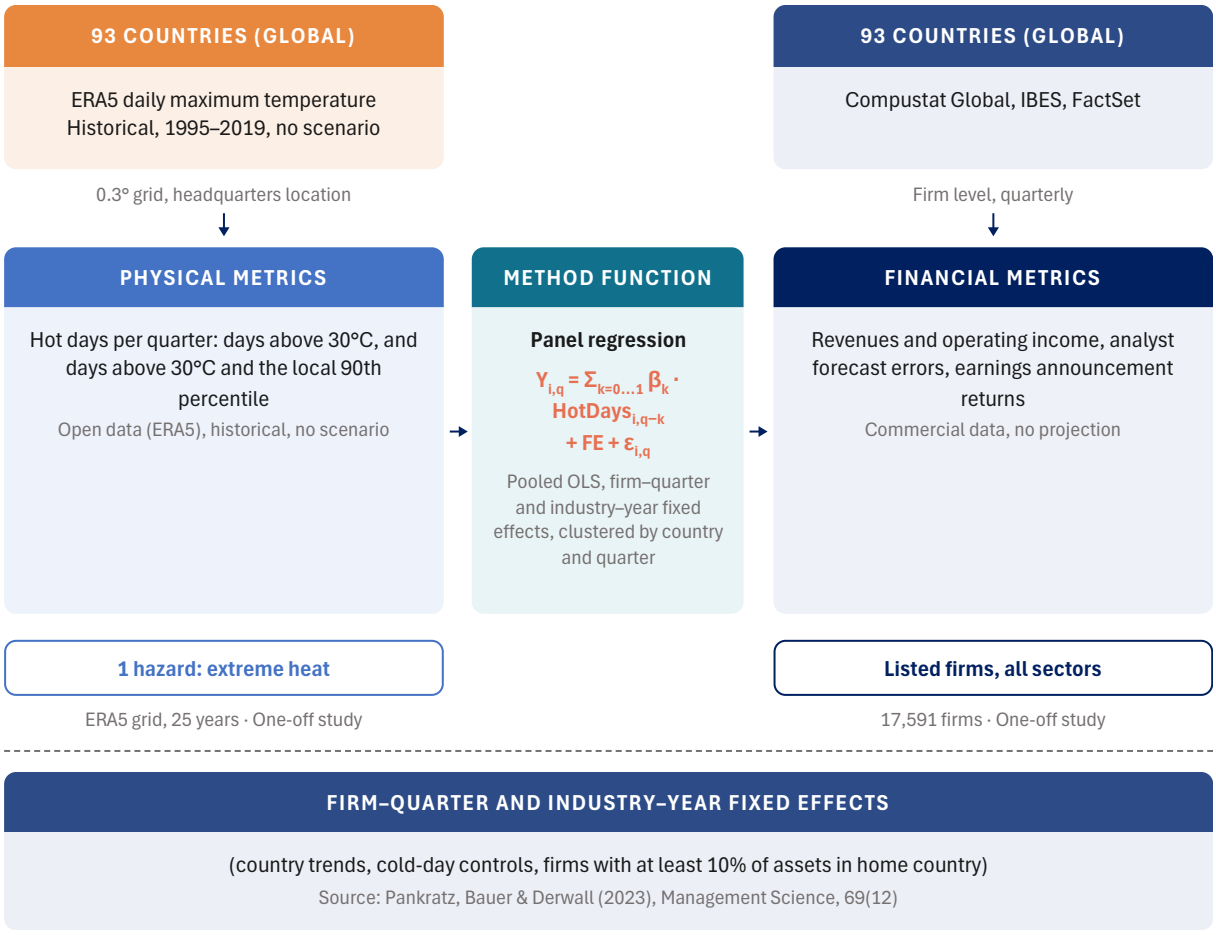


FIGURE 45. Model representation card – Enterprise value impact (extreme heat and firm performance, Pankratz et al., 2023)

What it does

The effect of extreme heat on firm performance and market valuation is estimated for 17,591 listed firms in 93 countries between 1995 and 2019 (Pankratz et al., 2023). Each firm's heat exposure is measured as the number of hot days per financial quarter at its headquarters, using ERA5 daily maximum temperatures. A day counts as hot when it exceeds 30°C, or 30°C and the local 90th percentile. Quarterly revenues and operating income are regressed on current and lagged heat exposure, with firm–quarter and industry–year fixed effects and country trends.

Heat is found to reduce both revenues and operating income: a one-standard-deviation increase in the number of hot days decreases revenues (operating income) by 0.6% (1.8%) of the average quarterly revenue. The analysis is then extended to markets. As heat exposure rises, analyst forecast errors and earnings announcement returns

become more negative, which shows that capital market participants do not fully anticipate the economic consequences of heat as a first order physical climate risk.

Why is this model useful?

This model can be used when you need empirical evidence that physical risk already affects the cash flows that drive enterprise value, and that markets misprice it. It fits a global, multi-sector portfolio, since it covers almost every listed firm with available data. Its open climate data and standard regression design also make it easy to replicate. Its estimates can calibrate the ΔEBITDA term of a climate-adjusted DCF for heat exposure.

What does this model teach us?

This model teaches us that heat hits operating income about three times harder than revenues, because costs do not fall when sales do. The EV effect therefore runs mainly through margins.

It also teaches us that analysts and investors are surprised by these losses. If the market does not yet price heat, current valuations overstate the enterprise value of exposed firms, and the gap closes only as losses are reported.

The trade-off is precision of location. Exposure is measured at the headquarters, not at the plants where heat actually disrupts production, which the authors expect to understate the true effect. The model also stops at historical data and carries no climate scenario.

Model Representation Card – Enterprise Value Model – Valuation under Climate Uncertainty



FIGURE 46. Model representation card – Enterprise value model (valuation under climate uncertainty, Wang et al., 2025)

What it does

The way climate-related exposures and disclosure practices are reflected in firm value is examined for non-financial firms across the Eurozone between 2016 and 2023 (Wang et al., 2025). Firm value is measured in three ways: Tobin's Q, EV/EBITDA and residual income, so the result does not rest on a single valuation metric. Physical climate risk, carbon intensity and structured climate disclosure are tested as drivers of value.

Both physical climate risk and carbon intensity are found to be negatively associated with firm value. Firms that disclose climate information in a structured way receive a valuation premium, especially when they operate in high-risk contexts.

Why is this model useful?

This model can be used when you need evidence that physical climate risk is already reflected in enterprise value multiples, not only in cash flows or returns. EV/EBITDA is the metric analysts use to price firms, so the result speaks directly to valuation practice. It fits a European portfolio of listed non-financial firms and supports the case that disclosure itself carries financial value.

What does this model teach us?

This model teaches us that markets price physical risk at the level of valuation multiples. Exposed firms trade at a discount on several measures of value at once.

It also teaches us that disclosure partly offsets that discount. Investors seem to reward firms that make their exposure readable, especially where the risk is high, which suggests that part of the penalty comes from uncertainty more than from the exposure itself.

The trade-off is depth. The study shows an association between exposure and value, but the public material does not explain how physical risk is measured or which hazards drive the effect, so it cannot be used to calibrate a hazard-specific EV adjustment.

Key Limitations

- **Terminal value sensitivity dominates.** In most DCF models, 60–80% of the total EV is in the terminal value. A small change in g or $WACC$ in the terminal period swamps the near-term cash flow adjustments. This means that EV climate impact estimates are highly sensitive to assumptions about the very long-run behaviour of both the firm and the climate (the least known quantities). Sensitivity tables showing ΔEV as a function of $(g, WACC)$ pairs should always accompany the central estimate.
- **Market already prices some physical risk.** To the extent that equity and credit markets already incorporate physical climate risk into prices (through climate beta, coastal property discounts, insurance sector repricing), the ΔEV computed relative to a baseline that assumes no climate risk will overestimate the incremental information. The baseline DCF should use current market-implied parameters (where observable) rather than historical averages that may not embed climate risk.
- **Discount rate and climate scenario correlation.** Severe physical climate scenarios (SSP5-8.5) imply not only higher physical damage but also higher macro volatility, potentially higher risk-free rates (inflationary impacts of climate damages) and higher equity risk premiums. Running a high-damage scenario with the same WACC as the baseline is internally inconsistent. Full consistency requires a macro-financial scenario that co-determines physical damage, GDP path, interest rates, and equity risk premiums simultaneously.
- **Stranding probability is highly uncertain.** The probability that a specific asset becomes economically non-viable before end of planned life depends on future policy (zoning, insurance regulation, building codes), future physical hazard evolution, and future adaptation technology costs — all deeply uncertain at the asset level. Point estimates of stranded asset values should be presented as ranges spanning at least an order of magnitude.

2.3.3 FINANCIAL RATIOS

Why financial ratios matter for physical climate risk

Financial ratios are the operational language of credit analysis, equity valuation, regulatory supervision, and internal performance management. They are the quantities that trigger covenant breaches, credit rating migrations, IFRS 9 stage reclassifications, regulatory capital surcharges, and investment committee decisions.

Understanding how physical climate risk propagates through financial ratios is therefore understanding the mechanism by which physical risk becomes observable and actionable in capital markets.

The models in sections 2.3.1 and 2.3.2 produce the underlying P&L and EV changes. This section shows how those changes flow through to the specific ratios that practitioners monitor, and provides the mathematical machinery for computing ratio sensitivity to physical climate shocks.

Mathematical Representation: Climate-Adjusted Financial Ratios

RETURN ON EQUITY (ROE) AND ITS CLIMATE DECOMPOSITION

$$RoE = \frac{Net\ Income}{Equity} = \left(\frac{Net\ Income}{Rev}\right) \times \left(\frac{Revenue}{Assets}\right) \times \left(\frac{Assets}{Equity}\right)$$

The DuPont decomposition separates RoE into three components: net profit margin, asset turnover, and financial leverage. Physical climate risk affects all three:

$$\Delta RoE \approx (\Delta Net\ Margin\ /\ Net\ Margin_0) \times RoE_0 + (\Delta Turnover\ /\ Turnover_0) \times RoE_0 + (\Delta Leverage\ /\ Leverage_0) \times RoE_0$$

Net margin channel: EBITDA compression (from 2.3.1) flows through depreciation and interest to reduce net income. For a 20% EBITDA compression:

$$\frac{\Delta NI}{NI_0} \approx \left(\frac{\Delta EBITDA}{EBITDA_0}\right) \times \frac{EBITDA_0(1-\tau)}{NI_0}$$

where NI is net income, $\Delta EBITDA/EBITDA_0$ is the relative change in EBITDA (−20% in this example), and τ is the tax rate.

It is amplified by the ratio of EBITDA to net income, which averages about 1.5 after tax for US non-financial firms (EBITDA margin of 17.4% against a net margin of 8.6%) but varies widely across sectors, rising in capital-intensive and highly indebted industries (Damodaran, 2026).

Asset turnover channel: Physical damage reduces the productive asset base temporarily (assets decline as write-downs occur) while revenue falls. The net effect on turnover depends on whether asset write-downs or revenue loss dominates.

Leverage channel: Capex escalation funded by debt increases the asset/equity ratio, increasing leverage and amplifying RoE movements in both directions.

RETURN ON ASSETS (ROA) AND RETURN ON INVESTMENT (ROI)

$$RoA = \frac{Net\ Income}{Total\ Assets}$$

$$\Delta RoA \approx \frac{\Delta\ Net\ Income}{Assets_0} - Net\ Income_0 \times \Delta\ Assets\ /\ {Assets_0}^2$$

Physical damage reduces the numerator (lower net income) and may reduce or increase the denominator (asset write-downs reduce assets, but adaptation capex increases them). The net direction is not always obvious: a large adaptation programme can inflate the asset base faster than income recovers, compressing RoA below baseline for multiple years post-event.

$$ROIC = \frac{EBIT(1-\tau)}{Equity + Debt - Cash}$$

where $EBIT$ is earnings before interest and taxes, τ is the tax rate, and $Equity + Debt - Cash$ is the capital invested in operations (equity plus net debt).

RoI (or $ROIC$) is the preferred metric for assessing whether a firm's return on its total capital base exceeds its WACC (the condition for value creation).

Physical climate risk compresses the Net Operating Profit After Taxes (NOPAT) via EBIT reduction and simultaneously inflates Invested Capital via adaptation capex. The result is a double compression of ROIC, and when $ROIC < WACC_s$ the firm is destroying value under the climate scenario.

LEVERAGE RATIOS: DEBT/EQUITY AND DEBT/EBITDA

Net Leverage = Net Debt / EBITDA

$$\Delta\left(\frac{ND}{EBITDA}\right) = \frac{\Delta ND \cdot EBITDA_0 - ND_0 \cdot \Delta EBITDA}{{EBITDA_0}^2}$$

where ND is net debt.

Physical events increase leverage in two simultaneous ways: debt rises (reconstruction and adaptation capex financed by new borrowing if insurance recovery is insufficient) and EBITDA falls (the P&L bridge).

The combined effect is super-additive: a 15% EBITDA compression combined with 20% debt increase raises $Net\ Debt/EBITDA$ by about 41% ($1.20/0.85$), against 35% under the first-order approximation. This is the most watched covenant metric in corporate credit and the most likely trigger for rating migration following a major physical event.

INTEREST COVERAGE RATIO (ICR) AND COVENANT BREACH

$$ICR = EBITDA / Interest\ Expense$$

$$\Delta ICR = \frac{\Delta EBITDA}{Interest_0} - \frac{EBITDA_0 \cdot \Delta Interest}{Interest_0^{\,2}}$$

The interest coverage ratio is, with debt-to-EBITDA, one of the two most common financial covenants in loan agreements: debt-to-EBITDA is the most frequent (Bordeman & Demerjian, 2022), and interest coverage covenants appear in 84% of the firms with covenants in the DealScan database (Greenwald, 2019). Thresholds are negotiated loan by loan.

An EBITDA compression of 20% moves ICR from 3.0× to 2.4×, potentially triggering covenant review. If adaptation capex is funded by new debt at higher rates (post-event market hardening), ΔInterest > 0 further compresses ICR.

Covenant breach probability under physical climate scenario s:

$$P(breach \mid s) = \Phi\left(\frac{Threshold - ICR_0 - E[\Delta ICR_s]}{\sigma_{\Delta ICR,s}}\right)$$

where Φ is the standard normal cumulative distribution function, $Threshold$ is the covenant level, ICR_0 is the current ratio, $E[\Delta ICR_s]$ is the expected change in the ratio under scenario s, and $\sigma_{\Delta ICR,s}$ is its standard deviation, reflecting uncertainty in the $EBITDA$ path.

$P(breach)$ can inform the IFRS 9 assessment of a significant increase in credit risk (Stage 2), where covenant breaches are listed among the indicators, and can be reported as a climate stress test output.

LOAN-TO-VALUE (LTV): THE COLLATERAL DIMENSION

$$LtV = Loan\ Balance\ / Property\ Value = L / V$$

$$\Delta LtV / LtV_0 \approx -\Delta V / V_0$$

(for a fixed loan balance, first-order approximation)

- LtV is the dominant risk metric for mortgage lending and real estate finance. Physical climate risk affects LtV through collateral value deterioration. The value of a climate-exposed property V declines when:

Insurance becomes unavailable or unaffordable: in jurisdictions where lenders require hazard insurance, such as the United States, loss of insurability can block refinancing and resale, pushing V towards distressed values.

Flood risk becomes priced by the market: homes exposed to sea-level rise already sell at a discount of about 7% compared with similar unexposed properties in the United States (Bernstein et al., 2019).

Physical damage from an acute event: repair costs reduce the net value of the collateral. For a property with:

V = €400k,

LOAN-TO-VALUE (LTV): THE COLLATERAL DIMENSION
a flood damage DR = 0.15 implies gross repair cost = €60k, reducing V to €340k (absent insurance). LtV rises from 75% to 88%. $$LtV_s(t) = L_t / (V_0 \times (1 - D_s(t)) \times (1 - \pi_s(t)))$$ • $D_s(t)$ is the cumulative physical damage fraction to the property by time t under scenario s, • $\pi_s(t)$ is the climate risk market discount applied by buyers at time t under scenario s (the stigma effect: market repricing independent of actual physical damage) The regulatory implication: under Basel III/CRR, mortgages with LtV > 80% attract higher risk weights. Under IFRS 9, a marked rise in LtV can indicate a significant increase in credit risk and support a move to Stage 2 (lifetime expected credit loss).

RATIO	FORMULA	CLIMATE IMPACT CHANNEL	DIRECTION	COVENANT / REGULATORY TRIGGER
RoE	Net Inc / Equity	EBITDA compression × DuPont leverage amplifier	↓ (compression)	Below hurdle rate: investor redemptions, management KPI breach
RoA	Net Inc / Assets	Net income ↓; assets ↑ (capex) or ↓ (write-down)	Ambiguous, usually ↓	Below peer median: rating pressure, activist risk
RoI / ROIC	NOPAT / InvCap	NOPAT ↓; InvCap ↑ (adaptation capex)	↓ (double compression)	ROIC < WACC: value destruction signal; impairment trigger
Net Debt/EBITDA	Net Debt / EBITDA	Debt ↑ (capex); EBITDA ↓ (damage)	↑ (leverage rises)	Breach at 4.0×–6.0× (sector): rating downgrade; covenant waiver
ICR	EBITDA / Interest	EBITDA ↓; Interest ↑ (new debt at higher rates)	↓ (coverage falls)	Breach at 2.0×–3.5×: lender acceleration; IFRS 9 Stage 2
LtV	Loan / Property Value	Property value ↓ (damage + stigma discount)	↑ (LtV rises)	Breach at 80%: higher risk weight (Basel); Stage 2 trigger; margin call

RATIO	FORMULA	CLIMATE IMPACT CHANNEL	DIRECTION	COVENANT / REGULATORY TRIGGER
Debt/Equity	Total Debt / Equity	Equity ↓ (write-downs); Debt ↑ (capex)	↑ (leverage rises)	Balance sheet covenant; equity raise trigger
Free Cash Flow Yield	FCF / Mkt Cap	FCF ↓ (lower EBITDA, higher capex)	↓	Below debt service: distress signal; dividend cut

TABLE 19. The ratio impact summary

IFRS 9 Stage Migration: what if a loan moves from stage one to stage 2?

The most operationally significant consequence of ratio deterioration from physical climate events is IFRS 9 stage migration for banks. Under IFRS 9, a loan moves from Stage 1 (12-month ECL) to Stage 2 (lifetime ECL) when there is a significant increase in credit risk since origination. The criteria for 'significant increase' include: PD doubling, ICR falling below covenant threshold, LtV rising above 80%, or the borrower's credit rating deteriorating by 2+ notches.

Physical climate events can trigger Stage 2 migration for an entire geographic concentration of loans simultaneously (a regional flood affecting all properties in a catchment area; a drought affecting all agricultural loans in a province). The provisioning impact is then not a single-loan issue but a portfolio event: ECL provisions jump from 12-month to lifetime basis for potentially thousands of loans at once, creating a cliff-edge P&L impact for the bank that bears no relation to the underlying default timing.

It can be formalized as follows:

$$ECL_{Stage\,2} = PD_{lifetime} \times LGD \times EAD \times DF$$

$$ECL_{Stage\,1} = PD_{12m} \times LGD \times EAD \times DF$$

$$\Delta Provision = (PD_{lifetime} - PD_{12m}) \times LGD \times EAD \times DF$$

For a €1m mortgage with

$$LGD = 25\%,$$

$$PD_{12m} = 0.5\%,$$

$$PD_{lifetime} = 3.5\%,$$

The discount factor $DF = 0.85$ is reflecting discounting over the 25-year remaining life:

$$\Delta Provision = (3.5\% - 0.5\%) \times 25\% \times €1m \times 0.85 = €6,375$$

Multiplied across a €5bn mortgage book with 15% of loans migrating to Stage 2 following a major flood event:

$\Delta\, Total\ Provision \approx €4.8m$ in the quarter of the event.

Model Representation Card – Financial Ratio Model – Floods and Bank Solvency

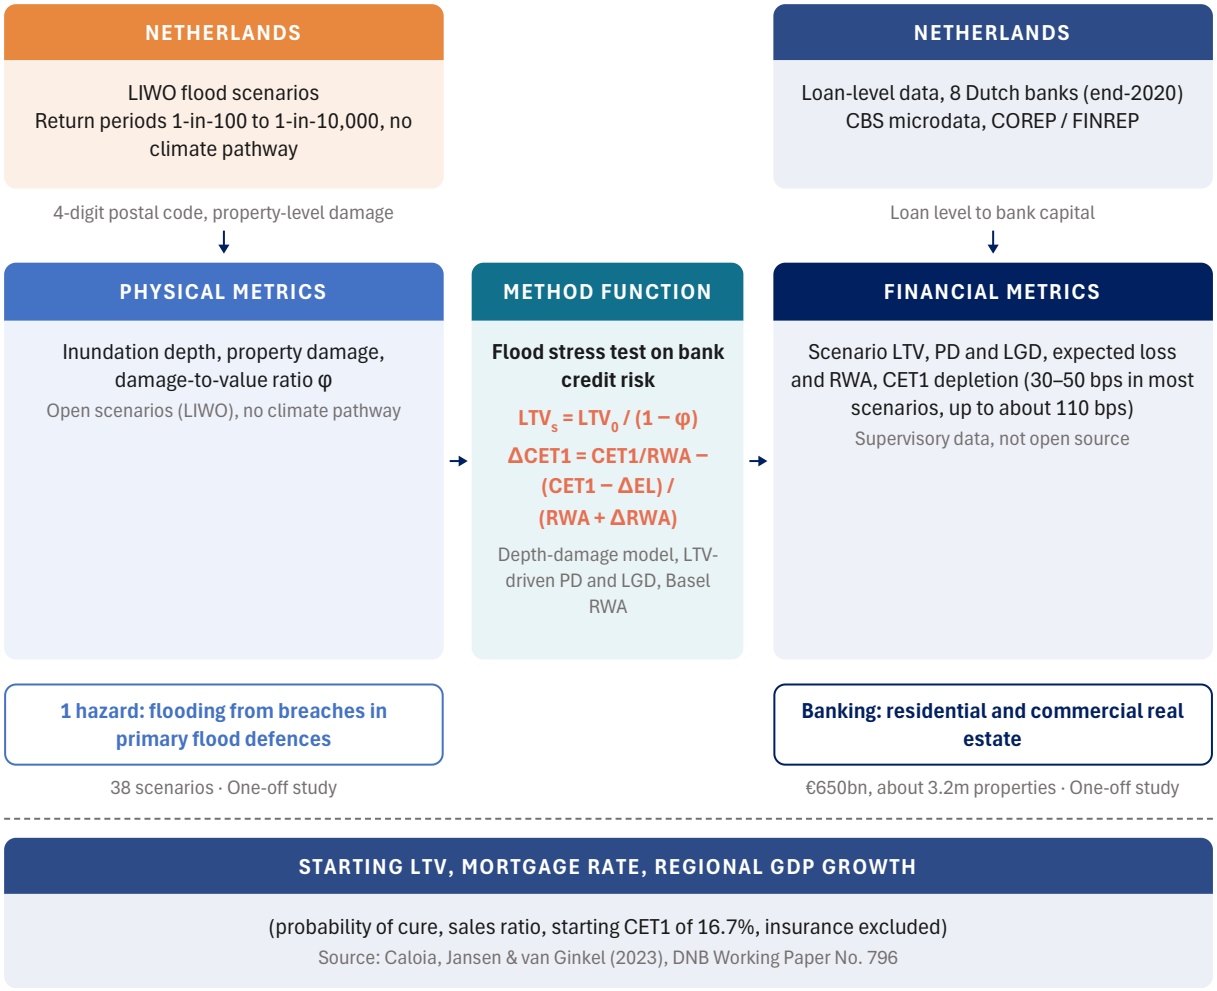


FIGURE 47. Model representation card – Financial ratio model (floods and bank solvency, Caloia et al., 2023)

What it does

The effect of floods on bank solvency is assessed through the credit risk channel, using the Netherlands as a case where part of the country lies below sea level and where insurance policies often exclude damage from breaches in primary flood defences (Caloia et al., 2023). Thirty-eight adverse flood scenarios are taken from the open LIWO database, with return periods between 1 in 100 and 1 in 10,000 years. Inundation depths are matched at postal-code level to loan-level data covering about €650 billion of residential and commercial real estate exposures held by eight Dutch banks.

Property damage is computed with a depth-damage model and expressed as a share of property value. That share raises each loan's LTV, which feeds a PD model estimated on LTV, mortgage rates and regional growth, and an LGD model based on collateral recovery. The resulting changes in expected loss and risk-weighted assets are translated into a change in the CET1 ratio. In most scenarios, capital falls by 30 to 50 basis points. The most extreme scenario brings the decline close to 110 basis points. Most of the effect comes through higher LGD, while PD matters more in the multiple-breach scenarios.

Why is this model useful?

This model can be used by a bank or a supervisor that wants to translate a physical hazard into the ratios that drive provisioning and capital: LTV, PD, LGD, expected loss and CET1. It fits mortgage and real estate books, where

collateral is the main exposure to physical risk. Its flood scenarios are public, and its LTV formula is simple enough to be applied to any loan book with geocoded collateral.

What does this model teach us?

This model teaches us that the starting LTV matters as much as the hazard itself. The same flood damage produces far larger capital losses when borrowers were already close to full leverage, which is common in the Netherlands.

It also teaches us that physical risk reaches bank capital mainly through collateral values, not through defaults. The first-order effect is a loss in what the bank recovers when a loan fails.

The trade-off is scope. Only direct property damage is modelled, without macroeconomic spillovers, repeated events or climate scenarios, so the results describe today's flood risk and probably understate future exposure.

2.3.4 MICRO AND MESO MODELS

What is it and why does scale matter?

The concept of scale has appeared throughout this section as an implicit constraint on what each model can and cannot say. The P&L bridge in 2.3.1 operates at the firm level. The DCF model in 2.3.2 operates at the entity or asset level. The ratio models in 2.3.3 operate at the firm or loan level. These are all micro models: they characterise the impact of physical climate risk on a single economic unit.

Meso models occupy the intermediate level between micro (single firm) and macro (national economy). They characterise the physical climate risk impacts on sectors, local economies, clusters of firms, urban agglomerations, or regional supply chains. They are necessary because physical climate risk does not affect firms in isolation: a flood in an industrial district simultaneously damages dozens of firms whose products and services are interconnected, whose workers share local infrastructure, and whose recovery depends on the same restoration resources. Understanding the physical risk impact on a sector or cluster requires a model that captures these interactions, which neither a single-firm DCF nor a national-level regression can provide.

The scale hierarchy in physical climate risk modelling

SCALE	UNIT OF ANALYSIS	TYPICAL MODEL	KEY INTERACTIONS CAPTURED	FINANCIAL OUTPUT
Micro (asset)	Single building, equipment unit	Damage function (section 2.1.1)	None. This is when you consider an isolated asset	Repair cost, replacement value
Micro (firm)	Single company	P&L bridge + DCF (sections 2.3.1–2)	Internal. Consider the fixed/variable cost structure, insurance	Revenue loss, EBITDA, EV, ratios
Meso (sector)	Industry cluster, value chain segment	Sectoral I-O, meso CGE	Inter-firm: supply chain, shared labour, common infrastructure	Sectoral GVA loss, supply-chain multiplier, employment impact

SCALE	UNIT OF ANALYSIS	TYPICAL MODEL	KEY INTERACTIONS CAPTURED	FINANCIAL OUTPUT
Meso (local)	City, district, river basin, coastal zone	Urban climate risk models, regional flood models	Geographic: shared infrastructure, local labour market, fiscal base	Property value gradient, local tax revenue, migration pressure
Meso (network)	Supply chain, financial network	Network contagion models (section 2.5)	Network: contagion, cascading failure, second-order dependencies	System-wide loss amplification, network resilience
Macro	National economy	Macro regression, CGE (section 2.2)	Economy-wide: trade, fiscal, monetary, international spillovers	GDP impact, sovereign risk, financial stability

TABLE 20. The scale hierarchy in physical climate risk modelling

Micro models: firm-level physical risk in full detail

A rigorous micro model of physical climate risk at the firm level integrates four components that are often modelled separately but must be combined to produce an internally consistent financial impact estimate:

- **Physical exposure assessment:** the geo-location of each of the firm's material assets (facilities, warehouses, data centres, ports, agricultural fields) and the hazard intensity at each location under the relevant scenarios. This requires asset registers with geographic coordinates — data that most firms do not maintain in machine-readable form and that must often be reconstructed from property records, satellite imagery, or third-party databases.
- **Vulnerability characterisation:** the damage function for each asset type at each location (construction class, age, flood defence level, seismic rating, fire-resistant design). The interaction between location and vulnerability is non-linear: an asset with a high-quality flood defence in a high-risk zone may have lower expected loss than an older unprotected asset in a moderate-risk zone.
- **Business interruption modelling:** the translation from physical damage at an asset to revenue loss and cost escalation at the firm level (the P&L bridge from 2.3.1). This requires firm-level data on revenue concentration by production location, supply chain alternatives, inventory buffers, and business continuity plans.
- **Financial statement integration:** the mapping of the physical and operational impacts into the specific line items of the firm's financial statements under the applicable accounting standard (IFRS or GAAP), including the timing of recognition, the classification of expenditure as capital or operating, the impairment testing of affected assets, and the insurance recovery accounting.

Meso models: cluster and sector dynamics

Meso models extend micro analysis to the level of industry clusters, urban agglomerations, or regional economies. They are particularly relevant for physical climate risk because extreme events are inherently local: a flood affects a district, not a country; a hurricane hits a coastal zone, not a nation. The economic consequences of that local event depend critically on the density and interconnectedness of economic activity in the affected area.

Three types of meso models are standard in physical climate risk analysis:

- **Sectoral input-output (I-O) models:** extend the national input-output (I-O) framework to the sector level within a single geography. For a flood in an industrial cluster, the sectoral I-O captures how the damage to the chemical sector propagates to the plastics sector (which uses chemical feedstocks), to the automotive sector (which uses plastic components), and to the packaging sector. The indirect loss multiplier at the sector level is typically larger than the national multiplier because sector-level I-O tables capture finer-grained inter-industry dependencies.
- **Urban and regional economic models:** model the physical risk impacts on local economies, including:
- the destruction of local tax base (property tax revenues fall with property values); the loss of local employment (firms relocate or downsize);
- the deterioration of local public services (fiscal stress reduces school and infrastructure quality);
- the induced migration of households away from high-risk areas.

These effects create a self-reinforcing dynamic: property values fall, reducing tax revenues, reducing service quality, which further depresses property values and accelerates outmigration. This spiral is already observable in US municipalities in high-flood-risk zones (see King et al. 2021 on managed retreat).

- **Value chain mapping:** explicitly trace the physical risk exposure along a specific supply chain from raw material extraction through processing, manufacturing, distribution, and retail. Value chain mapping is the most granular meso approach: it identifies the specific nodes in a supply chain that are both highly exposed to physical risk and highly critical (high indegree in the supply network), and quantifies the financial impact of their disruption. For a food company sourcing from a drought-prone agricultural region, value chain mapping reveals the exposure of specific crop sourcing locations, the availability of alternative suppliers, the premium for switching, and the revenue impact of supply shortfall at each stage of the value chain.

The micro-to-meso aggregation problem

Aggregating from micro (firm-level) physical risk estimates to meso (sector or portfolio level) is not simply a matter of summing firm-level impacts. Three effects prevent simple aggregation:

- **Correlation:** firms in the same geography and sector face correlated physical risk. A flood in an industrial zone damages multiple firms simultaneously. The portfolio loss distribution is not the sum of independent firm losses; it has a fatter tail because the worst cases co-occur. Section 2.1.2 on EVT provides the technical framework: the tail dependence coefficient λ_U determines how much portfolio diversification actually reduces tail risk.
- **General equilibrium feedbacks:** when multiple firms in the same sector are simultaneously damaged, market prices adjust. Surviving competitors gain market share and may raise prices. Input costs (labour, materials) rise due to reconstruction demand. These general equilibrium effects mean that the total sector loss is not the sum of firm losses: it is partly offset by competitive reallocation and partly amplified by input cost inflation. I-O and CGE models capture these effects; pure micro aggregation does not.
- **Common infrastructure:** firms in the same location share roads, electricity grids, water systems, and telecommunications. Damage to shared infrastructure creates a common disruption factor that affects all firms simultaneously, regardless of whether their own assets are physically damaged. A flood that leaves

factories structurally intact but destroys the access road, the power substation, and the water treatment plant can cause as much production disruption as direct physical damage to the factories. Micro models that focus only on the direct physical damage to firm assets systematically underestimate the total economic impact in areas with high shared infrastructure dependence.

Connecting micro and meso to the financial system: the transmission channels

THE THREE TRANSMISSION CHANNELS FROM MICRO/MESO PHYSICAL RISK TO FINANCIAL SYSTEM STRESS

Channel 1 – Credit channel. Physical damage reduces the creditworthiness of borrowers (PD uplift from 2.2) and collateral values (LtV deterioration from 2.3.3). Banks holding loans to damaged firms or secured by damaged properties face simultaneous increases in provisions (IFRS 9 Stage 2 migration), capital requirements (higher risk weights), and non-performing loan ratios. If the affected portfolio is large relative to the bank's capital base, the bank's own creditworthiness deteriorates, potentially triggering a credit supply contraction in the affected region.

Channel 2 – Market channel. Equity prices of firms with material physical climate exposure fall when a major event materialises or when new information about future exposure is revealed (e.g. new flood zone maps, updated insurance pricing). Insurance company stocks fall when catastrophe losses exceed reserving assumptions. Real estate investment trusts (REITs) with concentrated coastal or wildfire exposure reprice. These market movements can be rapid and large, creating mark-to-market losses for investors holding these securities and potentially triggering margin calls, forced sales, and liquidity stress if positions are leveraged.

Channel 3 – Insurance channel. The insurance sector is the primary risk absorption mechanism for acute physical events. If insured losses exceed insurance and reinsurance capacity (a scenario that EIOPA and APRA have both flagged as plausible under 2°C+ warming for specific perils and regions), the gap falls back on the uninsured: firms, households, governments, and ultimately the financial system through non-performing loans, sovereign fiscal stress, and reduced economic activity. The progressive withdrawal of insurance from high-risk markets (already occurring for wildfire in California, flood in Florida, cyclone in parts of Australia) is the most direct early signal of this channel activating.

Example: How flood impacts can go from Micro to Meso

A Category A flood event strikes the Rhine valley industrial corridor in Germany.

Its physical footprint is 280 km^2 inundated for 5 days, with a further 3 weeks of access disruption from road and rail damage.

Your asset register comprises 847 industrial facilities, 23,000 employees and $48bn of fixed assets.

Micro layer: firm-level P&L bridges

To assess the firm-level P&L, you need to consider 3 steps: the direct damage, the business interruption and the final impact on revenue:

1) Applying damage functions to each facility (construction type, flood depth from hydrodynamic model, equipment elevation), you manage to average a damage ratio (DR) of 0.12 across all facilities. Hence the total physical damage will be

$$48 \times 0.12 = \$5.76bn.$$

2) Using your historical data, you also average the downtime (including infrastructure disruption): 31 days. You consider this to be an additional $1.8bn loss.

3) Using the estimated business interruptions and impacts on your assets. Insurance recovery estimated at 55% of physical damage and 40% of BI losses. You will manage to recover 55% of $5.76 bn and 40% of your $1.8bn loss

The net uninsured loss will be $3.7bn across 847 firms (45% × $5.76bn + 60% × $1.8bn ≈ $3.67bn).

Meso layer: supply chain propagation

The Rhine valley corridor supplies automotive (23% of revenue concentration), chemicals (19%), and machinery (15%) sectors.

Applying the regional I-O (input-output) multiplier from an open-source database like the EXIOBASE Germany table:

$$\textit{total economic loss} = \$1.8bn \times 2.1 = \$3.8bn$$

$1.8bn direct and $2.0bn indirect through the supply chain.

Consider there is a key bottleneck. Three specialty chemicals producers account for 31% of regional output and supply the rest of the corridor; their 28-day downtime propagates to 14 downstream users.

Financial system transmission

Consider the regional bank exposure is $4.2bn of loans to firms in the affected area.

Applying ratio deterioration models:

$$\frac{Net\ Debt}{EBITDA}$$

rises on average from 3.2× to 4.1× for affected borrowers.

The interest coverage ratio (ICR) falls from 4.1× to 3.1×.

18% of the loan book migrates to IFRS 9 Stage 2.

$$Additional\ provision = \$4.2bn \times 18\% \times (PD_{lifetime} - PD_{12m}) \times LGD \approx \$85m.$$

Regional property values in the inundated zone fall 6–8% (stigma discount: empirically calibrated from the 2002 and 2013 Elbe floods). Residential mortgage LtV rises on average from 68% to 72%.

Key Limitations of Micro and Meso Models

- **Data availability is the binding constraint.** Rigorous micro modelling requires asset-level exposure data with geographic coordinates and construction attributes. This data is available for listed companies' major facilities (from annual reports, regulatory filings, satellite imagery analysis) but is essentially inaccessible for private firms, SMEs, and supply chain upstream nodes. The quality of the meso analysis is bounded by the weakest data link in the supply chain map.
- **Model boundary choices are consequential.** Where you draw the boundary of the meso model determines what is inside and what is treated as exogenous. A model covering the Rhine valley corridor treats national demand as exogenous. A national model treats international trade as exogenous. Selecting a boundary that excludes important transmission channels leads to systematic underestimation of total economic impact.
- **Non-linear recovery dynamics not well captured (Burke et al., 2015).** Standard micro models assume a smooth recovery path: production returns to baseline over the downtime period at a constant rate. Real recovery is non-linear: the first 30% of capacity is restored quickly, but the last 10% (damaged specialized equipment, regulatory certification for restarted facilities, key personnel who did not return) can take many

times longer. Models that use a linear recovery assumption underestimate the cumulative revenue loss and overestimate cash flow recovery speed.

- **Adaptation investment creates model circularity.** If the physical risk model informs adaptation investment, and adaptation investment changes the vulnerability function, the model must be run iteratively to reach a consistent equilibrium. Most operational models run a single pass (hazard → damage → financial impact) without updating the vulnerability function for endogenous adaptation. This produces impact estimates that are pessimistic in the long run (ignoring adaptation) but potentially too optimistic in the short run (failing to account for the capex and disruption of the adaptation process itself).

2.4 NETWORK AND SUPPLY-CHAIN IMPACT ASSESSMENT MODELS

Section 2.4 described how physical climate risk damages individual firms, deteriorates their financial ratios, and compresses their enterprise value. Those models operate on a single economic unit in relative isolation. This section removes that isolation. It describes the frameworks that model what happens when physical damage to one node in an interconnected system propagates to all the others.

Networks are everywhere in modern economies: supply chains link thousands of firms across dozens of countries into tightly coupled production systems; financial networks connect banks, insurers, central counterparties, and sovereigns through webs of credit, collateral, and payment obligations; infrastructure networks carry electricity, water, data, and fuel across geographic spaces whose topology is as consequential as the physical hazard itself. When a climate event strikes a node in any of these networks, the consequences radiate outward along the network edges in ways that no single-firm or single-sector model can capture.

The economic significance of network amplification is far from being marginal. The 2011 Thailand floods caused direct physical damage of approximately USD 15 billion. The total economic loss, once supply-chain propagation through the global automotive and hard-disk-drive industries was accounted for, exceeded USD 45 billion (a 3x multiplier) (Haraguchi and Lall, 2015)

The 2021 Texas winter storm caused USD 20 billion in direct damages; cascading failures across the electricity, water, and natural gas network systems multiplied the economic impact severalfold. Network effects are not a second-order correction: they are often the dominant component of total economic loss from major physical climate events (Busby et al., 2021).

This section covers five interlocking topics: macroeconomic frameworks and supervisory architecture (2.4.1), shock propagation mechanics (2.4.2), supply chain impact assessment with particular reference to the Schneider Electric / Epelbaum framework (2.4.3), systemic risk (2.4.4), and lagged effects (2.4.5).

Evidence markers. To make the quantitative content easier to use for prudential purposes, numerical ranges in this section carry a uniform marker: [S] denotes an estimate taken from the cited source; [I] denotes an illustrative calibration or practitioner judgement that has not been sourced to a specific study and must be validated on the user's own data before any prudential use. Unmarked figures are either sourced in the surrounding text or purely expository.

2.4.1 PHYSICAL RISKS AND MACRO-ECONOMIC FRAMEWORKS

Supply chains are of growing importance to all decision makers in the financial system, whether they're part of a company or representing a financial institution or even part of the country's administration. Critical raw materials (i.e. materials important for production in very sensitive sectors and subject to disruptions) are an example of such interest from the governments as lists of such materials.

The supervisory architecture: FSB, NGFS, and the transmission to financial regulation

The supervisory architecture: FSB, NGFS, standard-setting bodies, and the transmission to financial regulation

The global supervisory response to physical climate risk has been built through two interlocking institutions: the Financial Stability Board (FSB) and the Network for Greening the Financial System (NGFS). Understanding their respective roles and the frameworks they have produced is essential for practitioners working within a regulatory context, because these frameworks define the scenarios, the transmission channels, and the disclosure requirements within which all the models in this cookbook must ultimately operate.

The FSB and the NGFS, however, provide coordination and analytical reference frameworks rather than binding requirements: for practitioners, binding prudential expectations derive from the sectoral standard-setting bodies (the Basel Committee on Banking Supervision for banks, the International Association of Insurance Supervisors for insurers, and IOSCO for securities markets and disclosure), as transposed by national and regional authorities.

These are summarized at the end of this box.

THE FSB FRAMEWORK AND TOOLKIT OF CLIMATE RELATED RISKS

Financial Stability Board (FSB). The FSB is the G20 body responsible for monitoring and making recommendations about the global financial system. Its role in climate disclosure was historically channelled through the TCFD (Task Force on Climate-related Financial Disclosures), which it established in 2015 and whose 2017 recommendations became the reference for corporate climate risk disclosure. Following the publication of IFRS S1 and IFRS S2 by the International Sustainability Standards Board (ISSB) in June 2023, the TCFD was disbanded in October 2023 and the FSB asked the IFRS Foundation to take over the monitoring of companies' climate-related disclosures: IFRS S2, which incorporates the TCFD recommendations, is now the global baseline. The FSB coordinates international work on climate-related financial risks under its Roadmap for Addressing Climate-Related Financial Risks (July 2021), which covers disclosures, data, vulnerabilities analysis, and regulatory and supervisory approaches, and from which its analytical work on climate-related vulnerabilities derives. This FSB Roadmap is distinct from the G20 Sustainable Finance Roadmap (October 2021), which was produced by the G20 Sustainable Finance Working Group.

In 2025, the institution provided a framework for understanding and assessing the role physical risks play in climate-related vulnerabilities (FSB, 2025). The FSB explicitly defines a vulnerability as a property of the financial system that (i) reflects accumulated imbalances, (ii) raises the likelihood of a shock doing damage, and (iii) can lead to systemic disruption when a shock hits. It can largely be considered as a *hidden* risk.

THE FSB FRAMEWORK AND TOOLKIT OF CLIMATE RELATED RISKS

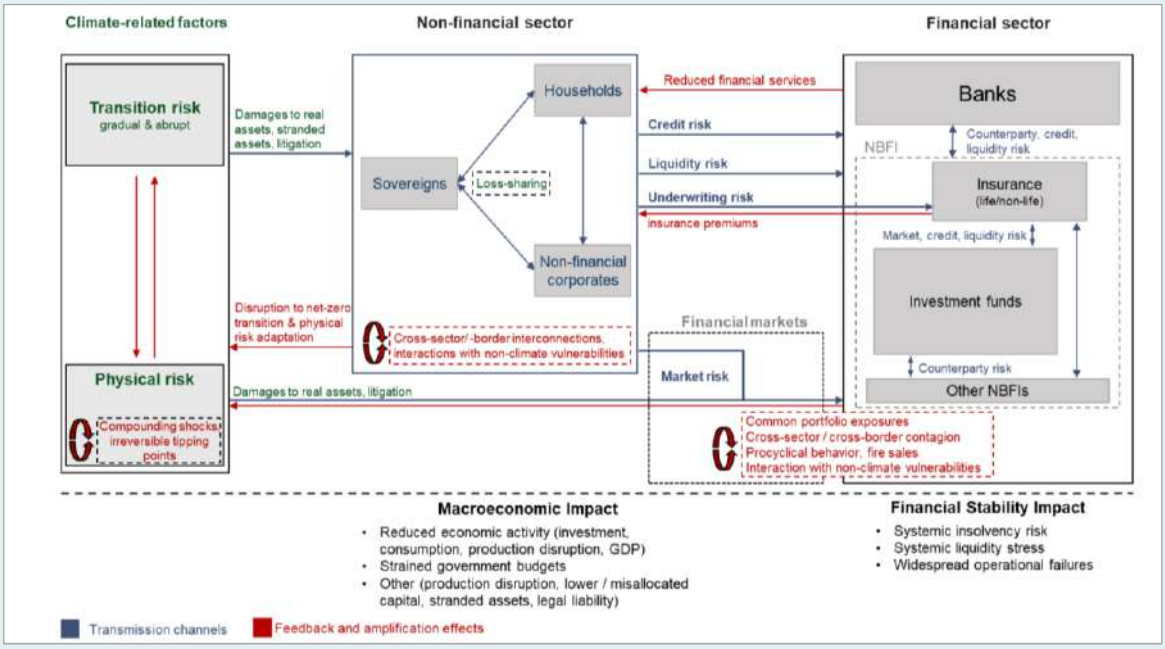


FIGURE 48. Framework for the assessment of climate-related vulnerabilities

Source: FSB, 2025

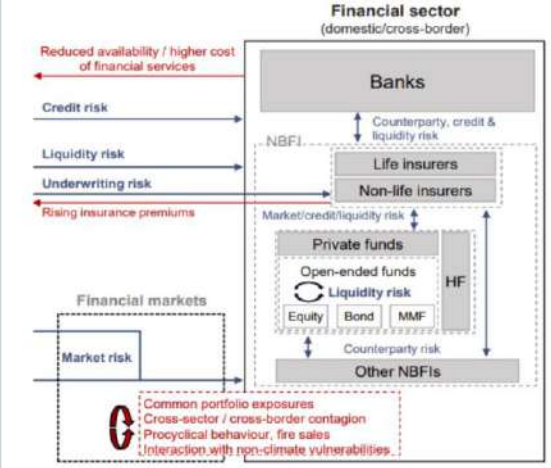
Physical risks propagate into financial risk through several distinct channels. Damage to real assets can impair a borrower's capacity to repay debt and reduce the value of the collateral backing that debt, generating credit risk for lenders. As market participants revise their expectations following the materialization of a hazard, asset prices can reprice abruptly, raising volatility and risk premia across affected sectors, which constitutes market risk.

Liquidity risk arises when depositors or investors withdraw funding from institutions perceived as overexposed to climate hazards, or when margin calls triggered by sudden repricing restrict access to funding. Insurers face underwriting risk directly, since acute and chronic physical hazards raise both asset damages and health-related claims.

Physical hazards can also generate operational losses when a bank's own infrastructure, such as branches or data centers, is damaged, and legal risk when due diligence or fiduciary obligations tied to climate exposure are called into question.

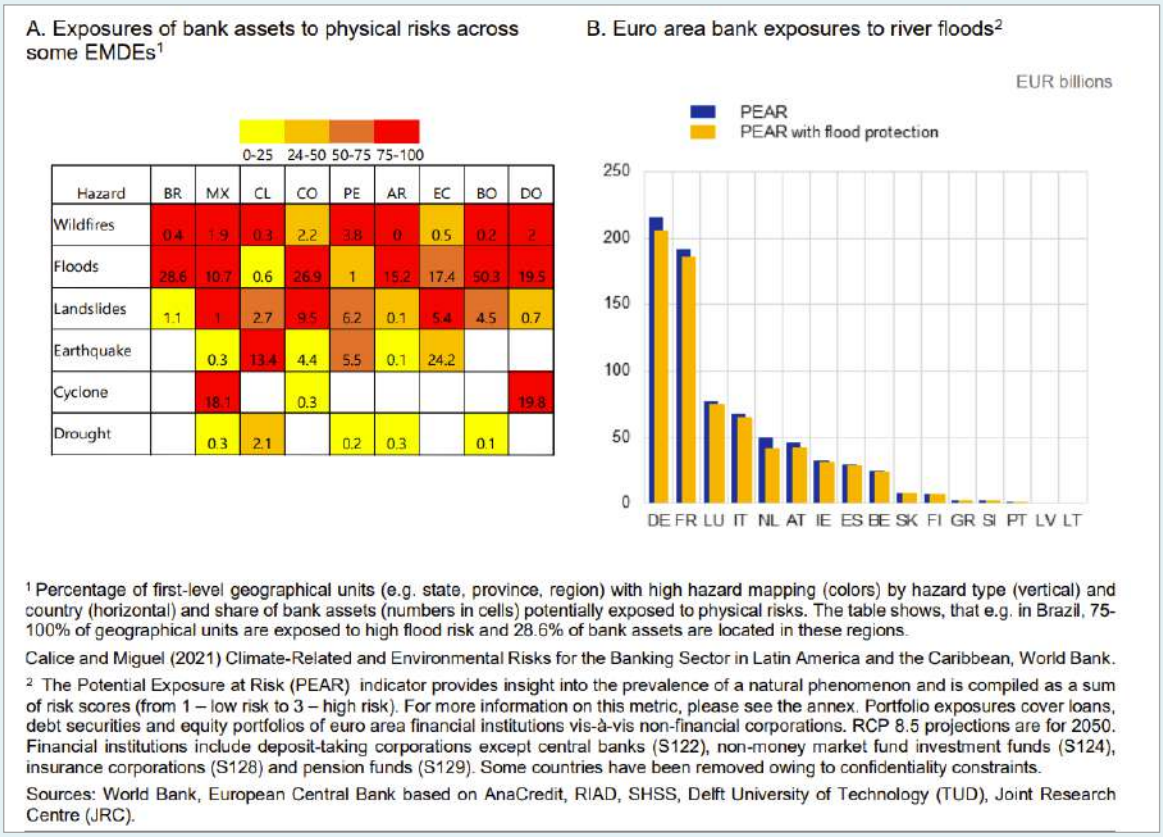
These channels are interconnected among financial institutions, particularly between banks, investment funds and reinsurers. This situation allows a shock that initially hits one sector to propagate to others through shared exposures, correlated selling behaviour, and funding relationships.

For example, a drop in a fund's asset valuations can raise both credit and liquidity risk for the banks that lend to it. At the system level, the framework points to three additional amplification mechanisms: spillovers and risk transfers between financial subsectors, market-wide amplification through interlinkages between institutions, and feedback loops in which reduced bank lending and shrinking insurance coverage feed back into the real economy, weakening the same borrowers whose exposure triggered the initial shock. In the most severe cases, hazards that become effectively uninsurable shift risk toward the public sector, which ends up absorbing the residual exposure as insurer or lender of last resort.



Risk channels	Transmission sources
Credit risk	Issuers of securities held in NBFI portfolios may default or face financial difficulties, especially for economic sectors with stranded assets Insurers may face credit risks through their direct lending and reinsurance contracts
Market risk	Exposed due to their portfolio holdings of equities and debt securities issued by non-financial corporates, financial institutions and sovereigns that are vulnerable to climate risks
Underwriting risk	Climate change may increase the cost of natural catastrophe claims as physical risks increase, which may lead to underwriting losses if insurers do not reflect this in their premiums or policy coverage
Liquidity risk	Abrupt repricing of some asset classes and sentiment shocks can trigger correlated redemption requests and increased collateral requirements across investors

FIGURE 49. NBFI block of the FSB framework and NBFI risk transmission channels



[1] Percentage of first-level geographical units (e.g. state, province, region) with high hazard mapping (colors) by hazard type (vertical) and country (horizontal) and share of bank assets (numbers in cells) potentially exposed to physical risks. The table shows, that e.g. in Brazil, 75-100% of geographical units are exposed to high flood risk and 28.6% of bank assets are located in these regions.

Calice and Miguel (2021) Climate-Related and Environmental Risks for the Banking Sector in Latin America and the Caribbean, World Bank.

[2] The Potential Exposure at Risk (PEAR) indicator provides insight into the prevalence of a natural phenomenon and is compiled as a sum of risk scores (from 1 – low risk to 3 – high risk). For more information on this metric, please see the annex. Portfolio exposures cover loans, debt securities and equity portfolios of euro area financial institutions vis-à-vis non-financial corporations. RCP 8.5 projections are for 2050. Financial institutions include deposit-taking corporations except central banks (S122), non-money market fund investment funds (S124), insurance corporations (S128) and pension funds (S129). Some countries have been removed owing to confidentiality constraints.

Sources: World Bank, European Central Bank based on AnaCredit, RIAD, SHSS, Delft University of Technology (TUD), Joint Research Centre (JRC).

FIGURE 50. Bank exposures to physical risk across jurisdictions

Source: FSB, 2025

They showed that risk exposure varies sharply across jurisdictions and hazard types.

Panel A shows the share of bank assets located in high-hazard regions across a set of Latin American and Caribbean emerging economies, disaggregated by hazard. Flood risk stands out as the dominant exposure channel in several countries, with Brazil, Colombia, and Ecuador each showing more than a quarter of bank

assets situated in regions where a large majority of geographical units face high flood hazard, while cyclone exposure is concentrated almost entirely in Mexico and the Dominican Republic.

Panel B applies the Potential Exposure at Risk (PEAR) indicator described earlier in this report to euro area banks' exposure to river floods, and ranks countries by the combined loan, debt security, and equity exposure of non-financial corporates to flood-prone borrowers under an RCP8.5 projection for 2050. Germany and France carry the largest absolute exposures, and the comparison between the plain PEAR bars and the flood-protection-adjusted bars show that existing flood defences reduce, but do not eliminate, this exposure for most countries in the sample.

The NGFS framework and toolkit of climate related

Network for Greening the Financial System (NGFS). The NGFS is a coalition of over 130 central banks and supervisors (as of 2025) committed to sharing best practices on climate risk. Its best-known output is the NGFS Climate Scenarios. The long-term scenarios (Phase IV, November 2023; Phase V, November 2024) comprise seven pathways grouped into four categories (Orderly, Disorderly, Hot house world, Too little too late), with end-of-century warming ranging from about 1.1–1.4°C to about 3°C. They have been complemented since 2025 by short-term scenarios covering a five-year horizon. NGFS scenarios are widely used as a starting point for supervisory climate scenario exercises (for example by the ECB, the Bank of England, DNB, the ACPR, APRA and MAS), generally after adaptation to the jurisdiction. They are analytical tools, not regulatory requirements in themselves.

In a 2024 survey, the NGFS offers a panoramic view on how financial institutions use climate risk data.

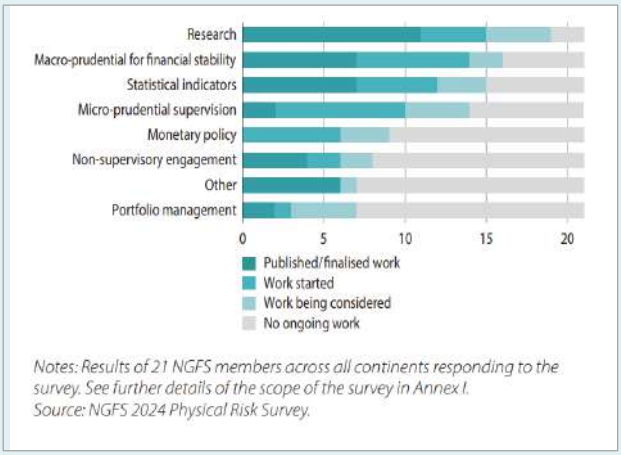
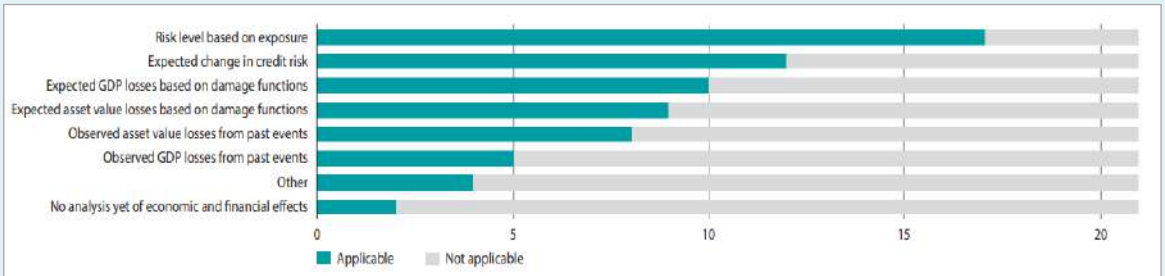


FIGURE 51. Different use cases for physical risk data being considered by surveyed institutions

If the FSB analytical framework describes how physical climate risk transmits into financial risk through channels such as credit, market, and underwriting risk, the NGFS's 2025 information note on leveraging physical climate risk data complements this framework with a typology of six metrics that central banks and supervisors use to measure this transmission in practice, shown in Figure 3.

THE FSB FRAMEWORK AND TOOLKIT OF CLIMATE RELATED RISKS



Note: Results of 21 NGFS members across all continents responding to the survey. See further details of the scope of the survey in Annex I.
Source: NGFS 2024 Physical Risk Survey.

FIGURE 52. Economic and financial impacts assessed by central banks and supervisors

These metrics span two dimensions: backward-looking versus forward-looking, and asset or credit level versus macroeconomic level. Risk level based on exposure identifies whether an asset or counterparty sits in a hazard-prone area, without a monetary loss estimate attached. Expected change in credit risk converts that exposure into probability of default or loss given default adjustments that supervisors can act on. Expected GDP losses and expected asset value losses apply damage functions under a given scenario to project forward-looking losses at the macroeconomic and asset levels respectively. Observed asset value losses and observed GDP losses draw on historical events to calibrate and validate the damage functions and credit risk models feeding the forward-looking metrics.

The typology tells us that physical risk assessment is layered (from exposure identification) through financial risk translation, to scenario-based loss projection, grounded throughout in empirical loss data. This progression maps directly onto the transmission pathway the FSB describes, with each NGFS metric type corresponding to a distinct stage from physical hazard to financial impact.

→The NGFS provided a summary of metrics used in analysis (NGFS, 2025). The FSB also has its own three-tier metric taxonomy, proxies, exposure metrics and risk metrics designating respectively.

1. Aggregate proxy metrics

Country or macro level. No portfolio or asset granularity.

INDICATOR	WHAT IT MEASURES	USE CASE	EXAMPLE
Historical aggregate losses from climate hazards	Records of past economic and insured losses, per event	Assessing past protection gaps, setting a starting point for exposure at risk	EM-DAT database
Forward-looking aggregate losses from climate hazards	Forward-looking acute risk estimates under different climate scenarios (for example, flood damage under various RCPs)	Assessing future increases in risk to infrastructure	Climate Analytics Climate Impact Explorer
Synthetic exposure at risk scores	Country-level climate vulnerability and readiness	Prioritizing adaptation investment	ND-GAIN Index

Hazard and exposure cross-analysis	Combines hazard data with socio-economic exposure	Assessing risk for multinational operations or sovereign investments	OECD Country Risk Analysis

TABLE 21. Aggregate proxy metrics

2. Portfolio and asset-level exposure metrics

Financial exposure to hazard-prone locations. Vulnerability and adaptation not yet factored in.

INDICATOR	WHAT IT MEASURES	USE CASE	EXAMPLE
Portfolio potential exposure at risk	Financial exposure to debtors located in hazard-prone areas, without adjusting for vulnerability or adaptation	Assessing risk distribution across portfolios and regions on a harmonized basis	Potential Exposure at Risk (PEAR), part of the ECB analytical physical risk indicators
Portfolio exposure to different hazard intensities	Value and share of a portfolio tied to debtors located in areas of varying hazard intensity, adding intensity information but still no vulnerability adjustment	Comparing relative risk levels across countries, scenarios and hazard types	ECB analytical physical risk indicators, scored 0 (no risk) to 3 (high risk)
Regional hazard-exposure proxy	Combines hazard data with asset exposure at a regional scale	Evaluating real estate portfolio risk for banks	S&P Physical Risk Heat Map
Asset-level potential exposure	Exposure of individual corporate assets to regional hazards	Corporate credit rating adjustments	Moody's Climate Risk Data

TABLE 22. Financial exposure to hazard-prone locations

3. Financial sector loss metrics

Translations of exposure into financial risk parameters: probability of default, loss given default, spreads, underwriting losses, expected loss. We already defined them in previous sections.

INDICATOR	WHAT IT MEASURES	USE CASE	EXAMPLE
Direct losses to physical assets	Asset-level loss estimates from extreme events via catastrophe models,	Underwriting risk assessment, premium pricing, reinsurance needs, or	Munich Re NatCatSERVICE

	potentially forward-looking with evolving exposure	as input to other financial risk metrics	
Climate-adjusted PD and LGD	Adjusts probability of default and loss given default for physical risk effects on borrower creditworthiness (for example, disrupted operations, reduced repayment capacity)	Integrating physical risk into credit risk models and stress tests	PD adjustments in NGFS short-term scenarios, particularly the Disasters and Policy Stagnation scenario
Climate-adjusted equity prices and corporate or sovereign spreads	Adjusts market risk metrics for the effect of physical risk on counterparty value, through higher indebtedness or reduced future cash flows	Integrating physical risk into market risk models and stress tests	Equity and spread adjustments in NGFS short-term scenarios, MSCI Climate Value at Risk
Underwriting risk	Impact of physical risk on insured people and assets	Climate stress testing and risk management for insurers	Evolution of loss ratios linked to forward-looking insured loss estimates and premium pricing assumptions, as in the ACPR 2023-2024 climate exercise
Other input financial risk metrics	Additional parameters needed across a portfolio, beyond PD, LGD and spreads	Feeding broader financial risk models	Shocks to risk-free rates in NGFS short-term scenarios
Portfolio-level expected losses	Expected loss across a portfolio at a given horizon and scenario, for credit and market risk, accounting for physical and financial collateral	Climate-induced credit or market risk estimation	ECB Normalised Exposure at Risk (NEAR) and Collateralised Exposure at Risk (CEAR); Expected Credit Loss as PD times LGD times Exposure at Default, adjusted for physical risk under a given scenario

TABLE 23. Financial sector loss metrics

Standard-setting bodies and sectoral supervisors. Binding prudential expectations are set by the sectoral standard-setting bodies and transposed by national and regional authorities.

Banking. The BCBS Principles for the effective management and supervision of climate-related financial risks (June 2022) set out eighteen principles, twelve for banks and six for supervisors. The revised Core Principles for Effective Banking Supervision (April 2024) explicitly integrate climate-related financial risks, and the BCBS published a voluntary Pillar 3 disclosure framework for climate-related financial risks in June 2025. These are

transposed regionally and nationally, for instance through the ECB Guide on climate-related and environmental risks (2020) and the EBA Guidelines on the management of ESG risks (January 2025) in the EU, or the PRA's supervisory statement SS3/19 and subsequent updates in the United Kingdom. Supervisory climate stress tests (e.g. ECB 2022, Bank of England CBES 2021) have so far been exploratory or learning exercises rather than inputs to capital requirements.

Insurance. The IAIS Application Paper on the supervision of climate-related risks in the insurance sector (April 2025, replacing the 2021 paper) and the December 2024 changes to ICP guidance (notably ICP 15 on investments and ICP 16 on enterprise risk management) set the international baseline. As a rule, insurance supervisors do not conduct regulatory solvency stress tests based on NGFS scenarios. In the EU, the Solvency II review (Directive (EU) 2025/2) now requires insurers to analyse climate change scenarios within their ORSA: this is a forward-looking risk-management tool owned by the insurer, distinct from a supervisory stress test. Supervisory exercises such as the ACPR's 2020 and 2023–2024 climate exercises remain exploratory.

Securities markets and disclosure. IOSCO endorsed the ISSB standards in July 2023 and has called on its members to consider them in their regulatory frameworks. The EU Sustainable Finance Disclosure Regulation (SFDR) and the FCA's Sustainability Disclosure Requirements (SDR) and investment labels are disclosure and labelling regimes, not applications of the NGFS framework. In the United States, the SEC climate disclosure rule adopted in March 2024 was stayed by the SEC in April 2024 pending litigation, and in March 2025 the SEC voted to end its defence of the rule.

A note for supervisors with limited data: a progressive approach

Supervisors in jurisdictions with limited granular data need not start with full scenario analysis. The NGFS metric typology presented above supports a progressive, three-step approach in which each step builds on the previous one:

1. **Exposure mapping.** Combine aggregate proxy metrics (EM-DAT historical losses, ND-GAIN, Climate Impact Explorer) with whatever portfolio data exist (credit registers, large-exposure reports, supervisory questionnaires) to locate exposures by sector and region, then by hazard zone where borrower or collateral location is available. Output: risk-level-based-on-exposure metrics such as PEAR-type indicators.
2. **Sensitivity analysis.** Apply simple, transparent shocks to the mapped exposures (collateral haircuts in flood-prone areas, PD or LGD add-ons for exposed sectors, single-hazard shocks calibrated on past events) without a full macroeconomic scenario. Output: expected change in credit risk and portfolio-level expected losses under stylised shocks.
3. **Scenario analysis.** Once exposures and sensitivities are understood, use NGFS scenarios, starting with the short-term scenarios for acute shocks and then long-term scenarios for chronic risk, to produce forward-looking loss estimates, possibly in a top-down exercise run by the supervisor. Output: expected GDP and asset value losses consistent with the forward-looking NGFS metrics.

NGFS physical risk scenarios: a reference framework

The NGFS scenarios are the starting point for any macroeconomic physical risk assessment conducted within a regulatory or supervisory context. They define the climate trajectories (temperature paths, sea-level rise, precipitation changes) and pair them with macro-financial impact estimates. The seven long-term core scenarios of Phases IV (2023) and V (2024) are summarised below; Phase IV replaced the former Divergent Net Zero scenario

with Low Demand and Fragmented World, and Phase V kept the same scenario set. Warming levels are approximate medians and vary slightly across vintages.:

NGFS SCENARIO	WARMING BY 2100	PHYSICAL RISK LEVEL	FINANCIAL RISK PROFILE	KEY MODELLING IMPLICATION
Net Zero 2050	~1.5°C	Low–moderate: limited additional physical risk above baseline	Transition risk dominant; physical risk contained	Physical damage functions applied at +1.5°C; low chronic drag
Below 2°C	<2°C	Moderate: some chronic impacts, moderate increase in extreme events	Balanced transition and physical risk	Damage functions at +1.8°C; moderate frequency uplift
Delayed Transition	~1.8°C	Moderate: similar to Below 2°C	Transition risk spike mid-century; physical comparable	Short-run physical similar; long-run financial disruption higher
Low Demand	~1.1–1.4°C	Low: lowest physical risk of all scenarios	Orderly transition driven by behavioural change and lower energy demand; limited transition risk	Physical damage at or below Net Zero 2050 levels; lower bound for physical risk
Nationally Determined Contributions (NDC)	~2.5°C	High: material increase in extreme event frequency and chronic damage	Physical risk dominant in developing countries	Damage functions at +2.5°C; significant frequency and intensity uplifts
Fragmented World (Too little, too late)	~2.3–2.4°C	High: physical risk approaching Hot House World levels	High transition and high physical risk: delayed, divergent policies with only partial achievement of net-zero targets	Combine transition shocks with damage functions at about +2.3°C; most adverse joint scenario
Current Policies (Hot House World)	~2.9–3.0°C	Very high: severe chronic and acute physical damages globally	Physical risk dominant; financial system severely stressed	Full non-stationary EVT; stranded asset write-downs; sovereign stress in tropical economies

TABLE 24. Physical and financial risk under the NGFS scenarios

Important caveats to consider for users. Because the long-term scenarios capture physical risk mainly through a damage function applied to gradual changes in temperature and precipitation, they essentially measure chronic physical risk; the NGFS itself acknowledges that acute physical risk (floods, cyclones, heatwaves, droughts) is largely underestimated in these pathways. Users interested in acute shocks should turn to the NGFS short-term scenarios (2025), which combine near-term transition and physical shocks over a five-year horizon, notably the Disasters and Policy Stagnation scenario, in which a sequence of severe natural disasters hits the economy in the absence of strengthened climate policy.

The damage function is also the main determinant of the projected GDP losses and has changed materially from one vintage to the next: Phase V replaced the Kalkuhl and Wenz (2020) function used in earlier phases with one based on Kotz, Levermann and Wenz (2024), which multiplied chronic physical losses under Current Policies several times over. That paper was retracted by Nature in December 2025, and the NGFS has asked users to take the retraction into account when interpreting the Phase V physical-loss variables, which do not affect the short-term scenarios; a revised methodology is expected in the next vintage. Results from different vintages are therefore not directly comparable, and the damage function should be treated as a key control variable (see the Model Representation Card below).

Transmission channels from physical risk to macroeconomic outcomes

The NGFS framework identifies four canonical transmission channels from physical climate hazards to macroeconomic outcomes and financial stability risks. Each channel has a different speed, scope, and model architecture:

TRANSMISSION CHANNELS: FROM PHYSICAL RISK TO MACROECONOMIC OUTCOMES

Channel 1: Direct capital destruction

Physical events destroy capital stock K, reducing productive capacity and the asset base of the financial system:

$$\Delta K_t = K_t \times -DR(I_t) \times \eta_{exposure}$$

Where:

- $\boldsymbol{K_t}$ is the capital stock at time t, the stock of productive capital exposed to the hazard before the event.
- Therefore, ΔK_t is the change in capital stock at time t. It's under 0 by construction as it represents capital destroyed by the physical event.
- $-\boldsymbol{DR}(\boldsymbol{I_t})$ is the damage ratio (or damage function), a function mapping hazard intensity *I* to the fraction of exposed capital destroyed. Bounded on $[0, 1]$.

$\eta_{exposure}$ is the exposure fraction, i.e. the share of K_t physically located within the hazard footprint. It's bounded on $[0, 1]$.

GDP impact

The GDP impact can be modelled with the standard production function: $\Delta Y_t = \Delta K_t \times MPK$

Where ΔY_t is the resulting change in gross domestic product at time t, driven by the capital stock destruction ΔK_t through the production function.

MPK is the marginal product of capital. MPK can be derived from national accounts as $\alpha \cdot Y/K$, using capital stocks and labour shares from the Penn World Table (Feenstra, Inklaar and Timmer, 2015). Naive estimates are higher in emerging than in advanced economies, but Caselli and Feyrer (2007) show that the gap narrows

substantially once natural capital and the relative price of capital are accounted for. The ranges below are therefore indicative and should be replaced by country-specific calibrations [I]:

For developed economies, consider $MPK \approx 0.08 - 0.12$.

For emerging markets, can be estimated to $0.12 - 0.20$.

For emerging markets, can be estimated to $0.12 - 0.20$.

Since ΔK is expressed in currency units, a 1% destruction of the capital stock corresponds, for a capital-output ratio of about 3, to roughly 3% of GDP, and therefore implies a 0.25–0.6% $GDP\ loss$ immediately, before reconstruction demand is accounted for.

Channel 2: Productivity loss from chronic climate change – the example of heat

TFP stands for Total Factor Productivity. At a time t, it is obtained by compounding the baseline TFP_0 with the climate drag accumulated over every period up to t.

$$TFP_t = TFP_0 \times \Pi_s(1 - \delta_{climate}(T_t, s))$$

Total Factor Productivity (TFP) is reduced by the cumulative climate drag $\delta_{climate}$ calibrated from the reduced-form damage functions in section 2.2.1.

At time 0, it acts as a baseline, a factor productivity absent of any climate related drag.

- $\boldsymbol{T_t}$ is the Prevailing annual average temperature at time t in a given country.
- $\boldsymbol{\Pi_s}$is the product operator applied across all periods s from the initial period up to t. This compounds the climate drag multiplicatively across every intervening period instead of applying it once.
- $\boldsymbol{s}$ is an index running over each period between the initial period and tt, over which the drag term is applied and compounded.

The climate drag function at period s, expressed as a function of temperature. It gives the fractional reduction in TFP applied at that period, bounded between 00 and 11.

Calibration of $\delta_{climate}$ requires care. A non-linear relationship between annual average temperature and the growth rate of GDP per capita (not TFP) has been estimated, with productivity peaking at around 13°C (Burke, Hsiang & Miguel, 2015).

Projected under unmitigated warming (RCP8.5), this relationship implies a global average income about 23% lower in 2100 than without climate change, with much larger losses in hot, mostly lower-income countries. No per-degree TFP growth effect is provided in that study, so $\delta_{climate}$ must be derived by the user. Whether temperature shocks have permanent growth effects or only level effects also remain debated (e.g. Newell, Prest & Sexton, 2021). Results are highly sensitive to this choice, which is why the damage function is listed among the control variables below. In this specification, $T*$ is the optimal temperature of approximately 13°C, at which economic productivity peaks.

Channel 3: Fiscal stress and sovereign risk

The change in the government debt-to-GDP ratio $\frac{\Delta D}{GDP}$ is a way to understand how the hazard at a national level can be considered, combining emergency spending, lost revenue, and the debt service cost of higher sovereign risk premiums.

$$\frac{\Delta D}{GDP} = \frac{\Delta G_{disaster}}{GDP} + \frac{\Delta Rev_{loss}}{GDP} + \frac{CL_{realised}}{GDP} + \Delta r_{climate} \cdot \frac{D_0}{GDP}$$

TRANSMISSION CHANNELS: FROM PHYSICAL RISK TO MACROECONOMIC OUTCOMES

Consider a very intuitive idea: climate events increase government spending.

- $\boldsymbol{\Delta G_{disaster}}$ is the additional primary spending triggered by the event: emergency response, reconstruction and social transfers.
- $\boldsymbol{\Delta Rev_{loss}}$ is the revenue shortfall, entered as a positive number: lower output reduces income, corporate and VAT receipts, so a larger shortfall raises the debt ratio.

The interest rate $\Delta r_{climate}$ term captures the effect of higher sovereign risk premiums ($\Delta r_{climate}$) on the debt service burden.

$CL_{realised}$ captures contingent liabilities that crystallise on the public balance sheet: payouts by public or publicly backed insurance and reinsurance schemes (e.g. the NFIP in the United States, the CCR-backed Cat Nat regime in France, Flood Re in the United Kingdom), calls on government guarantees, transfers to or rescues of state-owned enterprises (utilities, transport, housing), and bank recapitalization or deposit-insurance costs. In practice this is often the largest fiscal channel and should be modelled explicitly rather than folded into $\Delta G_{disaster}$.

The specification is deliberately partial: a complete debt-dynamics equation also includes the denominator effect of lower nominal GDP growth (the (r − g) 'snowball' term) and stock-flow adjustments. For calibration, the IMF–World Bank Debt Sustainability Framework for Low-Income Countries (LIC-DSF) includes a tailored natural-disaster stress test and a contingent-liabilities stress test, and the IMF Sovereign Risk and Debt Sustainability Framework for market-access countries (MAC SRDSF) allows tailored natural-disaster and contingent-liability scenarios. Both can be used to benchmark $\Delta G_{disaster}$, ΔRev_{loss} and $CL_{realised}$.

For small island developing states (SIDS) and highly exposed emerging markets, the fiscal impact of a single major event can add 5–15% of GDP to the debt stock [I], potentially triggering debt sustainability concerns and sovereign spread widening.

Channel 4: Financial sector amplification

For a bank, the relevant prudential metric is the CET1 capital ratio. Physical risk affects both its numerator (losses) and its denominator (risk-weighted assets):

$$\Delta CET1\ ratio_{bank} = \frac{CET1_0 - \Delta ECL_{climate} - \Delta FV_{securities}}{RWA_0 + \Delta RWA_{climate}} - \frac{CET1_0}{RWA_0}$$

- $\boldsymbol{\Delta ECL_{climate}}$ is the change in IFRS 9 expected credit loss allowances attributable to climate risk, including the move from 12-month to lifetime ECL when exposures migrate to Stage 2. Under IFRS 9 the ECL allowance is the loan-loss provision: this single term captures both, and adding a separate 'provisions' term would double count.
- $\boldsymbol{\Delta FV_{securities}}$ is the fair-value loss on climate-exposed securities. Its capital impact depends on accounting classification: losses on instruments measured at fair value through profit or loss (FVTPL) or through other comprehensive income (FVOCI) reduce CET1, whereas securities held at amortized cost produce no immediate capital loss but carry unrealized losses that crystallize if the bank has to sell them to meet liquidity needs.
- $\boldsymbol{\Delta RWA_{climate}}$ is the increase in risk-weighted assets as PDs and LGDs migrate under internal-ratings-based models (or as ratings are downgraded under the standardized approach). It lowers the ratio even before any loss is realized.

TRANSMISSION CHANNELS: FROM PHYSICAL RISK TO MACROECONOMIC OUTCOMES

Liquidity must be assessed alongside solvency: deposit outflows from affected regions, drawdowns of committed credit lines by affected corporates, and margin calls following abrupt repricing weigh on the LCR and NSFR and can force asset sales.

For insurers, the relevant metric is the Solvency II ratio (eligible own funds / SCR) or its local equivalent, and the main effects sit on the liability side: higher claims and technical provisions, higher cost or reduced availability of reinsurance, and a larger natural-catastrophe component of the SCR. Asset-side repricing is usually second-order for non-life insurers.

The change can be written:

$$\Delta SR = \frac{OF_0 + \Delta A - \Delta TP}{SCR_0 + \Delta SCR} - \frac{OF_0}{SCR_0}$$

- where $\boldsymbol{\Delta A}$ is the change in asset values and ΔTP the change in technical provisions.

Banks and insurers transmit physical risk shocks back to the real economy through three mechanisms:

Credit tightening: capital erosion from physical risk losses forces banks to reduce lending, amplifying the initial GDP shock through a credit multiplier.

Mark-to-market losses: repricing of climate-exposed securities (equities, bonds, real estate) creates valuation losses on bank and insurer balance sheets, further compressing capital.

Insurance coverage withdrawal: as insurers reprice or withdraw from high-risk markets, the uninsured share of physical losses rises, shifting more of the gross economic loss to uninsured firms and households whose balance sheets deteriorate faster.

Available databases on transmission channels

A significant body of standardized data infrastructure has been built (NGFS, 2025) over the past decade to support macro physical risk modelling. The following table maps the key databases to the transmission channels they support:

DATABASE / TOOL	PROVIDER	COVERAGE	TRANSMISSION CHANNEL SUPPORTED	ACCESS
NGFS Climate Scenarios v4	NGFS / IIASA	Global; 6 scenarios; 2020–2100	All channels; reference macro scenario	Open: ngfs.net/ngfs-scenarios-portal
NiGEM macro model	NIESR	40+ countries; quarterly	Capital destruction → GDP; fiscal → sovereign	Licensed (NIESR / Oxford Economics)
Oxford Economics Global Model	Oxford Economics	80+ countries; annual	Productivity → GDP; financial amplification	Licensed
EXIOBASE v3.8	EXIOBASE Consortium	43 countries; 163 sectors; 2011–2022	Supply-chain propagation (I-O network)	Open: exiobase.eu

DATABASE / TOOL	PROVIDER	COVERAGE	TRANSMISSION CHANNEL SUPPORTED	ACCESS
WIOD 2016	World Bank / EU	43 countries; 56 sectors	Supply-chain propagation	Open: wiod.org
EM-DAT	CRED	Global; 1900–present; 22,000+ events	Direct capital destruction; fiscal stress	Open: emdat.be
Swiss Re sigma	Swiss Re Institute	Global; 1970–present; insured + econ losses	Insurance gap; financial amplification	Open summary; full data licensed
BIS consolidated banking stats	BIS	30 reporting countries; bilateral bank claims	Financial network; credit contagion	Open: bis.org/statistics
ECB AnaCredit	ECB	Loan-level data; EU banking system	Credit channel; IFRS 9 stage migration	Restricted: central bank access
JRC PESETA IV	European Commission JRC	EU28; sector-level	Productivity → GDP; sector impacts	Open: joint-research-centre.ec.europa.eu
CLIMADA	ETH Zürich	Global; asset-level	Direct capital destruction; damage functions	Open: climada.ethz.ch
Aqueduct Floods	WRI	Global; 1km resolution	River flood hazard; capital at risk	Open: wri.org/aqueduct
Four Twenty Seven / Moody's	Moody's ESG	Asset-level; global	Physical risk score → credit impact	Licensed
MSCI Climate VaR	MSCI	Security-level; global	EV impact → portfolio risk	Licensed

TABLE 25. Available databases on transmission channels

Provider Box – Probabilities for scenarios and quantifying any type of assets – The Scientific Climate Ratings Ambition

Origins and architecture. Scientific Climate Ratings (SCR) emerged from the applied research of the EDHEC Climate Institute (ECI), where 16 PhD-level researchers oversee the methodology. SCR operates as a rating agency and will be regulated by ESMA. The offer consists of three products that share one scenario framework, one hazard and damage library, and one A-to-G rating scale.

- **ClimateMetrics (CMS)**, a self-service platform for any real asset worldwide, together with a rated universe of more than 6,000 infrastructure assets (Climate Exposure Rating, CERinfra, and Climate Risk Rating, CRRinfra).

- **Corporate Climate Risk Rating (Corporate CRR)**, covering more than 4,000 large-cap non-financial listed issuers, with launch targeted for the end of 2026.
- **Sovereign Climate Risk Rating (SovCRR)**, covering 191 countries built from more than 3,400 sub-national regions.

In each product the hazard is converted into a loss measure before any rating is assigned. This measure is damage or disruption as a share of asset value or revenue for real assets, an enterprise-value impact for companies, and a loss of GDP per capita for sovereigns. No product uses machine learning, and none relies on a live or near-real-time feed. Hazard, scenario and methodological datasets are refreshed annually, in line with major scientific releases.

At the heart of this method is the consideration of probability-weighted scenarios. Mainstream scenario sets, the NGFS included, describe plausible pathways without assigning them a likelihood. ECI extends the NGFS set with two high-warming tail scenarios, "Climate Breakdown" and "Climate Destabilization", and assigns an explicit probability to every scenario through maximum-entropy elicitation. The two tail scenarios carry 43.3% of the probability weight, so severe warming enters the central estimate. The probability-weighted result is called the Expected scenario.

$$X^E = \sum_{s=1}^{S} \pi_s \, X_s$$

- X^E is the output under the Expected scenario
- X_s is the output under scenario s
- π_s is the probability assigned to scenario s, with $\sum_s \pi_s = 1$
- S is the number of underlying scenarios

The scenario set depends on the product.

- **Real assets.** Physical metrics use SSP1-2.6, SSP2-4.5, SSP3-7.0 and SSP5-8.5, at five-year intervals from 2025 to 2100. Each of the eight transition scenarios is paired with one physical pathway.
- **Listed equities.** Six NGFS scenarios (Low Demand excluded), the two tail scenarios and the Expected scenario, at five horizons from 2030 to 2050.
- **Sovereigns.** Seven NGFS scenarios, the two tail scenarios and the Expected scenario, reported at 2035 and 2050 on a complete annual grid.

For listed equities and sovereigns, each scenario runs under three integrated assessment models (GCAM, MESSAGE and REMIND). Emissions are converted into warming with the MAGICC emulator, and the resulting impacts are averaged with equal weights.

Real assets: ClimateMetrics

Hazards. ClimateMetrics covers eight climate hazards (flood, wind, wildfire, heat, extreme precipitation, drought, landslide and subsidence), plus earthquake as a non-climate peril. Flood, wind, wildfire and heat carry monetised damage functions and feed the climate rating. Earthquake also carries damage functions and enters financial metrics, but stays outside the climate rating. A drought damage function is under development. The remaining hazards are reported as hazard indicators.

HAZARD	PHYSICAL PROXY	MAIN SOURCE	RESOLUTION
Flood (fluvial, pluvial, coastal)	Water depth (m)	Fathom (2024), 8 return periods from 5 to 1,000 years	30 m
Tropical cyclone	Maximum wind speed (m/s)	STORM, 10,000 synthetic years; TCHA18 for Australia	11 km (2 km Australia)
Extratropical storm	Maximum wind gust (m/s)	ERA5 since 1940; CORDEX and CNRM-CM6-1-HR	0.25° historical; 11–44 km projected
Heat	Wet bulb globe temperature (°C)	NEX-GDDP-CMIP6, median of five models	0.25°
Wildfire	Annual burn probability	MODIS burnt area (2001–2024); Copernicus land cover	500 m
Earthquake	Peak ground acceleration, 475-year return period	GEM Global Hazard layer	8 km

TABLE 26. Hazard table – ClimateMetrics

.

Footprint and damage. Each asset is drawn as a point, a line or a polygon. The hazard layer is reduced to a single intensity over that footprint by zonal statistics, and the damage function is applied at that stage. SCR shows in one published example that a generic radius overstates flood depth almost fivefold and damage more than threefold.

Damage curves come from a versioned library of 185 source curves (139 for flood, 46 for wind) drawn from FEMA Hazus, the GEM Global Vulnerability Model and DamageScanner, mapped to 133 asset types. Each curve assignment records its source and a confidence level (high, medium or low).

$$DF_{h,k} = f_{h,a,c}(\bar{H}_{h,k})$$

- $DF_{h,k}$ is the damage factor for hazard h at return period k, as a share of replacement value
- $f_{h,a,c}$ is the damage function for hazard h, asset type a and country c
- $\bar{H}_{h,k}$ is the hazard intensity over the asset footprint (e.g. flood depth) at return period k

Damage factors are integrated across return periods into an expected annual damage, using the area under the exceedance probability curve.

$$EAD_h = \sum_{k=1}^{K-1} (P_k - P_{k+1}) \frac{DF_{h,k} + DF_{h,k+1}}{2}$$

- EAD_h is the expected annual damage for hazard h, as a share of asset value
- P_k is the annual exceedance probability of return period k (e.g. 0.01 for 1-in-100 years)
- $DF_{h,k}$ is the damage factor at return period k
- K is the number of return periods (eight for flood)

Two hazards follow a different logic. Wildfire applies a conservative binary damage factor, so expected damage equals the annual burn probability. Heat runs on a dose-response chain from wet bulb globe temperature to loss of work ability and then to revenue, scaled by the work intensity of the asset type, its dependence on on-

site labour and a 50% global labour share. Only the deterioration relative to 1990–2014 counts, since average historical conditions are assumed to be priced in.

Business interruption. Downtime follows from structural damage. FEMA sector data give the interruption after complete destruction, scaled down in proportion to the share of the asset damaged (Mandel et al., 2025). The number of days lost gives the revenue impact.

Financial metrics. The Climate Exposure Rating reports damage as a share of replacement value, Average Annual Expected Loss per hazard, Physical Value at Risk, a Probable Maximum Loss (PML) at the 1% annual exceedance probability, the Cost of Inaction and the Adaptation Benefit. For wildfire the PML comes from a Monte Carlo simulation of burn probability, and for heat from the 99th percentile of revenue impact across the five climate models.

$$COI = \sum_{t=t_0}^{T} \frac{L_t}{(1+r)^{t-t_0}}$$

$$AB = COI^{\text{excl}} - COI^{\text{incl}}$$

- COI is the Cost of Inaction, the present value of expected climate losses
- L_t is the expected loss in year t (physical damage and direct carbon cost)
- r is the discount rate
- t_0 is the base year and T the chosen horizon
- AB is the Adaptation Benefit
- COI^{excl} and COI^{incl} are the Cost of Inaction without and with declared measures

The Climate Risk Rating goes one step further. It applies the climate-conditioned valuation and credit model described for listed equities to each asset, which yields an Enterprise Value impact, a probability of default, a loss given default and an expected loss. Insurance coverage is excluded on purpose, since annual contracts cannot guarantee protection over several decades.

Adaptation. Declared measures are quantified through ClimaTech, an ECI database of more than 100 standardized resilience strategies drawn from more than 200 peer-reviewed publications and reviewed by a committee of more than 20 external experts (including the EIB, the World Bank, WWF and PRI). Each measure is linked to the hazard and the return period it protects against. Its effectiveness is supplied by the user, or taken from ClimaTech by default.

$$DF_{h,k}^{\text{net}} = DF_{h,k}^{\text{gross}} \times \left(1 - e_{h,k}\right)$$

- $DF_{h,k}^{\text{net}}$ is the damage factor after declared measures
- $DF_{h,k}^{\text{gross}}$ is the damage factor before any declared measure
- $e_{h,k}$ is the effectiveness of the measure for hazard h at return period k, between 0 and 1

SCR quantifies implemented measures and grades no adaptation maturity. The net result is rated on the same scale as the gross one. Rated entities must provide verification evidence.

Rating scale. Each metric is scored against its distribution in the rated universe of more than 6,000 infrastructure assets. Each asset is compared with a peer group of the same infrastructure class in the same

climate zone, with at least 30 assets per group. The resulting 1-to-100 percentile score is mapped to the A-to-G scale, so a self-assessed asset on the platform can be compared with rated assets.

Listed equities: Corporate CRR

From sites to company. A company is assessed as a collection of production sites, each with its own location. Site results are combined using each site's share of productive capacity, measured with Global Energy Monitor data for energy assets and revenue as a proxy elsewhere.

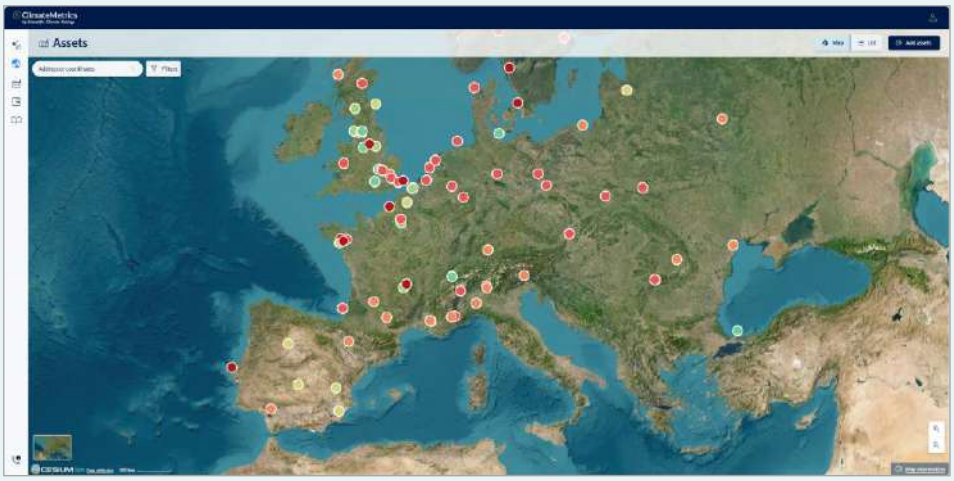


FIGURE 53. A quick look at an asset portfolio – ClimateMetrics

Six impact channels. The climate impact is split into six channels, each reported per scenario and for the Expected scenario.

- Two physical channels. The acute channel covers flood, wind and wildfire (asset damage and interruption) and heat waves (interruption without damage), site by site, with the same curves as ClimateMetrics. The chronic channel covers gradual productivity loss from warming, by sub-national region and sector.
- Three transition channels, namely direct carbon cost, market demand shift and supply-chain cost.
- One macroeconomic channel, the climate-driven change in cost of capital by sector and country.

Channel impacts are added without interaction effects, so the total can be attributed to each channel.

Valuation. A climate-consistent discounted cash-flow model translates annual impact factors into enterprise value. The counterfactual cash flow is normalised to one unit per year, so each shock reads as a relative deviation and no firm-specific forecast enters the rating.

$$\frac{\Delta EV}{EV} = \frac{\sum_{t=1}^{T} \frac{1+f_t}{(1+r_s)^t}}{\sum_{t=1}^{T} \frac{1}{(1+r_0)^t}} - 1$$

- $\Delta EV / EV$ is the climate impact on enterprise value, in percent
- f_t is the annual cash-flow impact factor in year t, summed across channels
- r_s is the discount rate under scenario s, including the macroeconomic channel
- r_0 is the climate-neutral baseline discount rate, based on present-day market conditions and sector and country risk premia

- T is the valuation horizon



FIGURE 54. Selected example of a Climate Exposure Rating – Horizon 2050 for a Gas-Fired Power Generation asset in the USA

Credit. A structural credit model carries the value impact to equity, debt and default risk. Current asset value, drift and volatility are inferred from the equity price series by inverting the Black-Scholes call formula. Under each scenario, the face value of debt, the one-year maturity, the risk-free rate and asset volatility stay at baseline, so the scenario moves the firm relative to its default barrier. Cash is carved out before the climate impact and added back afterwards.

$$DD = \frac{\ln\left(\frac{V_A}{F}\right) + \left(\mu - \frac{\sigma_A^2}{2}\right)\tau}{\sigma_A\sqrt{\tau}}$$

$$PD = N(-DD)$$

$$EL = PD \times LGD$$

- DD is the distance to default
- V_A is the climate-adjusted asset value of the firm
- F is the face value of debt
- μ is the asset drift
- σ_A is the asset volatility
- τ is the maturity, set to one year
- PD is the probability of default
- N is the standard normal cumulative distribution function
- EL is the expected loss, as a share of the face value of debt
- LGD is the loss given default, read by SCR as a lower bound

Rating. The Physical Impact Rating ranks the impact's percentile score within the rated universe, under a fixed reference condition (Expected scenario, 2035 horizon, average of the three integrated assessment models). Outputs cover 180 conditions (nine scenarios, five horizons, four model views), of which 120 are independent

valuations. Publicly disclosed or company-contributed measures are reflected after verification, which produces a net view alongside the gross one.

Sovereigns: SovCRR

Scope. The current version covers chronic temperature risk only. Rising temperatures are modelled as a persistent and compounding loss of GDP per capita, measured against a no-climate-change counterfactual. Adaptation is outside its scope, so all outputs are gross.

Stage I, historical estimation. A panel regression estimates how regional growth in output per capita responds to year-to-year variations in local temperature and precipitation. The panel covers 1,661 regions in 88 countries over 1979–2019, with output from the DOSE database and climate from NASA's GLDAS, aggregated with population-density weights.

$$g_{r,t} = f(T_{r,t}, P_{r,t}) + \mu_r + \delta_t + \varepsilon_{r,t}$$

- $g_{r,t}$ is the growth of output per capita in region r and year t
- f is the estimated response of growth to climate
- $T_{r,t}$ and $P_{r,t}$ are population-weighted temperature and precipitation
- μ_r is a region fixed effect, which isolates within-region variation
- δ_t is a year fixed effect
- $\varepsilon_{r,t}$ is the residual

Stage III, projection. The estimate is re-expressed as a region-specific quadratic function of the global mean temperature anomaly, fitted without an intercept, so damages read as deviations from a no-climate-change baseline. Parameters are estimated once per region and held constant across scenarios, models and years. Regional warming draws on 15 of 29 CMIP6 models, after implausibly hot models are screened out.

$$D_{r,t} = a_r \, \Delta T_t + b_r \, \Delta T_t^2$$

$$D_{c,t} = \sum_{r \in c} w_r \, D_{r,t}$$

- $D_{r,t}$ is the projected damage to output per capita in region r and year t
- ΔT_t is the global mean temperature anomaly relative to 1980–2010
- a_r and b_r are the region-specific damage parameters
- $D_{c,t}$ is the damage for country c
- w_r is the time-invariant population weight of region r within country c

The Expected-scenario impact on GDP per capita is converted into a 1-to-100 score and an A-to-G rating at 2035 and 2050. Ratings also exist for each scenario. Aggregate damages are consistent with NGFS (2024) NiGEM results on 49 economies and with the pattern found by Burke et al. (2015). Thirty-nine territories are excluded, mainly for insufficient climate-model resolution or contested sovereignty.

Roadmap. From Q4 2026, damages will feed a macro-fiscal layer that translates them into sovereign rating and credit spread changes, without an explicit default probability. From Q1 2027, coverage will extend to heatwaves and precipitation patterns, following NGFS Phase V, with GDP impacts split across 15 NACE sectors.

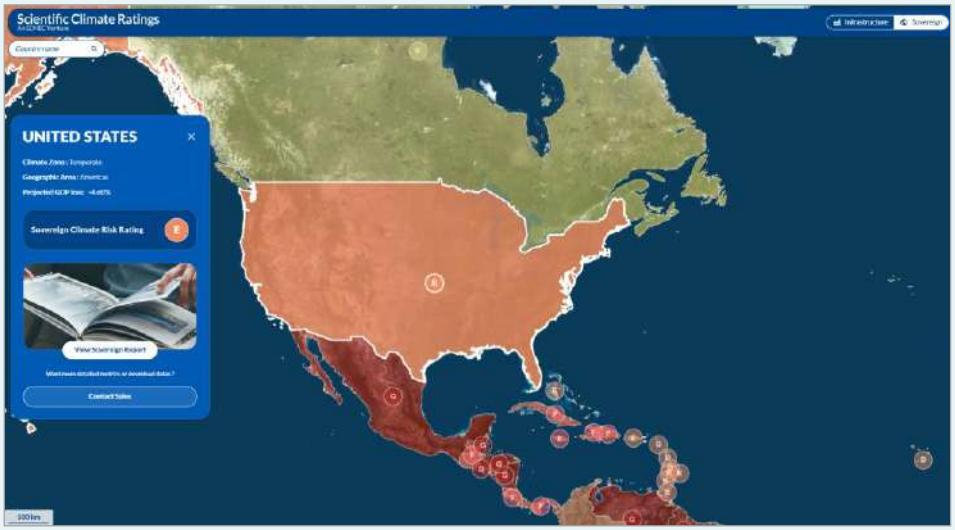


FIGURE 55. Selected example: 2050 Horizon Sovereign Risk Rating for the United States of America

Uncertainty, validation and transparency

Uncertainty. SCR distinguishes four sources of uncertainty.

- Model uncertainty, addressed through multi-model ensembles (CMIP6 models and three integrated assessment models).
- Scenario uncertainty, addressed through the full probability-weighted scenario set and the Expected scenario.
- Climate variability, which is not modelled explicitly.
- Damage estimation, addressed through documented curve sources and confidence levels.

Data Quality Scores on a 1-to-5 scale inspired by PCAF flag the provenance of inputs, independently of the robustness of the methods. No confidence intervals are produced. SCR rates its overall confidence as high (80–90%), materially higher at 2035 than at 2050 and beyond, since scenarios remain in a narrow band until around 2050 and diverge sharply afterwards.

Validation. No formal backtesting against realized losses has been performed, and no skill metric (RMSE, Brier score, hit rate) is published. Hazard layers are calibrated on observed records and validated against national reference maps (FEMA National Risk Index, Australia's TCHA18, UK Environment Agency and SEPA flood data). The corporate credit layer is calibrated to market data. Governance relies on three committees and an annual methodology review. Outcomes observed by clients are not tracked, except for sovereign spreads, which are compared with observed market ranges.

Tipping points. Discrete earth-system tipping points are not modelled as explicit thresholds. Their consequences are partly captured through the two tail scenarios, non-linear damage functions and return periods up to 1-in-1,000 years.

Transparency. Methodologies, hazard sources and scenario probability tables are public, and every output can be traced to its data and model versions. The proprietary layer covers the ClimaTech effectiveness values,

the asset footprints, the sector-country breakdown of NGFS pathways, the regional damage parameters and the rating calibration.

Note. Equations are simplified representations of the methodology as described by SCR in its survey answers. The Stage I functional form is stylized, and the exact specifications are available in SCR's published methodologies.

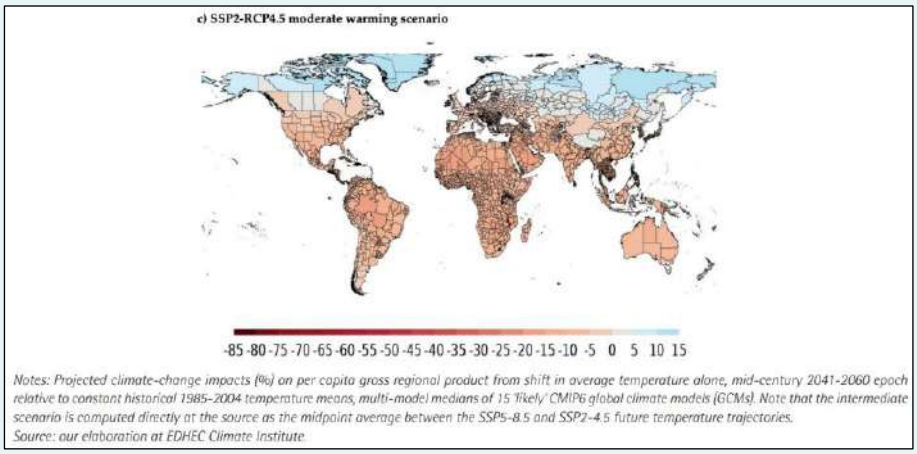


FIGURE 56. Clirmap SSP2-RCP4.5 moderate warming scenario (Edhec Clirmap, 2025)
Source: Edhec Clirmap – EDHEC-Climate-Induced Regional - Macroimpacts Projector - The Macroeconomic and Econometric Background, 2025

Model Representation Card – NGFS macro model – NGFS physical risk framework

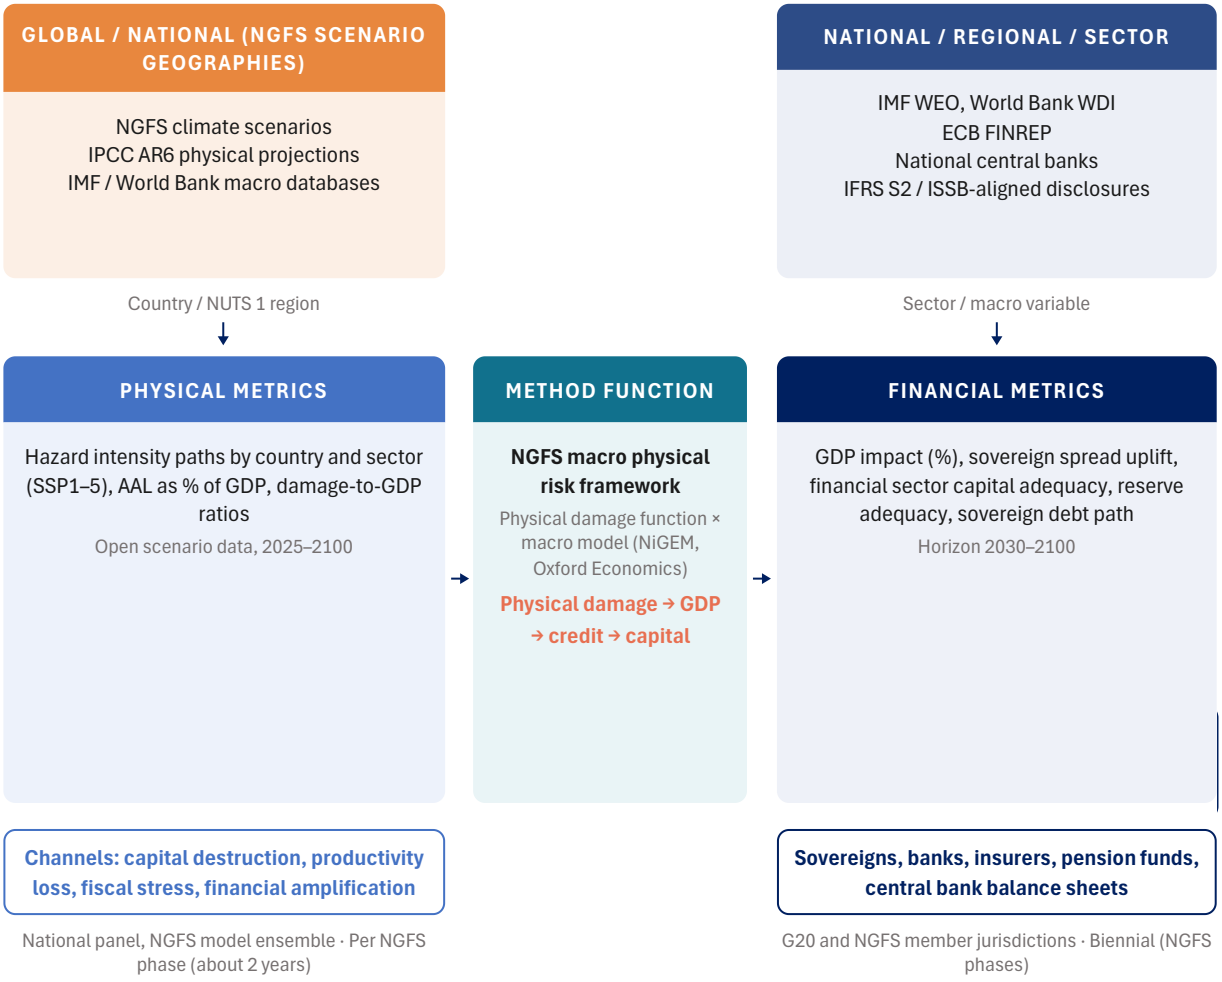


FIGURE 57. Model representation card – NGFS macro model (NGFS physical risk framework, NGFS 2023–2025)

USER	DATA ACCESS AND FEASIBILITY	TYPICAL PRUDENTIAL USE CASES
Supervisor / central bank	Full: open NGFS scenario data plus confidential supervisory data (FINREP/COREP, credit registers such as AnaCredit); can run the whole chain top-down.	Macroprudential: system-wide climate stress tests, sovereign and fiscal risk analysis, design of scenarios for bottom-up exercises.
Financial institution	Partial: open NGFS variables and own portfolio data; macro models usually licensed (NiGEM) or replaced by scenario paths provided by the supervisor.	Microprudential: ICAAP/ORSA scenario analysis, forward-looking macro overlays for IFRS 9 provisioning, strategic planning.
Researcher / market participant	Limited: open NGFS, IMF and World Bank data; no confidential exposures.	Disclosure: IFRS S2 scenario analysis; independent assessment and sovereign credit analysis.

TABLE 27. Implementation and Use Cases: NGFS Macro Model

2.4.2 SHOCK PROPAGATION

What is it?

Shock propagation models describe the mechanics by which a localized physical damage event travels through an economic or financial network to affect agents that were not directly struck. Throughout this section, an agent denotes a node of the network, which, depending on the context, may be a firm, a bank or another economic actor (insurer, fund, sovereign). The core insight is that modern economies are not collections of independent agents: they are networks of interdependent agents whose outputs are each other's inputs, whose debts are each other's assets, and whose production depends on shared infrastructure.

When a node fails, whether through physical damage, financial distress, loss of access to inputs or, as a second-order effect, loss of demand, the failure propagates along the edges of the network according to well-defined mathematical laws.

Shock propagation operates across three distinct network types that interact with each other in the real world, although in practice most models (particularly agent-based models) treat them relatively separately:

- **Production networks:** firms are nodes; supply relationships are directed edges (i → j means firm i supplies firm j). A shock to supplier i propagates forward to customer j (supply disruption) and backward from customer j to alternative suppliers (demand reallocation). The I-O (input-output) framework is the linear approximation of production network propagation; network models extend this to capture non-linearities, bottlenecks, and graph topology.
- **Financial networks:** banks, insurers, funds, and sovereigns are nodes; loans, bonds, equity stakes, derivatives, and interbank exposures are edges. A shock to bank A's capital (from physical risk losses on its loan book) propagates to bank B if B holds A's debt or has derivatives exposure to A. The propagation mechanism is counterparty credit risk: financial network contagion models map this propagation.
- **Infrastructure networks:** electricity grids, water systems, road and rail networks, telecommunications, and port networks are physical graphs. Physical damage to a critical node (a substation, a bridge, a port) propagates through the infrastructure network by reducing service availability to all downstream users. The

economic consequences of infrastructure network failure are distinct from and compounding with the production and financial network effects.

Mathematical Representation: Four Propagation Channels

MATHEMATICAL REPRESENTATION: FOUR PROPAGATION CHANNELS

CHANNEL A: BACKWARD (UPSTREAM) PROPAGATION OF A FINAL-DEMAND SHOCK: THE LEONTIEF MODEL (DEMAND-DRIVEN I-O, LINEAR)

The Leontief (demand-driven) I-O model gives the total output change across all sectors following a change Δf in final demand for sector k:

$$\Delta x = (I - A)^{-1} \cdot \Delta f = L \cdot \Delta f$$

The propagation from a final-demand shock in sector *k* back to all supplying (upstream) sectors *j* is given by column *k* of the Leontief inverse L: $L_{\{jk\}}$ is the total output of sector j required per unit of final demand for sector k.

For example, a €1bn fall in final demand for chemicals (for instance because flood-hit customers cancel orders) reduces the output of each upstream sector j by €$L_{j,chem}$ billion: energy by €$L_{energy,chem}$ bn, mining and basic inorganic inputs by €$L_{mining,chem}$ bn, and so on across all N sectors.

The Leontief model thus propagates demand shocks upstream, towards suppliers; the downstream effect on chemical users (automotive, plastics) can be captured by the Ghosh model of Channel B.

Channel B: Forward (downstream) propagation of a supply shock: Ghosh model (supply-driven I-O, linear)

When a supplier loses production capacity, its downstream customers lose inputs. The Ghosh model (supply-driven I-O) propagates this loss forward:

$$\Delta x' = \Delta v' (I - B)^{-1} = \Delta v' G$$

- $\Delta x'$ is the row vector of output changes across sectors.
- $\Delta v'$ is the row vector of changes in primary inputs (value added) by sector.
- $B_{ij} = \frac{z_{ij}}{x_i}$ is the Ghosh output coefficient, the share of the output of sector i sold to sector j Here z_{ij} is the flow from sector i to sector j and x_i is the total output of sector i.
- $G = (I - B)^{-1}$ is the Ghosh inverse. G_{ij} is the total output of sector j generated per unit of primary input in sector i.

Hurricane Harvey (2017) illustrates this mechanism. A network model of the US petrochemical industry shows that the capacity disruptions along the storm track changed the production level of 170 chemical units in 26 states, far beyond the directly affected area (DeRosa et al., 2019). At a broader scale, Middelanis et al. (2022) simulate Harvey on over 7000 regional economic sectors with more than 1.8 million supply chain connections. They find that while the United States were able to balance regional production failures by the original 2017 hurricane, this option becomes less viable under future warming.

The Ghosh model is rarely used as a causal model. It assumes fixed output allocation coefficients, so buyers would keep absorbing a fixed share of each supplier's output whatever happens to their own demand, which economists consider implausible (Oosterhaven, 1988). Dietzenbacher (1997) showed that it is better interpreted as a price model. In practice it serves mainly as a measure of forward linkages, that is, of how exposed downstream sectors are to a supplier.

MATHEMATICAL REPRESENTATION: FOUR PROPAGATION CHANNELS

Channel C: Non-linear propagation — the ARIO model

The Adaptive Regional I-O (ARIO) model (Hallegatte 2008) introduces a critical non-linearity: a production bottleneck constraint. Firms cannot produce beyond their current capacity even if orders exist, and capacity takes time to rebuild:

$$x_i(t) = min(x_i^{demand}(t), \quad x_i^{\text{capacity}}(t))$$

$$x_i^{\text{capacity}}(t+1) = x_i^{capacity}(t) + \mu_i \times (x_i^{max} - x_i^{capacity}(t))$$

- $x_i^{demand}(t)$ = I-O derived demand for sector i at time t: what customers want
- $x_i^{capacity}(t)$ = physically available capacity of sector i at time t: what the sector can deliver
- μ_i = reconstruction rate: fraction of damaged capacity restored per period (calibrated from post-event data: typically 0.05–0.20 per month for complex industrial facilities)

The capacity constraint creates a bottleneck amplification: when a critical sector (e.g. specialty chemicals, semiconductors, port operations) loses capacity, all downstream sectors simultaneously experience supply shortfalls even if their own facilities are intact. The total loss from the constrained ARIO model is generally greater than the unconstrained I-O model because bottlenecks prevent demand reallocation from clearing excess demand.

Several models derived from ARIO extend this framework, notably the dynamic production network models developed by Pichler, Farmer and co-authors (e.g. Pichler et al. 2022, applied to the COVID-19 shock), which add explicit inventory dynamics, production functions distinguishing critical from non-critical inputs, and endogenous demand responses.

Channel D: Financial network propagation — DebtRank

DebtRank (Battiston et al., 2012) is a standard model of financial network contagion. Each bank i has equity E_i and holds claims on other banks. A physical risk shock reduces the equity of exposed banks, and the loss propagates through bilateral exposures. Following Bardoscia et al. (2015), each bank passes on only its new distress at each step:

As written, each bank passes on its full cumulative distress at every step, so the same loss is counted repeatedly. The fix is to propagate only the new distress (Bardoscia et al., 2015):

$$h_i(t+1) = \min\left(1,\ h_i(t) + \sum_j W_{ji}\left(h_j(t) - h_j(t-1)\right)\right)$$

- $h_i(t)\ is the$ relative equity loss of bank i at time t: 0 = no loss, 1 = full default
- $W_{\{ji\}} = \frac{e_{\{ji\}}}{E_i}$ is the impact weight: fraction of bank i's equity that would be wiped out by a full default of bank j (where $e_{\{ji\}}$ is the bank *i*'s exposure to bank *j*)
- $h_j(t) - h_j(t-1)$ is the new distress of bank *j* during the last step.

The total systemic impact of a physical risk shock is the sum of equity losses across all banks:

$$Systemic\ Loss = \Sigma_i h_i(\infty) \times E_i.$$

The multiplier $\Sigma_i h_i(\infty) \times E_i\ /\ initial\ shock$ measures the financial network amplification.

MATHEMATICAL REPRESENTATION: FOUR PROPAGATION CHANNELS

Calibrations of European and US banking networks find amplifiers of 1.5–3.0× for moderate shocks and up to 5–8× for shocks near systemic tipping points [I]. These ranges are indicative: published DebtRank applications report results for specific networks and dates (e.g. Battiston et al., 2012, on US institutions receiving emergency Federal Reserve funding in 2008–2010; Battiston et al., 2017, for a climate-policy application to the euro-area financial system), and amplification must be recomputed on the supervisor's own exposure network.

Integrated propagation: combining channels

$$Total\ Loss = Direct_{cat} + Indirect_{I-O} + Capacity_{ARIO} + Financial_{DebtRank} + Infrastructure_{network}$$

In practice, these channels are partially additive and partially substitutable.

The insight given by this specific approach is that ignoring any single channel systematically underestimates total economic impact.

For the Thailand 2011 floods, total damage and losses amounted to about USD 46.5 billion, and the manufacturing sector bore roughly 70% of this total after the flooding of six industrial estates, which then disrupted supply chains well beyond Thailand (World Bank, 2012).

Graph-theoretic measures: identifying critical nodes

Beyond the propagation models, graph theory provides metrics to identify which nodes are most systemically important, that is, whose failure would generate the largest propagation loss. The metrics below apply to firm-level networks, such as real supply chains or interbank exposures. Sector-level I-O matrices have their own centrality measures, such as PageRank-type indices or forward and backward linkage indices. Three metrics are useful for physical climate risk analysis:

- **Betweenness centrality:** the fraction of all shortest paths in the network that pass through node *i*. A node with high betweenness is a chokepoint: many supply or financial flows route through it. Its physical damage propagates to a disproportionate share of the network.

$$BC(i) = \Sigma_{s \neq i \neq t} \frac{\sigma_{st}(i)}{\sigma_{st}}$$

- where σ_{st} is the number of shortest paths from s to t and $\sigma_{st}(i)$ is those passing through i.

For the global semiconductor supply chain, TSMC (Taiwan Semiconductor Manufacturing Company) has the highest betweenness centrality of any single manufacturer globally, making it the single most consequential node for physical climate risk network propagation.

In the global semiconductor supply chain, TSMC is widely regarded as a critical chokepoint, which makes it a prime candidate for this kind of analysis.

- **Exposure-weighted centrality (EWC):** betweenness centrality multiplied by the physical hazard exposure at the node's location.

$$EWC(i) = BC(i) \times Hazard_{score}(i).$$

- This is the quantity that directly determines which nodes should be prioritised for adaptation investment or alternative sourcing strategies: high EWC nodes are both critical to the network and highly exposed to physical risk.

Systemic importance of financial institutions (G-SIB methodology):

In financial networks, the closest analogue of centrality measures is the indicator-based methodology used to identify systemically important global banks (G-SIBs). This methodology is defined by the Basel Committee on Banking Supervision (BCBS, 2011, revised in 2013 and 2018), not by the FSB, and relies on five equally weighted categories: size, interconnectedness, substitutability/financial institution infrastructure, complexity, and cross-jurisdictional activity. The FSB publishes the resulting list of G-SIBs each November. For insurers, the FSB discontinued the annual identification of global systemically important insurers (G-SIIs) in December 2022, following the IAIS's adoption of the Holistic Framework for systemic risk in the insurance sector (2019). G-SIB scores are indicator-based measures of the impact that a bank's failure would have on the global financial system, not of its probability of failure. Physical climate exposure therefore does not raise a bank's G-SIB score as currently measured: it raises the probability that a systemically important institution experiences distress. For climate analysis, the relevant combination is systemic importance (G-SIB or D-SIB score) together with the climate-driven probability of distress, which is how exposure-weighted measures such as the EWC above should be read in a financial context.

Model Representation Card – Shock propagation – Tropical cyclones (De Maximy et al., 2024)

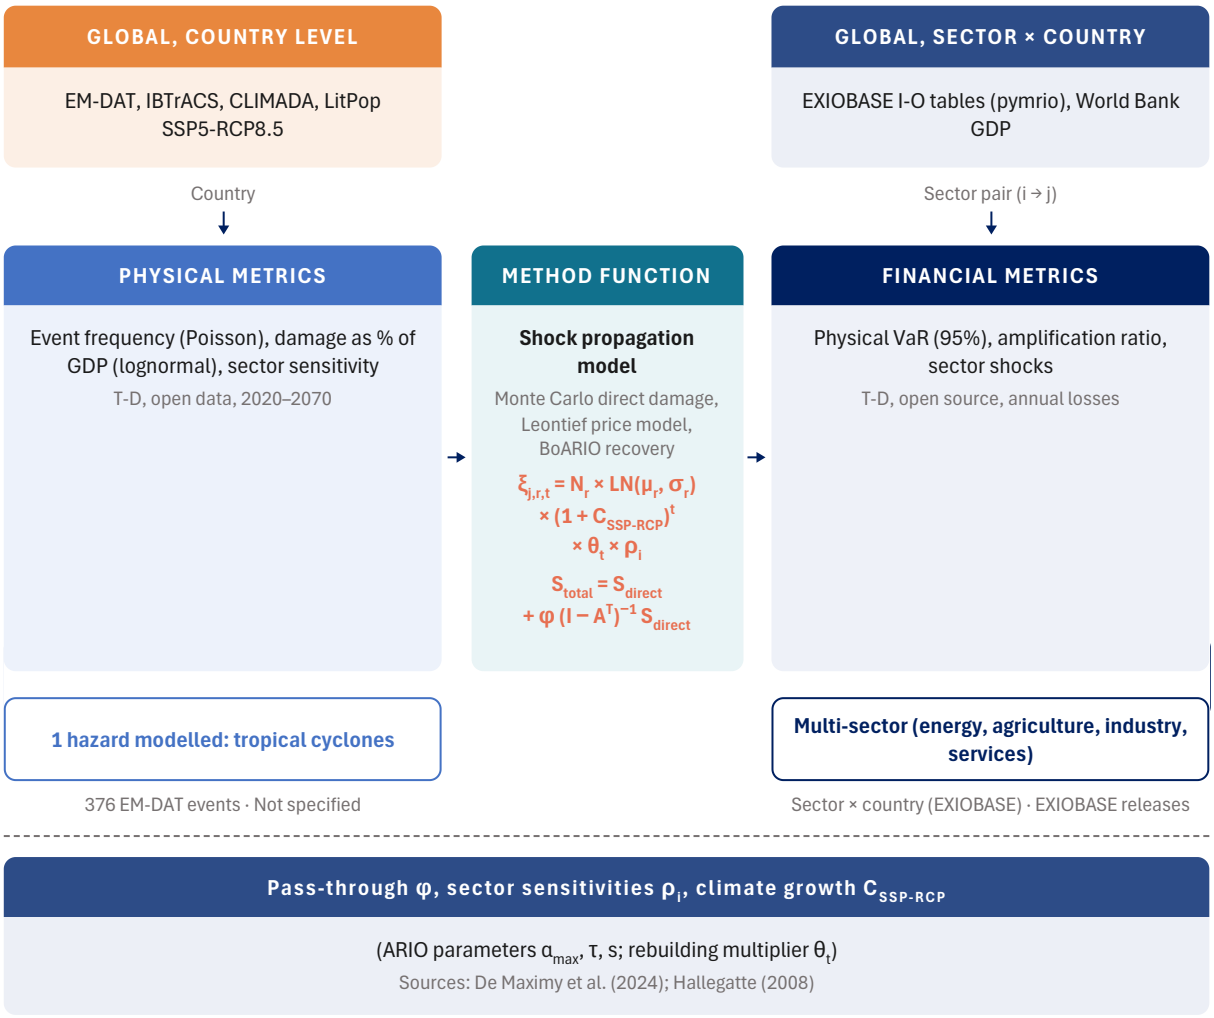


FIGURE 58. Model representation card – Shock propagation model (tropical cyclones, De Maximy et al., 2024)

Why is this model useful?

The model is useful if you need a portfolio-ready estimate of indirect physical losses without asset-level geolocation data. All of its inputs are public (EM-DAT, IBTrACS, EXIOBASE), and its tools are open source (pymrio, BoARIO). It turns a direct damage distribution into sector shocks that can be weighted by market capitalisation and read as a physical VaR, which is the format asset managers and risk teams already use.

What does this model teach us?

Indirect losses are of the same order as direct losses. In the authors' illustrative 20% direct shock, the amplification ratio reaches 67% inside the United States and 74% inside Australia, with further spillovers to the rest of the world. Frequency matters as much as severity. Under their assumptions, an event with a 20-year return period could recur about every four years by 2050, and the rebuilding multiplier θ_t captures the losses that compound when a new event hits before recovery is complete. The authors also show the limits of the approach. ARIO parameters are event-specific, the pass-through of localised physical shocks is uncertain, and the numerical examples are presented as illustrative only.

2.4.3 SUPPLY CHAIN IMPACT

Network theory and the operational supply chain risk

Section 2.4.2 described the mathematics of propagation through abstract networks. This section operationalises those mechanics at the level of the firm's extended supply chain: the specific web of Tier-1, Tier-2, and deeper-tier suppliers whose physical exposure and operational continuity determine the firm's own revenue at risk from physical climate events.

Supply chain physical climate risk assessment is a distinct discipline from macro network modelling. Where the I-O framework treats sectors as aggregated nodes in a national economy, supply chain models work with named suppliers at specific geographic locations with known product categories, volumes, and contractual relationships. This granularity enables a qualitatively different set of management actions: not just 'the electronics sector in Japan is exposed to typhoon risk' but 'Supplier X in Kumamoto, who provides 62% of our microcontroller supply, is in the projected track of Typhoon Nanmadol with a 34% probability of production shutdown exceeding 15 days'.

Empirical evidence: physical shocks propagate along supply chains

The propagation of natural disaster shocks through production networks has been widely documented in the US and Japan:

- **United States.** Major natural disasters over the past 30 years were used to identify firm-level shocks (Barrot & Sauvagnat, 2016). Affected suppliers imposed substantial output losses on their customers, especially when they produced specific inputs. These losses translated into market value losses and spilled over to other suppliers.
- **Japan.** The Great East Japan Earthquake of 2011 was studied with inter-firm transaction data (Carvalho et al., 2021). The disruption propagated upstream and downstream, reaching both direct and indirect suppliers and customers of the firms hit. The aggregate effect was estimated at a 0.47 percentage point decline in Japan's real GDP growth in the following year.
- **Japanese affiliates in the United States.** The same event was used to estimate input substitutability (Boehm et al., 2019). Output of these affiliates fell roughly one-for-one with the decline in imports from Japan. This is consistent with a near-Leontief relationship between imported and domestic inputs. In the short run, firms barely substitute away from a disrupted input.

- **Climate-specific evidence.** Heat at supplier locations was shown to reduce the operating income of both suppliers and their customers (Pankratz & Schiller, 2024). Customers also responded. When a supplier's realised exposure exceeded expectations, the probability of ending the relationship increased. Replacement suppliers had lower expected exposure. Similar results were found for floods.

A FIRM-LEVEL SUPPLY CHAIN PHYSICAL RISK FRAMEWORK

The firm-level approach was formalized and tested at Ford Motor Company (Simchi-Levi et al., 2015). It is usually carried out in five steps.

Step 1: Mapping and geo-location. Suppliers are located at their production sites, since a billing address can be far from where goods are made. Coverage is usually high for Tier-1 and partial beyond it. Lack of data on lower-tier suppliers was identified as a major obstacle at Ford (Simchi-Levi et al., 2015).

Step 2: Hazard overlay and exposure scoring. Each supplier facility is overlaid with standardized physical hazard datasets for the six most material perils: river flood (JBA Global Flood Maps), tropical cyclone wind (IBTrACS best-track data), earthquake (GEM Global Earthquake Hazard Map), water stress (WRI Aqueduct), heat stress (Copernicus ERA5) and wildfire (GlobFire database). Heat and floods at supplier locations were shown to reduce the operating income of both suppliers and their customers (Pankratz & Schiller, 2024). A composite physical risk score is assigned to each facility with a weighted scoring approach, calibrated to the vulnerability of the firm's sector. Electronic components, for instance, are particularly sensitive to flood inundation.

Step 3: Criticality. Exposure alone is a poor guide. A highly exposed supplier of a commodity part with many alternatives poses less risk than a moderately exposed sole source of a specific component. Disaster shocks were found to propagate mainly through specific inputs (Barrot & Sauvagnat, 2016). At Ford, exposure was unrelated to spend, and some of the largest exposures came from low-cost parts (Simchi-Levi et al., 2015). A criticality index CI_i is therefore built from substitutability, qualification lead time and revenue dependence, not from spend.

Step 4: Revenue at risk. Exposure and criticality are combined into an expected revenue loss per supplier. The key comparison is between how long a supplier takes to recover and how long the firm can operate without it (Simchi-Levi et al., 2015). The formulation is given below.

The Revenue at Risk (RAR) for each supplier i is computed as:

$$\boldsymbol{RAR_i = P(disruption_i \mid hazard) \times Severity_i \times CI_i \times (Revenue\ from\ supplier_i / 365)}$$

Where:

- $\boldsymbol{RAR}$ is the $Revenue\ at\ risk_i$, the expected annual revenue loss caused by the disruption of supplier i. Aggregated across all suppliers, the portfolio RAR gives the firm's total exposure to physical climate supply chain disruption.
- $P(disruption)$ is the annual probability that a hazard stops production at supplier i for longer than the firm's inventory buffer.
- $\boldsymbol{Severity_i}$ is the expected number of days of lost production beyond the buffer, given such a disruption
- $\boldsymbol{CI_i}$ is the criticality index mentioned before
- $\boldsymbol{Rev_i}$ or $Revenue\ from\ supplier_i$ is the annual revenue of the product lines that cannot be produced without supplier i, so Rev_i / 365 is the revenue lost per day of disruption.

A FIRM-LEVEL SUPPLY CHAIN PHYSICAL RISK FRAMEWORK

Step 5: Resilience investment optimisation.

Revenue at risk is used to allocate a resilience budget across diversification (qualifying alternative suppliers), buffer inventory (increasing safety stock at critical nodes), supplier adaptation support (co-investing in site hardening) and contractual risk transfer (requiring business continuity plans as a procurement condition).

The allocation minimises the total RAR subject to the budget. Diversification only helps if the alternative supplier lies outside the same hazard footprint. Such adjustments were observed in practice. Firms replaced suppliers whose realised climate exposure exceeded expectations with less exposed ones (Pankratz & Schiller, 2024).

Mathematical Representation: Revenue at Risk (RAR) Model

REVENUE AT RISK: THE CORE FORMULA

Consider a small extension of the formula we have just given, the $Revenue\ at\ Risk$:

$$RAR_i = P(D_i > B_i) \times E[D_i - B_i \mid D_i > B_i] \times CI_i \times \rho_i \times \frac{Rev_i}{365}$$

- $\boldsymbol{P(D_i > B_i)}$ is the probability that disruption duration D_i exceeds the inventory buffer B_i (days).
- This is derived from:

$$P(D_i > B_i) = P(\textbf{hazard event}) \times P(\textbf{production halt} \mid \textbf{hazard}) \times P(D_i > B_i \mid \textbf{halt})$$

The first term comes from the cat model OEP; the second from the damage function; the third from a disruption duration distribution calibrated to historical supply chain events

- $\boldsymbol{E[D_i - B_i \mid D_i > B_i]}$ is the expected number of days of disruption beyond the buffer, given that the buffer is exceeded (Embrechts et al., 1997): the tail expectation of the downtime distribution above B_i (analogous to TVaR for financial losses)
- $\boldsymbol{CI_i}$ is the Criticality Index. It ranges from 0 (fully substitutable) to 1 (sole source, no alternative).

$$CI_i = f(substitutability, lead\ time, revenue\ impact)$$

- $\boldsymbol{\rho_i}$ is the revenue impact rate, the share of daily revenue lost per day of supply disruption from supplier i (from 0 to 1), accounting for partial recovery from alternative sourcing at premium cost and inventory drawdown.
- $\boldsymbol{Rev_i}$ is the total revenue at risk from this supplier's component (the revenue of product lines that cannot be produced without this supplier's component).

Portfolio aggregation with correlation

$$E[L] = \sum_i RAR_i \qquad \mathrm{Var}(L) = \sum_i \mathrm{Var}(L_i) + \sum_{i \neq j} \mathrm{Cov}(L_i, L_j)$$

- $\boldsymbol{L_i}$ is the annual revenue loss caused by the disruption of supplier i, so that $RAR_i = E[L_i]$

REVENUE AT RISK: THE CORE FORMULA

- $\boldsymbol{L = \sum_i L_i}$ is the total annual revenue loss of the supplier portfolio.
- $\boldsymbol{Var(L_i)}$ is the variance of the loss from supplier i, and $\mathrm{Cov}(L_i, L_j)$ is the covariance between the losses from suppliers i and j.

The covariance term is critical because suppliers in the same geography face correlated physical risk from the same events. $\mathrm{Cov}(L_i, L_j)$ for two suppliers in the same flood zone is strongly positive, since they both fail in the same event. The expected loss is unchanged, but the tail of the portfolio loss widens. Concentration risk is therefore measured with a tail metric:

$$RAR_\alpha = \mathrm{VaR}_\alpha(L)$$

- $\boldsymbol{RAR_\alpha}$ is the portfolio revenue at risk at confidence level α (for example 95%), and $\mathrm{VaR}_\alpha(L)$ is the α-quantile of the annual loss distribution.
- This is the geographic concentration risk that naive single-supplier analysis misses. The covariance is estimated from the event-level correlation in the cat model stochastic event set. If both suppliers are in the footprint of the same simulated event, their failures are perfectly correlated for that event.

Resilience optimization: minimizing RAR subject to budget

$$\min_x \; RAR(x) = \sum_i RAR_i(x_i) \quad \textbf{s.t.} \quad \sum_i \mathbf{Cost}(x_i) \leq \mathbf{Budget}$$

Where x_i is the resilience investment choice for supplier i (dual-source qualification, inventory buffer increase, facility hardening, contractual protection).

- $\boldsymbol{RAR_i(x_i)}$ is a decreasing function of x_i: more resilience investment reduces the probability and duration of disruption. $\mathrm{Cost}(x_i)$ is the cost of that investment

Typical resilience levers and their RAR reduction efficiency (€of RAR reduced per €of investment) are:

Dual sourcing: reduces CI_i from 1.0 to 0.5–0.7 (depending on geographic correlation of alternative source). Investment: qualification cost + price premium on diversified volume. RAR reduction: 30—60% per sole-source relationship addressed [I].

Safety stock increase: increases B_i from 14 days to 45 days, reducing $P(D_i > B_i)$ for events of moderate duration. Investment: inventory carrying cost (typically 15–25% of inventory value p.a.). RAR reduction: 40–70% for events with duration < 45 days; lower for longer events, since the extra 31 days of buffer then cover only a small share of the disruption [I].

Supplier adaptation co-investment: reduce $P(production\ halt \mid hazard)$ by supporting supplier facility hardening (flood barriers, backup power, elevated equipment). Investment: 50–80% of adaptation cost shared. RAR reduction: 20–45% depending on hazard type and adaptation measure effectiveness [I].



FIGURE 59. Model representation card – Supply chain business interruption (ClimVaR, Epelbaum, 2026)

What does the model do?

Supply chain physical risk is modelled inside a discounted cash flow valuation called ClimVaR (Epelbaum, 2026). The physical impact is expressed as business interruption days (BID), the share of the year during which production stops.

First, the direct interruption of a site or country-sector is obtained by combining the interruption caused by each hazard, assuming the hazards are independent:

$$\boldsymbol{\beta}^{(t_0)} = 1 - \prod_{h \in H} \left(1 - \boldsymbol{\beta}^{(t_0 h)}\right)$$

- $\boldsymbol{\beta}^{(t_0)}$ is the total direct interruption share of the site or country-sector.
- $\boldsymbol{\beta}^{(t_0 h)}$ is the direct interruption share caused by hazard h
- $\boldsymbol{H}$ is the set of hazards covered.

Second, the interruption is propagated upstream through the multi-regional input-output (MRIO) table. A firm stops producing when its own site stops, or, on the days it is running, when its suppliers stop:

$$\boldsymbol{\beta}^{(t)} = \boldsymbol{\beta}^{(t_0)} + \left(\mathbf{1} - \boldsymbol{\beta}^{(t_0)}\right)\left[A^{\top}\boldsymbol{\beta}^{(t)}\right] \quad \Leftrightarrow \quad \boldsymbol{\beta}^{(t)} = (I - D\,A^{\top})^{-1}\boldsymbol{\beta}^{(t_0)}$$

- $\boldsymbol{\beta}^{(t)}$ is the total interruption share, including the whole value chain.
- $\boldsymbol{A}$ is the matrix of technical coefficients from the OECD ICIO tables (45 ISIC sectors), and $\boldsymbol{A}^{\top}$ is its transpose.
- $\boldsymbol{DD} = \boldsymbol{diag}\left(\mathbf{1} - \boldsymbol{\beta}^{(t_0)}\right)$ counts supplier interruption only on days when the firm itself is running.
- $\boldsymbol{I}$ is the identity matrix.

Three levels of information are handled:

1. country-sector benchmarks;
2. a geolocated site with an unknown supply chain;
3. a site whose rank-1 suppliers are known, where the coefficients become
4. $\boldsymbol{A_{aa'}} = \boldsymbol{P_{aa'}}/\boldsymbol{Y_a}$ purchases from supplier a' divided by the site's revenue.

The total interruption is then split by supply rank (direct operations, rank 1, rank 1 and beyond) and by hazard, using the Shapley value. The interruption share reduces revenue in the cash flows. Inputs are drawn from lognormal distributions in a Monte Carlo simulation, which yields a ClimVaR and an expected shortfall at 95%.

What does it teach us?

- A firm can lose revenue from climate events while none of its own sites is hit. Upstream interruptions add to the direct ones, and the propagation term $\boldsymbol{A}^{\top}\boldsymbol{\beta}^{(t)}$ can dominate for sectors that depend heavily on imported inputs.
- Decomposition by rank shows where in the value chain the risk sits. A firm exposed mainly at rank 2 and beyond gains little from hardening its own sites or auditing its direct suppliers.

Supplier data has measurable value. Replacing sector averages with actual purchase shares at rank 1 changes the estimate. In the paper's illustrative example, precise financial and geospatial information reduces the CVaR at 95% by about 45%.

The framework is fully transparent. Each formula is published, so the model can be replicated and audited, unlike most vendor tools.

Why is it useful?

For corporates: it links procurement, operations and finance in one valuation. Adaptation actions (relocation, dual sourcing, site hardening) are ranked by return on investment against the reduction in ClimVaR.

For investors and banks: the same formalism applies to a portfolio, a company, an entity, a site or a product. It yields a supply chain component of climate risk that is consistent across scales and can feed a valuation or a credit assessment.

It is one of the few published models that treats supply chain business interruption as a distinct, quantified channel. Most physical risk models stop at the asset's own location.

The Tier-N visibility problem

The single most intractable challenge in supply chain physical climate risk assessment is Tier-2 and deeper visibility. Most firms have reasonable knowledge of their Tier-1 suppliers. Very few have reliable knowledge of their Tier-2 suppliers (the suppliers' suppliers), and essentially none have systematic visibility of Tier-3 and beyond. Yet physical climate events most frequently cause catastrophic supply chain disruption not through Tier-1 damage but through Tier-2 and Tier-3 bottlenecks (e.g. a single component manufacturer six tiers deep in the automotive supply chain whose sole-source chemical feedstock supplier is in a flood zone).

The 2011 Japan earthquake and tsunami illustrated this vividly: the automotive supply chain disruption was ultimately caused not by assembly plant damage (Tier-1, rapidly identified) but by the destruction of a single plant in Miyagi prefecture that manufactured 60% of the world's supply of a specific automotive-grade pigment used in black paint, which took nine months to restore. No mainstream supply chain risk management tool had mapped this dependency before the event.

- **Commercial supply chain intelligence databases** (Resilinc, Riskmethods, Dun & Bradstreet Supply Chain Analytics): these platforms use spend data, trade flow data, and supplier self-disclosure to construct probabilistic Tier-2 and Tier-3 maps. Coverage is best for electronics and automotive; weakest for agriculture and pharmaceuticals.
- **Trade flow reconstruction** using customs data (COMTRADE, UNCTAD, Panjiva/S&P Global Market Intelligence): import/export flows at the HS code level reveal supplier dependencies that firms do not disclose directly.
- **Input-output imputation** using EXIOBASE or WIOD: the national I-O table for each country provides probabilistic estimates of which sectors supply which other sectors, allowing Tier-2 dependencies to be inferred from Tier-1 sector classifications even without firm-level data.
- **Satellite imagery and alternative data** (Kayrros, Planet Labs, Maxar, Descartes Labs): satellite-derived signals of facility activity (parking lot occupancy, thermal signatures, nighttime light emission) can detect supply disruptions in real time even when suppliers do not disclose them, providing event-triggered intelligence for Tier-N nodes.

2.4.4 SYSTEMIC RISK

What is systemic risk and how does physical climate change it?

Systemic risk can be defined in several ways. It has been defined as "a risk of disruption to financial services that is caused by an impairment of all or parts of the financial system and has the potential to have serious negative consequences for the real economy" (IMF, BIS & FSB, 2009).

A slightly different definition is adopted here. Systemic risk is understood as the risk that the failure of one institution, market or infrastructure node triggers a cascade of failures across the financial system as a whole, with consequences that cannot be contained by the normal mechanisms of market discipline, insurance or central bank intervention. It differs from the sum of individual risks, since it is an emergent property of the network (Haldane & May, 2011; Battiston et al., 2016).

Physical climate change increases systemic risk through three distinct mechanisms that have no precedent in the historical calibration of systemic risk models:

- **Common exposure concentration:** climate hazards affect broad geographies, simultaneously impairing the balance sheets of many financial institutions with correlated exposures. A severe Atlantic hurricane season damages all banks with mortgage, corporate, and infrastructure loan books in the affected coastal

regions simultaneously. This common factor structure eliminates the diversification benefit that makes individual risks manageable at the system level.

- **Non-stationarity of the risk environment:** systemic risk models (VaR, stressed VaR, SRISK) are calibrated to historical financial data. Physical climate change alters the underlying distribution of shocks (e.g. higher frequency, higher intensity, shifted geography) making historical calibration increasingly unreliable. The systemic risk models are trained on a world that is progressively less representative of the world they are being used to manage.
- **Insurance capacity constraints:** the insurance sector is the primary absorption mechanism for physical risk losses at the firm and household level. The retreat of insurers operates mainly through pricing and non-renewal at annual renewals, rather than through an annual exhaustion of reinsurance capacity: as expected losses and reinsurance costs rise, premiums increase, terms tighten (higher deductibles, exclusions) and policies in the riskiest areas are not renewed. The uninsured share then cascades into the financial system as non-performing loans, sovereign contingent liabilities, and real economy demand destruction. Risk also concentrates in insurers of last resort: public or quasi-public residual-market and pooling schemes (Citizens Property Insurance in Florida, the California FAIR Plan, the US National Flood Insurance Program, Flood Re in the United Kingdom, the CCR-backed Cat Nat regime in France) grow as private insurers retreat and become fiscal and financial nodes in their own right, whose deficits fall on taxpayers, on policyholders through post-event assessments, or on the State. The progressive retreat of insurance from high-risk markets (coastal flood in Florida, wildfire in California, cyclone in parts of Australia and the Pacific) is transferring physical risk from the balance sheets of insurers (where it is priced and reserved for) to the balance sheets of banks, households, governments, and ultimately central banks.

MATHEMATICAL REPRESENTATION: SRISK AND CLIMATE-ADJUSTED SYSTEMIC RISK

Srisk: the canonical systemic risk measure (Brownlees & Engle, 2017)

$$SRISK_i = \max(0,\ k \cdot D_i - (1-k) \cdot (1 - LRMES_i) \cdot W_i)$$

$SRISK_i$ is the capital shortfall of institution i in a systemic stress event: the amount by which its equity would fall below the prudential minimum

- $\boldsymbol{W_i = MV_i\ (equity)}$, it is the current market capitalisation of institution i
- $\boldsymbol{D_i = Book\ Debt_i}$, it is the book value of debt obligations
- $\boldsymbol{k}$ is the prudential capital ratio applied to total assets, a quasi-leverage ratio. It is 8% in the baseline V-Lab specification and 5.5% for European institutions reporting under IFRS.
- $\boldsymbol{LRMES_i}$ is the Long-Run Marginal Expected Shortfall: the expected fractional loss of institution i's equity value conditional on a broad market decline of 40% over six months. It is commonly approximated from the daily MES (Acharya, Engle & Richardson, 2012) as:

$$LRMES_i \approx 1 - \exp\left(-18 \times MES_i^{daily}\right)$$

Total system SRISK is $\sum_i \max(SRISK_i, 0)$.

Without climate adjustment, $SRISK$ aggregates the capital shortfall across the financial system in a severe market downturn. Climate-adjusted $SRISK$ extends this by conditioning $LRMES$ on a physical climate shock, instead of or in addition to a market shock.

MATHEMATICAL REPRESENTATION: SRISK AND CLIMATE-ADJUSTED SYSTEMIC RISK

The $SRISK$ formula follows from the definition. After a market crash, equity falls to $((1 - LRMES_i)\, W_i$. The required capital is k times total assets $D_i + (1 - LRMES_i)\, W_i$

The capital shortfall is the required capital minus the remaining equity, which gives $k \cdot D_i - (1 - k) \cdot (1 - LRMES_i) \cdot W_i$

Climate-adjusted SRISK: incorporating correlated physical losses

A following extension is proposed here:

$$MES_i^{climate}(s) = E\left[\frac{\Delta W_i}{W_i} \,\middle|\, \text{Physical shock} = s,\ \text{System in distress}\right]$$

The physical shock s is a realisation of the climate scenario (e.g. a 1-in-50-year correlated flood event across a major financial center).

- $\boldsymbol{\Delta W_i}$ is the equity loss of institution i resulting from: (a) direct credit losses on its climate-exposed loan book; (b) mark-to-market losses on climate-exposed securities; (c) business interruption losses on its own operations; and (d) reputational and market confidence losses if climate risk disclosure reveals concentrated exposure.
- $\boldsymbol{MES_i^{climate}(s)}$ replaces $LRMES_i$ in the SRISK formula for each institution, and the system-level measure is obtained by summing the resulting capital shortfalls

Network default cascade: the threshold model

In the threshold model of financial contagion (Gai & Kapadia, 2010), each bank defaults when its equity buffer is exhausted. When a bank j defaults, each of its creditors i loses its exposure e_{ji} times the loss given default on bank debt. The cascade propagates as follows:

$$E_i(t+1) = E_i(0) - L_i^{climate}(t) - \sum_{j \in \mathcal{D}(t)} e_{ji} \times LGD_j$$

$$\textbf{Bank i defaults if } E_i(t+1) < 0$$

- $\boldsymbol{E_i}$ is the equity of bank i at step t and $E_i(0)$ is its initial equity.
- $\boldsymbol{L_i^{climate}(t)}$ is bank i's cumulative climate-related credit loss by step t, from the PD model in section 2.2.
- $\boldsymbol{\mathcal{D}(t)}$ is the set of banks that have defaulted by step t.
- $\boldsymbol{e_{ji}}$ is the exposure of bank i to bank j
- $\boldsymbol{LGD_j}$ is the loss given default on the debt of bank j.

The term $L_i^{climate}(t)$ is the direct physical channel. The term $\sum e_{ji} \times LGD_j$ is th financial contagion channel, which propagates losses from failed counterparties. Simulations on actual interbank networks find that contagious defaults are unlikely, but that when contagion does take place, it can bring down a substantial fraction of the banking system (Upper, 2011). Contagion in such networks has been shown to be "robust-yet-fragile", rare but potentially very large (Gai & Kapadia, 2010). Amplification must therefore be computed on the actual exposure network, and cannot be transposed from another system.

Non-linearity of systemic risk

MATHEMATICAL REPRESENTATION: SRISK AND CLIMATE-ADJUSTED SYSTEMIC RISK

Note that the relationship between network structure and systemic risk is non-linear. When shocks are small, a densely connected financial network is more stable, because losses are shared among many counterparties. Beyond a critical shock size, the same interconnections become the channel through which losses propagate, and the network turns fragile (Acemoglu et al., 2015). Small increases in shock severity near this threshold therefore produce disproportionately large increases in systemic risk. This non-linearity explains why physical climate change is a systemic risk concern at warming levels that look modest in terms of average GDP impact. The distribution of outcomes has a fat tail at the system level even when the expected value of individual losses is moderate.

The FSB systemic risk assessment framework for climate

The FSB's 2025 framework highlights several amplification mechanisms through which physical climate risk can become systemic (FSB, 2025):

- **Cliff-edge effects from insurance withdrawal:** when insurers withdraw from a market, properties become unmortgageable (most mortgage lending requires insurance), triggering forced sales, price collapse, and bank collateral deterioration simultaneously. This is not a gradual process: it is a discrete cliff-edge once insurers collectively determine a market is uninsurable, which can happen within a single underwriting cycle. Empirical work on the insurance–mortgage nexus supports this mechanism: Blickle and Santos (2022) find that mandatory flood insurance requirements following flood-map updates reduce mortgage lending in newly mapped areas; Sastry, Sen and Tenekedjieva (2023) show that, in Florida, climate losses push large national insurers out and that lenders respond to the resulting reliance on fragile local insurers by transferring more mortgages to the government-sponsored enterprises; and Keys and Mulder (2024) document sharp increases in homeowners' insurance premiums in high-risk areas using mortgage escrow data.
- **Fire sale dynamics in climate-exposed asset markets:** when market participants simultaneously update their climate risk assessments (triggered by a new flood map, a new regulatory requirement, or a major climate event), the simultaneous repricing of climate-exposed assets can trigger liquidity-driven price spirals. The 2020 Australian bushfire season triggered a measurable repricing of wildfire-exposed residential property in Victoria within three months of the peak fire season.
- **Central bank balance sheet exposure:** quantitative easing programs have placed climate-exposed assets (sovereign bonds, corporate bonds, covered bonds backed by mortgages) on central bank balance sheets. Climate-related credit events therefore have the potential to impair central bank equity, a concern explicitly raised by the ECB, Bank of England, and Banque de France in their climate risk assessments.

The case for considering macroeconomic variables

The systemic consequences of a natural catastrophe go beyond direct damage. A macroeconomic model can show how a single severe event spreads to growth, prices, interest rates and financial markets. In the scenario shown below, a magnitude 7.7 earthquake hits Los Angeles in 2018 and causes about USD 850 billion of damage (Mahalingam et al., 2018). The shock was simulated with the Oxford Economics global macroeconomic model and compared with a baseline without the event.

US GDP growth falls sharply in the year of the event and rebounds strongly during reconstruction. Bond yields and the central bank policy rate rise above the baseline, equity prices drop and recover only gradually, and inflation and the exchange rate fluctuate for several years. The trade balance also deteriorates. This figure shows that a physical event becomes a financial shock through the macroeconomy. Interest rates, asset prices and inflation are the

variables that feed into banks' and insurers' balance sheets. Systemic risk assessment should therefore include these macro-financial paths alongside direct losses.

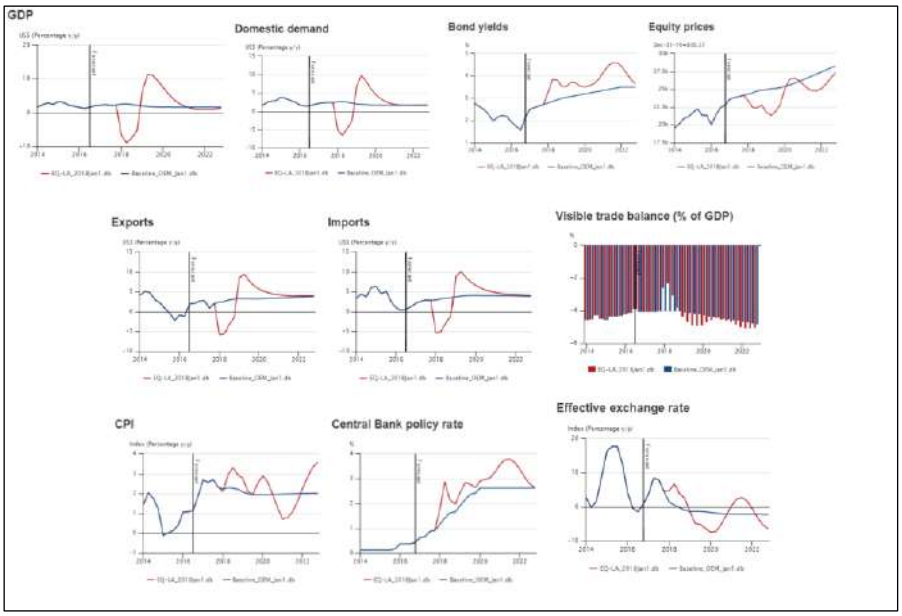


FIGURE 60. The case for considering macroeconomic variables – Impact of EQ-LA scenario on key macroeconomic variables. Source: Mahalingam et al., 2018

Model Representation Card – Flood-induced systemic risk – Financial networks (Mandel et al., 2021)

GLOBAL, 109 COUNTRIES (95% OF WORLD GDP)

GLOFRIS (river), DIVA (coastal)
SSP1-RCP2.6 / SSP5-RCP8.5

Country (30″ grid, 12,000 coastal segments)

109 COUNTRIES, 4 INSTITUTIONAL SECTORS

JRC Finflows, BIS locational banking statistics, World Bank GFDD, IMF debt data

Country × institutional sector

PHYSICAL METRICS

Yearly direct damage to physical capital (% of world GDP), return periods 2 to 1,000 years

T-D, open models, 2030 / 2050 / 2080

METHOD FUNCTION

Flood shock propagation in financial networks

Direct damage → insurance, government, private sector → domestic financial sector → DebtRank across countries

$$A_{j,n}(t) = [E_n(t-1) / E_n(t-2)] \, A_{j,n}(t-1)$$

FINANCIAL METRICS

CVaR 95% vs financial sector capital, financial amplification ratio, country-induced VaR

T-D, open data, 50,000 Monte Carlo draws

2 hazards: river and coastal floods

Flood maps, 109 countries · One-off study

Insurance, government, private and financial sectors

Cross-border exposures, 109 countries · One-off study

Adaptation (constant vs declining protection), insured share σ_n, fiscal space β

(state cover τ, insurer pass-through ρ, financial leverage and network centrality, static network, independent hazards)

Source: Mandel et al. (2021)

FIGURE 61. Model representation card – Flood-induced systemic risk through financial networks (Mandel et al., 2021)

What does the model do?

The model estimates how floods can threaten the stability of the global financial system. River and coastal flood damages are projected for 109 countries up to 2080, with and without adaptation. These losses are first shared between insurers, the State and private companies. They then reach the banks that lend to them, and finally spread to banks in other countries through cross-border financial links.

What does it teach us?

Financial links double the cost of floods on average, and can multiply it by ten in some cases. Adaptation is decisive. Without new flood defences, global risk in 2080 would be about ten times today's level, against 50% higher with adaptation. The most exposed countries are highly indebted financial hubs such as the United Kingdom, France, Germany, the Netherlands and Switzerland. In a severe scenario, their potential losses could approach the capital of their banking sectors, so a climate-driven financial crisis cannot be ruled out.

Why is it useful?

It is one of the few published models that links physical climate risk to systemic financial risk with real cross-border exposure data. It shows supervisors which countries generate the risk and which ones are most exposed. It also shows that adaptation protects financial stability as well as physical assets.

USER	DATA ACCESS AND FEASIBILITY	TYPICAL PRUDENTIAL USE CASES
Supervisor / central bank	Full: supervisory balance-sheet data, bilateral exposures and market data.	Macroprudential: system-wide stress tests, calibration of systemic or sectoral buffers, concentration monitoring, G-SIB/D-SIB analysis.
Financial institution	Partial: market data for SRISK/CRISK on itself and peers; no network data.	Microprudential: reverse stress testing in ICAAP/ORSA, counterparty risk management.
Researcher / market participant	Partial: SRISK and CRISK computable from public market data (e.g. NYU V-Lab); no network data.	Market monitoring and research; no direct disclosure use.

TABLE 28. Implementation and Use Cases: The systemic risk model

2.4.5 LAGGED EFFECTS

Why timing matters: the multi-speed nature of physical climate impacts

Physical climate risk does not produce all of its financial consequences in the quarter in which the hazard event occurs. Bank stability deteriorates for several years after a weather-related disaster (Noth & Schüwer, 2023), and national incomes may remain below their pre-disaster trend for two decades (Hsiang & Jina, 2014)

The full economic and financial impact of a major climate event unfolds over months, years, and sometimes decades, through mechanisms that operate at very different speeds. Understanding the timing structure of physical climate impacts is essential for: IFRS 9 provisioning (which requires lifetime ECL for Stage 2 loans); long-horizon stress testing (supervisory climate stress tests, with horizons from 3 to 30 years); infrastructure investment appraisal; and sovereign debt sustainability analysis.

The lagged effects of physical climate risk can be categorized into four temporal phases:

PHASE	TIME HORIZON	PRIMARY MECHANISM	FINANCIAL MANIFESTATION	MODEL REQUIRED	EVIDENCE
Immediate	0–3 months	Direct physical damage, business interruption	Insurance claims, collateral revaluation, equity repricing, first IFRS 9 Stage 2 transfers	Catastrophe model, P&L bridge (2.3)	Damaged properties lose 17–22% of their value right after the event (Ortega & Taşpınar, 2018); significant abnormal returns of P&C insurers around most hurricane landfalls (Feria-Domínguez et al., 2017); output losses of hit suppliers pass to their customers (Barrot & Sauvagnat, 2016)
Short-run	3–18 months	Reconstruction demand, insurance market hardening, scarcity of materials and labour	Higher premiums, capex surge, input cost inflation, local labour market tightening	ARIO and input-output models (2.4.2)	Reconstruction wages and costs rise after catastrophes (“demand surge”) (Döhrmann et al., 2017); reinsurance premiums are high relative to expected losses, especially after catastrophes (Froot, 2001); floods and storms raise inflation in the short run, mainly in developing economies (Parker, 2018)
Medium-run	18 months–5 years	Capital reallocation, firm relocation, population migration	Property repricing, rising non-performing loans, rating migration, change in bank portfolio composition	Regional economic models, survival models (2.2), DCF (2.3)	Bank stability weakens for several years after disasters (Noth & Schüwer, 2023); a flood-zone price discount emerges gradually and persists (Ortega & Taşpınar, 2018)
Long-run	5–30 years	Chronic climate change, asset stranding, widening insurance gap,	Progressive LtV deterioration, higher cost of capital, stranded asset write-downs, sovereign	Non-stationary EVT (2.1.2), macro scenarios (2.4.1), EV model (2.3)	National incomes stay below trend twenty years after cyclones (Hsiang & Jina, 2014)

PHASE	TIME HORIZON	PRIMARY MECHANISM	FINANCIAL MANIFESTATION	MODEL REQUIRED	EVIDENCE
		regulatory response	debt accumulation		

TABLE 29. The four temporal phases

Mathematical representation: distributed lag models for physical risk

DISTRIBUTED LAG MODEL: CAPTURING THE FULL TEMPORAL IMPACT PROFILE

For a discrete climate event, the impact on a variable Y (GDP, credit losses, property values) over the following periods can be estimated with a distributed lag model (Hsiang & Jina, 2014):

$$Y_t = Y_t^{baseline} + \sum_{k=0}^{K} \beta_k \times Shock_{t-k} + \varepsilon_t$$

- Y_t is the value of the variable at time t
- $Y_t^{baseline}$ is its value without the event.
- $Shock_{t-k}$ is the intensity of the event that occurred k periods earlier.
- β_k = the lag-k impulse response coefficient: the marginal effect of a unit shock at time t-k on Y at time t
- K is the lag horizon, the number of periods over which the effects persist.
- ε_t is an error term.

The sequence $\{\beta_0, \beta_1, \dots, \beta_K\}$ is the impulse response function (IRF) and the cumulative effect $\sum_k \beta_k$ is the total impact. For a major flood, a stylised IRF for GDP shows the following pattern:

an immediate contraction from destruction $\beta_0 < 0$;

a reconstruction rebound $\beta_1 > 0$;

a recovery plateau $\beta_2, \dots, \beta_4 \approx 0$;

possibly a lasting loss $\beta_5, \dots, \beta_{20} < 0$ if reconstruction is financed by debt or productive workers emigrate.

Applied to 6,700 cyclones, this approach finds that national incomes fall relative to their pre-disaster trend and do not recover within twenty years, in both rich and poor countries (Hsiang & Jina, 2014). A synthetic-control study finds no significant long-run effect of even extremely large disasters once political changes are controlled for (Cavallo et al., 2013).

Credit loss lag structure: IFRS 9 and the delayed provisioning problem

$$ECL_t = \sum_{k=0}^{T} PD_{t+k|t} \times LGD_{t+k|t} \times EAD_{t+k|t} \times DF_k$$

- ECL is the lifetime expected credit loss at time t

DISTRIBUTED LAG MODEL: CAPTURING THE FULL TEMPORAL IMPACT PROFILE

- $\boldsymbol{PD_{t+k|t}}$ is the marginal probability of default in period $t+k$, conditional on the information available at t
- $\boldsymbol{LGD_{t+k|t}}$ is the loss given default in that period.
- $\boldsymbol{EAD_{t+k|t}}$ is the exposure at default in that period.
- $\boldsymbol{DF_k}$ is the discount factor at the effective interest rate.
- $\boldsymbol{T}$ is the remaining life of the loan.

The IFRS 9 lifetime ECL integrates the expected loss over the full remaining life of the loan, discounting each future period's loss at the effective interest rate. Physical climate risk enters through $PD_{t+k|t}$: the forward-looking probability of default in period t+k, conditional on current information.

For a 25-year mortgage on a coastal property, $PD_{t+k|t}$ increases over the life of the loan as sea-level rise, frequency of flooding, and insurance availability deteriorate.

A bank that provisions only for current-period default risk therefore systematically underestimates the lifetime ECL of climate-exposed loans.

Property value deferred stigma effect: the information cascade

$$V_t = V_0 \prod_{k=1}^{t} (1 + g_k) \times \left(1 - \pi_t(I_t)\right)$$

- $\boldsymbol{V_t}$ is the property value at time t, and V_0 its initial value.
- $\boldsymbol{g_k}$ is the market growth rate in period k
- $\boldsymbol{\pi_t}$is the climate stigma discount at time t between 0 and 1, which depends on the information I_t available to the market about the property's physical climate risk. π_t increases when new information arrives: new flood zone maps, insurance repricing, distressed sales of neighbouring properties, or high-profile climate events that update risk perceptions.

Property value deterioration from physical climate risk is therefore neither smooth nor proportional to the underlying risk. It is driven by information. After Hurricane Sandy, damaged properties in New York lost 17–22% of their value immediately. Undamaged flood-zone properties, priced like the rest of the city before the storm, developed a gradual discount that reached 8% five years later, with no sign of recovery (Ortega & Taşpınar, 2018). Such repricing can produce cliff-edge LtV deterioration in mortgage books. A portfolio well within covenant at 70% LtV can breach an 80% threshold for many loans at once when the information shock hits.

Long-run population and economic structure effects

$$N_{region}(t) = N_{region}(0)\ \exp\left(\int_0^t \left(g_{nat} - m_{climate}(s)\right) ds\right)$$

- $\boldsymbol{N_{region}}$(t) is the regional population at time t and $N_{region}(0)$ the initial population.
- $\boldsymbol{g_{nat}}$ is the natural population growth rate.
- $\boldsymbol{m_{climate}(s)}$ is the climate-driven emigration rate at time s the share of the regional population leaving per period because of physical risk (recurrent flooding, extreme heat, loss of the employment base).

DISTRIBUTED LAG MODEL: CAPTURING THE FULL TEMPORAL IMPACT PROFILE

The emigration rate can be estimated from observed migration flows after major events. After Hurricane Katrina, the population of Orleans Parish fell by about half between 2005 and 2006. After Hurricane Maria, Puerto Rico lost about 4% of its population in one year.

Population loss has cascading economic consequences. Tax base erosion reduces public services, which accelerates outmigration and further erodes property values and bank collateral. This self-reinforcing dynamic can turn physical climate risk into a lasting fiscal and financial concern in exposed regions. Sovereign credit models for small island developing states should therefore include this population channel in long-run debt sustainability projections.

Lagged effects and model horizon mismatch

The most practically important consequence of lagged physical climate effects is the mismatch between the horizon over which effects materialize and the horizon of standard financial models.

- 1-year horizon (Basel III credit risk, Solvency II SCR). It captures only the immediate phase. Long-run property value deterioration, migration dynamics and structural regional economic effects lie entirely outside the model. Regulatory capital calibrated to a 1-year horizon therefore systematically underestimates the solvency risk from physical climate change for long-duration books (mortgages, infrastructure loans, life insurance liabilities).
- 3-year horizon (regular supervisory stress tests, such as the EU-wide EBA exercise). It captures short-run reconstruction and credit migration effects, but neither medium-run regional restructuring nor long-run chronic damage. Climate stress tests have started to address this. The Bank of England's 2021 Climate Biennial Exploratory Scenario used a 30-year horizon, and the ECB's 2022 climate stress test combined short-term shocks with a 30-year scenario. Most regular stress tests nevertheless remain at around 3 years.
- 30-year horizon (mortgage books, infrastructure). This is the appropriate horizon for capturing the full lagged effect profile. It requires forward-looking climate scenarios with substantial uncertainty in the tails. The non-stationary EVT framework of section 2.1.2 is one way to extend physical risk distributions over this horizon.
- Generational horizon (pension funds, sovereigns, life insurers). Some physical climate effects operate over 50 to 100 years: managed coastal retreat, irreversible biodiversity loss, permanent shifts in agricultural productivity zones. NGFS long-term scenarios extend to 2100, but standard financial models do not. Integrating these effects into capital market frameworks requires intergenerational discounting and valuation models that use scenario paths up to the end of the century.

Key Limitations of Network and Supply-Chain Models

- **Data on network topology is incomplete and proprietary**. Production network data (WIOD, EXIOBASE) are available at sector level but not at firm level. Bilateral exposures between financial institutions are available to supervisors under strict confidentiality, but not to researchers or market participants, who must reconstruct them from aggregate statistics. Supply chain topology beyond Tier-1 is only partially known, even by the focal firm itself (Simchi-Levi et al., 2015). Network models therefore run on approximations of the true network, and contagion simulations rest on strong assumptions that can bias their results (Upper, 2011). The approximation matters most in the tail of the distribution, where rare but severe cascades occur.

- **Linear I-O models ignore the time path of recovery**. The Leontief and Ghosh models are static. They give the total effect of a shock under fixed input shares, but say nothing about how long recovery takes, and they ignore capacity constraints, inventories and substitution. Fixed input shares are a reasonable approximation in the short run, since output was found to fall roughly one-for-one with disrupted imported inputs (Boehm et al., 2019). Over longer horizons, supplier qualification, contractual obligations and logistics constraints make recovery slow and non-linear. The ARIO model partly addresses this (Hallegatte, 2008), but its reconstruction rate is calibrated on past events and may not represent future events under changed climate conditions.
- **Systemic risk models rely on current structures**. SRISK is estimated from historical market data (Brownlees & Engle, 2017). DebtRank applications use today's balance sheets and exposure networks, even when driven by forward-looking climate scenarios (Mandel et al., 2021). Neither captures the gradual deterioration of collateral quality, the progressive withdrawal of insurers from high-risk areas (Sastry et al., 2023), or the slow build-up of climate-related non-performing loans in the years after disasters (Noth & Schüwer, 2023). Physical climate systemic risk builds up slowly, and market-based systemic risk measures are poorly suited to detecting it early.
- **Lagged effect duration is highly uncertain**. The coefficients of lag models are estimated on historical events, and the evidence is contradictory. National incomes were found not to recover within twenty years after cyclones (Hsiang & Jina, 2014). A synthetic-control study finds no significant long-run effect of even extremely large disasters once political changes are controlled for (Cavallo et al., 2013). Future events may also differ from historical analogues in severity, adaptation capacity, policy response and the pre-existing fragility of regional economies. Multi-decade lag estimates should therefore be treated as orders of magnitude.
- **Behavioural tipping points are outside the models**. The models in this section capture mechanical cascades, including discontinuous default cascades (Gai & Kapadia, 2010) and non-linear network effects (Acemoglu et al., 2015). They take the behaviour of insurers, rating agencies and governments as given. They do not capture abrupt shifts in that behaviour, such as a coordinated exit of insurers from a coastal market, a sudden sovereign downgrade after a catastrophic event, or a political decision to abandon a coastal area. These tail events are precisely what systemic risk analysis should illuminate, and they remain beyond the current model frontier.

2.5 SCENARIO-BASED AND PROJECTIVE MODELS

A note on structure: how this section relates to the rest of Part 2

The plan for this cookbook originally posed the question of whether section 2.5 should 'dissolve into other parts'. Having now written sections 2.1 through 2.4, that question has a clear answer: yes, substantially — but not entirely.

SSP-RCP projections have already been introduced as the scenario backbone of every model in Part 2. Damage functions use them to adjust vulnerability curves for future climate. Actuarial models use them to construct non-stationary $\lambda(t)$ parameters. The NGFS macro framework in section 2.4.1 is built on them. Extreme Value Theory in section 2.1 uses them for non-stationary GEV parameter estimation. Repeating the full SSP/RCP taxonomy here would be redundant.

Cat models (OEP, AEP, PML, EP curves) are already the subject of section 2.1.1 (damage functions) and section 2.1.3 (loss and frequency curves), where the full mathematical architecture is laid out. Reproducing that here would be a wholesale duplication.

What this section therefore does is different and complementary: it provides the integrating framework that connects all the individual model families to the scenario machinery that makes them forward-looking and policy-relevant. Section 2.5.1 treats SSP-RCP projections as a modelling discipline in themselves (i.e. how scenarios are constructed, how uncertainty is structured, and how they are paired with impact models). Section 2.5.2 treats stress testing as an applied framework that assembles scenario inputs and model outputs into regulatory and management deliverables. Section 2.5.3 gives the definitive reference treatment of cat model output metrics (OEP, AEP, PML, EP, ELT, YLT) as a consolidated glossary and technical reference. Section 2.5.4 briefly introduces the specificities of some particular Machine-Learning based models. Section 2.5.5 addresses the predictive limits of all the models in Part 2 — the honest frontier of what physical climate risk models can and cannot know.

2.5.1 SSP-RCP PROJECTIONS

The architecture of climate scenarios

Every forward-looking physical climate risk model in this cookbook requires a climate scenario: a self-consistent representation of future greenhouse gas concentrations, resulting atmospheric warming, and the cascade of physical changes that follow. The standard architecture for producing these scenarios is the coupled SSP-RCP framework, which pairs socioeconomic narratives with radiative forcing pathways to produce a structured matrix of possible futures.

The framework has two components that must be understood separately before they can be used together correctly:

Shared Socioeconomic Pathways (SSPs): five narratives describing how human society might evolve over the 21st century in terms of population, economic development, urbanisation, technology, governance, and inequality. They are not climate scenarios: they describe the world that drives emissions, not the climate that results from them. SSP1 describes a world of sustainability, low population growth, and rapid technological progress. SSP3 describes fragmentation, high population growth in developing regions, and slow technology diffusion. SSP5 describes fossil-fuel-led development with high energy demand and rapid economic growth.

Representative Concentration Pathways (RCPs): four forcing trajectories defined by their 2100 radiative forcing level in watts per square metre: RCP2.6 (peaking before 2100 and declining), RCP4.5 (stabilising at 4.5 W/m^2), RCP6.0, and RCP8.5 (rising throughout the century). Each RCP produces a specific warming trajectory when run through a climate model, but the same RCP can be achieved by different SSP socioeconomic pathways.

In CMIP6 (the current generation of climate model experiments underlying IPCC AR6), the SSPs and RCPs are paired into combined ScenarioMIP experiments. The five most commonly used pairings in physical climate risk modelling are:

SCENARIO	FORCING (2100)	BEST-ESTIMATE WARMING (2081–2100)	DESCRIPTION	FINANCIAL RISK PROFILE	CLOSEST NGFS SCENARIOS
SSP1-1.9	~1.9 W/m^2	~1.4°C	Very low emissions, net zero CO_2 around 2050	Minimal additional physical risk; transition risk dominant	Net Zero 2050 (~1.5°C), Low Demand (~1.1–1.4°C)
SSP1-2.6	~2.6 W/m^2	~1.8°C	Low emissions, net zero CO_2 around 2075	Low to moderate physical risk; manageable with adaptation	Below 2°C (<2°C), Delayed Transition (~1.8°C)
SSP2-4.5	~4.5 W/m^2	~2.7°C	Intermediate emissions, roughly in line with current pledges	Moderate to high: material increase in catastrophe losses; coastal stress	NDCs (~2.5°C), Fragmented World (~2.3–2.4°C)
SSP3-7.0	~7.0 W/m^2	~3.6°C	High emissions, CO_2 emissions doubling by 2100	High: severe chronic damage; rising systemic risk	Above all NGFS scenarios; Current Policies (~2.9–3.0°C) lies between SSP2-4.5 and SSP3-7.0
SSP5-8.5	~8.5 W/m^2	~4.4°C	Very high emissions, CO_2 emissions tripling by 2075	Very high: extreme physical risk in exposed regions	No NGFS equivalent

TABLE 30. The different IPCC scenarios and their NGFS equivalent

How scenarios are translated into physical hazard projections

A climate scenario (SSP-RCP combination) is an input to a global climate model (GCM). The GCM simulates atmospheric, oceanic, and land surface dynamics to produce gridded outputs of temperature, precipitation, wind, sea level, soil moisture, and dozens of other variables at 0.5°–1° spatial resolution and annual–daily temporal resolution. These GCM outputs are then fed into impact models (hydrological models for flood, wind field models for cyclone, fire behaviour models for wildfire) to produce the hazard intensity metrics used by the damage functions and actuarial models in earlier sections.

The chain can be represented.

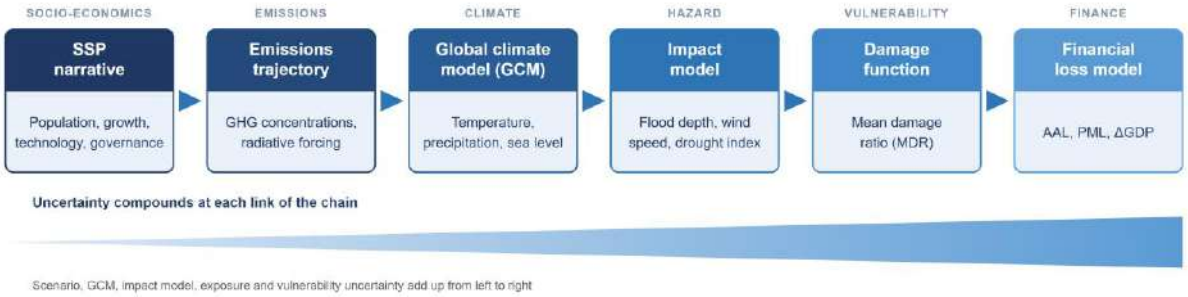


FIGURE 62. From climate scenario to financial loss: the modelling chain

Each link in this chain introduces uncertainty, and the uncertainties compound. The ISIMIP (Inter-Sectoral Impact Model Intercomparison Project) framework systematically quantifies this full chain uncertainty by running multiple GCMs through multiple impact models and comparing results.

THE UNCERTAINTY CASCADE: FROM EMISSIONS TO FINANCIAL LOSS

Total uncertainty in a financial loss estimate σ^2_{total} decomposes into five additive components:

$$\sigma^2_{total} = \sigma^2_{scenario} + \sigma^2_{GCM} + \sigma^2_{impact} + \sigma^2_{exposure} + \sigma^2_{vulnerability}$$

$\boldsymbol{\sigma^2_{scenario}}$ = scenario uncertainty: which SSP-RCP pathway the world follows. Irreducible at any given time; reducible over time as emissions trajectories are observed. Typically the dominant source of uncertainty beyond 2050

$\boldsymbol{\sigma^2_{GCM}}$ = model uncertainty: different GCMs produce different climate responses to the same forcing. Reducible in principle by improving physical understanding; the IPCC AR6 likely range for equilibrium climate sensitivity is 2.5–4°C (best estimate 3°C)

$\boldsymbol{\sigma^2_{impact}}$ = impact model uncertainty: different hydrological/fire/cyclone models produce different hazard intensities from the same GCM output. Often 30–60% of variance in flood depth or wind speed projections for a given location

$\boldsymbol{\sigma^2_{exposure}}$ = exposure uncertainty: asset registers are incomplete; future asset locations are unknown; land use changes alter exposure over the model horizon

$\boldsymbol{\sigma^2_{vulnerability}}$ = vulnerability model uncertainty: damage functions may not hold under unprecedented climate states; construction standards and materials change over time

IPCC AR6 WG1 projects a best estimate of +2.7°C for SSP2-4.5 in 2081–2100 relative to 1850–1900, with a very likely range of 2.1–3.5°C.

For regional precipitation change, uncertainty is much larger: some GCMs project a 20% increase in summer precipitation in central Europe under SSP5-8.5; others project a 15% decrease. This regional uncertainty dominates financial risk estimates for flood and drought much more than the global temperature uncertainty.

THE UNCERTAINTY CASCADE: FROM EMISSIONS TO FINANCIAL LOSS

Interpreting the scenario spread as a risk distribution

Scenario analysis is not the same as probability-weighted expected value computation. The five SSP-RCP scenarios are not assigned probabilities; they represent distinct plausible futures designed to span the range of policy and social outcomes. The correct interpretation for financial risk purposes is:

For disclosure (TCFD/IFRS S2): present results across all scenarios without probability weighting. Show the range. This allows readers to assess exposure under both optimistic and pessimistic futures.

For risk management: identify the scenarios under which the institution is most vulnerable. Design management actions that are robust across scenarios (i.e. reduce risk under SSP2-4.5 and SSP5-8.5, not just the base case).

For capital allocation: use the worst-case credible scenario (typically SSP5-8.5 or the NGFS Hot House World) for tail risk capital and provisioning, consistent with the conservatism principle in prudential regulation.

Downscaling: from global models to asset-level risk

Global climate models operate at 50–200 km resolution — far too coarse for asset-level physical risk assessment. A single GCM grid cell may contain a mountain range, a coastal zone, an urban heat island, and a river floodplain with very different physical risk characteristics. Two downscaling approaches bridge the gap:

Statistical downscaling: establishes empirical relationships between large-scale GCM variables (temperature, pressure patterns) and local observed climate variables (rainfall at a specific rain gauge, river flow at a specific gauge station). These relationships are then applied to GCM projections to produce local-scale climate projections. Advantages: computationally cheap; well-calibrated to historical data. Limitations: assumes the statistical relationships are stationary — valid for moderate climate change but potentially invalid in novel climate states where large-scale circulation patterns change qualitatively.

Dynamical downscaling (Regional Climate Models, RCMs): runs a high-resolution regional climate model (12–25 km) nested within the GCM output, simulating physical processes at finer scales. The EurO-CORDEX ensemble (12 km resolution over Europe) and NA-CORDEX (North America) are the standard frameworks. Advantages: physically consistent; captures mesoscale processes (convective precipitation, orographic enhancement) that statistical downscaling misses. Limitations: computationally intensive; only available for a subset of GCMs and regions; still too coarse for urban-scale flood modelling.

For asset-level applications, a third step is required: hyper-local modelling using asset-specific data. Flood risk models use 1–10m digital elevation models (DEMs) and hydrodynamic simulations to produce asset-level flood depth estimates from the RCM precipitation projections. Wildfire models use 30m land cover and fuel moisture data. These hyper-local models sit at the bottom of the uncertainty cascade and are the most data-intensive component of the full chain.

Provider Box – EY Climate Analytics Platform (EY CAP)

EY CAP (Climate Analytics Platform) is an internal tool used by EY's Climate Change and Sustainability consultants. It is not sold as a SaaS product, and clients receive its outputs through consulting missions. End clients include corporates across all sectors, public entities, asset managers (private equity and real estate)

and development finance institutions. Financial institutions usually prefer to build their own models, and EY supports that work directly. The methodology and scripts have been independently reviewed by academics at Imperial College.

Hazard modelling. The platform provides more than 92 climate indicators covering the main physical hazards, including floods, extreme heat, drought, wind and avalanches. Hail is not covered at this stage. The historical baseline relies on ERA5 and ERA5-Land reanalysis, which is used to bias-correct and downscale CMIP5 and CMIP6 projections. Projections cover 2030, 2035, 2040, 2050 and 2085 under SSP1-2.6, SSP2-4.5 and SSP5-8.5.

Hazard resolution is about 10 km after downscaling, 30 m for flooding, and 1 km for landslides and urban temperature across more than 135 cities. Urban heat and storm surge indicators use machine learning, and several other indicators rely on simulation. Indicators are updated through two releases a year.

Hazard and financial impact. EY CAP use damage-based models. Physical indicators are mapped to financial impacts through more than 35 hazard-related damage functions, each with asset-type specific versions, covering more than 60 asset types. These functions draw on engineering models, observed losses and expert judgement, and they are calibrated through client missions. Sector-specific functions are developed when a mission requires them, for example for motorway maintenance costs or alcohol ageing. Econometric models, such as OLS or SVAR, are used to estimate yield-price elasticities. Outputs include financial loss in percentage or absolute terms, revenue at risk and additional operating costs such as cooling. A transition risk module uses NGFS trajectories for carbon and energy prices.

Scoring and adaptation. The platform produces a gross risk score, from very low to very high, together with a gross financial impact. It does not compute a net risk score or a resilience score. Residual risk and resilience are assessed case by case during missions, for example from the client's enterprise risk management practices. An adaptation module ranks the most relevant actions for the main hazards through a cost-effectiveness optimisation.

Uncertainty and validation. Uncertainty is communicated through intervals around financial estimates and through scenario ranges. Climate models are debiased with historical reanalysis, and some indicators, such as tropical cyclones, rely directly on historical events. Ongoing validation is qualitative. For high-priority sites, EY visits the site and compares the modelled gross risk with past impacts and existing measures, such as business continuity plans or asset elevation, which yields a residual risk estimate. No quantitative performance metric is tracked.

Positioning. EY sees its main strength in the vulnerability and financial impact layer, where many data providers stop at hazard indicators. The tool supports a rapid screening of thousands of sites, followed by bespoke modelling for priority assets or sectors. Its methodology is not published in open access, and clients receive hazard, vulnerability and damage outputs together with the mission deliverables.

Assigning probabilities to scenarios? The EDHEC approach

The EDHEC Climate Institute challenges the convention that scenarios carry no probability.

Its Climate Scenarios programme attaches probabilities to climate outcomes by quantifying model uncertainty, structural uncertainty in economic growth, and policy uncertainty measured by the gap between promised policies and observed actions. In "How to Assign Probabilities to Climate Scenarios".

Rebonato, Melin and Zhang draw on 5,905 social cost of carbon estimates from 207 academic sources and combine an expert-elicitation method with a maximum-entropy method that makes minimal assumptions. They find a 35 to 40% chance that warming exceeds 3°C by the end of the century.

The Institute also extends the NGFS set with two hotter scenarios, Climate Destabilisation and Climate Breakdown, consistent with both the SSP-RCP and NGFS frameworks, since the authors note that NGFS takes Middle of the Road as its single narrative. Scientific Climate Ratings applies these probabilities in its ratings. It computes an enterprise value under each scenario, weights each value by its probability, and obtains a single expected impact that considers all scenarios at once.

A bank running a stress test keeps the scenario-dependent values, while an investor allocating capital uses the expected value. Investors need these probabilities to compute the expected discounted cash flows at the heart of asset pricing models. The approach complements the disclosure practice described above, which presents the full range without weighting.

Model Representation Card – Probabilistic climate scenarios – EDHEC (Rebonato, Melin & Zhang, 2025)

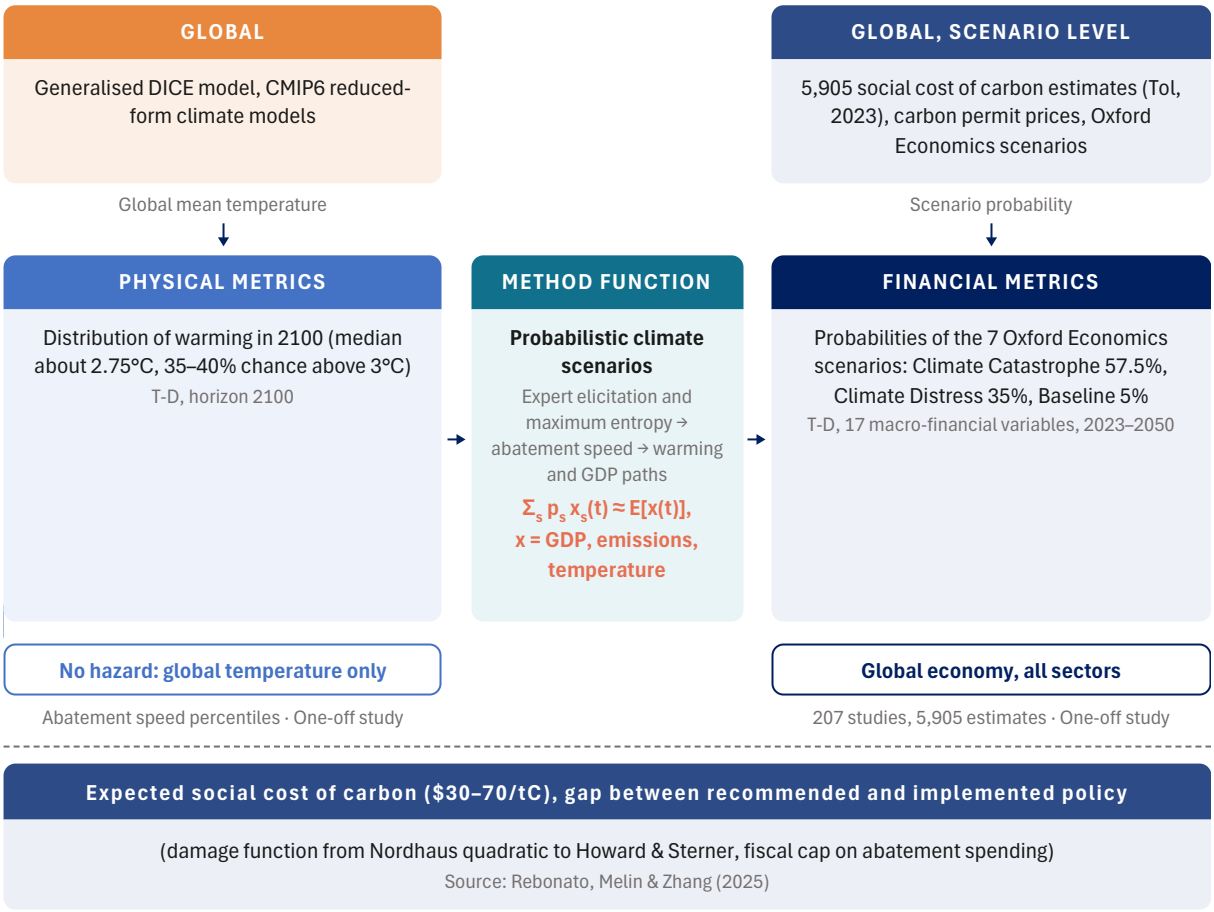


FIGURE 63. Model representation card – Probabilistic climate scenarios
Source: Rebonato et al., 2025

What does the model do?

The model attaches probabilities to climate scenarios, which are usually presented without any. It estimates how fast emissions are likely to be cut, in two ways. The first starts from economists' estimates of the cost of carbon and corrects them for the gap between what they recommend and what governments actually do. The second makes as few assumptions as possible. Both give the likely warming by 2100. The probabilities are then attached to the seven Oxford Economics climate scenarios.

What does it teach us?

Limiting warming to 1.5°C is very unlikely, and even 2°C has a very low probability. The median warming by 2100 is about 2.75°C, with a 35 to 40% chance of exceeding 3°C. Applied to the Oxford Economics scenarios, more than 90% of the probability goes to the two scenarios with the slowest emissions cuts: Climate Catastrophe (57.5%) and Climate Distress (35%). The Net Zero scenarios receive about 1% each, because they would require a carbon tax rising faster than any government has accepted so far.

Why is it useful?

Without probabilities, an investor cannot compute an expected value, and all scenarios look equally plausible. With them, each scenario's result can be weighted by its likelihood. The model also shows that the scenarios most often used as references are among the least likely, while high-warming scenarios carry most of the probability. Physical risk therefore deserves more weight than transition risk.

2.5.2 STRESS-TEST ANALYSIS

What stress testing is and what it is not

A stress test is a structured analysis of the financial impact of an adverse scenario on an institution or system, designed to identify vulnerabilities, quantify resilience, and inform capital and risk management decisions. It is not a forecast of what will happen: it is an exploration of what could happen under a defined set of adverse conditions, and what the financial consequences would be if they did.

Physical climate stress testing differs from standard macroeconomic stress tests in three critical ways that shape the model architecture:

Scenario horizon: standard stress tests cover 3–5 years. Physical climate stress tests must extend to 30 years (Bank of England CBES) or beyond to capture the full lagged effect profile documented in section 2.4.5. This creates a fundamental tension with the 1-year capital horizon of Basel III and Solvency II: the regulatory capital framework was not designed for horizon-spanning physical risks that accumulate gradually rather than arriving as sudden shocks.

Non-stationarity: standard stress tests apply adverse shocks to a stationary baseline. Physical climate stress tests must apply shocks whose distribution is itself changing over time — the 2050 baseline is already more severe than the 2025 baseline, before any adverse scenario is applied. The stress is therefore layered on top of a deteriorating trend, not a stable equilibrium.

Physical-financial interaction: the physical damage causes financial losses; but financial losses (credit tightening, insurance withdrawal, fiscal stress) feed back into the real economy and amplify the physical impact. Standard stress tests model only the financial side; physical climate stress tests must model the two-way interaction, which requires coupling the macro-financial model to the physical impact model.

EXERCISE	AUTHORITY	HORIZON	APPROACH TO PHYSICAL RISK	KEY OUTPUT	STATUS
CBES (Climate Biennial Exploratory Scenario)	Bank of England / PRA	30 years	Bottom-up: banks and insurers use their own models. Three NGFS-based scenarios: Early Action, Late Action,	Projected credit and insurance losses by scenario; explicitly not used to set capital	Run in 2021, results published May 2022

EXERCISE	AUTHORITY	HORIZON	APPROACH TO PHYSICAL RISK	KEY OUTPUT	STATUS
			No Additional Action (3.3°C by 2050)		
ECB climate risk stress test	ECB Banking Supervision	1 year (flood; drought and heat), 3 years (disorderly transition), 30 years (transition paths)	Bottom-up for 41 banks, using ECB templates	Credit and market losses of about €70bn in the short-term and physical scenarios; learning exercise with qualitative input to SREP and no direct capital impact	2022
ACPR climate exercises	ACPR / Banque de France	2020–2050 (pilot); 2050 plus a short-term 2027 horizon (second exercise)	Bottom-up. The pilot covered banks and insurers; the second exercise covered insurers only, with a sequence of droughts followed by floods	Cost of risk, claims and solvency under transition and physical scenarios	Pilot results May 2021; insurance exercise results May 2024
"Fit-for-55" one-off climate scenario analysis	EBA, EIOPA, ESMA and ECB, with ESRB scenarios	8 years (2022–2030)	Top-down; transition risk only, with no physical risk	First-round losses for 110 banks, 2,331 insurers, 629 pension funds and about 22,000 investment funds	Results published

TABLE 31. The major regulatory stress test frameworks

THE STRESS TEST MODEL ARCHITECTURE: A STEP-BY-STEP FRAMEWORK

Step 1: Scenario definition and translation

The regulatory or management scenario set is selected first. Regulatory exercises use scenarios set by the supervisor, usually adapted from NGFS scenarios. Internal stress tests can add proprietary scenarios reflecting the institution's specific geographic concentrations. The climate scenario is then translated into hazard intensity paths at each location. Most applications take these paths directly from climate and impact models. A simple alternative scales current hazard intensity with global warming:

$$I_{j,t}(s) = I_{j,0} \times (1 + \alpha_{j,peril}(s) \cdot \Delta T_t(s) + \beta_{j,peril}(s) \cdot (\Delta T_t(s))^2)$$

where $\boldsymbol{I_{j,t}(s)}$ is the hazard intensity at location j at time t under scenario s; α and β are the linear and quadratic response coefficients (calibrated from the GCM-impact model chain);

$\boldsymbol{\Delta T_t(s)}$ is global warming above pre-industrial at time t. This translates the abstract temperature path into the hazard metrics used by the damage functions.

Step 2: Exposure mapping

Every material asset in the portfolio is geocoded: mortgages by postcode, corporate loans by main facility location, sovereign bonds by country, infrastructure loans by asset location. Each asset is matched to the hazard intensity path $I_{j,t}(s)$ at its location. For assets without precise location data, sector-level averages or conservative assumptions are used.

Step 3: Physical impact → financial loss translation

$$\boldsymbol{\Delta ECL_{i,t}(s) = EAD_i \times [\Delta PD_{i,t}(s) \times LGD_i + PD_{i,t}(s) \times \Delta LGD_{i,t}(s)]}$$

This is the central formula of every physical climate stress test. ΔPD is the climate-driven PD uplift from the logit/Merton model (section 2.2). ΔLGD is the collateral value deterioration. The first term captures the increased frequency of default; the second captures the lower recovery given default on climate-damaged collateral.

The threshold is the Basel III CET1 minimum of 4.5%, raised to 7% with the capital conservation buffer, plus Pillar 2 requirements and other buffers, or an institution-specific ICAAP target. If $CET1_{stress}(s)$ falls below the threshold under a material scenario, the institution is identified as vulnerable. It must then hold additional capital or reduce its climate-exposed positions.

Step 4: Aggregation and capital impact

$$\boldsymbol{\Delta CET1(s) = -\frac{\sum_i \Delta ECL_{i,t}(s)}{RWA_{total}}}$$

$$\boldsymbol{Capital\ adequacy: CET1_{stress}(s) = CET1_{baseline} + \Delta CET1(s) \geq Threshold?}$$

The threshold is typically is typically the Basel III minimum CET1 ratio of 4.5%, or 7% including the capital conservation buffer, increased by the institution's Pillar 2 requirement and applicable buffers, or set at an institution-specific ICAAP target.

If $CET1_{stress}(s)$ falls below the threshold under any material scenario, the institution is identified as vulnerable and must either hold additional capital buffers or reduce climate-exposed positions.

Step 5: Sensitivity and scenario comparison

$$\boldsymbol{Climate\ VaR(s) = CET1_{baseline} - CET1_{stress}(s)}$$

Report Climate VaR across all scenarios in basis points of CET1. The scenario spread ($Climate\ VaR_{SSP5-8.5} - Climate\ VaR_{SSP1-2.6}$) quantifies the financial benefit of climate mitigation for the institution: the capital that would be freed by a world that limits warming to 1.5°C rather than 4.4°C. This is the metric that makes climate stress tests decision-relevant for management: it translates climate policy into the language of capital efficiency.

Bottom-up vs. top-down stress tests

Stress test designs fall on a spectrum between fully top-down, where the authority applies a common model to all institutions, and fully bottom-up, where each institution uses its own models to compute losses. Each approach has advantages and limitations.

Top-down (Fit-for-55 exercise, central bank exercises on their own balance sheets): results are comparable across institutions, faster to produce and independent of the quality of institutions' models. However, they rely on simplified risk scores and average damage functions and miss institution-specific concentrations and granular collateral data.

Bottom-up (CBES, ACPR): each institution uses its own models, data and expert judgement, which captures granular exposures. However, results are less comparable across institutions, and their quality depends on the maturity of each institution's climate risk models. Supervisory challenge and model validation are therefore required.

Hybrid (ECB 2022 climate stress test): banks project their own losses using common scenarios and templates set by the supervisor. This combines institution-level granularity with a degree of comparability.

2.5.3 CAT MODELS

Catastrophe model output metrics are the standard vocabulary of property catastrophe (re)insurance, and they are increasingly adopted in banking and sovereign risk contexts as physical climate risk becomes embedded in mainstream financial analysis. This section provides the definitive reference treatment of the key metrics, their mathematical definitions, their relationships to each other, and their practical interpretation and limitations. It consolidates and extends the treatment in sections 2.1.1 and 2.1.3.

The stochastic event catalogue: the foundation

Every cat model output metric is derived from a stochastic event catalogue: a large database of simulated events, each characterised by physical parameters (track, intensity, footprint) and an annual rate of occurrence. The catalogue is the probabilistic representation of the full range of possible hazard realizations in any given year. Understanding the metrics requires first understanding the catalogue structure:

CATALOGUE COMPONENT	DEFINITION	TYPICAL SCALE	ROLE IN OUTPUT METRICS
Event	A specific realised hazard scenario (e.g. a hurricane with defined track, intensity, and footprint)	50k–500k events per peril	Unit of the catalogue; each event has a rate and a loss
Annual rate $\boldsymbol{\lambda_e}$	Expected number of events per year identical to event e (inverse of return period). For rare events, $\lambda_e << 1$	Typically 0.001–0.10 per year	Determines the probability weight of each loss in the OEP/AEP computation
Ground-up loss $\boldsymbol{L_e}$	The physical loss generated by event e to the exposed portfolio, before application of	USD thousands to billions	Primary input to EP curves; represents total economic damage

CATALOGUE COMPONENT	DEFINITION	TYPICAL SCALE	ROLE IN OUTPUT METRICS
	insurance policy conditions		
Insured loss I_e	L_e after application of deductibles, limits, co-insurance, and treaty structures	Always ≤ L_e	Used for insurance pricing; determines insurer-specific EP curves
Event Loss Table (ELT)	The full catalogue expressed as $\{e, \lambda_e, L_e, \beta_e\}$ where β_e is the coefficient of variation of loss given event e	One row per event	The primary computational object from which all metrics are derived
Year Loss Table (YLT)	A simulation of N years, each year containing the sum of all events occurring in that year (sampled from the Poisson event set)	N = 10,000–100,000 years	The primary input for AEP and multi-year aggregate metrics

TABLE 32. The components and their role

THE FIVE "CANONICAL" OUTPUT METRICS

1. Average Annual Loss (AAL)

$$AAL = \sum_e \lambda_e \times E[L_e] = \int_0^\infty (1 - F_S(x))\, dx$$

The AAL is the rate-weighted expected loss across all events. It is the pure premium (the actuarially fair cost of risk before loadings, expenses, and profit). The second equality shows that AAL equals the area under the aggregate exceedance probability curve (i.e. a geometric identity that provides a useful cross-check).

For a USD 1bn property portfolio in a moderate-risk area: AAL typically ranges from 0.2–2.0% of TIV (USD 2–20m per year). For a high-risk coastal portfolio: 1.5–5.0% of TIV.

2. Occurrence Exceedance Probability (OEP)

$$OEP(L) = P(\exists\, e \text{ in year}: L_e > L) = 1 - P(\text{all } L_e \leq L)$$

$$OEP(L) = 1 - exp(-\sum_{e:\, L_e > L} \lambda_e)$$

THE FIVE "CANONICAL" OUTPUT METRICS

OEP answers: what is the probability that the single worst event in a year exceeds loss L? It is the per-occurrence view: relevant for per-risk reinsurance treaties, for regulatory stress tests asking about the single worst event, and for rating agency PML stress tests.

The OEP curve decreases from $1 - exp(-\Sigma_e \lambda_e)$ at L = 0, the probability that at least one event occurs in the year, to 0 at the maximum possible loss.

3. Aggregate Exceedance Probability (AEP)

$$AEP(L) = P(S_{year} > L)$$

AEP answers: what is the probability that total losses accumulated over a full year exceed L? AEP is always to the right of OEP at every probability level: for the same loss level, the aggregate is more likely to be exceeded than any single event because multiple events accumulate. The gap between AEP and OEP quantifies the multi-event contribution. The gap is largest at short return periods and for portfolios exposed to many frequent events, such as geographically diversified books. At long return periods, a single large event dominates the annual total and AEP converges towards OEP.

4. Probable Maximum Loss (PML) at return period T

$$PML_T^{OEP} = OEP^{-1}\left(\frac{1}{T}\right)$$

$$PML_T^{AEP} = AEP^{-1}\left(\frac{1}{T}\right)$$

PML is the loss level exceeded with probability 1/T per year.

At T = 200: PML_{200} is exceeded with 0.5% annual probability (the Solvency II SCR calibration point).

At T = 250 and T = 100, it is exceeded with 0.4% and 1.0% annual probability, return periods commonly used by rating agencies and in reinsurance purchasing. $PML_{AEP} \geq PML_{OEP}$ always. The difference is larger for diversified portfolios with many moderate events accumulating; smaller for concentrated single-peril books.

5. Tail Value-at-Risk (TVaR) / Conditional Tail Expectation (CTE)

$$TVaR_\alpha(S) = E[S \mid S > VaR_\alpha(S)] = VaR_\alpha(S) + E[S - VaR_\alpha \mid S > VaR_\alpha]$$

$$= VaR_\alpha(S) + \varepsilon_\alpha(S)$$

where ε_α is the mean excess function at VaR_α

TVaR is considered by many a "coherent" risk measure (unlike VaR/PML): it satisfies sub-additivity, meaning the TVaR of a diversified portfolio is always less than or equal to the sum of individual TVaRs.

It is the capital metric of the Swiss Solvency Test (99%) and of Basel market risk rules (97.5%), while Solvency II relies on a 99.5% VaR. It is better suited than VaR to heavy-tailed losses. $TVaR_{99.5}\%$ corresponds to the average loss in the worst 0.5% of years: always larger than PML_{200}, often by 30–60% for typical cat loss distributions with $\xi \approx 0.3$.

The relationships between metrics: A navigation guide

HOW THE METRICS RELATE TO EACH OTHER AND TO TREATY STRUCTURES

AAL vs. PML: AAL is the expected cost; PML is the tail. For an example cat portfolio,

$$PML_{250} \,/\, AAL \;\approx\; 15 - 30\times .$$

This ratio measures tail heaviness: higher ratios indicate more extreme event dominance. Catastrophe-exposed books have ratios of 20–40× [I]; diversified books: 8–15× [I].

OEP vs. AEP: The OEP-AEP spread measures the multi-event contribution to annual losses. The spread is largest at short return periods and for diversified books, where many moderate events from different regions accumulate in the same year. It narrows at long return periods, where a single large event dominates the annual total, and is small for single-region, single-peril books.

PML vs. TVaR: TVaR > PML always. TVaR captures the severity of losses in excess of PML. For ξ = 0.3:

$$TVaR_{99.5} \;\approx\; 1.4\times\; PML_{200}$$

For $\xi \;=\; 0.5\text{:}\; TVaR \;\approx\; 2\times\; PML.$

Capital set at $PML_{\{200\}}$ is systematically inadequate for heavy-tailed portfolios: the average loss when the 200-year PML is breached can be 40–100% above the PML itself.

Treaty attachment points: Per-occurrence catastrophe excess of loss (XL) treaties attach at the OEP curve: the treaty pays when a single event exceeds the attachment. Aggregate stop-loss treaties attach at the AEP curve: the treaty pays when annual aggregate losses exceed the attachment. The price (rate on line) of each treaty is driven by the expected loss in its layer.

Divided by the limit U − A, this gives the expected loss on line: $E[layer\ loss] = \int_A^U EP(L)\,dL.$

A is the attachment point and U the exhaustion point of the layer, so $U - A$ is the limit. $EP(L)$ is the exceedance probability (OEP for a per-occurrence layer, AEP for an aggregate layer).

Climate-adjusting cat model outputs

Standard cat models are calibrated to current climate conditions. Incorporating physical climate change requires adjusting the event catalogue for projected changes in hazard frequency and intensity under each scenario:

CLIMATE-ADJUSTED OEP: FREQUENCY AND INTENSITY UPLIFT

For each event e in the current catalogue, the climate-adjusted rate and loss are:

$$\boldsymbol{\lambda_e(t,s) \;=\; \lambda_e^{\text{current}} \;\times\; m_{freq}(peril, region, t, s)}$$

CLIMATE-ADJUSTED OEP: FREQUENCY AND INTENSITY UPLIFT

$$L_e(t,s) = L_e^{current} \times m_{int}(peril, region, t, s)$$

where $\boldsymbol{m}_{freq}$ and m_{int} are climate change multipliers calibrated from GCM-impact model chains.

Example values: for North Atlantic tropical cyclone under SSP5-8.5 by 2050: m_{freq} = 0.85 (Knutson et al., 2020) (frequency slightly decreases but severe events increase), m_{int} = 1.15–1.25 (intensity increases for Category 4–5 storms).

For European riverine flood under SSP2-4.5 by 2050: m_{freq} = 1.15–1.30, m_{int} = 1.05–1.15. (Knutson et al., 2020; Emanuel, 2005, 2011)

Climate-adjusted AAL uplift

$$AAL(t,s) = \sum_e \lambda_e(t,s) \cdot L_e(t,s)$$

The frequency and intensity multipliers interact multiplicatively on AAL. A 15% frequency increase and 15% intensity increase combine to a 32% AAL uplift (not 30%, because of the interaction). At the tail, the frequency effect is damped.

For a generalized Pareto tail with shape ξ, a frequency multiplier m raises long-return-period losses by about m^{ξ}, while the intensity multiplier scales them one-for-one. With $\xi \approx 0.3$, a 15% frequency and 15% intensity uplift raise the 250-year PML by about

$$1.15 \times 1.15^{0.3} \approx 1.20$$

Around 20%, less than the 32% AAL uplift (for example, an application of this shows that 100-year flood peaks double in frequency within three decades under RCP8.5. (Alfieri et al., 2015))

Non-stationarity and the moving baseline problem

$$OEP_{eff}(L,t,s) = 1 - \exp\left(-\sum_{e:\, L_e(t,s)>L} \lambda_e(t,s)\right)$$

The effective return period of a fixed loss level L changes over time as the climate shifts. A flood defence designed for the current 100-year event provides 100-year protection in 2025 but only 60–70-year protection in 2050 under SSP2-4.5.

Model Representation Card – Climate-adjusted flood catastrophe model – Contiguous US (Wing et al., 2022)



FIGURE 64. Model representation card – Climate-adjusted flood catastrophe model (contiguous US, Wing et al., 2022)

What does the model do?

It simulates floods from rivers, rainfall and the sea for every building in the contiguous United States. For each building, it estimates water depth for floods of different rarity, from the 5-year to the 500-year event. These depths are turned into repair costs through damage curves built from over two million insurance claims. The losses are then averaged into a single yearly expected cost. The exercise is repeated for today's climate and for the 2050 climate.

What does it teach us?

US flood losses already average about 32 billion dollars a year. Climate change alone raises this figure by about a quarter by 2050. Population growth into flood-prone areas raises the risk about four times more than climate change does. Exposure therefore matters more than hazard over this horizon.

Why is it useful?

It shows a complete catastrophe model chain in an open, peer-reviewed form: hazard, exposure, vulnerability, then average annual loss. It also shows the simplest way to climate-adjust a cat model. The event probabilities are changed and the rest of the chain is kept fixed. Results are reported by building, which suits mortgage, real estate and insurance portfolios.

2.5.4 ML-BASED MODELS

Several providers already use machine learning at the margin of their modelling chain. We mentioned some already, refinq uses a large language model only to interpret risk metrics and generate adaptation recommendations, once losses have been computed. EY CAP applies machine learning to two climate indicators, urban heat and storm surge. Scientific Climate Ratings uses none in any product. In these cases, the core of the model remains a damage function applied to a hazard projection.

The model descriptions above show that a scenario is a structured assumption about the future, and that any statement about future losses rests on relationships observed in historical data. Machine learning follows the same logic. It learns patterns from past observations and uses them to produce projections. Two uses can be distinguished.

Direct prediction. Algorithms such as gradient boosting or random forests estimate a loss, a damage ratio or a risk class directly from hazard, exposure and asset characteristics, without imposing a functional form on the relationship.

Data generation. Generative models, such as generative adversarial networks (GANs), learn the joint distribution of climate variables and simulate a large number of plausible trajectories. These synthetic series can then feed a classical damage model.

Three examples illustrate these approaches.

Wasserstein GAN for drought (Heranval et al., 2026). The authors train a conditional Wasserstein GAN on Météo-France monthly Soil Wetness Index maps for the Grand Est region (1960–2024, 8 km grid), conditioned on eleven DRIAS weather covariates. The model simulates 1,000 trajectories of soil moisture over 2025–2050 under RCP 4.5 and RCP 8.5. The authors then reproduce the CatNat eligibility rules for drought-induced subsidence and derive a distribution of insured losses. On the test period, the model identifies about 90% of eligible communes, with weaker performance in mountainous areas.

Kayrros. The Paris-based company combines satellite imagery with machine learning and data fusion to monitor physical risks, with a focus on wildfire. Its Wildfire Risk Monitor produces pre-fire risk maps, active-fire detections and post-fire damage assessments for investors, insurers and public authorities. The model architecture is not published.

ScyAI. The Zurich-based start-up launched its platform in February 2026, after a pre-seed round of about €2 million. It combines a company's operational data with external hazard models to build quantified risk profiles of real assets, presented in terms close to those used by underwriters. Its target users are manufacturers, energy producers and other owners of large asset portfolios. The methods used are not disclosed.

These approaches share three limits. A model trained on historical data may extrapolate poorly beyond the range of conditions it has observed, which is precisely the situation under strong warming. Its outputs are harder to trace back to a physical mechanism than those of a damage curve. Validation also requires long, consistent loss records, which remain scarce for most hazards. When a provider states that it "uses AI", the key question is therefore where the algorithm sits in the chain, in the generation of hazards, in the damage function, or in the interpretation of results.

Provider Box – Riskthinking.ai

RiskThinking.AI is a Toronto-based provider serving banks, insurers, asset managers and corporates. Its Climate Digital Twin supplies the physical climate risk data on the Bloomberg Terminal. Canada's banking supervisor, OSFI, mandated it for its 2024 Standardized Climate Scenario Exercise, in which more than 100

federally regulated institutions assessed their portfolios with it. It covers eleven hazard families (63 indices) and provides every output as a full probability distribution rather than a point estimate or score.

Two design choices distinguish it.

First, scenario uncertainty is carried into the model rather than chosen in advance. The model draws on 1,378 CMIP6 experiments from 44 climate models. Its "Stochastic View" weights all SSP pathways at once, and individual SSP, RCP and NGFS pathways remain available.

Second, a single bias-corrected climate ensemble drives every hazard, so rainfall, heat, wind and humidity at a site come from one consistent simulated atmosphere. Hazards can therefore be combined on a common probability basis instead of being added up from separate models. A patented "Spanning Set" method builds the joint multi-hazard distribution for each asset. For flood, the climate ensemble drives a hydrological and hydrodynamic chain (PCR-GLOBWB2, LISFLOOD-FP) directly rather than rescaling historical flood maps.

The model keeps the loss components apart. Direct damage is expressed as a damage ratio, the share of the replacement cost of a building or piece of infrastructure destroyed by the hazard. Business interruption is estimated separately, as the share of operational output lost. Both become accounting amounts once multiplied by the replacement value of the asset and by the revenue it generates. Because full distributions are retained, portfolio tail measures such as VaR and CVaR can be aggregated by convolution rather than by averaging asset-level figures.

The provider maintains a database of more than 3.5 million geolocated physical assets, linked to about 130,000 subsidiaries and more than 14,000 listed parent companies. It is updated continuously with monthly change logs and is built by machine-learning extraction verified by people. Damage computed at the site can therefore be rolled up to the subsidiary and then to the listed parent.

The credit module, currently in controlled release, models physical damage as discrete shocks to the borrower's asset value, arriving according to a Poisson process.

$$\frac{dV_t}{V_{t-}} = \mu\, dt + \sigma\, dW_t - J\, dN_t$$

- where $\boldsymbol{V_t}$ is the asset value of the borrower,
- $\boldsymbol{\mu}$ its expected growth rate,
- $\boldsymbol{\sigma}$ its volatility in normal times,
- $\boldsymbol{W_t}$ a standard Brownian motion,
- $\boldsymbol{N_t}$ a Poisson process counting climate events and J the fraction of asset value destroyed by each event.

The probability of default then enters the IFRS 9 expected credit loss, $ECL = PD \times LGD \times EAD$.

Validation is carried out at the hazard level. Each newly observed record-breaking event is checked against the modelled distribution for that location. Overall, 97% fell within it; in 2025, 96% of heat records and 97% of cold records did, against 77% and 85% for the best single deterministic SSP scenario. Misses feed into the annual re-calibration. Damage functions rest on published engineering sources (JRC, HAZUS) and are not calibrated to insured losses. Clients can substitute their own curves, and raw hazard intensities and intermediate outputs are supplied so results can be checked independently.

2.5.5 PREDICTIVE LIMITS

The honest frontier: what physical climate risk models cannot do

This section closes Part 2 of the cookbook. It is the most important section to read last, because it provides the epistemic frame within which all the preceding technical content must be interpreted. Every model described in sections 2.1 through 2.5.3 has predictive limits, which are boundaries beyond which its outputs are no longer reliable representations of reality. Knowing where those limits lie is not a confession of inadequacy: it is a precondition for using these models responsibly.

The predictive limits of physical climate risk models operate at four levels: structural limits (what the model architecture cannot represent), empirical limits (what the data cannot support), horizon limits (how far forward projections remain valid), and epistemic limits (what is genuinely unknown and cannot be reduced by better models or more data). Each is addressed in turn.

Structural limits: what the model architecture cannot represent

Models are linear where reality is non-linear. Most of the models in Part 2 are linear or log-linear: I-O models assume proportional propagation; damage functions assume smooth MDR-intensity relationships; panel regressions assume constant marginal effects of temperature. Physical reality is full of thresholds, tipping points, and phase transitions: a levee either holds or fails; a crop either survives or fails; an insurance market either remains open or withdraws. Smooth models systematically underestimate the probability and severity of threshold crossings.

Models treat agents as passive. Economic models assume that firms, households, and governments respond to physical risk in predictable, historically-calibrated ways. In reality, agents adapt: they move, rebuild better, change crops, insure differently, lobby for regulation. Adaptation changes the damage function, the vulnerability curve, and the insurance penetration rate. Models that do not endogenise adaptation overstate long-run physical risk. Models that assume full optimal adaptation understate it. Neither extreme is correct, and the correct intermediate is unknown.

Models do not represent genuine novelty. All calibrated models are trained on historical data. They can extrapolate beyond the historical range (as EVT does for return levels) but they cannot represent qualitatively novel outcomes: the first time a major financial center floods, the first time a G20 sovereign default because of climate damages, the first time a major reinsurer fails due to climate losses. These events have no historical analogues on which to calibrate and must be treated as scenario narratives rather than probability estimates.

Models handle one system at a time. Most physical climate risk models operate on either the physical system (climate → hazard) or the economic system (hazard → financial loss) but not the interaction between them. The physical system affects the economic system (damage → loss); but the economic system also affects the physical system (capital available for adaptation → reduced future damage). This feedback loop is absent from virtually all operational physical climate risk models, and its absence means that the long-run impacts under severe scenarios are systematically underestimated.

Empirical limits: what the data cannot support

Historical records are short relative to the return periods of interest. A 100-year historical record contains at most 100 realizations of annual maximum losses. Estimating a 500-year return level from 100 observations requires extrapolation by a factor of 5 into the unobserved tail. The uncertainty in that estimate (as quantified in section 2.1.2 on EVT) typically spans a factor of 2–3 around the point estimate. All PML figures beyond the 100-year return period should be understood as carrying this irreducible sampling uncertainty.

Geographic coverage is uneven. Physical hazard databases, insurance claims records, and economic loss data are rich in North America and Western Europe, moderate in East Asia and Australia, and sparse to absent in Sub-Saharan Africa, South Asia, and most of the Pacific. Damage functions calibrated on US hurricane data are applied to cyclones in Bangladesh or Madagascar with minimal empirical support. Vulnerability functions for informal settlements, traditional construction, and non-engineered buildings are largely uncalibrated against claims data. The models are most precise where the risk is already well-understood and best-managed; they are weakest precisely where the human need is greatest.

Financial impact data is even sparser than physical data. The chain from physical damage to financial consequence — the PD uplift from a flood, the LGD change from property damage, the EV impact of chronic heat stress — requires matching financial records (loan performance, equity returns, credit spreads) to precisely geo-located physical risk observations. Very few institutions have performed this matching systematically, and almost none have done so across the full range of perils and geographies. The β_2 coefficients in the PD logit models are estimated from small, potentially unrepresentative samples; the standard errors are large; and the causal identification is fragile. The financial impact chain is the empirically weakest link in the entire physical climate risk modelling framework.

Horizon limits: how far forward projections remain valid

1–5 year horizon: high confidence, well-calibrated. Cat model outputs for current-climate OEP/AEP curves, AAL, and TVaR are well-calibrated against historical data for return periods up to approximately 50–100 years. Stress tests at this horizon using current-climate models are reliable within the EVT uncertainty bounds described in section 2.1.2. Physical climate change is a minor correction at this horizon for most perils.

5–20 year horizon: moderate confidence, scenario-dependent. Climate change begins to materially alter hazard frequencies and intensities. The NGFS scenarios diverge from each other as emissions trajectories separate. Model uncertainty (GCM spread) is the dominant source of variance, and it cannot be reduced by better calibration of the impact models: it reflects genuine uncertainty in the climate system's response to forcing. Financial impact models at this horizon are more reliable for chronic effects (heat stress, water stress) than for extreme event frequency changes.

20–50 year horizon: low confidence, high scenario dependence. At this horizon, the scenario spread dominates all other sources of uncertainty. The choice of SSP pathway produces factor-of-two differences in projected physical damages for most perils and regions. Non-linear tipping points (ice sheet instability, permafrost carbon feedback) become potentially relevant but are not captured in standard CMIP6 models. Asset-level projections at this horizon should be treated as scenarios (possible futures) rather than forecasts (probable futures).

50–100 year horizon: exploratory, not predictive. At century-scale horizons, the compounding of scenario uncertainty, model uncertainty, and structural uncertainty (will the economic system, institutional framework, and technology environment be recognisable?) renders quantitative loss projections essentially meaningless as predictions. They retain value as exploratory scenario analysis: exploring the range of possible futures and identifying which strategic options remain viable across that range. But they should never be presented as forecasts of likely financial outcomes, and capital or provisioning decisions should not be made on their basis without explicit acknowledgement of the exploratory nature of the analysis.

Epistemic limits: what is genuinely unknown

Beyond the structural, empirical, and horizon limits lies a category of uncertainty that is irreducible: genuine epistemic uncertainty about the behaviour of complex systems in unprecedented states. These are the 'known unknowns' and 'unknown unknowns' of physical climate risk modelling, and they deserve explicit treatment because they represent the limits of model-based knowledge.

Tipping point proximity: whether the climate system is close to tipping points (Amazon dieback, Atlantic Meridional Overturning Circulation collapse, West Antarctic Ice Sheet instability) that would cause non-linear acceleration of physical risk is genuinely unknown. The best scientific assessment is that some tipping points could be reached at 1.5–2.0°C of warming; others require 4°C or more. The probability distributions over tipping point thresholds are wide and not well-constrained by available evidence. Any financial model that assumes the absence of tipping points is making a non-trivial assumption about the tail of the physical risk distribution.

Adaptation effectiveness at scale: what happens to physical risk when adaptation investment scales from the current level (roughly 5–10% of what is needed, per UNEP adaptation gap reports) to the level required to maintain tolerable risk under 2°C+ warming is genuinely unknown. Physical protection measures (sea walls, building codes, early warning systems) reduce risk; but their effectiveness at scale, their interactions with each other, and their political sustainability are uncertain. Models that assume either no adaptation or full optimal adaptation are both wrong; the correct assumption lies in between and is not well-characterised.

Distributional effects and social tipping points: physical climate risk is profoundly unequal in its distribution. Low-income households in tropical countries face orders-of-magnitude higher exposure than high-income households in temperate countries. The financial models in this cookbook largely operate on the assumption that economic agents are homogeneous within their modelled categories. The social and political consequences of extreme inequality in physical climate risk — mass migration, political instability, conflict over water and agricultural resources — are not captured in any of the financial models described here. These consequences could generate financial instability that dwarfs the direct physical damage from the climate events themselves.

Model-induced complacency: perhaps the most insidious predictive limit is the risk that having a model produces false confidence. For example consider a PML_{250} estimate with a 95% confidence interval of [USD 300m, USD 900m] is technically correct and intellectually honest. But the act of producing a number, even a range, induces a false sense of control over something that may be fundamentally uncontrollable. The appropriate epistemic posture for physical climate risk is not ‘we have modelled this, therefore we understand it’ but rather ‘we have modelled the parts we can quantify; the parts we cannot quantify may be larger than the parts we can’

FOUR PRINCIPLES FOR USING PHYSICAL CLIMATE RISK MODELS WITHIN THEIR LIMITS

Principle 1: Range over point. Never present a single-point physical climate risk estimate without an uncertainty range.

The uncertainty range is essential information about the quality of the knowledge. A decision-maker who understands that the 250-year PML estimate ranges from USD 300m to USD 900m is better positioned than one who believes it is precisely USD 600m. For regulatory disclosures under IFRS S2 and TCFD, presenting scenario ranges is required.

Principle 2: Stress the model, along with the scenarios.

A scenario stress test that uses a single model and a single set of model parameters does not reveal model uncertainty. Every physical climate risk analysis should include a model stress: what does the output look like if the tail index ξ is 0.1 higher? If the damage elasticity α is doubled? If the recovery rate ρ is halved? Model stress reveals the assumptions that dominate the output and allows decision-makers to focus validation effort where it matters most.

Principle 3: Distinguish what you know from what you have assumed.

A well-designed physical climate risk model makes its assumptions explicit and distinguishes empirically calibrated parameters from expert judgments and default values. The AAL is calibrated from data. The 2050 climate multiplier is assumed from a model chain. The LtV stigma discount is assumed from a single academic study. These have very different epistemic statuses and should be communicated differently. Treating all model parameters as equally reliable is a common and consequential error.

Principle 4: Use models to narrow the action space, not to determine it.

Physical climate risk models should be used to eliminate clearly inferior decisions (investing in an unprotected infrastructure in a 1-in-10-year flood zone is a wrong move regardless of model precision) and to identify the option space worth exploring (adaptation investments that reduce risk under all plausible scenarios are more valuable than those that reduce risk only in the most likely scenario). The models should narrow the range of defensible choices, not substitute for judgment about which choice to make.

A final word about the models

As seen in Parts 1 and 2, models are a simplification of the world, they do not reflect the "truth" nor even an exact fact about the reality, they are an approximation, and in the best case, a signal using what already happened showing us something likely to happen in the future.

"All models are wrong but some are useful" as George E.P. Box once said.

Part 3 of this cookbook will address data: the databases, the data quality challenges, the emerging alternative data sources, and the governance frameworks for managing physical climate risk data across a financial institution. The models described in Part 2 are only as good as the data fed into them. Part 3 will make that constraint explicit and provide the practitioner tools to navigate it.

3 THE INGREDIENTS – DESIGN CHOICES AND DATA

The previous section catalogued the main families of physical climate risk financial impact models. Knowing which families exist is necessary for model selection. The second question is harder. Given a specific institution, portfolio and objective, which data sources and modelling choices will produce the most defensible result? Statistical tests of model discrimination compare models already fitted to a common dataset (Atkinson, 1970). Borowiak (1989) extended this approach with likelihood-ratio tests for nested and non-nested models.

Such tests separate the quality of a model's fit from the appropriateness of the model as a representation of the underlying physical and financial process. In physical climate risk, the decisive choices are made earlier, and a model can fit historical losses well while being structurally mis specified for the forward-looking problem (see section 3.4.4).

Part 3 organizes the design choices a practitioner must make before running any model.

These choices determine what data enters the model, at what resolution, over what time horizon and through what financial lens. Poor upstream choices cannot be corrected by sophisticated downstream estimation. We decompose each model along four axes.

The first is hazard representation (section 3.1), which covers the hazards, the physical resolution and the timeline chosen.

The second is exposure and vulnerability inputs (section 3.2), which covers tangible assets, socio-economic variables and the adaptation and resilience factors considered.

The third is financial proxy selection (section 3.3), which covers accounting, market and actuarial variables. The fourth is aggregation architecture (section 3.4).

3.1 THE HAZARD CHOICES AND THEIR REPRESENTATION

The hazard is the physical event that initiates the loss chain. Before any vulnerability function or financial model is applied, the practitioner must decide which physical phenomenon is modelled, how its intensity is measured, and at what spatial and temporal resolution.

These choices become visible in the outputs, in the risk maps. They constrain every downstream step and determine whether the resulting model can be validated against observed losses, projected into future climate scenarios and replicated across institutions.

Some experts in the field described a clear tension in this design space with characteristic directness: "RMS and AIR Worldwide will never show you their high-resolution maps." The institutions with the most refined hazard representations are also those with the greatest commercial incentive to keep them proprietary. The practitioner must therefore make explicit choices about which publicly available or licensable hazard data to use. This is a trade-off between transparency and expected performance. A financial institution that pays little attention to hazard choices may implicitly accept a lower spatial resolution and a lower accuracy in its assessment of physical climate risk.

This choice is difficult in practice. Climateurope2, a European research project that runs until 2027, aims at "supporting and standardizing climate services in Europe and beyond". The project interviewed around thirty public and private climate service providers. According to their accounts, the first difficulty a practitioner meets is the abundance of sources. Some private data providers were abandoned because clients preferred national public providers. Other organizations declined to buy external expertise because they judged its added value too low. Guidance on source selection often relies on personal networks.

The quality of documentation is also uneven across providers. The level of detail on methods, temporal coverage, metadata and processing steps varies from one supplier to another, and the same data product can change from one version to the next. Two needs follow from this situation. Users lack clear guidance to navigate the available datasets, and access to data remains difficult, with formats that seldom work well together. The hazard choices described in this section should be read against this background. A model inherits the version, the format and the documentation gaps of the hazard dataset it relies on.

3.1.1 SINGLE VS MULTIPLE HAZARD

The first and most consequential hazard design choice is whether to model one physical process at a time or to represent the simultaneous or sequential occurrence of multiple hazards at the same location. This choice has direct implications for data requirements, model tractability, and the adequacy of the loss estimate.

The single-hazard assumption and its uses

The dominant architecture in both academic research and commercial practice is the single-hazard model (i.e. consider only one peril at a time).

Single-hazard in damage-based models can be calibrated against historical claims data for a defined peril. The outputs are comparable across studies and can be validated against observed events.

For many applications, and especially for portfolios with concentrated peril exposure such as a portfolio of coastal properties subject to Atlantic hurricane risk, a single-hazard model captures the dominant driver of loss adequately (Montgomery & Kunreuther, 2018).

Single-hazard data takes two principal forms, each with distinct epistemological properties that determine what type of model it can support.

Physics-based hazard data

Physics-based data is generated by numerical simulation of the atmosphere, ocean, or land surface. It represents the physical process directly, using the governing equations of fluid dynamics, thermodynamics, or hydrology.

The example of ERA5, Chelsea, Copernicus, and others

Among the most common examples of physics-based models, you will find ERA5 reanalysis wind fields (ECMWF), CHELSA precipitation grids, Copernicus flood return period maps, and USGS ShakeMap ground motion products. The key characteristics are:

- Resolution is bounded by computational cost. ERA5 operates at 31km global resolution. Hydrodynamic models for river flood can reach 10m but only at the catchment scale, not globally. This creates a fundamental resolution tension: the models with the best physical realism cover the smallest areas.
- Uncertainty is physically interpretable. Because the model is grounded in physical law, its uncertainty structure is transparent: uncertainty in boundary conditions (sea surface temperature, soil moisture initial

state), parametric uncertainty in sub-grid processes, and structural uncertainty from numerical scheme choices. This interpretability is a major advantage when communicating model confidence to risk managers and regulators. To measure this uncertainty in flood models for example, you can use QUALYPSO, an open-source algorithm in R (see section 3.4.2 for more information and a detailed discussion about uncertainty).

- Historical record length is limited. Most reanalysis products extend back 40 to 70 years. For rare tail events with return periods above 100 years, this is structurally insufficient. This is why physics-based data is typically used as calibration input for stochastic event sets rather than directly as the loss driver.
- Climate change signal is directly embeddable. Physics-based data from climate model projections (CMIP6, CORDEX) carries a physically consistent climate change signal that can be passed directly to vulnerability functions. This forward-looking property is absent in purely actuarial approaches.
- High resolution leaves physical processes out. Even the 30 m × 30 m standard of the Bristol flood model (Wing et al., 2024) may miss processes that drive local damage. A wildfire, for example, reduces soil permeability, so heavy rain falling on a burned slope can trigger a flash flood that a flood-only model does not anticipate. Local topography matters as well, since the shape of a valley can channel runoff towards a village downstream. A fine grid improves the representation of the hazard, but it captures these interactions only if the model includes the underlying processes.

Provider Box – More hazards more problems: the Munich Re approach to Climate Physical Risk

Munich Re's approach is a good illustration of the advance reinsurance companies have when it comes to considering climate physical risks. As mentioned in section 2.1.3, Munich Re covers a lot of hazards on its Location Risk Intelligence platform. However, the platform directly provides financial metrics as outputs for only four hazards.

PERIL	PRESENT-DAY HAZARD	RETURN PERIODS	RESOLUTION	CLIMATE PROJECTION	SCENARIOS AND HORIZONS
River flood	JBA global flood maps on bare-earth terrain, with flood defences included through an adjusted standard of protection	6, from 25 to 1,000 years	30 m	Changes in peak run-off at basin scale from CMIP5 and land surface models, used to scale current flood depths	RCP 4.5 and 8.5; 2030, 2050, 2100
Tropical cyclone	NATHAN zoning; wind fields of all historical storms simulated on a grid, then a frequency analysis per cell	More than 20, from 2 to 10,000 years	0.1° grid	Changes in frequency and intensity from the GFDL HiFLOR high-resolution model	RCP 4.5 and 8.5; 2030, 2050, 2100
Storm surge	Global surge model with ocean connectivity through creeks, rivers and estuaries, attenuated by land cover,	10, up to 1,000 years	About 30 m	IPCC AR6 sea level rise combined with projected changes in tropical cyclone risk	SSP1-2.6 to SSP5-8.5; 2030, 2040, 2050, 2100

	and validated against observed events				
Extratropical storm	Reanalysis downscaled and calibrated with gust data from six national weather services; tail fitted with a generalized Pareto distribution with an upper bound	More than 25, from 0.5 to 100,000 years	0.01° (about 1 km)	Current climate only	None

TABLE 33. Four examples of perils

In total, Munich Re covers 28 hazards. Fifteen Natural Hazard Scores describe the present-day hazard at a location, and thirteen Climate Hazard Scores project the change in hazard under climate scenarios up to 2100. The resolution differs from one hazard to another (see section 3.1.3).

HAZARD	SCORE FAMILY	RESOLUTION	SCENARIOS	HORIZONS	BASIS
River flood	Natural Hazard	30 m	Current climate	Current	JBA flood extents for return periods of 50, 100 and 500 years, defended and undefended
Flash flood	Natural Hazard	250 m	Current climate	Current	Meteorological data, soil sealing, slope and flow accumulation
Storm surge	Natural Hazard	30 m	Current climate	Current	Inundation areas for return periods of 100, 500 and 1,000 years, defended and undefended
Tropical cyclone	Natural Hazard	5 km	Current climate	Current	Probable maximum wind for a 100-year return period, from basin-specific models
Extratropical storm	Natural Hazard	1 km	Current climate	Current	Probable maximum wind for a 100-year return period, between 30° and 70° latitude
Hail	Natural Hazard	Polygons	Current climate	Current	Meteorological data, elevation and lightning distribution
Tornado	Natural Hazard	Polygons	Current climate	Current	Annual frequency per 10,000 km^2, interpolated from meteorological data
Lightning	Natural Hazard	Polygons	Current climate	Current	Strikes per km^2 and year from satellites and ground networks
Earthquake	Natural Hazard	Polygons	Current climate	Current	Probable maximum intensity on the Modified Mercalli scale for a 475-year return period
Volcano	Natural Hazard	Polygons	Current climate	Current	Volcanic Explosivity Index and annual return periods

Tsunami	Natural Hazard	90 m	Current climate	Current	Inundation areas for return periods of 100, 500 and 1,000 years
Wildfire	Natural Hazard	1 km	Current climate	Current	Climatological and land cover data
Subsidence	Natural Hazard	500 m	Current climate	Current	Clay shrink-swell, groundwater depletion, sinkholes and mining
Avalanche	Natural Hazard	250 m	Current climate	Current	Starting zones and flow paths from elevation and land use
Landslide	Natural Hazard	1 km	Current climate	Current	Rainfall-triggered and earthquake-triggered hazard, from the World Bank
River flood	Climate Hazard	30 m	RCP 4.5, 8.5	Current, 2030, 2050, 2100	Current flood model scaled with an aggregate of climate and hydrological models (Munich Re and JBA)
Storm surge	Climate Hazard	30 m	SSP1-2.6, 2-4.5, 3-7.0, 5-8.5	Current, 2030, 2040, 2050, 2100	Global surge model with sea level rise and projected cyclone changes
Sea level rise	Climate Hazard	30 m	SSP1-2.6, 2-4.5, 3-7.0, 5-8.5	2030, 2040, 2050, 2100	Permanent inundation and surge intensification from IPCC AR6 projections
Tropical cyclone	Climate Hazard	5 km	RCP 4.5, 8.5	Current, 2030, 2050, 2100	Proprietary cyclone model combined with a high-resolution climate model
Subsidence	Climate Hazard	500 m	SSP1-2.6, 2-4.5, 3-7.0, 5-8.5	Current, 2030, 2040, 2050, 2100	Clay shrink-swell, groundwater depletion, sinkholes and mining
Heat stress	Climate Hazard	11 km	SSP1-2.6, 2-4.5, 3-7.0, 5-8.5	Current, 2030, 2040, 2050, 2100	Rising temperatures, extreme heat and heat waves
Heat-humidity stress	Climate Hazard	11 km	SSP1-2.6, 2-4.5, 3-7.0, 5-8.5	Current, 2030, 2040, 2050, 2100	Wet-bulb temperature
Cold stress	Climate Hazard	11 km	SSP1-2.6, 2-4.5, 3-7.0, 5-8.5	Current, 2030, 2040, 2050, 2100	Several temperature parameters
Fire weather stress	Climate Hazard	11 km	SSP1-2.6, 2-4.5, 3-7.0, 5-8.5	Current, 2030, 2040, 2050, 2100	Fire Weather Index (ignition, spread, fuel)
Precipitation stress	Climate Hazard	11 km	SSP1-2.6, 2-4.5, 3-7.0, 5-8.5	Current, 2030, 2040, 2050, 2100	Duration, intensity and frequency of precipitation

Drought stress	Climate Hazard	11 km	SSP2-4.5, 5-8.5	Current, 2030, 2040, 2050, 2100	SPEI index and dry spells
Water scarcity	Climate Hazard	Polygons	SSP1-2.6, 3-7.0, 5-8.5	Current, 2030, 2040, 2050, 2100	Freshwater supply deficit, with the World Resources Institute
Permafrost	Climate Hazard	1 km	RCP 2.6, 4.5, 8.5	Current, 2050, 2100	Northern hemisphere extent from the ESA Climate Change Initiative

TABLE 34. The 28 MunichRe Hazards

For the four hazards with financial outputs, Munich Re applies vulnerability functions to each type of asset. Combining the vulnerability with the respective hazard information results in an exceedance probability curve specific to the location, and CEL is its expected value, the area under that curve.

The provider also treats these curves as non-stationary, since climate change shifts the hazard side while exposure growth increases the amount of assets at risk. These effects need to be separated before any trend is drawn. Recent loss experience shows why. In 2025, insured losses passed 100 billion dollars, driven primarily by wildfires and severe convective storms in the US (Munich Re, 2026). Neither peril is among the four hazards for which the platform provides financial metrics, which brings the box back to its title. The more hazards a provider covers, the harder it becomes to turn each of them into a financial figure.

Downscaling: or how to turn a low resolution grid into a high resolution footprint

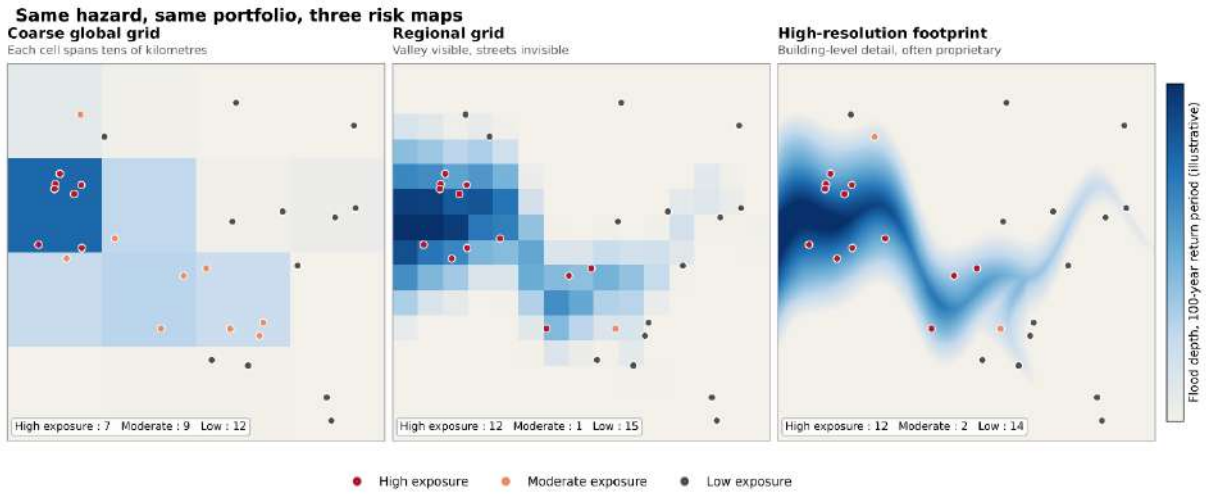


FIGURE 65. A three-step example on how downscaling operates
Source: Author

We mentioned downscaling in Part 2. Downscaling is the step that moves a hazard layer from the coarse grid of a global climate model to the resolution at which assets are located. It follows two main approaches.

Dynamical downscaling runs a regional climate model inside the global one, as in the EURO-CORDEX simulations at about 12 km over Europe. Statistical downscaling links large-scale variables to local observations through empirical relationships such as quantile mapping, and it usually includes a bias correction.

In its simplest form, quantile mapping replaces each simulated value by the observed value that has the same rank in the historical distribution:

$$\hat{x}_p = F_{o,h}^{-1}\left(F_{m,h}(x_p)\right)$$

- where $\boldsymbol{x_p}$ is the value simulated by the climate model for a future date, for instance a daily rainfall total, $\hat{x}_p$is the corrected local value,
- $\boldsymbol{F_{m,h}}$ is the distribution function of the model over the historical period, which gives the rank of x_p among simulated values, and $F_{o,h}^{-1}$ is the inverse distribution function of local observations over the same period, which returns the observed value at that rank.

A day at the 95th percentile of the model's rainfall is actually a day at the 95th percentile of the observed local rainfall.

When observations are available on a finer grid than the model, the same operation corrects the bias and will add local detail. However, please note that it can distort the change in extremes projected by the model, which is why variants such as quantile delta mapping preserve the projected change in each quantile (Cannon et al., 2015).

Both approaches sharpen the picture. But they do not create information absent from its inputs.

The building-level detail of the third panel depends on a fine elevation model, local hydraulic data and knowledge of flood defences. For this reason, a downscaled projection and a high-resolution footprint answer different questions, and a financial institution should know which of the two its provider delivers.

Provider Box – Mitiga Solution – EarthScan and Bayesian inference

Mitiga Solutions, a Barcelona company spun out of the Barcelona Supercomputing Center (BSC-CNS), Spain's national supercomputing center, offers physical climate risk assessment at global scale. Its EarthScan platform starts from an ensemble of CMIP6 and CORDEX projections. Models are selected on several criteria, including their transient climate sensitivity, their ability to capture local extremes and their skill in reproducing key modes of variability, such as ENSO and its teleconnections. Model output is bias-corrected with a trend-preserving method and validated against more than 40 years of Earth observation from over 100,000 weather stations across 180 countries, together with reanalysis, satellite and LiDAR data. It is then brought to asset scale (~1 km) using statistical-analogue, deep-learning and dynamical downscaling.

Extremes are then estimated with Mitiga's proprietary Multiple Futures Model (MFM), which combines extreme value theory with Bayesian inference in a single statistical framework. Instead of assuming a stable climate, it captures how extremes evolve over time, including changes in both their severity and their frequency, which do not always move in the same direction.

It estimates full probability distributions rather than single best-fit values. As a result, uncertainty from the physics (natural variability, differences between climate models and the emissions pathway) and from the maths and statistics (estimating trends and tails from a limited number of extreme events) is carried through to the results, which describe multiple plausible futures rather than a single projection.

EarthScan covers 13 hazards: heat stress, heavy precipitation, extreme precipitation, flooding coastal, riverine and pluvial, hail, drought, cold waves, wildfire, extreme wind, and changing climate patterns in wind, temperature and precipitation, together with a combined physical risk view. Spatial resolution depends on the hazard, reaching 90 m for terrain-dependent hazards such as flooding. Results span 1970 to 2100 in five-year steps under three emission scenarios (business-as-usual, 2040 emissions peak and Paris-aligned). For each scenario and time horizon, the MFM yields return periods and return levels (for example, the severity of the 1-

in-10, 1-in-100 or 1-in-500-year event) and annual exceedance probabilities, each reported as a range (5th, 50th and 95th percentiles). These are summarised in A-to-F EarthScan Ratings, which make exposure comparable across assets and portfolios. Combined with asset vulnerability, they also feed financial metrics, including Climate Value-at-Risk, average annual loss, return-period losses, and capex and opex impacts.

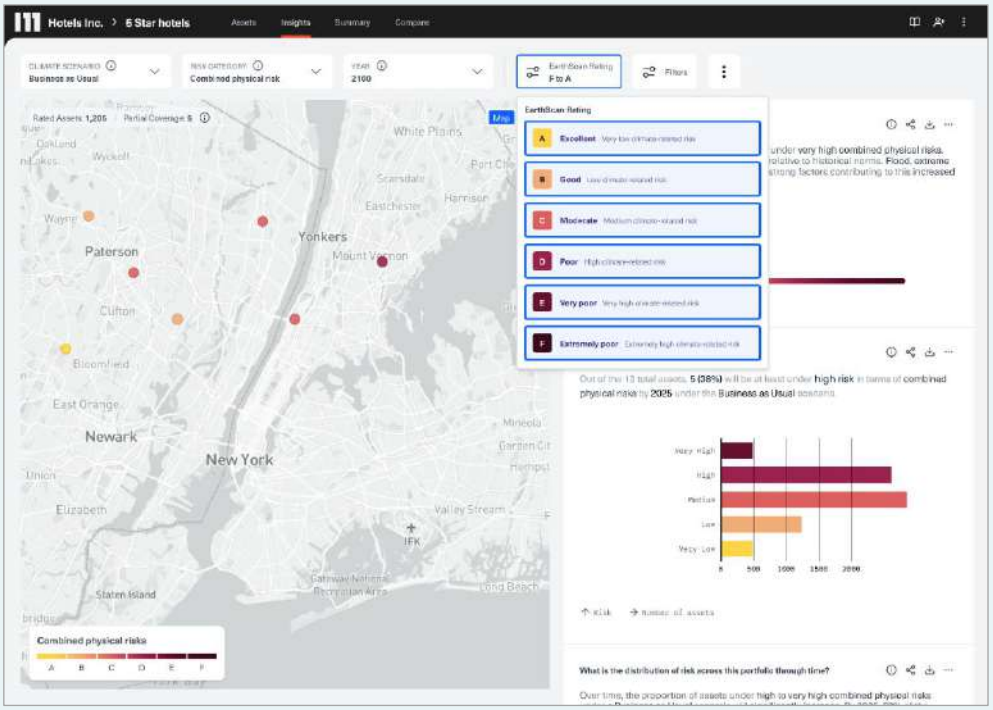


FIGURE 66. EarthScan platform interface

Source: EarthScan

The interface is showing the combined physical risk assessment of a hotel portfolio ("Hotels Inc., 5 Star hotels") in the New York metropolitan area under the Business as Usual scenario for the year 2100. The map places each asset according to its EarthScan rating, a six-level scale from A (Excellent, very low climate-related risk) to F (Extremely poor, extremely high climate-related risk). The right panel shows how assets are distributed across five risk levels (Very Low to Very High) and gives a written summary of portfolio exposure. Users can change the climate scenario, risk category and time horizon with the selectors at the top of the interface.

Provider Box – The case for Fathom Flood models

Fathom is a spin-off of the University of Bristol. Its hydrology group has produced global flood hazard models for over a decade, and the methodology behind the current Global Flood Map is published in Wing et al. (2024). The paper is the reference used by the provider in its answers to our survey. The model runs at a 30 m × 30 m resolution worldwide, and at 10 m wherever high-resolution terrain data exists. Earlier versions ran at about 900 m and then 90 m. This progression illustrates the gap between a downscaled climate grid and a true hazard footprint described above.

Fathom treats flooding as three distinct physical processes with three dedicated models. Fluvial (river) flooding arises when a channel overflows. Pluvial flooding arises when intense rainfall exceeds the absorption

and drainage capacity of the ground. Coastal flooding arises from the combination of tide, storm surge and waves. The three share one hydraulic core, and each has its own boundary conditions.

The modelling chain can be separated down into four steps:

Terrain. The model uses FABDEM, a version of the Copernicus elevation model with trees and buildings removed by machine learning, so the water flows over the bare ground. Where national LiDAR data is available, it complements FABDEM.

River network and bathymetry. River channels follow the MERIT-Hydro hydrography. Channel depth (bathymetry) is rarely observed at global scale, so Fathom estimates channel capacity with a gradually varied flow solver.

Boundary conditions. Gauging station records train a machine-learning regionalisation that produces flood frequency curves for ungauged rivers. In cross-validation against gauges, the mean absolute error of these curves fell from 126% to 57%. Coastal boundaries come from a regional frequency analysis of tide, surge and waves. Pluvial inputs come from a statistical rainfall model based on ERA5 and sub-daily rain gauges, which yields Intensity-Duration-Frequency curves globally.

Hydraulic engine. LISFLOOD-FP solves a simplified form of the shallow water equations to propagate water across the terrain. Flood defences are applied afterwards as a post-process, and so are climate change factors, which allow the maps to represent any 21st-century scenario (SSP, RCP or NGFS).

The physical proxy delivered to the user is flood depth per return period. Fathom then converts depth into ground-up loss through engineering vulnerability functions. In the Global Flood Cat model, these losses can be refined into gross (insured) and net (of reinsurance) losses using the open-source Oasis Loss Modelling Framework.

The model is validated against observed flood extents with the Critical Success Index:

$$CSI = \frac{H}{H + M + F}$$

- Where **H** is the number of cells flooded both in the model and in the observation (hits),
- **M** is the number of cells flooded in the observation and missed by the model (misses)
- **F** is the number of cells flooded in the model and dry in the observation (false alarms).

A CSI of 1 means a perfect match. Against engineering flood maps, satellite images and post-event surveys, Fathom reports a CSI of about 0.75 for flood extent and an average error of about 0.6 m for water levels. The authors note that these scores approach the ceiling set by errors in the observations themselves.

Two features make Fathom an interesting case.

The whole methodology is published in peer-reviewed journals, and the provider sells data products for other platforms to distribute. A financial institution can therefore trace each hazard layer back to its sources, which is rare among commercial providers. Published validation scores such as the CSI also let users judge the model's uncertainty for themselves.

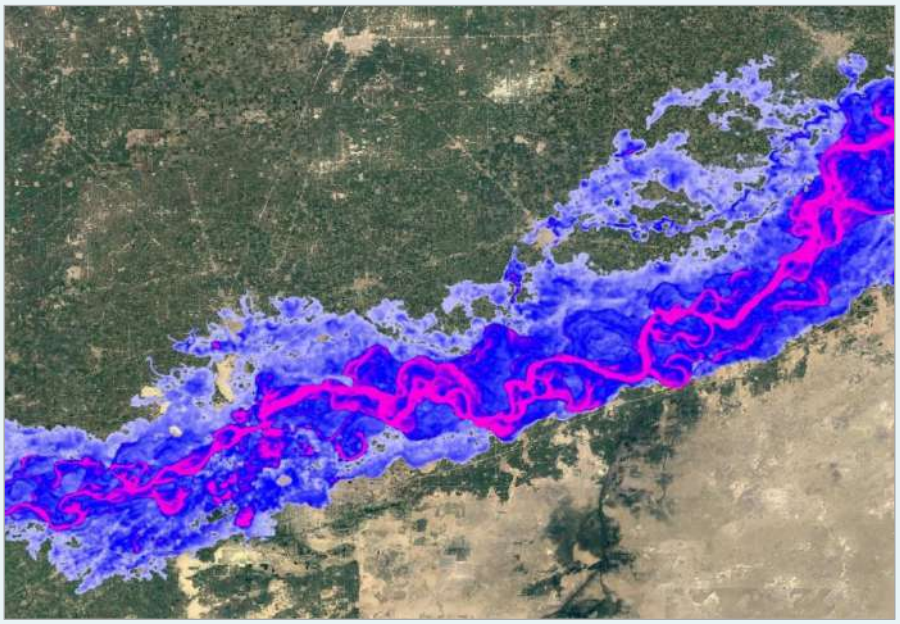


FIGURE 67. 200-year fluvial flood depth map in Pakistan.
Darker blues indicate greater depths, and permanent water is shown in pink. Source: Fathom

Provider Box – Hydroclimat, a hazard-focused physics-based approach

Hydroclimat is a French provider specialising in quantitative assessment of physical climate and hydrological hazards. It produces location-specific indicators, such as flood depth, precipitation intensity, extreme temperatures, wind hazards, drought and coastal flooding, from asset coordinates or spatial datasets. Its work focuses on hazard and exposure assessment, and it does not directly estimate damage or financial losses.

Its clients span insurance, banking, infrastructure, energy, industry and public authorities. The indicators can feed third-party damage, vulnerability and actuarial models, and they also support climate adaptation strategies. The hazards cover those relevant to the EU Taxonomy and the CSRD. Historical reference periods vary by dataset, and projections extend to 2100. These follow SSP scenarios, global warming levels or the French reference warming trajectory for adaptation (TRACC), depending on the client's framework.

The modelling combines physically based approaches with statistical methods. Standard products rely on CMIP6 ensembles, climate reanalyses, Copernicus and Météo-France data, together with proprietary hazard datasets. Bespoke studies add 1D and 2D hydraulic simulations, and resolution ranges from kilometre-scale climate datasets to metre-scale hydraulic simulations for selected engineering studies.

Machine learning supports bias correction, downscaling, and hydrological modelling and projections.

The model is using machine learning for downscaling, hydrological forecasting and hybrid physics-informed components.

Uncertainty is characterised through multi-model climate ensembles and statistical distributions, including quantile ranges (Q05, Q50 and Q95) where available. Validation compares the hazard outputs with historical observations and benchmark datasets. Validation against losses falls outside its scope, since the model stops at the hazard.

> The provider considers stochastic event-based approaches better suited to retrospective insurance analysis than to forward-looking projection. Its outputs are therefore scenario-based, without event catalogues.
>
> Hydroclimat shows a division of labour found across the market. A specialized hazard provider brings resolution and physical expertise, and the financial institution keeps control of the damage and financial layers. The quality of the final financial figure then depends on the client's damage functions as much as on the hazard data.

Actuarial weather indices

Raw hazard data are hard for a financial institution to use directly. A reanalysis grid gives temperature, rainfall or wind every hour and for every cell. A weather index condenses these variables into one number per place and per period. The number is easier to read, track and contract on. For single-hazard models, indices are often the most practical entry point to hazard data. They come in two families. Monitoring indices track how climate extremes drift over time. Contract indices trigger a payout in parametric (index-based) insurance.

Monitoring indices: the Actuaries Climate Index

The Actuaries Climate Index (ACI) was developed jointly by the North American actuarial societies for the United States and Canada. It works like a consumer price index for climate extremes. Its basket holds six components, each a monthly series starting in 1961. Each component is expressed as a standardized anomaly against the 1961–1990 reference period, which works like a z-score:

$$ACI = 1/6\ (T90_{std} - T10_{std} + P_{std} + D_{std} + W_{std} + S_{std})$$

- With $T90_{std}$ is the anomaly in the frequency of temperatures above the 90th percentile
- $\boldsymbol{T10_{std}}$ is the anomaly in the frequency of temperatures below the 10th percentile. It enters with a negative sign because fewer cold extremes mean warming
- $\boldsymbol{P_{std}}$ is the anomaly in the monthly maximum rainfall over five consecutive days
- $\boldsymbol{D_{std}}$ is the anomaly in the annual maximum number of consecutive dry days (precipitation below 1 mm)
- $\boldsymbol{W_{std}}$ is the anomaly in the frequency of wind speeds above the 90th percentile
- $\boldsymbol{S_{std}}$ is the anomaly in mean sea level, measured at tide gauges

An index value of 0.5 means that extremes have increased by half a standard deviation on average compared with the reference period.

Actuarial weather indices are constructed from observed meteorological records to characterise the statistical behaviour of weather-related losses at a given location or region.

They summarize the historical frequency and severity distribution of weather events in a form directly usable by actuarial loss models. Examples include catastrophe loss indices (PCS in the US, PERILS in Europe), parametric trigger indices (wind speed exceedance at a named station, cumulative rainfall over a basin), and standardized weather event counts (number of days exceeding a temperature threshold, annual 99th percentile precipitation).

Actuarial indices have several properties:

- They are directly calibrated to observed loss experience, which makes them tractable for pricing and reserving.
- They are parsimonious: a single index value can trigger a parametric insurance payment without requiring a full damage calculation.

- They can characterise the tail of the distribution with greater statistical confidence than physical models alone for well-observed perils in data-rich geographies because they are constructed from long meteorological records.

The method has travelled. Australia built its own version in 2018. Zhou et al. (2023) computed an Iberian Actuarial Climate Index (IACI) from ERA5-Land reanalysis over 6,526 cells of 0.1° × 0.1°. The North American ACI uses 12 sub-regions. In the Iberian case, sea level and warm extremes mainly drive the sign of the index. In France, CNP Assurances and the DIALog Chair adapted the ACI to French data as the Indice Climatique Actuariel Français (ICAF) in 2024.

Note also that the value of the ACI family lies in its transparency. Each component is public and replicable. Its limit is that it describes hazard trends only and has no link to losses.

Contract indices: weather index insurance

In index insurance, the payout depends on a measured weather variable instead of a loss assessed after the event. Joseph et al. (2026) reviewed 87 studies and found 49 distinct indices. The most common are the Standardized Precipitation Index (SPI), the Normalized Difference Vegetation Index (NDVI) and degree-day indices.

Two recent applications show how an index is chosen and priced. Our first example provides a better understanding a drought and flood cover for the Cantareira water supply system in São Paulo (Gesualdo et al., 2024). They tested SPI, SPEI, consecutive dry and wet days, streamflow and precipitation against reservoir levels. Only two-year accumulated precipitation passed their correlation threshold of 0.7, with a value of 0.76. Another designed a frost cover for Australian wheat producers based on heating degree days below a temperature threshold (Barratt et al., 2025):

$$HDD_{T_{min}}([t_1, t_2]) = \sum_{t=t_1}^{t_2} max\{0, T - T_{min}(t)\} - \varphi$$

- $[t_1, t_2$ is the risk period, from flag leaf emergence to 10 days after flowering
- $T_{min}(t)$ is the minimum temperature on day t
- T is the frost threshold (0, −1 or −2°C)
- φ is the strike, the index level at which the contract starts to pay (set to zero in the study).

The payout equals the index multiplied by a fixed amount per degree.

Every index carries basis risk, the gap between the payout and the actual loss. We can distinguish four types (Joseph et al., 2026). Spatial basis risk comes from distance between the measuring point and the asset. Temporal basis risk comes from a mismatch between the index period and the sensitive period of the exposed asset. Design basis risk comes from a poorly specified index. Technological basis risk comes from the estimation model.

They have three limitations:

- First, they are backward-looking by construction (i.e. calibrated on historical observations), they cannot represent a non-stationary climate in which the frequency and severity of extreme events is shifting beyond historical bounds, for example when tipping points occur (Lenton et al., 2008; Dietz et al., 2021; Armstrong et al., 2022).
- Second, they are geographically and perennially uneven: dense observational networks in Western Europe and North America contrast with sparse coverage in Sub-Saharan Africa, South-East Asia, and the Arctic, precisely the regions where climate risk is projected to grow fastest. The actuarial index for a parametric

drought product in the Sahel rests on a fraction of the data that would support an equivalent European product.

- Third, an underappreciated limitation is basis risk: the index may not track the actual losses of the specific portfolio being protected. If a parametric wind insurance product triggers at wind speeds above 28 m/s at a single reference station, and the insured assets are 60km inland where wind speeds were 24 m/s during the same event, the insured sustains a loss that the index does not capture. Basis risk is the central design problem of all parametric products and determines the extent to which actuarial indices can serve as proxies for physical damage in portfolio risk models.

HAZARD	INDEX	WHAT IT MEASURES	TYPICAL DATA	EXAMPLE OF USE
Drought (clay soil subsidence)	SWI, Soil Wetness Index	Soil water content on a 0 to 1 scale, relative to usable reserve	Météo-France SIM model, 8 km grid	French CatNat drought recognition (return period ≥ 25 years). Olivier Lopez's CARE programme on future droughts in the Grand Est (Heranval et al., 2026)
Drought	PDSI, Palmer Drought Severity Index	Soil water balance from precipitation and temperature	Station or gridded series	Long-run US drought monitoring
Meteorological drought	SPI	Precipitation anomaly over 1 to 48 months, standardized	Station or gridded rainfall	Agricultural index insurance
Drought under warming	SPEI	SPI with evapotranspiration added	Rainfall and temperature	Hazard screening under climate scenarios
Vegetation stress	NDVI, VCI	Vegetation health from satellite reflectance	MODIS, Sentinel	Pasture and livestock insurance
Frost	HDD below threshold	Cumulative degrees below a frost threshold	Daily minimum temperature	Wheat frost cover (Barratt et al., 2025)
Heat	CDD, GDD	Cumulative degrees above a threshold	Daily temperature	Energy demand, crop heat stress
Multi-hazard trend	ACI, IACI, ICAF	Composite anomaly of six extremes	Stations, reanalysis, tide gauges	Climate monitoring for insurers

TABLE 35. Climate indices commonly used to measure physical hazards in insurance and financial risk assessment

Compound event data

Compound events are configurations in which two or more physical drivers co-occur, interact, or occur in rapid succession such that the resulting impact exceeds what would be expected from each driver acting independently (Edgeley et al., 2024; Zscheischler et al., 2018). The IPCC AR6 WG1 chapter on extreme events (Zscheischler et al., 2020; IPCC, 2021) formalized a taxonomy with four main compound types:

COMPOUND TYPE	DEFINITION	FINANCIAL RISK EXAMPLE
Preconditioned	A weather or climate precondition aggravates the impact of a hazard	Wildfire followed by heavy rain. Burned slopes absorb less water, so moderate rainfall triggers flash floods and debris flows. Flood damage after a fire is often uninsured, because take-up drops within a few years (Edgeley et al., 2024)
Multivariate	Several drivers or hazards occur together at the same place and time	Storm surge, high river discharge and heavy rain at an estuary. No driver alone would breach the coastal defences, but together they do (Bevacqua et al., 2019)
Temporally compounding	A succession of hazards hits the same area before recovery is complete	Successive hurricanes in one season, such as Irma and Maria in the Caribbean in 2017. Damage from the first event raises vulnerability to the second, so season losses exceed the sum of independent-event estimates
Spatially compounding	Hazards in several connected locations cause an aggregated impact	One weather system flooding several countries at once, as in Germany, Belgium and the Netherlands in July 2021. A portfolio diversified across regions loses that diversification, and aggregate losses exceed what independent regional models predict

TABLE 36. Compound Climate Events and Their Financial Risks

There is a challenge in gathering data for compound events. Standard single-hazard datasets are statistically independent by construction: flood return periods are computed by fitting a univariate extreme value distribution to river flow; wind return periods are fitted to wind speed records at stations. Combining two such independent datasets to estimate the joint probability of simultaneous high-wind and high-flood events will structurally underestimate the true co-occurrence probability if the two hazards are driven by a shared atmospheric precursor such as an extratropical cyclone (Dupuis & Jones, 2006; Hillier et al., 2015; Martius et al., 2016; Li et al., 2023; Kerkhofs et al., 2025).

That's why compound risk requires multivariate extreme value analysis. We already mentioned the extreme value theory in Part 2. But we will typically need to consider how two hazards are related through copula models. A copula separates the marginal distribution of each hazard (fitted independently) from the dependence structure between them.

The choice of copula family determines the tail dependence properties of the resulting joint distribution and has first-order effects on the estimated probability of simultaneous extreme events.

COPULA FRAMEWORK FOR COMPOUND HAZARD MODELLING

How can a practitioner tell whether two hazards are related? Three signs suggest dependence. The hazards share a physical driver, such as an extratropical cyclone producing both wind and heavy rain. Their annual maxima are correlated in the historical record. Their extremes coincide in past events more often than chance would predict. When these signs are present, the joint behaviour of the hazards should be modelled explicitly.

Let X_1 and X_2 be two hazard intensity random variables, for example flood depth and wind speed. By Sklar's theorem (Sklar, 1959), their joint distribution can be written as

$$F(x1, x2) = C(F1(x1), F2(x2))$$

- Where $\boldsymbol{F}$ is the joint cumulative distribution function of X_1 and X_2.
- $\boldsymbol{F_1}$ and F_2 are the marginal cumulative distribution functions, fitted to each hazard separately.
- $\boldsymbol{C}$: $[\mathbf{0}, \mathbf{1}]^{\mathbf{2}} \rightarrow [\mathbf{0}, \mathbf{1}]$ is the copula, which carries the whole dependence structure.

The copula family matters most in the tails. It is characterized by the upper tail dependence coefficient

$$\lambda_U = \lim_{u \to 1^-} P\left(U_2 > u \mid U_1 > u\right)$$

- $\boldsymbol{\lambda_U}$ is the probability that one hazard is extreme given that the other is extreme, in the limit.
- $\boldsymbol{U_1 = F_1(X_1)}$ and $U_2 = F_2(X_2)$ are the hazards expressed on a uniform scale.
- $\boldsymbol{u}$ is the probability level, which tends to 1.

The lower tail dependence coefficient λ_U is defined in the same way for joint low values. Four families are common in practice (Joe, 1997; Nelsen, 2006).

Gaussian copula. It has no tail dependence for any correlation below 1 ($\lambda_U = \lambda_L = 0$). Joint extremes become asymptotically independent, so the Gaussian copula underestimates their probability when hazards tend to peak together (Embrechts et al., 2002).

Student-t copula. It has symmetric tail dependence ($\lambda_U = \lambda_L > 0$). It suits hazards whose joint extremes and joint low values are both more frequent than under a Gaussian model.

Gumbel (extreme value) It has upper tail dependence only ($\lambda_U = 2 - 2^{1/\theta}$ with $\theta \geq 1$). It suits hazards whose extremes tend to co-occur, such as tropical cyclone wind and storm surge.

Clayton copula is used with lower tail dependence only ($\lambda_L = 2^{-1/\theta}$ with $\theta > 0$). suits hazards defined by jointly low values, such as low river flow and low soil moisture during a drought.

The joint exceedance probability

The probability that both hazards exceed given thresholds is

$$P(X_1 > x_1, X_2 > x_2) = 1 - F_1(x_1) - F_2(x_2) + C\big(F_1(x_1), F_2(x_2)\big)$$

- $\boldsymbol{x_1}$ and x_2 are the intensity thresholds for each hazard, for example the 100-year flood depth and the 100-year wind speed.

COPULA FRAMEWORK FOR COMPOUND HAZARD MODELLING

The key variable to assess the financial impact will be the return period of the joint event:

$$T = \frac{1}{P(X_1 > x_1, X_2 > x_2)}$$

Consider two hazards each at their 1-in-100-year level, so that

$$F_1(x_1) = F_2(x_2) = 0.99$$

Under independence,

$$T = 1/(0.01 \times 0.01) = 10,000$$

Under a Gumbel copula with $\theta = 3$

the diagonal simplifies to $C(u,u) = u^{2^{1/\theta}}$ which gives a joint probability of about 0.0074 and a return period of about 135 years. The resulting tail dependence $\lambda_U \approx 0.74$ will shorten the joint return period by a factor of about 75. The probability-weighted annual loss contribution of this event configuration rises by the same factor.

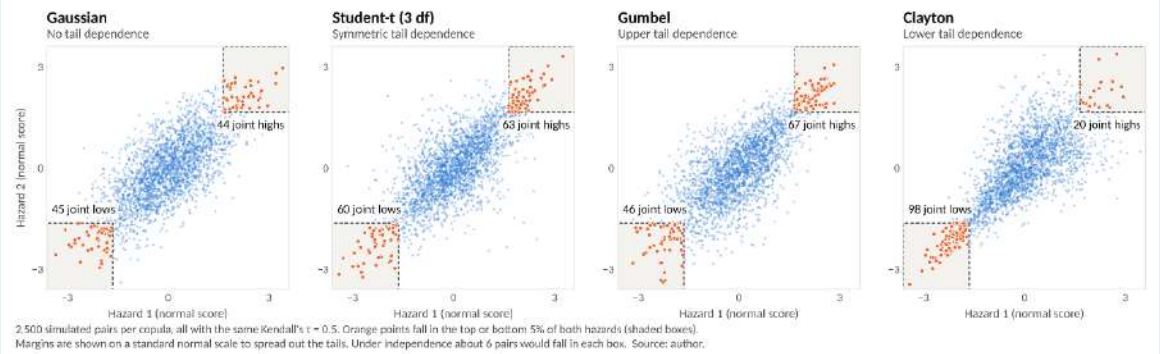


FIGURE 68. Tail dependence of four copulas with identical Kendall's tau

Four copulas with the same rank correlation (Kendall's $\tau = 0.5$) and different tail behaviour. Orange points are pairs in the top or bottom 5% of both hazards. The Gumbel copula places more pairs in the joint-extreme corner than the Gaussian, which is why it shortens joint return periods for hazards such as cyclone wind and surge.

In practice, compound event data is rarely available in directly usable form. The practitioner's options are ordered by data availability and statistical tractability:

- Full joint simulation: physics-based climate models can generate multivariate hazard footprints that preserve the physical co-occurrence structure. CMIP6 large ensemble runs provide thousands of years of synthetic climate history from which compound event statistics can be estimated. This is computationally intensive but physically consistent. ERA5 reanalysis can serve as a shorter but observation-validated alternative.
- Copula fitting to historical observations: fit marginal distributions to each hazard independently, estimate the copula from historical co-occurrence records. This requires long, concurrent time series of both

hazards at the same locations, which are available for some peril pairs (coastal flood and wind) but not others (wildfire and drought).

- Conditional scenario analysis: fix one hazard at a given return period and estimate the conditional distribution of the second given the first. This is a practical compromise widely used in the insurance industry for riverine and coastal flood combined assessments.
- Structured expert elicitation: where data are insufficient for statistical fitting, expert knowledge about the physical co-occurrence mechanisms can be formalized through a Bayesian prior. This is the approach recommended in the actuarial literature for perils with limited historical records.

Why single-hazard models systematically underestimate tail risk

The financial consequence of the single-hazard assumption is asymmetric. At moderate return periods (1-in-10 to 1-in-50 years), ignoring the dependence between hazards produces a small error in loss terms. The error grows in the tail, at return periods of 1-in-100 to 1-in-250 years, which is the range relevant for regulatory capital.

The Solvency II SCR, for instance, corresponds to a 99.5% VaR over one year, or a 1-in-200-year loss. When hazards share a physical driver and show upper tail dependence, their extremes tend to occur together. A flood at the 99th percentile is then more likely than an average flood to come from an atmospheric configuration that also produces storm surge or extreme wind.

Institutions that use single-hazard models for capital calculations therefore underestimate their tail loss exposure whenever the hazards in their portfolio are positively dependent. The size of the bias depends on the peril mix and the geography of the portfolio.

In Europe, the co-occurrence of storm surge and heavy precipitation makes compound coastal flooding more likely than independent models suggest, and this probability is projected to rise along parts of the northern European coast (Bevacqua et al., 2019). In Great Britain, assuming independence underestimates the probability of co-occurring extreme wind and river flow episodes by a factor of 2 to 4 (Hillier et al., 2025).

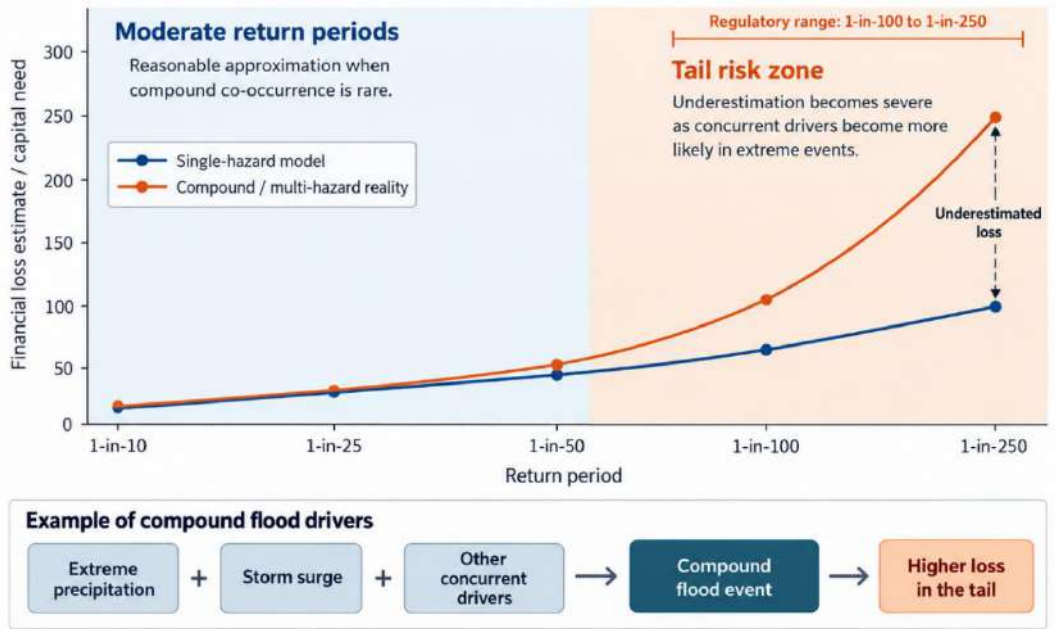


FIGURE 69. Underestimation of tail losses by single-hazard models compared with compound hazard events

DIMENSION	SINGLE-HAZARD	COMPOUND / MULTI-HAZARD
Data requirement	Univariate time series for one intensity metric. 30 to 50 years are usually sufficient to fit the marginal distribution	Concurrent multivariate time series at the same locations. Much longer records, often a century or more, are needed for stable estimation of tail dependence
Statistical method	Univariate extreme value theory (GEV, GPD). Well established, tractable and standard in actuarial practice	Multivariate extreme value theory and copulas. Computationally demanding and sensitive to the choice of copula family
Vulnerability function compatibility	Direct. Most vulnerability libraries (e.g. Hazus, FLEMOps, JRC depth-damage curves) are calibrated for single-hazard inputs	Limited. Hazus combines wind and surge damage for hurricanes, but most other combinations require custom multi-hazard damage functions or a combination rule applied to single-hazard losses
Regulatory treatment	Standard in Solvency II natural catastrophe modules and in supervisory climate stress tests	Encouraged in supervisory guidance on climate risk (e.g. BCBS, 2022; EIOPA, 2021) but not yet standard in Pillar 1 capital calculations
Tail bias	Underestimates joint tail losses when hazards are positively dependent	Consistent if the copula is correctly specified, but adds model uncertainty from the choice of family, especially in the upper tail where data are sparse
Primary use cases	Pricing and reserving for single-peril products, single-peril scenario analysis, physical risk heat maps	Reinsurance aggregate covers, sovereign catastrophe risk assessment, portfolio AEP curves for multi-peril geographies

TABLE 37. Comparison: single-hazard vs compound hazard data requirements

Provider Box – The case of Climafin multi-hazard approach

We already mentioned Climate Finance Alpha (CLIMAFIN) in section 2.2, a Paris-based firm founded by academics from Paris 1 Panthéon-Sorbonne, the University of Zurich and Utrecht University. Its structural credit-risk model, CLIMACRED-PHYS, is one of the few published models that connects several physical hazards to the price of bonds and equities (Mandel et al., 2025).

The model follows the hazard, exposure and vulnerability chain in five steps. First, it geolocates the production facilities behind each security, with the help of the data provider Sequantis. Second, it samples hazard

trajectories at each location. Third, it converts these into capital destruction and business interruption. Fourth, it passes the impacts to the firm's cash flow and balance sheet. Fifth, it derives the probability of default and the value of the firm's securities, in the tradition of Merton (1974).

The model covers six hazards with a common climate forcing, the ISIMIP3b protocol, which draws on an ensemble of CMIP6 climate models. Each hazard has its own physical proxy. Tropical and winter storms use maximum wind gust, from IBTrACS synthetic tracks and the Copernicus windstorm events. River and coastal floods use flood depth, from GLOFRIS maps at about 1 km. Wildfires and landslides use an occurrence indicator, from MODIS burned area and a global landslide atlas. The combined impact on a facility in a given year is driven by the strongest hazard.

$$D_{i,t}^{S} = g_i \cdot \max_{h \in H} f_{h,c_i,s_i}(x_{h,i,t}^{S})$$

- $D_{i,t}^{S}$ is the share of total capital of facility i destroyed in year t under scenario S
- g_i is the share of tangible capital of the facility, taken from EU KLEMS sector averages
- H is the set of hazards (tropical storm, winter storm, river flood, coastal flood, wildfire, landslide)
- f_{h,c_i,s_i} is the vulnerability function of hazard h , specific to the country c_i and sector s_i of the facility
- $x_{h,i,t}^{S}$ is the intensity of hazard h at the location of facility i in year t

The authors argue that hazards rarely strike the same facility in the same year, so adding their impacts gives results close to taking the maximum. This choice treats hazards as independent at the facility level and leaves out the compound interactions described above. Business interruption days are then derived from $D_{i,t}^{S}$ with FEMA sector coefficients.

Applied to 2,934 firms of the MSCI ACWI and about 130,000 facilities, the model gives a mean annual loss of 2.22% of productive capital under RCP 6.0 (published ISIMIP2b version) in 2050. The 99th percentile reaches 3.80%. Tropical cyclones and coastal floods drive most of the impact. Financial leverage amplifies these direct losses.

The mean equity loss ranges from 10.1% to 12.9%, about four to six times the direct capital loss, because financial leverage amplifies the impact on shareholders. The model also shows cross-border transfers of risk. Losses at facilities in Asia land on the balance sheets of firms headquartered in Europe and North America. The main lesson of CLIMACRED-PHYS is that the leverage of the issuer and the location of its facilities can matter as much as the hazard data. Its treatment of hazards as independent also shows how far multi-hazard coverage is from compound modelling.

Key Limitations

The stationarity problem is more acute for compound events than for single hazards (Zscheischler et al., 2020). Single-hazard marginal distributions change under climate change in ways that can be projected from physical model ensembles: flood frequencies shift with precipitation intensification; wind speed distributions shift with changes in storm track and extratropical cyclone activity. The dependence structure between hazards, encoded in the copula, is a more ambiguous object under climate change. There is limited theoretical or empirical guidance on whether the copula family and parameters estimated from historical records will remain stable in 2 degrees or 3 degrees warmer climate. This uncertainty propagates directly into compound event return period estimates and represents a structural limit of current multi-hazard modelling.

Data asymmetry between perils and geographies constrains copula estimation. Copula parameter estimation for compound events requires long concurrent records of both hazard variables at the same locations. For European

Atlantic coastlines, where wind and surge observations extend back a century, this is feasible. For compound drought and heat in Sub-Saharan Africa, or compound typhoon and landslide in South-East Asia, the observational record is far shorter and spatially sparse. Institutions with portfolio exposure in these geographies face a structural data gap that cannot be resolved by statistical sophistication alone.

Actuarial indices suppress the compound signal. Annual loss indices aggregate across all events in a year and across all perils contributing to those events. A year in which a portfolio suffers a moderate wind loss and a moderate flood loss will produce the same annual index value as a year with a single large wind event of equivalent total cost. The compound structure is invisible in the aggregated index. This is not a problem for pricing aggregate annual loss products but means that actuarial indices cannot be used to estimate the compound event contribution to tail risk. Physics-based or copula-based approaches are required for that purpose.

Vulnerability functions for compound hazards remain an open research problem. As noted in section 2.1.1, no widely deployed vulnerability model captures multi-hazard interactions in a physically consistent way. The practical consequence for multi-hazard data users is that even a correctly specified joint hazard distribution cannot be passed directly to a standard damage function. The joint intensity must be decomposed, each component run through a single-hazard vulnerability function, and interaction effects estimated through structural adjustments that are themselves poorly calibrated. Until multi-hazard vulnerability libraries are developed and validated against compound event claims data, the compound hazard data chain remains broken between the hazard module and the loss module.

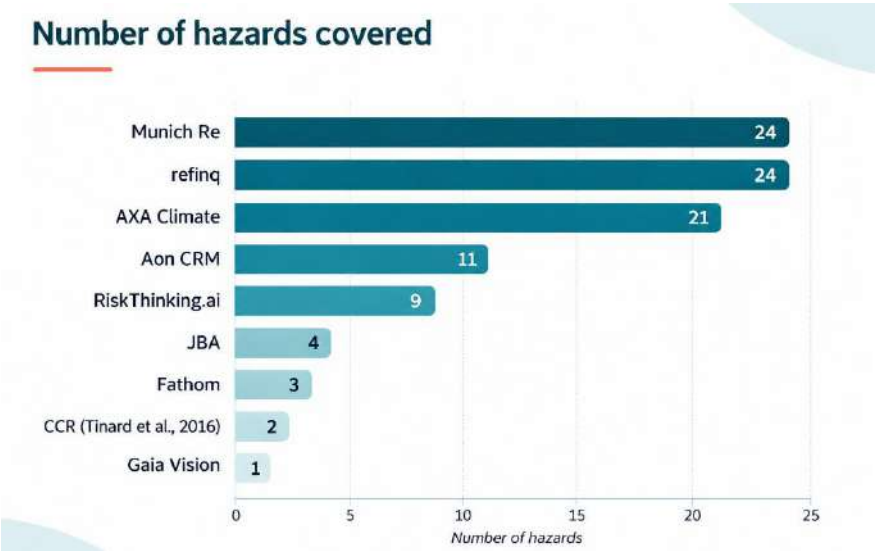


FIGURE 70. An illustration of the Breadth vs. Depth dilemma – Number of hazards by provider

3.1.2 HAZARD TYPE

Hazard type is the second axis, together with the single- versus multiple-hazard architecture discussed in the previous section, that determines how physical risk data must be structured. Where a given hazard sits within a classification scheme, whether it is acute or chronic, whether it acts through fire, water, wind, or ground, and whether it falls within or outside the conventional boundary of physical climate risk, determines which datasets are relevant, which vulnerability functions apply, and which regulatory disclosure regime governs its treatment.

Classification frameworks for hazard type

We have already written about the different perspectives of framing physical climate risks in Part 1.

The most widely used classification distinguishes acute from chronic hazards. Acute physical hazards are discrete, event-driven perils with a defined onset and duration, such as tropical cyclones, floods and wildfires. Chronic physical hazards are slow-onset, long-duration shifts in average conditions, such as sea level rise, mean temperature increase and changing precipitation patterns.

The TCFD recommendations introduced this distinction (TCFD, 2017), and Appendix A of the EU Taxonomy Climate Delegated Act later formalized it (European Commission, 2021). The distinction has direct data implications. Acute hazard data is typically expressed as an event-based intensity metric with an associated return period. Chronic hazard data is typically expressed as a time series or a projected trend under a given warming scenario.

A second, complementary classification groups hazards by the physical medium through which they act, independently of their temporal signature. Fire-based hazards include wildfire. Water-based hazards include river, pluvial and coastal flood, as well as drought. Wind-based hazards include tropical and extratropical cyclones, hailstorms and tornadoes. Ground-based hazards include landslide, subsidence and permafrost thaw.

The EU Taxonomy itself combines both logics, since its hazard table crosses the acute and chronic categories with temperature-, wind-, water- and solid mass-related hazards (European Commission, 2021). This peril-based grouping is closer to how catastrophe modelling vendors and reinsurers structure their hazard libraries, and it is generally the organising principle behind commercial hazard datasets (IRDR, 2014; Mitchell-Wallace et al., 2017).

Other classification frameworks in use across the industry include the acute and chronic split applied in the NGFS climate scenarios (NGFS, 2024), the ISO 14091 guidelines for climate vulnerability and risk assessment (ISO, 2021), and the climatic impact-driver framework of the IPCC Sixth Assessment Report (Ranasinghe et al., 2021). The IPCC framework sorts hazards into seven types, namely heat and cold, wet and dry, wind, snow and ice, coastal, oceanic, and other.

None of these frameworks is universally adopted. Insurance companies frequently maintain an internal mapping table that reconciles their chosen regulatory taxonomy with the peril categories used by their data vendors and catastrophe models. Maintaining this mapping is itself a data governance task.

Boundary cases: geophysical hazards and the definition of physical climate risk

A recurring practitioner question concerns the place of seismic and volcanic hazards within a physical risk data architecture. Are earthquakes and volcanic eruptions physical risks? They are. Both cause direct physical damage to assets, and both are routinely included in natural catastrophe risk frameworks (IRDR, 2014). Are they physical climate risks? The answer is more qualified. In the conventional sense used by the TCFD, the EU Taxonomy and most regulatory physical risk frameworks, they fall outside the scope of physical climate risk. The TCFD defines physical risks as risks resulting from climate change (TCFD, 2017). The list of climate-related hazards in Appendix A of the EU Taxonomy Climate Delegated Act contains no seismic or volcanic hazard (European Commission, 2021).

Physical climate risk, as the term is used throughout this cookbook, refers to hazards whose frequency or severity is expected to change, and in most cases intensify, as a function of anthropogenic greenhouse gas emissions and the resulting warming trend. The causal chain runs from rising emissions to a hotter and more variable climate, then to a higher intensity or frequency of the physical hazard, and finally to a larger expected financial impact. Earthquakes and volcanic eruptions do not fit this chain at the time horizons and spatial scales relevant to financial risk management. The geoscience literature does identify mechanisms through which climate can influence geophysical activity. Glacier retreat unloads the crust and can raise eruption rates in glaciated volcanic regions such as Iceland, and seasonal changes in surface water loads can modulate local seismicity (McGuire, 2010; Johnson et al., 2017;

Swindles et al., 2018). These effects operate over centuries to millennia, or remain too small and too local to enter a financial risk model. No established causal link allows anthropogenic emissions to be translated into a measurable change in global seismic or volcanic hazard over the horizon of a climate stress test.

The relationship between volcanic activity and climate runs largely in the opposite direction. Large explosive eruptions inject sulphate aerosols into the stratosphere. These aerosols reflect incoming solar radiation and produce a measurable, though temporary, global cooling effect, as observed after the 1991 Mount Pinatubo eruption (Hansen et al., 1992; Robock, 2000). Volcanic activity can therefore affect both climate and financial outcomes. The reverse channel, from human activity or climate trends to volcanic activity, remains negligible for practical risk assessment.

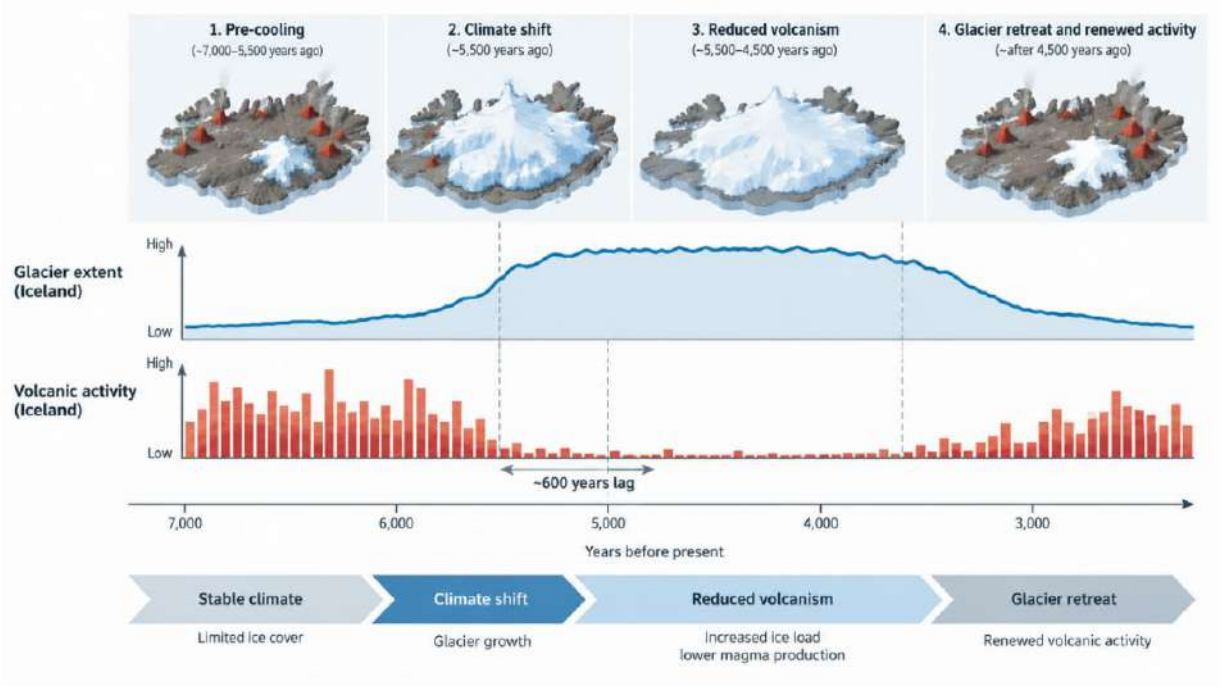


FIGURE 71. Illustration of the volcanic activity and Glacier extent.
Source: Adapted from Swindles et al., 2018

The table below summarises the contrast in causal structure between climate-driven hazards and geophysical hazards.

CAUSAL LINK	CLIMATE-DRIVEN HAZARDS	GEOPHYSICAL HAZARDS
Human activity → hazard intensity	Established: rising GHG emissions increase intensity or frequency via a warming climate	Not established at financially relevant scales. Proposed mechanisms such as glacial unloading act over centuries or locally (McGuire, 2010; Swindles et al., 2018).
Hazard → climate trends	Not applicable: the hazard does not itself drive climate	Large explosive eruptions can cause short-term global cooling

CAUSAL LINK	CLIMATE-DRIVEN HAZARDS	GEOPHYSICAL HAZARDS
		through stratospheric aerosols (Robock, 2000).
Hazard → financial impact	Direct	Direct

TABLE 38. Causal links between human activity, hazards, climate trends and financial impact

Earthquakes show no causal channel of practical relevance in either direction. Seasonal water storage can modulate the timing of small earthquakes on individual faults (Johnson et al., 2017). The literature, however, documents no detectable effect of anthropogenic climate change on global earthquake rates, and seismic activity has no measurable effect on climate (McGuire, 2010). Most regulatory taxonomies consequently classify earthquakes as a non-climate physical hazard. Many institutions nevertheless include them in their broader physical risk or natural catastrophe perimeter for capital adequacy and disclosure purposes.

This distinction has practical consequences for data governance as well as for taxonomy. An institution's physical risk data architecture is generally built around the natural catastrophe perimeter as a whole, which is wider than the physical climate risk perimeter defined by the TCFD or the EU Taxonomy. Earthquake and volcanic hazard data therefore typically sit within the same data infrastructure, vendor contracts and governance processes as wildfire, flood and windstorm data, even though only the latter group is expected to change with the climate. Practitioners should document this scope decision explicitly, since disclosure frameworks that refer to physical climate risk may require the geophysical component to be reported separately or excluded.

Data requirements by hazard type

Once a hazard perimeter has been defined, whether restricted to climate-driven perils or extended to the full natural catastrophe set, the practitioner must translate each hazard type into the specific physics-based variables required to characterise it, and identify where those variables can be sourced. The table below summarises this mapping for the hazard types most commonly represented in financial institution physical risk perimeters.

HAZARD TYPE	PHYSICS-BASED VARIABLES REQUIRED	EXAMPLE DATA SOURCES
Wildfires	Fuel load and vegetation dryness, fire weather index (temperature, humidity, wind speed, drought codes), historical burned area	Canadian Forest Fire Weather Index System (Van Wagner, 1987); Copernicus European Forest Fire Information System (San-Miguel-Ayanz et al., 2012); NASA MODIS burned area and VIIRS active fire products (Giglio et al., 2018; Schroeder et al., 2014) Mandel et al. (2025), MODIS burned area
Flood, river (fluvial)	River discharge, return period flood extent and depth, digital elevation model	Copernicus Global Flood Awareness System (Alfieri et al., 2013); JRC global flood hazard maps (Dottori et al., 2016); national flood hazard maps; Copernicus Digital Elevation Model

HAZARD TYPE	PHYSICS-BASED VARIABLES REQUIRED	EXAMPLE DATA SOURCES
		Jewson et al. (2023), JRC PESETA IV / LISFLOOD; Mandel et al. (2025), JRC depth-damage functions
Flood, pluvial (surface water)	High-resolution precipitation intensity, urban drainage capacity, land surface permeability	Radar-derived precipitation nowcasting products; national meteorological service rainfall records Bourget et al. (2024), national precipitation records for the United States and Canada
Flood, coastal	Storm surge height, sea level rise projections, bathymetry, tide gauge records	IPCC AR6 sea level projections (Fox-Kemper et al., 2021); Permanent Service for Mean Sea Level tide gauge records (Holgate et al., 2013); coastal digital elevation model products Bernstein et al. (2019), NOAA sea level rise data; Mandel et al. (2025), coastal flood depth maps
Earthquake	Peak ground acceleration, seismic hazard return period, fault proximity, soil liquefaction susceptibility	USGS ShakeMap (Wald et al., 2005); Global Earthquake Model seismic hazard maps (Pagani et al., 2020) Maneerat et al. (2024), USGS earthquake catalogue and peak ground acceleration map
Tropical cyclone	Wind speed field, central pressure, storm track, coupled storm surge	NOAA International Best Track Archive for Climate Stewardship (Knapp et al., 2010); ECMWF ERA5 reanalysis (Hersbach et al., 2020) Bressan et al. (2024) and Mandel et al. (2025), IBTrACS
Hailstorm	Convective available potential energy, hail size distribution, radar reflectivity	European Severe Weather Database (Dotzek et al., 2009); national weather radar networks
Heat wave	Temperature anomaly and duration, humidity-adjusted heat index, exceedance thresholds	ECMWF ERA5 reanalysis (Hersbach et al., 2020); heat wave metrics and definitions (Perkins & Alexander, 2013)

TABLE 39. Physics-based variables and example data sources by hazard type

This list is not exhaustive: drought, landslide, subsidence, and permafrost thaw are increasingly represented in institutional physical risk perimeters and follow the same logic of mapping a hazard type to its governing physical variables and available data sources. The remainder of this Part addresses how these datasets are sourced, validated, and governed in practice.

Hazard model impacts in practice: examples

The following examples illustrate how the hazard-to-variable mapping in the previous table translates into published financial impact models. Each example is drawn from our literature review and is selected to show a distinct modelling approach.

Floods: coastal

Coastal flood models combine a surge or extreme sea level hazard layer with an exposure layer and a depth-damage function. The main modelling choice concerns how the surge hazard is generated and how far the model projects it into the future.

A probabilistic surge model can be used to price storm surge risk for residential properties, as shown for Pensacola, Florida. Risk-based premiums are compared with the uniform pricing of the US National Flood Insurance Program, and elevating homes proves cost-effective mainly for high-value, low-lying properties in the highest-risk zones (Montgomery & Kunreuther, 2018).

In France, a regional climate model has been coupled with a storm surge model and the insured portfolio of the national natural catastrophe scheme. Total flood and storm surge insured losses are projected to rise from about €855 million in 2015 to €1.3 billion in 2050, and growth in insured values contributes more to this increase than climate change (Moncoulon et al., 2016).

At the global scale, coastal flood depth maps have been translated into default probabilities and equity losses through a structural credit risk model, in which coastal floods rank as the second most damaging hazard after tropical cyclones (Mandel et al., 2025).

Floods: pluvial

Pluvial flooding is driven by short, intense rainfall over surfaces that cannot absorb or evacuate it. Models therefore need high temporal resolution precipitation data and a representation of urbanisation.

Precipitation dynamics over 12 to 48 hours and urbanisation indicators have been used to model pluvial flood occurrences in the United States and Canada. Generalized additive models are compared with tree-based machine learning methods and used to project insurance losses under climate model scenarios. Generalized additive models extrapolate better to future climates, and both climate change and urbanisation increase projected losses, with strong heterogeneity across states and provinces (Bourget et al., 2024).

Household survey evidence from Germany shows that surface water floods cause the lowest losses and the shortest recovery times among inland flood pathways. Warning systems also perform worse for pluvial and flash floods than for river floods, which limits the scope for short-notice mitigation (Thieken et al., 2022).

Floods: fluvial / river flood

River flood models are the most mature class of flood model. They typically rely on hydrological simulation of river discharge, return-period flood maps and depth-damage functions, often within a catastrophe model structure.

The JRC PESETA IV river flood simulations are used to derive how annual losses in each European NUTS2 region scale with global mean surface temperature. The logarithm of losses rises approximately linearly with warming up to about 1.2 °C, and some regions show close to a doubling of losses per degree. A covariance matrix is also provided so that insurers can adjust catastrophe models for spatially correlated changes across Europe (Jewson et al., 2023).

At the asset level, correlated river flood events can be simulated with Gaussian, Student-t and vine copulas, and the resulting losses propagated to firm value, probability of default, value at risk and expected shortfall. Ignoring spatial correlation or insurance structure distorts these tail risk measures by large margins (Kerkhofs et al., 2025).

Long-term flood insurance premiums in the Netherlands can rise several-fold in high-risk areas under climate change if no adaptation takes place (Aerts & Botzen, 2011).

Sea level rise and real estate prices

Sea level rise is a chronic hazard whose financial effect appears in asset prices long before any physical damage occurs. Most models in the literature are hedonic or difference-in-differences regressions that test whether markets discount properties exposed to future inundation. US coastal homes exposed to sea level rise sell at a discount of about 7% compared with equivalent unexposed homes.

Rents show no discount, which indicates that the price effect reflects long-horizon risk and leaves current use value unaffected, and the discount is concentrated among non-owner-occupied properties (Bernstein et al., 2019).

A comparable hedonic approach applied to Melbourne finds no significant discount once dwelling characteristics are controlled for, so the capitalisation of sea level rise risk varies across markets (Fuerst & Warren-Myers, 2021).

The same exposure also affects credit. Lenders charge higher interest rates on 30-year mortgages for properties exposed to sea level rise, while the premium is absent for 15-year loans, which mature before the risk materialises (Nguyen et al., 2022).

Storms

Storm models in the literature focus mainly on tropical cyclones, for which track archives such as IBTrACS allow probabilistic event sets to be generated. The choice of intensity variable and the spatial granularity of exposure both have a strong effect on results.

The CLIMADA platform has been used to assess tropical cyclone and chronic risk at the level of individual assets and to aggregate the results to portfolio value at risk. Replacing asset-level locations with headquarters locations underestimates portfolio value at risk by up to 71% and omitting the acute tail (250-year return period) underestimates losses by up to 82% (Bressan et al., 2024).

Minimum sea level pressure predicts normalized US hurricane damage better than maximum sustained wind speed, a result that matters for the design of parametric triggers (Klotzbach et al., 2020). In credit markets, loan spreads rise by about 17 to 19 basis points for borrowers indirectly exposed to hurricanes (Correa et al., 2022).

Hailstorms

Hail is one of the least represented hazards in financial literature. Damage depends on hail size and convective storm structure at very fine spatial scales, and observational records are short and inconsistent.

One of the most complete hail-specific impact model available is an open-source radar-based damage model for buildings and cars in Switzerland. The hazard is the maximum expected severe hail size (MESHS), derived from 1 km weather radar data.

It is combined with about 989,000 buildings insured by cantonal building insurers over 2002 to 2021 and with more than 50,000 car insurance claims from a private insurer. Sigmoidal impact functions link hail size to the share of assets damaged and to the mean damage ratio, and the model runs on the open-source CLIMADA platform. Across more than 100 large hail events, the model reproduces the correct order of magnitude of the number of damaged buildings in 91% of cases and of total building damage cost in 77% of cases.

Performance is lower for cars, at 74% and 60% respectively. Most of the remaining errors come from the limits of radar-based hail detection, with building vulnerability playing a smaller role (Schmid et al., 2024). The concentration of losses illustrates the tail behaviour of the peril, since the hailstorm of 28 June 2021 alone caused about CHF 483 million in building claims, roughly 35% of the 20-year total (Schmid et al., 2024).

At the global scale, a statistical model trained on hail reports and driven by ERA5 reanalysis reconstructs the probability of very large hail (5 cm or more) from 1950 to 2023, and the results are compared with insured hail losses from Munich Re's NatCatSERVICE.

Europe shows the sharpest and only widespread significant increase in very large hail frequency, led by northern Italy, and the rise in European losses is attributed mainly to more frequent severe hailstorms. In the United States and Australia, losses have also risen while hail frequency shows no significant upward trend, so growth in exposure and vulnerability drives losses there (Battaglioli et al., 2026). The distinction matters for financial institutions, because a hail loss trend can reflect climate change in one region and portfolio growth in another.

A multi-peril approach is also possible. Hail can be treated jointly with tornadoes, floods and winter storms at census tract level, with losses simulated by Monte Carlo, the tail fitted with a generalized Pareto distribution, and the results used to price multi-peril catastrophe bonds (Bahrami et al., 2025). Hail-specific models nevertheless remain scarce. Institutions with material hail exposure, such as motor and property insurers in central Europe, still depend largely on vendor catastrophe models, and open models calibrated on insurance claims offer one of the few independent benchmarks.

Earthquakes and volcanoes

Geophysical hazards are modelled with the same hazard, exposure and vulnerability structure as climate-driven hazards, without a climate scenario dimension (see the boundary cases discussion above). Their main use in this cookbook is as a benchmark for the robustness of catastrophe models.

A factorial sensitivity analysis of an earthquake catastrophe model shows that construction type and event catalogue length drive most of the uncertainty. Model outputs vary by 30% to 50% depending on assumptions, and misplacing buildings geographically can halve estimated losses (Born et al., 2023).

For insured earthquake losses in New Zealand, a mixed-effects model that accounts for spatial correlation between territorial authorities improves predictive power (Di Lascio et al., 2025). For volcanic hazard, an event study measures significant negative abnormal returns for European airlines after the 2010 Eyjafjallajökull ash cloud closed European airspace (Mazzocchi et al., 2010).

Earthquakes also serve as a non-climate placebo in several studies of the review. They do not move loan spreads in the way hurricanes do, which supports the separation between climate and non-climate physical risk (Correa et al., 2022).

Wildfires

Wildfire models face two specific difficulties. Fire spread depends on vegetation, slope and wind at the scale of tens of metres, and wildfire triggers cascading hazards such as post-fire flash floods.

Wildfire risk is priced in US municipal bond yields for maturities between 1 and 15 years. No effect appears for shorter maturities, where short-term predictability reduces risk for bondholders, or for longer ones (Berry-Stölzle & Hao, 2026). An event study across US sectoral indices finds significant negative abnormal returns for the insurance sector lasting up to 28 days after major wildfires (Tavor, 2024). The financial dimension of the wildfire and flash flood cascade is documented in northern Arizona. Flood insurance take-up rises after a fire and then falls from about 50% to 22% within a few years, which leaves most post-fire flood damage uninsured (Edgeley et al., 2024).

Provider Box – Image Recognition and Fires: the Kayrros approach

Kayrros covers about ten physical hazards at different levels of maturity, and wildfire is by far the most advanced. The approach is built on machine learning models trained on satellite imagery. The models learn from observed environmental conditions, such as vegetation state, drought, precipitation and slope, and they do not use stochastic event catalogues.

This approach keeps the explanation of each score visible at the level of the individual asset.

Data. ESA Sentinel-2 imagery is the main source, and more than 9 million images have been processed. Imagery from private providers (Airbus, Maxar, Planet) reaches a resolution of 50 cm and is used for calibration and specific local cases. Radar and hyperspectral satellites and public meteorological datasets complete the inputs. The historical baseline relies on about a decade of well-exploitable imagery since 2015, and future projections follow IPCC scenarios.

Before the fire. Risk maps are produced at 50 m resolution and refreshed monthly.

During the fire. Sentinel-2 provides full global coverage every 3 to 5 days. In France, geostationary satellites refresh active fire detection about every 15 minutes, and real-time detection reaches an average accuracy of about 70%, which varies with fire size.

After the fire. Damage assessment maps at 10 m resolution are delivered within days of containment.

Results are computed at the asset level and aggregated to company and portfolio level. The original clients were commodity traders, followed by insurers and more recently asset managers, banks and infrastructure funds .

How the images are read. Raw satellite images go through a processing chain before any risk model uses them, developed in partnership with the Centre Borelli image processing laboratory of ENS Paris-Saclay. The first stage produces analysis-ready data. Clouds are detected and masked, atmospheric effects are corrected, and images from different dates and satellites are co-registered so that each pixel covers the same ground location over time.

Images from several constellations, such as Sentinel-2 and Landsat-8, are then brought to a common radiometric scale. Reference images are selected automatically, and stable ground features serve as anchors, so that a change in a pixel's value reflects a change on the ground and not a change of sensor or lighting.

The second stage extracts hazard-specific indicators from these harmonized time series. Wildfire relies on vegetation indices that measure plant density and health and on thermal anomalies that flag active fires. Flood relies on radar imagery combined with terrain elevation models, which works through cloud cover, to estimate water extent and height.

Object detection methods can identify features as small as two pixels across in medium-resolution imagery (Kayrros, 2020, 2026). Machine learning models are finally trained on these indicators and calibrated against observed fire perimeters and very high resolution imagery.

Three limitations are worth noticing. First, the modelling chain stops at physical risk scores and activity indicators, so translating the scores into financial losses such as average annual loss or revenue loss remains the user's task. Second, tipping points and adaptation measures are not modelled, and NGFS scenarios are not used. Third, hail is not covered.

The trade-off of the satellite approach mirrors its strength. The current conditions are observed directly, which keeps exposure estimates up to date between annual model updates. The short historical record, however, limits the estimation of return periods for extreme fire seasons.

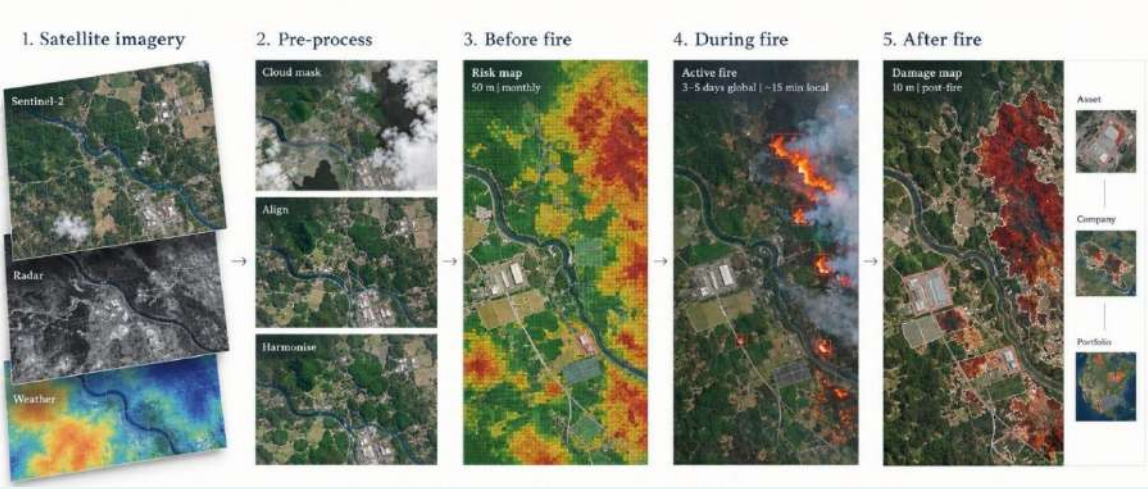


FIGURE 72. Satellite-based wildfire monitoring workflow

3.1.3 SPATIAL RESOLUTION

A good way to discriminate between models is to compare their covered area. Looking at the whole world sometimes mean you do not focus resources on assessing the most "valuable" sites (i.e. you waste resources going global when you should focus on one specific area).

To understand the tradeoff between physical resolution and impact assessment, we can focus on the grid size and its location. There is a wide range on providers focused on one region of the globe or even one country.

The resolution-hazard relationship? Are smaller grids better?

Hazards that vary sharply over short distances require fine-grained data. This is very intuitive. Flood depth can differ by more than a metre between two adjacent buildings depending on micro-topography. Wildfire spread depends on vegetation type and slope at the scale of tens of metres, and landslide susceptibility is governed by local soil and slope conditions. For these hazards, a resolution coarser than approximately 10 to 100 metres misclassifies individual assets as exposed or unexposed, which is generally unacceptable for asset-level underwriting or investment decisions. The cost of imprecise location is heavy. Replacing asset-level locations with company headquarters underestimates portfolio value at risk by up to 71% (Bressan et al., 2024), and misplacing buildings geographically in a catastrophe model can halve estimated losses (Born et al., 2023).

Hazards that vary smoothly over larger areas tolerate coarser resolution. Heat wave intensity, drought severity and the long-run sea level rise trend change gradually across tens to hundreds of kilometres. A global reanalysis grid at about 31 km, or a climate model grid at 25 to 100 km, therefore captures the relevant signal for portfolio-level assessment without material loss of information (Eyring et al., 2016; Hersbach et al., 2020). Tropical cyclone wind fields and coastal storm surge fall between these two extremes. They require an intermediate resolution of typically 1 to 5 km to resolve the interaction between the storm and the coastline, since tropical cyclone intensities are significantly under-represented in global reanalyses, with clear improvements in more recent, higher resolution products (Hodges et al., 2017).

Finer resolution also carries costs. Processing time and storage grow rapidly as grid cells shrink. A fine output grid can also convey false precision when the underlying climate signal comes from a coarser model and has simply been

downscaled. The appropriate resolution is therefore the coarsest one that still separates exposed from unexposed assets for the hazard and the decision at hand.

HAZARD TYPE	TYPICAL REQUIRED RESOLUTION	REPRESENTATIVE GLOBAL DATASET RESOLUTION
Flood (river / pluvial)	10-100m	Global flood hazard models: 30 m–1 km (Wing et al., 2024; Dottori et al., 2016)
Wildfire	10-30m (vegetation-dependent)	NASA VIIRS active fire: 375 m; MODIS burned area: 500 m (Schroeder et al., 2014; Giglio et al., 2018). ESA Sentinel-2 optical imagery: 10–20 m (Drusch et al., 2012)
Tropical cyclone / storm surge	1-5km	ECMWF ERA5: about 31 km (Hersbach et al., 2020)
Earthquake ground motion	100m-1km (soil-dependent)	Global seismic hazard map on a regular global grid (Pagani et al., 2020); Global site amplification proxy at 30 arc-seconds, about 1 km (Allen & Wald, 2007)
Heat wave / drought	10-50km	ERA5: about 31 km; CMIP6: 25–100 km (Eyring et al., 2016; Hersbach et al., 2020)

TABLE 40. Required spatial resolution vs. resolution of representative global datasets

The resolution-coverage-cost trade-off

As noted in Section 3.1.1, the resolution achievable with physics-based hazard data is bounded by computational cost, which creates a structural trade-off between resolution and geographic coverage. Global datasets such as ERA5 operate at about 31 km (Hersbach et al., 2020).

Global flood hazard models used to reach about 90 m at best (Sampson et al., 2015), however, they now reach about 30 m, and 10 m where high-resolution terrain data exists (Wing et al., 2024). River flood mapping at around 10 m requires hydrodynamic modelling that remains tractable for an individual catchment but not for a global portfolio, and global flood models are designed for large-scale screening without the precision needed for local decisions (Ward et al., 2015).

If you manage a globally diversified portfolio, you are facing a choice between a uniform coarse-resolution dataset everywhere and accepting that asset-level risk is systematically smoothed. Or you can also commission high-resolution modelling for a subset of material or concentrated exposures and accept that coverage will be incomplete.

This trade-off is compounded by an availability asymmetry. High-resolution hazard products are disproportionately available in high-income geographies with dense observational networks and mature national flood or wildfire agencies, such as the United Kingdom, the Netherlands and the United States. Observations are sparse in Africa and parts of South America, and global reanalysis products agree less on extreme precipitation in these regions (IPCC, Seneviratne et al., 2021).

Global flood models also diverge sharply where local data is lacking. Across Africa, six global flood models agree on flood extent only about 30% to 40% of the time (Trigg et al., 2016). Institutions operating globally should therefore treat resolution as a data quality dimension that varies by geography, and assess it region by region within each vendor dataset.

The appropriate resolution also depends on the unit of analysis. For a single material asset, such as a manufacturing plant that represents a large share of an insured value or a loan exposure, asset-level resolution is generally warranted regardless of hazard type, because a misclassification at that scale carries direct financial consequences. Replacing asset-level locations with company headquarters underestimates portfolio value at risk by up to 71% (Bressan et al., 2024). For a large, granular retail portfolio, an averaged or statistically representative resolution is often sufficient, because idiosyncratic error at the level of any single address is diversified away across the portfolio. Concentrated exposure, whether by value or by geography, is the condition under which coarse resolution data is most likely to understate true risk, since spatially correlated losses materially increase portfolio tail risk (Kerkhofs et al., 2025).

Provider Box – Hypermeteo: a practical use of downscaling

Hypermeteo illustrates the trade-off between resolution and coverage. The company comes from the meteorological firm Radarmeteo and supplies physical hazard data only. Its service, Climate Sentinel Scenario, covers about 20 hazards aligned with the EU Taxonomy over Italy, a choice of depth over breadth. Flood and landslide maps are produced at 100 m resolution, and the other hazards at 5 km, finer than most global providers offer for hazards other than flood.

Its inputs are regional simulations from a nine-member multi-model ensemble of the EURO-CORDEX programme. The ensemble combines five global climate models (CNRM-CM5, EC-EARTH, HadGEM2-ES, MPI-ESM-LR, NorESM1-M) with five regional climate models (ALADIN63, RCA4, REMO2009, REMO2015, HIRHAM5), at a starting resolution of 11 km.

Hypermeteo bias-corrects and statistically downscales these simulations in house up to 2099, under RCP 4.5 and RCP 8.5. The method is quantile mapping, with a variant selected for each variable. Quantile delta mapping is used for precipitation, wind speed and gusts, which preserves the projected change in each quantile, as discussed above (Cannon et al., 2015). Detrended quantile mapping is used for temperature. Hypermeteo is also exploring downscaling approaches based on deep learning.

The model then bias-corrects them and statistically downscales them in house up to 2099, under RCP 4.5 and RCP 8.5. The provider does not specify which statistical method it applies.

The downscaled fields are validated against in situ measurements and reanalysis datasets, with the root mean squared error, the mean absolute error, the bias and contingency tables.

Hypermeteo produces no damage estimate and no financial metric. Its clients, mainly infrastructure managers, insurers, infrastructure asset managers and banks, feed the hazard layer into their own models, and the company is looking for a financial-impact partner to complete the chain. For a portfolio concentrated in Italy, such a regional layer can replace the coarse global hazard grids used by most platforms. For an international portfolio, it covers only part of the exposure.

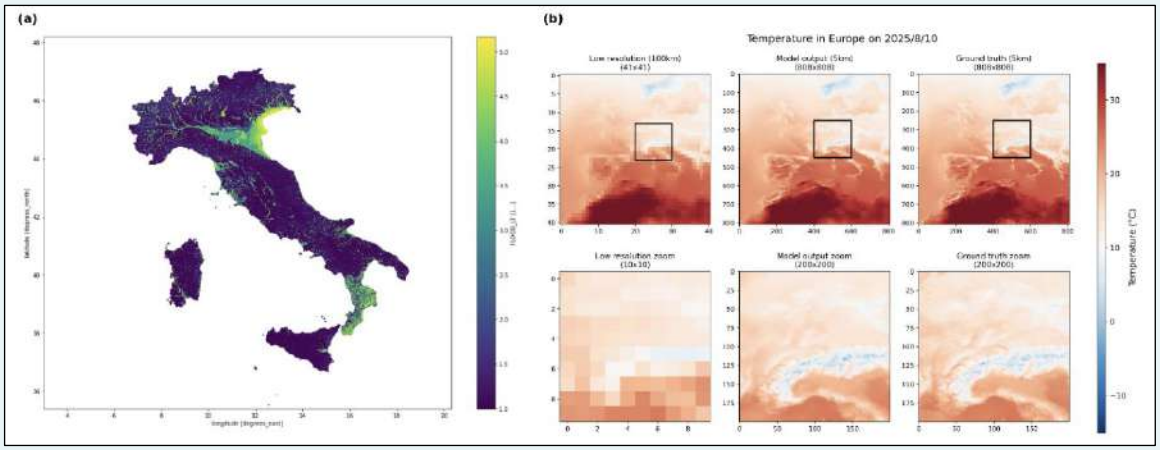


FIGURE 73. (a) Flood hazard classes across Italy; (b) downscaling of a European temperature field from 100 km to 5 km resolution
Source: Hypermeteo

Provider Box – Intensel – an Asia-focused model

Intensel shows how much a damage estimate gains from a detailed description of each building. The provider runs its own global hazard models, with a strong presence in Asia and Southeast Asia. It covers about 18 hazards, and its flood maps reach 0.5 to 5 m resolution at the asset. It keeps rainfall, river and coastal flooding as separate hazards, where many platforms merge them into a single flood layer.

Each asset is described by a two-dimensional footprint and a three-dimensional valuation, so the model knows which floors and contents lie below the projected water level. Vulnerability then draws on more than 300 damage curves, selected by asset type, use and construction features such as the presence of a basement. The model also accounts for the surrounding environment, such as the distance to drainage infrastructure, and for adaptation measures such as flood walls or raised floors.

Several machine learning and deep learning models support this chain, and satellite imagery helps map the local adaptation factors. Outputs are expressed in dollars per asset, as an average annual loss and as losses for return periods such as the 100-, 200- and 500-year events.

Intensel also reports operational loss, the number of days of disruption per hazard and location, and an adjustment of credit spread and loss given default for lenders. Results come under six scenarios with five return periods each and are refreshed every month, with annual projections available on request.

Flood outputs are validated against observed events, with an accuracy of about 85% measured with metrics that include the root mean squared error. This validation has not yet been published in an academic journal. Intensel also makes a point every user should keep in mind.

The final financial figure depends on client parameters, such as asset values, daily revenue, inventory and insurance cover, so it cannot be standardized across clients. The more the client documents its own assets, the closer the estimate gets to its real exposure.

Provider Box – Screening at basin scale: the WWF Risk Filter Suite

The WWF Risk Filter Suite is a free, web-based screening platform for companies and financial institutions. It brings together the Water Risk Filter, the Biodiversity Risk Filter and a Place-Based Impact Estimator. A user uploads the locations of its operations or suppliers, and the tool returns a risk profile for each site.

The Water Risk Filter is organized in four levels. It has 3 risk types (physical, regulatory and reputational), 17 risk categories and 67 indicators. Physical risk covers four categories, which are water scarcity, flooding, water quality and the status of ecosystem services. Each indicator is scored on an integer scale from 1 (very low risk) to 5 (very high risk). The scores are then aggregated with default weights for 25 industries (WWF, 2022).

The suite has no grid resolution in the usual sense. Its spatial unit is the river basin, drawn mainly from HydroBASINS Level 7 polygons, and some indicators use country boundaries. Every site inside the same basin receives the same basin score. Two plants a few kilometres apart in the same catchment therefore look identical, whatever their elevation or distance to the river. The basin score can be adjusted with an operational questionnaire on the water use of each site.



FIGURE 74. WWF Water Risk Filter interface showing basin in northern China and the Korean Peninsula
Source: WWF Water Risk Filter, 2026

The forward-looking layer covers 2030 and 2050 under three pathways that combine RCP and SSP scenarios. These projections exist only at category level, and WWF is phasing them out. The methodology describes average, recent conditions, and some of the underlying datasets are old (for example). The tool also stops at a risk score and does not translate it into financial loss.

It can be useful for a first screening of a large supply chain or portfolio, especially for water-related hazards. Please note that it cannot support an asset-level loss estimate.

INDICATOR	DATASET	PERIOD COVERED	UPDATE STATUS
Aridity Index	Global Aridity Index v2 (Trabucco & Zomer, 2019)	1970–2000	Infrequent
Water Depletion	WaterGAP3 (Brauman et al., 2016)	1971–2000	Infrequent
Blue Water Scarcity	Mekonnen & Hoekstra (2016)	1996–2005	No updating
AWARE	WULCA (Boulay et al., 2018)	2010	No updating
Projected change in drought	ISIMIP2b (Frieler et al., 2017)	1661–1860 vs. about 2050	No updating

Projected change in flood	ISIMIP2b with CaMa-Flood (Frieler et al., 2017)	1661–1860 vs. about 2050	No updating
Surface water quality	World Bank (Damania et al., 2019)	Not stated	No updating
River fragmentation	Grill et al. (2019)	Not stated	No updating
Freshwater biodiversity	Tedesco et al. (2013)	Projection to about 2090 (SRES A2)	No updating

TABLE 41. Table of datasets used in the WWF Risk Filter suite

Provider Box – Five metres and ten thousand years: JBA flood maps

The challengers to Fathom flood maps are scarce, and JBA is one of the few challengers credible on that field. JBA Risk Management is a UK flood specialist with an engineering and public sector background. Flood accounts for about 98% of its activity, across fluvial, pluvial and coastal flooding. JBA also provides wind data, mainly because wind fields drive coastal surge, and it offers a UK subsidence add-on from a partner.

The hazard layer comes from JFLOW, a two-dimensional hydraulic engine that solves the shallow water equations. Its inputs are gauging station records for rivers and rainfall, ERA5 reanalysis, licensed terrain data and coastal bathymetry. The resolution reaches 5 m in the UK, France, much of continental Europe and parts of Australia and the United States. Elsewhere it is 30 m, and it is never coarser. The local terrain data sets the limit.

JBA delivers two products. The first one is **static hazard maps** that give flood depth at fixed return periods, from 20 to 1,500 years in France.

A 10,000-year stochastic event catalogue captures the spatial correlation between flood events across a portfolio. Climate change and tipping points are kept outside this catalogue and treated through bespoke storylines based on RCP, SSP or warming levels. In JBA's view, adding them to the catalogue would break its internal logic.

Ground-up loss comes from depth-damage functions built on the UK Multi-Coloured Manual and USACE data, applied to the sum insured. Insured and reinsured losses are computed in the open-source Oasis Loss Modelling Framework, which returns average annual loss, return period losses, VaR and TVaR. Banks generally ask for flood scores such as the Floodability Index, since they already run their own probability of default engines.

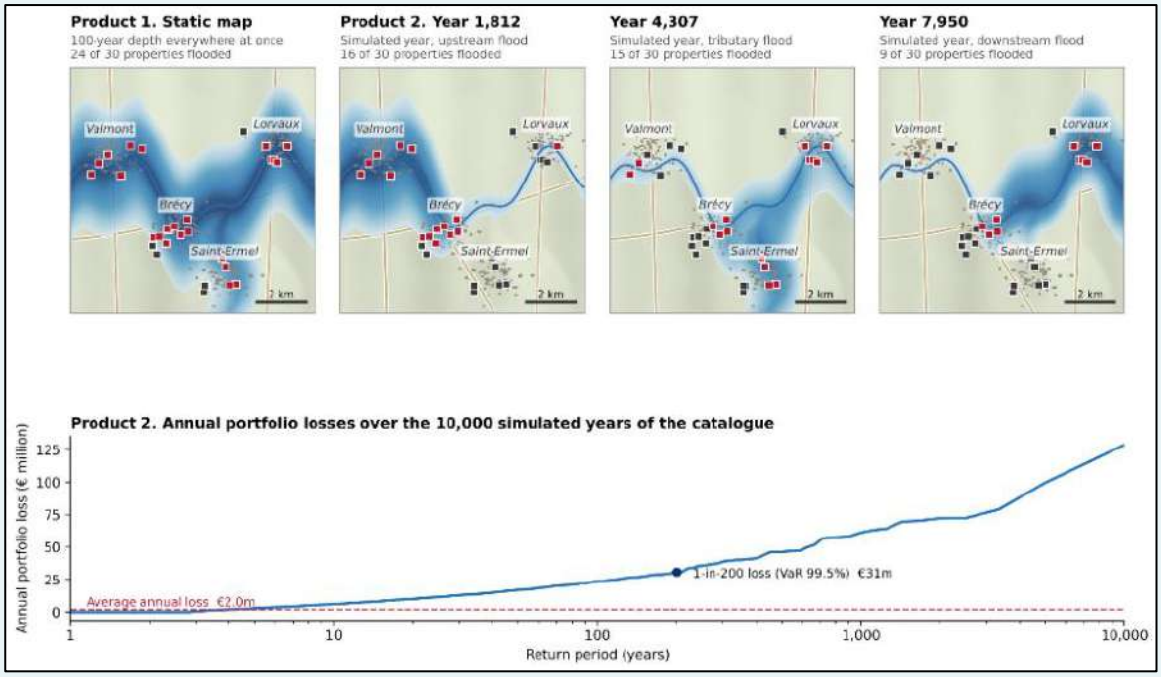


FIGURE 75. Static flood hazard map compared with a stochastic event catalogue for a portfolio of 30 properties Illustration example.

Validation against claims is occasional and mostly qualitative, because GDPR restricts the sharing of loss data. Client portfolios are also often coarser than the model, for instance at canton level, and JBA uses a disaggregation module to avoid false precision. JBA shows that the finest hazard map is only as useful as the location data of the portfolio it is applied to.

Provider Box – A digital twin of every building: First Street (MSCI)

First Street is a US provider of physics-based climate risk data. MSCI announced its acquisition in June 2026 and completed it in August 2026 for about USD 123 million, net of cash acquired. First Street's models now feed MSCI's physical risk solutions. Its digital twin covers about 2.5 billion structures worldwide, and MSCI's geospatial base layer attributes each asset to its owning or operating company.

The hazard suite is expanding, and MSCI plans to reach more than 30 hazards with global coverage by April 2027. As of October 2026, it includes river and coastal flooding, wildfire, wind (tropical cyclones, extratropical storms and severe convective storms), hail, drought and extreme temperatures. Flood, wildfire and wind come from physics-based models. Drought and extreme temperatures come from statistical models, and hail, tornadoes and cyclones combine both. Inputs combine public agencies (NOAA, NASA, USGS, FEMA, USDA, USFS and Copernicus) with commercial sources.

Resolution varies by hazard, depending on data inputs, model complexity and commercial needs. It is 10 m for all flood layers, 250 m for global subsidence, 300 m for global wildfire and about 6 km for wind, heat and cold. Projections follow CMIP6 SSP1-2.6, SSP2-4.5 and SSP5-8.5 at asset level, with horizons up to 2100. Additional scenarios are planned for April 2027.

Vulnerability rests on a digital twin of structures, which records building type, materials and foundation. Component-level engineering damage functions convert hazard into damage and downtime. Users can add site-level adaptation measures to estimate the avoided losses they bring. Outputs include average annual loss

and downtime per asset, reported at the 10th, 50th and 90th percentiles. A company module links firms to the footprint of their operations and value chain.

The provider validates its models against observed events. Its wildfire model reaches over 90% accuracy in California, and no building rated as low risk was exposed during the 2025 Los Angeles fires. The methods are published in peer-reviewed journals. Tipping points are not modelled explicitly. Overall, the combined offering pairs fine-resolution hazard modelling with an engineering vulnerability layer and an asset-to-issuer mapping. This makes it one of the most complete hazard-to-loss chains in the review.

Each asset in a portfolio is matched to a building-specific model that records its footprint, height and number of stories, construction materials, and foundation type or first-floor elevation. Engineering damage functions combine these attributes with the local hazard to give average annual loss (AAL) and expected downtime for each asset.



FIGURE 76. From localization to digital twin
Illustrative example.

Provider Box – Climate Risk Services – Strong models for assets located in developing countries

Climate Risk Services works mainly for development finance institutions, impact funds and commercial banks in emerging and frontier markets, although its platform covers the whole world. The method relies almost entirely on open-source data, and the client only provides location, sector and financial exposure. The current tool is a scoring framework combining hazard, exposure, sector sensitivity and adaptive capacity. It produces a gross risk score, an adaptation maturity assessment and a net risk score, but no resilience score.

Resolution reaches about 1 km when coordinates are available and falls back to province or country level otherwise. The tool covers a 1981–2010 baseline and three future periods (2010–2040, 2040–2070 and 2070–2100) under RCP 4.5 and RCP 8.5. Earthquakes and hail are not covered, tipping points are not modelled, and transition risk is assessed alongside physical risk.

With FMO, the Dutch entrepreneurial development bank, CRS compared its sector sensitivity scores with the non-performing loans observed in the portfolio. The comparison remains qualitative, without a single accuracy metric. An external consultant also reviewed the physical risk methodology during this partnership. The methodology is not open access, but every platform client can download it.

The provider is moving from scores to financial losses. Extreme value distributions are fitted to hazard data to build hazard curves. These are combined with damage curves drawn from frameworks such as those of FEMA and the EU Joint Research Centre and verified with development finance partners. For each hazard and asset type, several candidate curves are fitted and the one with the smallest error is kept. Uncertainty is expressed through model ensembles, since the hazard data comes from a multi-model average. The damage function and financial translation modules are still under development.

Provider Box – Gaia Vision – A Japanese-focused model

Gaia Vision is a Tokyo start-up founded in September 2021 by researchers from the University of Tokyo. Its technical adviser is Dai Yamazaki, an associate professor at the university's Institute of Industrial Science, who developed the global river flood model CaMa-Flood (Yamazaki et al., 2011) and the MERIT global terrain and hydrography datasets. The company runs its own global flood simulations with river flood modelling technology developed at the University of Tokyo, combined with climate projections.

Its Climate Vision platform turns these simulations into flood maps at 30 m resolution in Japan and 90 m elsewhere, for the present climate and for global warming levels of 1.5°C, 2°C and 4°C. Users register their sites and obtain the current and future flood risk for each one, together with a financial assessment based on the asset value of the site. The platform serves corporates, financial institutions and local governments, mainly for scenario analysis, TCFD disclosure and the choice of new sites. Gaia Vision describes these maps as suited to an initial screening of a portfolio.

The same hydrological expertise feeds two operational services. Water Vision, developed with JAXA's Today's Earth system, forecasts the extent and depth of river flooding up to 36 hours ahead for evacuation decisions and business continuity planning. A second tool simulates how levees and dams would reduce flood damage, including under climate change. Gaia Vision thus links weather forecasting, flood defence engineering and long-term climate risk, using the same expertise in river hydrology for a 36-hour alert and for a 2°C scenario.

Provider Box – Scientific Climate Ratings (SCR) – A polygon approach

We mentioned Scientific Climate Ratings (SCR) in section 2.4. It is an independent climate rating agency within the EDHEC Business School venture ecosystem, working with the EDHEC Climate Institute. It has submitted its application for registration to ESMA. SCR starts from a familiar premise: a climate score has little value for an investor until it is translated into financial terms. Its physical risk module for real assets and corporates therefore feeds a valuation model, which returns the impact on enterprise value, equity, debt and probability of default.

SCR starts from the shape of the asset, using precise asset boundaries. Physical risk is local, and two assets in the same street can face very different hazards. Each asset is drawn as a set of points, lines or polygons, and large assets are split into sub-elements. **An airport, for example, separates runways from buildings, since each component has its own damage function and its own value**. A point with a 500 m radius, the usual fallback when the exact location is unknown, produced a 100-year flood damage two to three times lower than the precise footprint.

The modelling chain then follows six steps.

STEP	CONTENT
1. Location	Identify where the asset sits
2. Geometry	Draw points, lines or polygons for the asset and its sub-elements
3. Hazard	Overlay historical and projected hazard maps for several return periods (flood depth, wind speed for tropical and extratropical cyclones, wildfire probability and burnt area)
4. Vulnerability	Select a damage function by asset type, from 0 (no damage) to 1 (total destruction)
5. Loss	Convert intensities over the footprint into damage, aggregate them at asset level, and weight each return period by its annual probability to obtain an average annual loss
6. Adaptation	Reduce damage according to the resilience measures in place

TABLE 42. The 6 SCR modelling steps

Step 5 feeds two ratings. The Climate Exposure Rating stops at the average annual loss. The Climate Risk Rating propagates climate shocks to cash flows, enterprise value and probability of default.

Scientific Climate Ratings start from the shape of the asset, using precise asset boundaries. Physical risk is local, and two assets in the same street can face very different hazards. Each asset is drawn as a set of points, lines or polygons, and large assets are split into sub-elements. An airport, for example, separates runways from buildings, since each component has its own damage function and its own value. Scientific Climate Ratings illustrated the stakes with a site in the United Kingdom. A point with a 500 m radius, the usual fallback when the exact location is unknown, produced a 100-year flood-damage two to three times lower than the precise footprint.

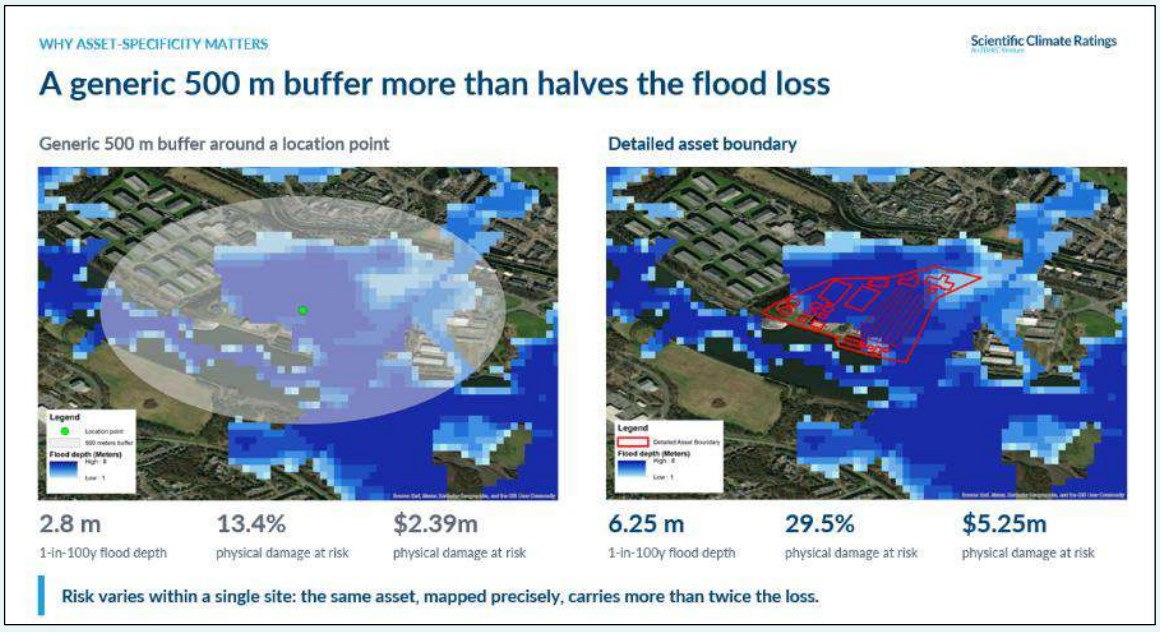


FIGURE 77. Why Asset Specific Boundaries Matter

The way the footprint meets the grid depends on the resolution of each hazard map. Flood depth comes from Fathom maps at 30 m, wildfire burnt area from MODIS at 500 m, tropical cyclone wind speed at 0.08°, and extratropical storm gusts and heat stress at 0.25°. At the flood and wildfire resolutions, the polygon simply covers a set of pixels. At coarser resolutions, a long asset such as a road may cross several cells, and each cell is weighted by the share of the asset it contains. Flood is the only hazard mapped at metre scale, and most other climate variables sit at about 25 km. The polygon approach therefore moves the resolution question from the hazard grid to the geometry of the asset, and its precision depends on the quality of the location data.

Damage functions come from peer-reviewed sources, mainly FEMA Hazus, DamageScanner and the GEM Global Vulnerability Model. They cover 94 infrastructure asset classes with documented flood and wind curves, and fragility curves are assigned to every TICCS infrastructure class and NACE sector for flood, wind, earthquake and heat. Financial translation is narrower. The ClimateMetrics platform assesses seven hazards at indicator level (flood, wind, wildfire, heat, precipitation, drought and landslide), of which five are quantified financially. Drought is measured as a hazard, and its financial translation is in development. Damage covers structure, contents and inventory, and business interruption through conversion factors from the literature and FEMA.

SCR relies on Scientific Climate Scenarios, which build on the NGFS set with three additions. Economic variables are projected at sector and country level. The warming range is extended beyond the 1.3°C to 2.9°C of NGFS with two scenarios developed by the EDHEC Climate Institute, Climate Destabilisation (3.6°C) and Climate Breakdown (4.4°C).

Each scenario receives a probability of occurrence. Hazards follow the matching SSP-RCP pathway: SSP1-2.6 for Net Zero 2050, Below 2°C and Delayed Transition; SSP2-4.5 for NDC, Fragmented World and Current Policies; SSP3-7.0 for Climate Destabilisation; and SSP5-8.5 for Climate Breakdown. A bank running a stress test would use the impact under one scenario. An investor allocating capital would use the probability-weighted expected value, a move from “what if” to “what’s likely”. The expected impact is finally mapped to a relative scale from A to G, based on where each asset falls in the distribution of the rated universe.

Adaptation draws on ClimaTech, an EDHEC Climate Institute database of more than 1,800 technological applications of adaptation and mitigation, compiled from more than 200 sources. SCR ran the quantitative analysis of a June 2026 report by IFC and AXA Climate on three Brazilian infrastructure assets. For the Mato Grosso transmission line, highly exposed to wildfire, climate risk reduced net asset value by more than 30% under a high-emission scenario. Adaptation measures with a benefit-to-cost ratio above 1 cut this loss by 25 percentage points, or 22 points net of their cost, roughly 8 dollars of asset value protected per dollar invested.

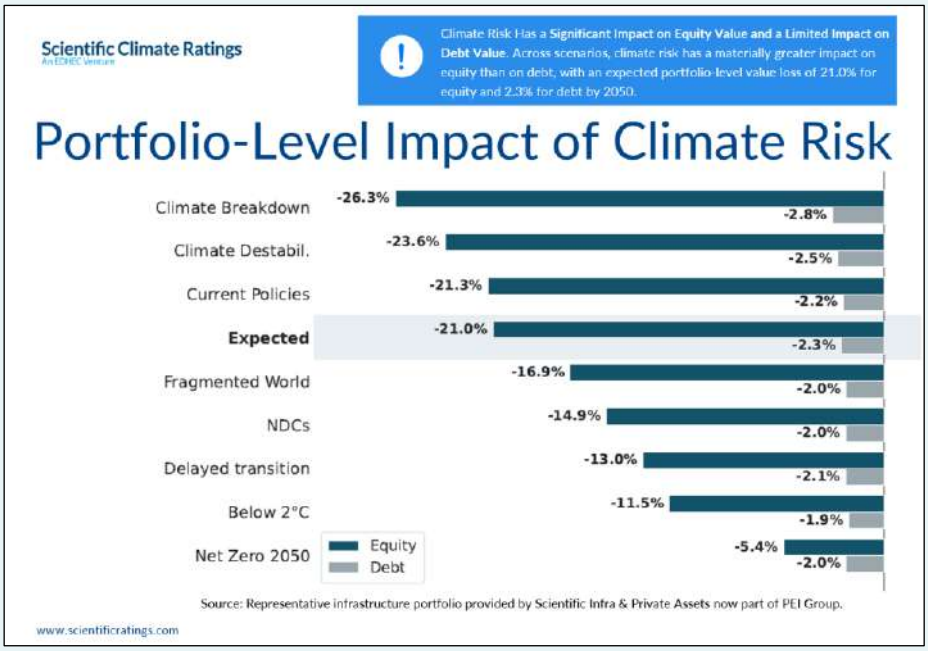


FIGURE 78. Climate Risk Hits Equity Broadly, Debt Narrowly
Source: SCR

On a sample of about 443 infrastructure firms, the equity impact by 2050 is 21.0% under the expected scenario, from 5.4% in Net Zero 2050 to 26.3% in Climate Breakdown.

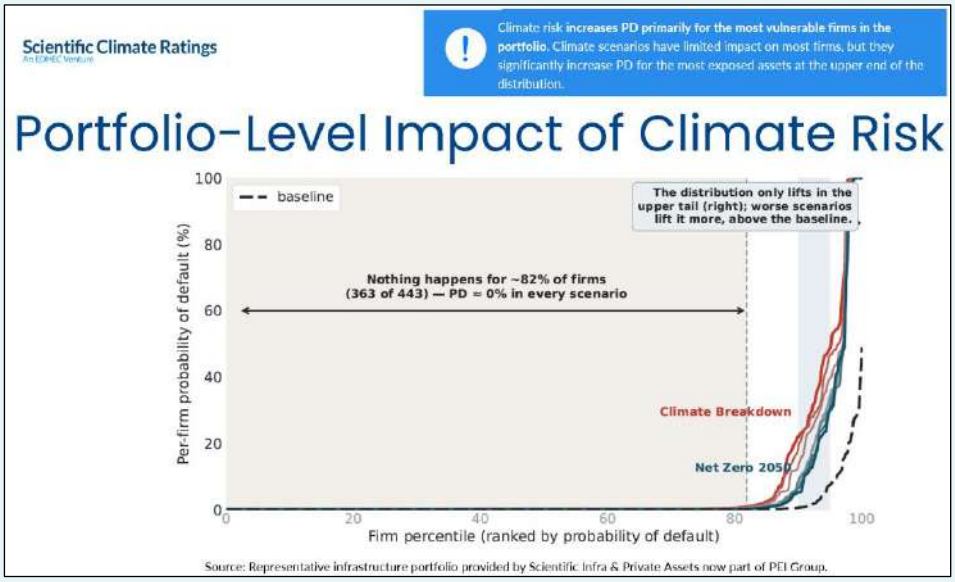


FIGURE 79. For Debt, the Risk Is Concentrated in a Small Set of Exposed Firms

The expected impact on debt is only 2.3%, and about 82% of the firms see no increase in probability of default in any scenario. Floods cause the heaviest damage per asset, while storms affect the largest number of assets. Disruption outside the footprint remains hard to capture. Brisbane airport raised its runways and built a flood wall, yet the roads leading to it remain exposed. SCR partly addresses this through regional chronic impacts on sector productivity, which also capture supply chain effects indirectly.

Coverage includes more than 6,000 infrastructure assets, 191 sovereigns and more than 3,400 sub-national regions, with about 4,000 listed equities announced for 2026. Each company is treated as a portfolio of production sites weighted by their share of revenue. Site locations come from sustainability reports, regulatory

filings and other public sources, and rated companies are invited to share more precise data, which can replace model estimates once reviewed. The full methodology is publicly available (Scientific Climate Ratings, 2025a, 2025b).

About resolutions

Flood is the only hazard mapped at metre scale (30 m). Wildfire inputs sit at 100 to 500 m. Every other climate variable sits at roughly 9 to 44 km, with most at 0.25° (about 25 to 28 km).

Full methodology documentation is publicly available, supporting transparency and independent scrutiny of the ratings. Details on resolutions are shared below.

HAZARD	INDICATOR / DATASET	SPATIAL RESOLUTION
Flood (fluvial, pluvial, coastal)	Flood depth, Fathom global flood maps (2024), return periods 5 to 1,000 years	Global 30 m × 30 m
Tropical cyclone (historical)	Max wind speed, STORM synthetic catalogue (Bloemendaal et al., 2020) based on IbTracks NOAA	Global 0.08° x 0.08°
Tropical cyclone (projections)	Max wind speed, CHAZ global wind-hazard dataset	Global 0.08° x 0.08°
Extratropical storm (historical)	Wind gust speed, ERA5 reanalysis (ECMWF)	Global 0.25° × 0.25°
Extratropical storm (projections)	Wind gust speed, NASA NEX-GDDP CMIP6	Global 0.25° × 0.25°
Wildfire	Burnt date, NASA MODIS MCD64A1	Global 500 m × 500 m
Wildfire	Tree cover and built-up cover density, Copernicus Global Land Service	Global 100 m × 100 m
Wildfire	Fire Weather Index, NASA NEX-GDDP-CMIP6	Global 25 km × 25 km
Heat stress	WBGT, NASA NEX-GDDP-CMIP6	Global 0.25° × 0.25°
Extreme precipitation	1-day and 5-day annual max precipitation, NEX-GDDP-CMIP6	Global 0.25° × 0.25°
Drought (water scarcity)	SPI-6 drought duration and magnitude, NEX-GDDP-CMIP6	Global 0.25° × 0.25°

	Water stress from Aqueduct	
Earthquake	Peak Ground Acceleration (PGA) with a 10% probability of being exceeded in 50 years, GEM global hazard map	Global 0.05° x 0.05°
Landslide	Rainfall-triggered landslide susceptibility, ARUP Global Landslide Hazard Map	Global 100m x 100m
Subsidence	Subsidence rate (in cm/year), from Herrera-García et al. 2021 (Science)	Global 0.02° x 0.02°

TABLE 43. Spatial resolutions used by SCR

Comparing resolutions

As a French saying goes, comparison is not reason. The figures below should not be read as a ranking. The four providers do not model the same hazards, do not serve the same clients and do not even use the same spatial unit. A basin polygon and a 5 m grid cell answer different questions. A "finest resolution" also applies only where high-quality terrain data exists, usually in Western Europe, North America and parts of Australia. Still, placing the providers side by side shows how each one has chosen to handle the same challenge, which is to represent a hazard at the scale where it causes damage.

PROVIDER	SPATIAL UNIT	FINEST RESOLUTION	FINANCIAL OUTPUT
WWF Risk Filter Suite	River basin (HydroBASINS L7)	Basin polygon	Risk score (1 to 5)
Fathom	Grid	30 m – (10 m locally)	Ground-up, gross and net loss
JBA	Grid	5 m - 30 m elsewhere	AAL, return period loss, VaR
First Street	Grid and building	10 m - US flood	AAL, downtime
Scientific Climate Ratings	Asset polygon (points, lines, polygons)	Asset footprint, on 30 m flood grids	AAL, impact on enterprise value, equity, debt and PD; A to G rating

TABLE 44. Comparing spatial unit and resolution by provider

The table shows a link between the two dimensions. The finer the spatial unit, the further the modelling chain goes toward a financial loss. A basin score supports a first screening of a supply chain. A 30 m global grid gives an investor a consistent view across countries. A 5 m map with a stochastic catalogue fits the pricing needs of an insurer, and a building-level digital twin brings the analysis down to the single asset.

The right provider therefore depends on the decision at hand and on the quality of the user's own data. A 5 m flood map brings little to a portfolio whose locations are known only at canton or state level. Before choosing, a financial

institution should ask each provider for the resolution actually achieved in the countries where its exposures sit, and check that its own location data can match it.

3.1.4 HISTORICAL VS REAL TIME VS PROJECTION

The choice of temporal reference, whether to use historical observed data or forward-looking climate projections, and over what time horizon, is not a modelling detail. It depends on what the analysis is trying to consider: the risk of the specific decision at hand, over the specific period for which that decision is binding.

Historical baseline versus forward-looking projection

Historical data characterises hazard frequency and severity as observed over the available record, typically 30 to 70 years for physics-based reanalysis products and considerably longer for select long-running actuarial series. Its principal strength, discussed in Section 3.1.1, is that it is directly validated against observed outcomes. Its principal weakness is that it is backward-looking by construction. It cannot represent a hazard distribution that shifts outside historical bounds under a warming climate. For this reason, insurers and reinsurers have coupled climate projections with catastrophe models to estimate how average annual losses and extreme losses could evolve under warming scenarios (Association of British Insurers, 2009). A climate model can also generate long synthetic series for both the current climate and a future climate, which historical records alone cannot provide (Tinard et al., 2016).

Forward-looking projections are generated by climate models run under defined emissions or warming scenarios. The most common are the CMIP6 Shared Socioeconomic Pathways (SSP1-2.6, SSP2-4.5, SSP5-8.5) and their predecessors, the Representative Concentration Pathways (Jewson et al., 2023; Mandel et al., 2025; Schneider, 2025). These projections embed a physically consistent climate change signal that historical data cannot provide. They do so at the cost of two additional sources of uncertainty that historical data does not carry. The first is which emissions pathway will actually be realized. Climate scenarios used in financial regulation and investment generally carry no probability, and estimates that assign them one place the median warming at about 2.7 °C by 2100, with a very low probability of limiting warming to 1.5 °C (Rebonato et al., 2025). The second is how faithfully a given climate model translates that pathway into local hazard intensity. Estimated impacts are sensitive to the choice of global climate model, and a subset of CMIP6 models runs warmer than the scientific consensus (Schneider, 2025).

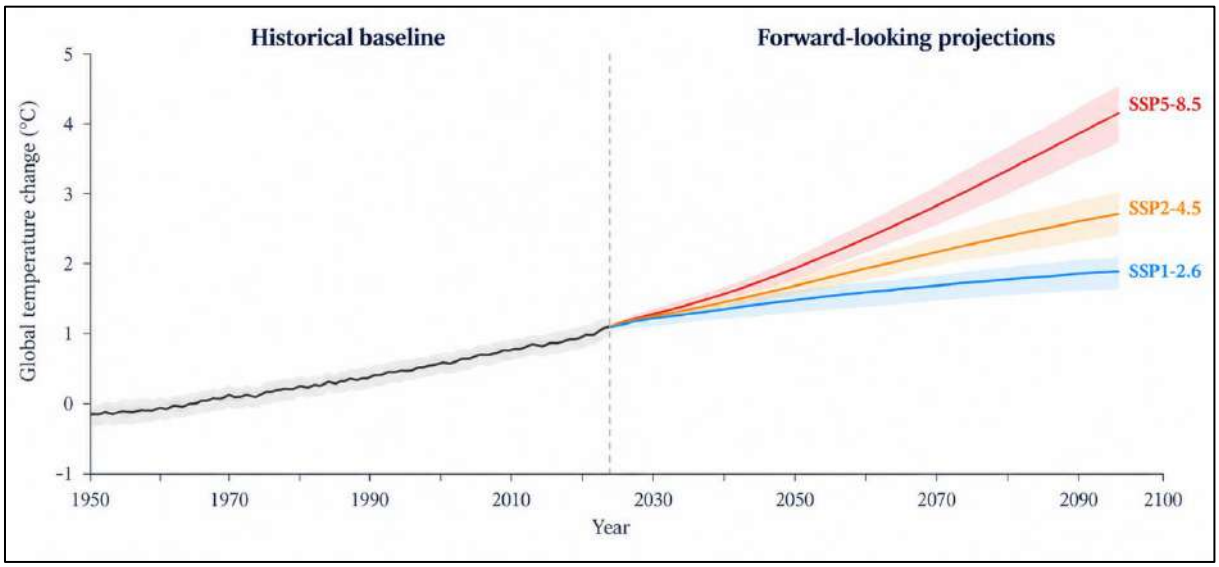


FIGURE 80. Historical global temperature change and projections to 2100 under three CMIP6 scenarios (SSP1-2.6, SSP2-4.5 and SSP5-8.5)

Real time assessments

While historical and projected data both describe a distribution of possible events, real-time data describes the event that is taking place. It comes from satellite imagery, weather station networks, river gauges and short-range forecasts, and it is refreshed at intervals of minutes to days.

A real-time assessment answers a different question. It does not estimate how often an event may occur. It estimates which assets are affected now, and what the event will cost.

Wildfire monitoring illustrates the first use. Kayrros detects active fire hotspots from satellite imagery every 15 minutes, sends alerts when a fire approaches a monitored asset, and produces damage maps at 10 m resolution within days of containment (Kayrros, 2025). One of the firm's biggest clients was Bloomberg. For an asset manager, this information matters before any loss is recorded. A fire close to a listed company's production site can move its share price within hours.

The French natural catastrophe regime illustrates the second use. The Caisse Centrale de Réassurance (CCR), the public reinsurer of this regime, relies on an operational network that includes Météo-France, the Copernicus Emergency Management Service and the national flood forecasting service Vigicrues. CCR's deterministic models combine the observed hazard footprint with its exposure database and produce a first estimate of insured losses about five days after the event, with flood outputs at 25 m resolution (CCR, 2020). Insurers and public authorities use this estimate to set reserves and prepare the response before claims are filed.

Real-time data has limits for portfolio risk modelling. It records events one at a time and carries no information on their frequency, so it cannot replace a historical record or a projection for pricing, capital or long-term strategy. Satellite archives also cover only a few decades, and most near-real-time products are proprietary. In practice, real-time data serves two roles in the models reviewed in this Cookbook. It supports monitoring and loss estimation after an event, and it feeds the observed footprints that later recalibrate historical event sets and vulnerability functions.

How to choose a scenario? The horizon depends on the institutional use case

As highlighted in the PARC Foundation's *Alignment Cookbooks* (PARC, 2024), scenario selection is a key methodological decision. Portfolio alignment methodologies combine current or projected climate performance with carbon budgets and low-carbon pathways, and their results depend materially on the pathway selected.

Different institutional users of physical risk data are exposed over structurally different horizons.

Insurance and reinsurance companies price and reserve against tail events defined over long return periods. Solvency II calibrates the capital requirement to a value-at-risk at a 99.5% confidence level over one year, which corresponds to a 1-in-200-year event (European Parliament & Council of the European Union, 2009). Catastrophe models routinely report probable maximum losses at 250-year and 1,000-year return periods for the same purpose (Born et al., 2023). The underlying policies are annually renewable, yet the capital held against them must reflect the full tail of the hazard distribution. Insurers therefore rely on the longest available historical records and on forward-looking event sets calibrated to multi-century return periods.

Asset managers face a different profile. Holding periods are typically measured in months to a few years, and mark-to-market valuation reacts to changes in perceived risk well before a hazard event materialises. Asset managers are therefore more exposed to volatility in the market's forward assessment of physical risk than to the realized tail event itself. Coastal homes exposed to projected sea level rise already sell at a discount of about 7% compared with equivalent unexposed homes, decades before any inundation (Bernstein et al., 2019). Market reactions to realized events are also short-lived and sector-specific, with insurance stocks showing negative abnormal returns for up to 28 days after major wildfires (Tavor, 2024). The data requirements of asset managers therefore emphasise shorter-horizon, scenario-sensitive projections and frequent re-assessment over long-run tail statistics.

Banks and long-term asset owners sit between these two positions. Mortgage and infrastructure loan tenors of 15 to 30 years require projections aligned with that specific horizon, which is longer than the one-year solvency horizon of an insurer and longer than the holding period of an actively traded fund. Lenders already price this horizon. A premium for sea level rise exposure appears on 30-year mortgages and is absent on 15-year loans, which mature before the risk materialises (Nguyen et al., 2022). Supervisory exercises typically reference a small number of standardized scenario-horizon combinations, most often NGFS scenarios run to 2050 (ACPR, 2021; APRA, 2021; FSB & NGFS, 2022).

INSTITUTION TYPE	RISK HORIZON	DATA EMPHASIS
Insurance / reinsurance	1 year (annual renewal); capital calibrated to 200-500 year tail events	Long historical record and forward-looking stochastic event sets (European Parliament & Council of the European Union, 2009; Born et al., 2023)
Asset managers	Months to a few years (holding period)	Short-horizon, scenario-sensitive re-pricing (Bernstein et al., 2019; Tavor, 2024)
Banks (mortgage / infrastructure lending)	15-30 years (loan tenor)	Medium-to-long horizon climate projections, NGFS scenarios (ACPR, 2021; Nguyen et al., 2022)
Long-term asset owners / infrastructure investors	30+ years	Long-horizon projections under multiple warming scenarios (CISL, 2019)

TABLE 45. Risk horizons and data emphasis of physical climate risk assessment by type of financial institution

Selecting a timeline is therefore inseparable from selecting a scenario. A ten-year projection under SSP1-2.6 and the same ten-year horizon under SSP5-8.5 can imply materially different hazard intensities, and estimated financial impacts increase with scenario severity (Mandel et al., 2025). A horizon specified without its pathway leaves the assessment incomplete. Practitioners should treat the pair, horizon and scenario, as a single design choice, addressed further in the discussion of aggregation choices later in this Part.

FRAMEWORK	SCENARIO	LOGIC	APPROX. LONG-TERM WARMING (RELATIVE TO 1850–1900)
IPCC / CMIP6	SSP1-2.6	Low-emissions pathway	~1.8 °C by 2081–2100
IPCC / CMIP6	SSP2-4.5	Intermediate pathway	~2.7 °C by 2081–2100

FRAMEWORK	SCENARIO	LOGIC	APPROX. LONG-TERM WARMING (RELATIVE TO 1850–1900)
IPCC / CMIP6	SSP5-8.5	Very high-emissions pathway	~4.4 °C by 2081–2100
NGFS	Net Zero 2050	Early, strong transition	~1.4 °C by 2100 (after a peak near 1.6 °C)
NGFS	Below 2°C	Gradual transition	Below 2 °C (67% chance)
NGFS	Current Policies	Only currently implemented policies	~3 °C by 2100
IEA (WEO 2024)	NZE	Net-zero energy emissions by 2050	~1.5 °C by 2100
IEA (WEO 2024)	APS	Announced climate pledges met in full and on time	~1.7 °C by 2100
IEA (WEO 2024)	STEPS	Stated policies, adopted or announced	~2.4 °C by 2100
TRACC (France)	Reference warming trajectory	French adaptation reference, based on current global policies	+4 °C in France by 2100 (~+3 °C global)

TABLE 46. Main climate scenario frameworks used in physical risk assessment

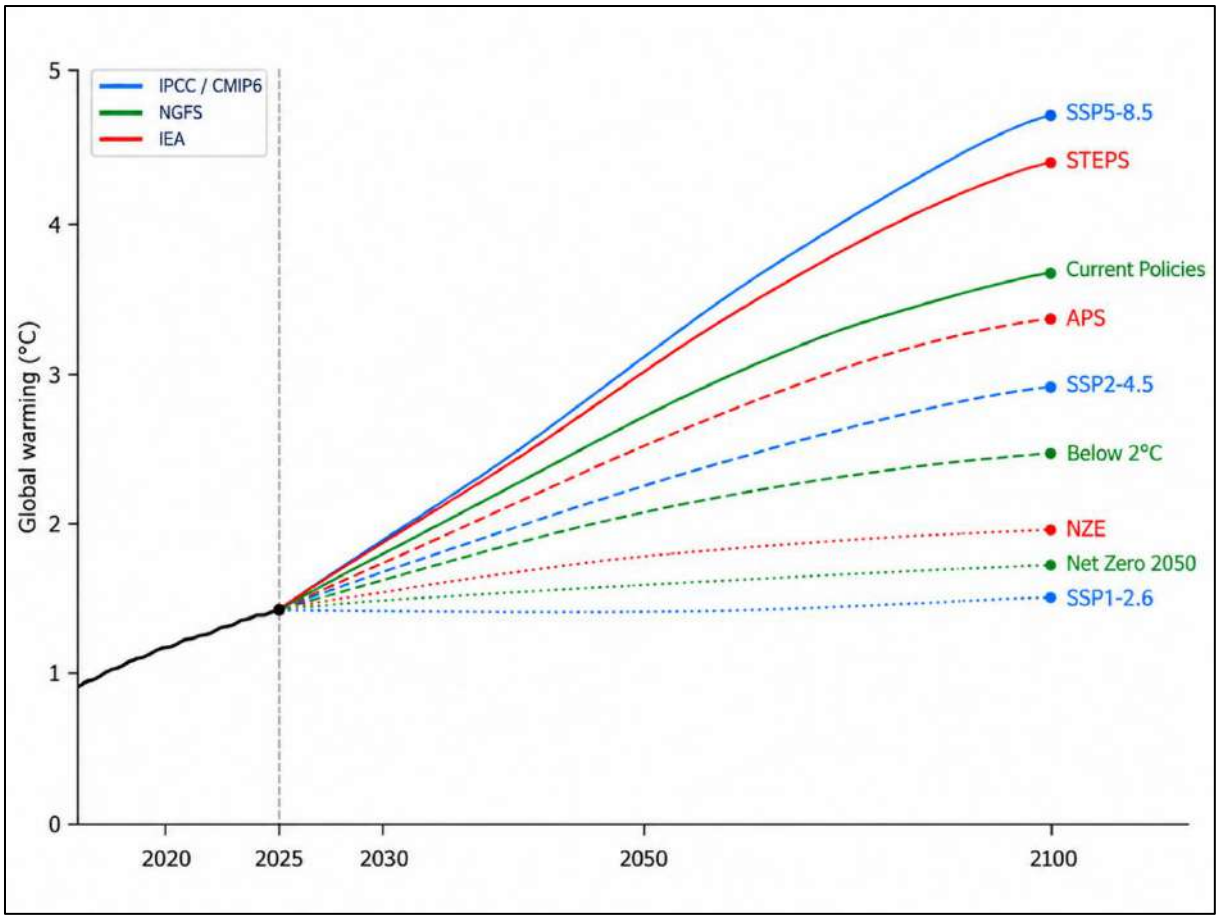


FIGURE 81. Global warming pathways to 2100 under IPCC/CMIP6, NGFS and IEA scenarios

Scenario frameworks are not directly interchangeable. The IPCC/CMIP6 Shared Socioeconomic Pathways are primarily climate and emissions pathways used to generate physical climate projections. SSP1-2.6 represents a low-emissions pathway, SSP2-4.5 an intermediate pathway, and SSP5-8.5 a very high-emissions pathway, with end-of-century warming estimates of roughly 1.8°C, 2.7°C and 4.4°C respectively relative to 1850–1900.

NGFS scenarios are designed for climate-related financial risk analysis and combine physical climate outcomes with assumptions on transition timing, policy and macroeconomic adjustment. Net Zero 2050 is aligned with a low-warming pathway, Below 2°C represents a more gradual transition, while Current Policies leads to substantially higher physical risk.

IEA scenarios are primarily energy-system scenarios. NZE describes an energy transition consistent with limiting warming to around 1.5°C, APS reflects announced policy commitments, and STEPS reflects stated or implemented policies. Because the three frameworks are constructed for different purposes, scenario names should not be treated as exact equivalents. Comparisons should therefore rely on the underlying emissions, temperature and policy assumptions rather than on labels alone.

3.2 EXPOSURE AND VULNERABILITY MODELLING CHOICES

Once the hazard has been characterized, the second major design axis concerns exposure and vulnerability. It asks what is at risk, and how sensitive it is to a given hazard intensity. Exposure defines where the assets subject to a hazard are located and what they consist of; vulnerability defines how much damage a given hazard intensity is expected to cause to that exposure. Both are shaped by choices that are as consequential as the hazard choices addressed in Section 3.1, yet are frequently given less attention in institutional risk frameworks.

3.2.1 ASSET LOCATION

The most basic exposure question is also the most consequential. It concerns the geographic precision at which the assets in a portfolio are located. Country-level or postal-code-level location is often the only information available for granular retail loan books or diversified equity holdings. Address-level geocoding or building footprint data is achievable for direct real estate or infrastructure holdings.

Several banks still rely on sector-level statistical approaches in place of client-level asset data, and each step towards a more tailored quantitative method requires significantly more granular location data (BCG, 2024). This precision determines which hazard data, at which resolution as discussed in Section 3.1.3, can meaningfully be intersected with the exposure.

The financial consequences of location precision are large. The accuracy of asset location affects both the estimated severity of a hazard and the resulting risk assessment, and the effect is strongest for perils that vary over short distances, such as flood, and weakest for heat (GARP, 2025). Replacing asset-level locations with company headquarters underestimates portfolio value at risk by up to 71% (Bressan et al., 2024), and misplacing buildings geographically in a catastrophe model can halve estimated losses (Born et al., 2023). A coarse location matched against a fine-resolution flood map produces a false sense of precision. The hazard output should never be reported at a finer grain than the underlying exposure data supports.

Asset location also conditions the characteristics of the exposure itself, beyond its position on a hazard map. Vulnerability to physical risk varies greatly by location and by business type (BCG, 2024). One channel is the mix of labour and capital in production. Heat exposure reduces the productivity of labour-intensive firms, which respond over time by substituting capital for through investment and automation (Xiao, 2021). High-temperature shocks raise energy costs and lower productivity for small plants, while large plants show little effect (Ponticelli et al., 2023).

A manufacturing facility in South and South-East Asia is, for example, more likely to be labour-intensive, with production continuity depending on workforce availability. The same asset class in Western Europe tends to be more capital-intensive, with continuity depending on equipment replacement time and specialized supply chains. Location further determines exposure through the supply chain, since heat and flood days at supplier locations reduce supplier revenues and propagate to their customers (Pankratz & Schiller, 2024). Identical physical damage to nominally identical asset classes can therefore translate into materially different total financial impact depending on region. A vulnerability function calibrated in one geography cannot be assumed to transfer to another without an explicit adjustment for labour and capital intensity.

Example

Consider an institution wanting to assess the financial impact of wildfire on its portfolio over the coming year. Two datasets are available. The first covers 100% of the assets at a resolution of 100 km by 100 km. The second covers only Asia, where 75% of the portfolio is located, at a resolution of 1 km by 1 km. Which dataset should be used?

The higher-resolution regional dataset is the better choice. Wildfire spread is governed by vegetation type and slope at the scale of tens of metres, as discussed in Section 3.1.3. Specialised wildfire products are built accordingly, with pre-fire risk maps at 50 m and post-fire damage maps at 10m. A 100 km grid misclassifies individual assets as exposed or unexposed, whatever its coverage. The 1 km dataset produces a decision-useful exposure estimate for 75% of the portfolio. The remaining 25% should be reported as an explicit data gap. Filling it with the coarse global proxy would add no reliable asset-level information for this hazard type.

Data for Asset Location

Location data comes either from administrative identifiers, which link a financial position to a registered entity or establishment, or from direct observation of the asset. Administrative identifiers vary widely in the precision they offer. Entity-level identifiers often point only to a registered headquarters. Even recent supervisory research that

matches metre-level flood maps with loan-level data acknowledges that headquarters-based geocoding can miss the sites of multi-establishment firms (Albertazzi et al., 2026). The table below summarises the main sources.

SOURCE	LEVEL IDENTIFIED	LOCATION INFORMATION	COVERAGE (EXAMPLE)
Legal Entity Identifier (LEI)	Legal entity	20-character ISO 17442 code linked to reference data and ownership structure. It identifies the entity and not its physical sites.	Global (GLEIF)
Commercial company and location databases	Company and address	About 693 million companies and 173 million physical addresses compiled from primary sources	Global (Veridion)
Vendor asset-mapping services	Asset	GPS coordinates or asset polygons supplied by the client. Where none are available, locations are retrieved from company websites with AI tools.	Global
Asset-level climate ratings	Infrastructure asset	Ratings for more than 6,000 infrastructure assets at launch, with listed companies to follow	Global (EDHEC Climate Institute (2025))
Satellite imagery (Sentinel-2)	Physical footprint	Direct detection of assets in optical imagery at 10 to 20 m resolution, used where no administrative identifier is available	Global (Drusch et al. (2012))
SIRET number (Sirene register)	Establishment	X/Y coordinates for all publicly listed establishments, published monthly. About 79% are located at street-number precision, and about 11% of establishments have opted out of publication.	France (Insee, 2026)
EPA Facility Registry Service	Regulated facility	Facility location and identification data combined from EPA programmes and federal, state and tribal records, updated monthly	United States (U.S. EPA (2025))

TABLE 47. Sources for identifying and locating assets, from legal entity to physical site

Mapping Exposure: examples from the field

Data providers locate assets in four distinct ways. The approaches differ in the input they require from the client, the precision of the location they return, and the information they attach to that location beyond coordinates.

Veridion - Resolving companies to locations from web data. The first approach starts from a company name or an address and resolves it to the business operating there. Veridion collects content from company websites, trade registries, map pins, satellite images and news. A language model turns this content into structured company records, and an entity-resolution model matches them to existing records.

Addresses must resolve to coordinates within the stated country. Data for regions with sparse information and for micro-companies is delivered at lower confidence. The location service accepts only addresses that point to a precise building, street address or facility access point. It rejects cities, postal areas and countries, and it distinguishes an operating site from a registered legal address. This approach recovers the operational footprint behind an entity identifier, which a headquarters address alone cannot provide.

Building digital twins from client-supplied sites. The second approach relies on the client's own location data. ScyAI is a startup that geocodes the sites uploaded by the client and generates a 3D digital twin of each one from terrain and elevation data, satellite imagery and on-site risk engineering reports. A live demonstration covered about 30 manufacturing and headquarters sites of a single automotive group.

Location matching reaches about 99% for very large companies. Quality is lower when a portfolio is built from company names alone, without confirmed site locations. Physical damage is then translated into property damage, business interruption and replacement cost from the property and machinery values the client holds for insurance placement.

Observing assets and their activity from satellites. The third approach detects the asset directly in imagery. Kayrros maps the physical footprint of industrial assets such as steel mills and cement plants from Sentinel-2 and higher-resolution commercial imagery, and links each footprint to the legal entity and the portfolio. The same imagery yields activity indicators.

Heat signals emitted by blast furnaces, measured in infrared bands, provide a proxy for production. One survey covered 62 blast furnaces in 24 of the largest integrated steel mills in Europe, close to 70% of regional capacity. Accuracy for steel mills exceeds 90%. This approach provides an exposure measure that is independent of corporate disclosure, and it shows whether an asset is operating at the moment a hazard strikes.

Pre-mapped asset universes described as geometries. The fourth approach delivers ratings on a universe of assets that the provider has already located and drawn. Scientific Climate Ratings, developed at the EDHEC Climate Institute, describes each asset as a geometry made of points, lines or polygons, stored as shapefiles. The choice of geometry depends on the asset and on its sub-elements. An airport, for example, is split into runways and buildings, and each component receives its own damage function.

The polygons are then overlaid on gridded hazard maps, such as flood depth, wind speed or burned area, for the historical period and for future climate scenarios. At fine resolutions, such as 10 to 30 m for flood, the asset is assessed on the set of pixels that its polygon covers. At coarser resolutions, around 10 km for many climate variables, an extended asset such as a road is apportioned across grid cells according to the share of its surface in each cell.

The precision of the geometry matters. For an asset in Great Britain, a precise polygon produced two to three times more damage from a 100-year flood than the common fallback of a point with a 500 m radius, which missed part of the risk. The client receives an asset-level result without supplying any location data. Coverage, however, can be limited to the provider's universe.

PROVIDER	CLIENT INPUT	HOW THE ASSET IS LOCATED	LOCATION OUTPUT	MAIN LIMIT
Veridion	Company name or address	Web crawling, language-model extraction and entity resolution across company websites, trade registries,	Operating business at a precise building or street address, distinguished from a registered legal address	Lower confidence in sparse-data regions and for

PROVIDER	CLIENT INPUT	HOW THE ASSET IS LOCATED	LOCATION OUTPUT	MAIN LIMIT
		map pins and satellite images		micro-companies
ScyAI	Client site list, property and machinery values, risk engineering reports	Automated geocoding and 3D digital twin built from terrain, elevation and satellite data	Individual site modelled in 3D, aggregated to company and portfolio level	Lower quality when only company names are available
Kayrros	GPS coordinates or asset polygons (optional)	Satellite detection of the physical footprint and of activity signals, such as furnace heat	Asset footprint linked to the legal entity, with production indicators	Limited to assets visible from space
Scientific Climate Ratings	None	Pre-mapped asset universe drawn as points, lines and polygons (shapefiles), split into sub-elements	Asset geometry overlaid on hazard grids, apportioned across cells by surface share	Limited to the covered universe

TABLE 48. Asset location methods of selected providers, from client input to location output

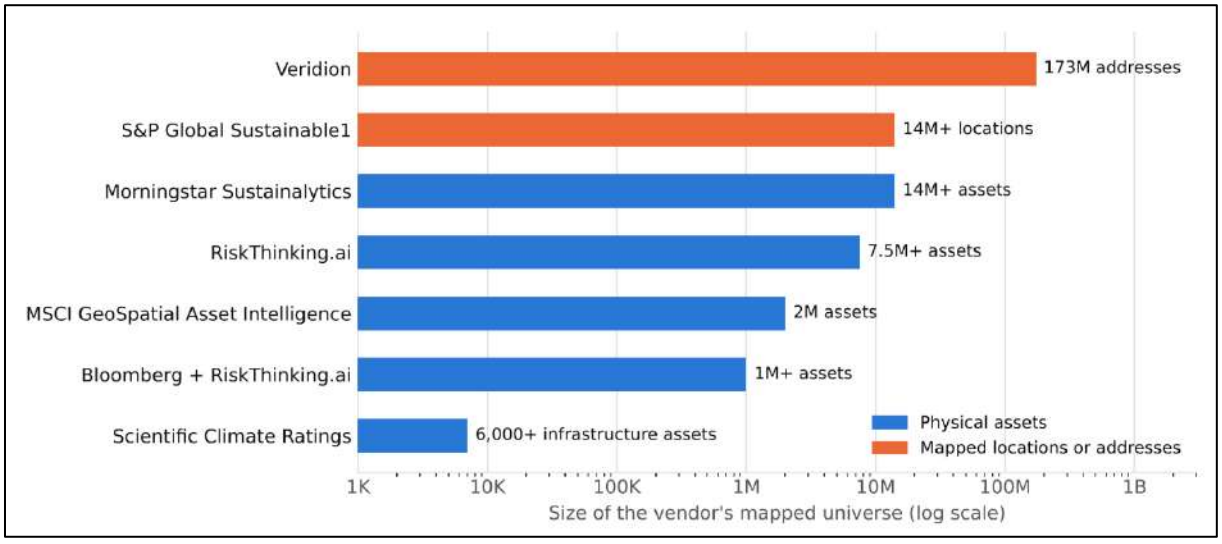


FIGURE 82. Size of the asset universes mapped by selected vendors, as declared by the vendors (log scale)

3.2.2 SOCIO-ECONOMIC VARIABLES

Hazard intensity and asset location do not fully determine financial impact. The socio-economic conditions around an asset shape how vulnerable the exposure is and how well the surrounding population can absorb and recover from a loss. Two assets with identical flood exposure can face different financial outcomes because of building standards, sector sensitivity and local adaptation (Wang et al., 2026). Recovery after a disaster also depends on income, employment and insurance coverage, and a severe event can leave uninsured households in financial distress for years (Prohaska, 2020; Wagner et al., 2022). At the macroeconomic level, uninsured losses drive most

of the lasting output cost of a catastrophe, because insurance payouts are what fund reconstruction (von Peter et al., 2024). Socio-economic conditions therefore deserve to be treated as a data category of their own, on the same footing as the hazard and location choices covered above.

Income level as a proxy for construction quality and recovery capacity

Physical risk data most often proxies socio-economic conditions with GDP per capita, or with the related split between emerging and developed markets. Income level correlates with several factors that matter for decisions but are rarely observed directly. These are how strictly building codes are enforced, how common informal or non-compliant construction is, how deeply insurance has penetrated, and how fast and how well reconstruction happens.

Lower-income countries face greater exposure to physical risk than higher-income ones (Castoldi et al., 2024). Building code adoption and enforcement vary widely between local governments, partly because of cost, and the insurance protection gap is widest where penetration remains limited, in Africa and Asia in particular (Golnaraghi & Vichev, 2025; EIOPA & ECB, 2023).

The global flood damage database of the European Commission's Joint Research Centre reflects these differences. Its country-level maximum damage values come from regressions on socio-economic World Development Indicators (Huizinga et al., 2017).

For a given hazard intensity, an asset in a lower-income region will therefore tend to suffer proportionally greater damage and take longer to recover. The hazard itself is the same, and the difference comes from a less resilient construction stock and weaker recovery infrastructure. Vulnerability functions calibrated on claims from high-income markets with high insurance penetration carry the construction standards and insurance conditions of those markets. Applied unadjusted to lower-income regions, they may misstate losses (Castoldi et al., 2024; Huizinga et al., 2017).



FIGURE 83. Illustration of differences in hazard impact in different areas

Social vulnerability and asset criticality indicators

Beyond income, practitioners increasingly use composite social vulnerability indicators. These combine population density, age structure, poverty headcount and access to essential services to identify the populations and assets that are most exposed, or least able to cope and recover. One example at national level is the INFORM Risk Index,

maintained by the European Commission's Joint Research Centre. It combines three dimensions, hazard and exposure, vulnerability, and lack of coping capacity, and now offers subnational models for 23 countries and regions (European Commission Joint Research Centre, 2026). Another example is the United States' CDC/ATSDR Social Vulnerability Index, which ranks every census tract on 16 American Community Survey variables grouped into four themes (Flanagan et al., 2011; CDC/ATSDR, 2024). Indices of this kind already inform physical risk models. A social vulnerability index has been used to steer hurricane recovery funding towards low-income assistance and elevation grants (Eid & El-Adaway, 2015). A census-tract dataset combining socio-economic, physical risk and housing data shows that FEMA risk scores correlate inversely with per capita income in Miami (Tedesco et al., 2022).

Asset or regional criticality is a related but separate concern. The loss of a dominant local employer, a hospital or a piece of critical infrastructure can amplify the financial and social impact of a hazard event well beyond the direct damage to that asset. Second-round effects spread through the local economy and reach other exposures that the institution holds in the same area. Input-output models of these cascades show that indirect losses can match or exceed direct damage, and that sectors such as finance can suffer most even when they are hit indirectly (De Maximy et al., 2025). The same propagation runs through supply chains, as shocks to suppliers in hazard-exposed locations pass on to their customers' sales and earnings (Pankratz & Schiller, 2024).

Resolution of socio-economic data

Socio-economic variables usually come at a coarser resolution than physical hazard data, most often for whole countries or first-level administrative subdivisions. Common sources are the World Bank's World Development Indicators, the UNDP Human Development Index and national census or household survey programs (UNDP, 2025). Subnational coverage is improving. Exposure indicators can now be built at national and subnational levels with global coverage (Maes et al., 2022). Tools such as EDHEC-CLIRMAP project climate damage to gross regional product for more than 3,660 subnational regions (Schneider, 2025). Because of this resolution mismatch, socio-economic adjustments are generally applied as a regional multiplier on a vulnerability function, and seldom as an asset-specific input. This treatment is coarser than what the hazard and location data of Sections 3.1 and 3.2.1 can support. Practitioners should record it explicitly as a data quality limitation.

Example: should two identical warehouses share one vulnerability function?

Consider two otherwise identical warehouses exposed to the same 1-in-100-year flood intensity. The first sits in a high-income region with dense insurance penetration and enforced building codes. The second sits in a lower-income emerging-market region with limited insurance and informal construction. Should both receive the same vulnerability function?

They should not. A vulnerability function calibrated on high-income claims would understate the expected damage ratio and the recovery time of the second warehouse, for the reasons set out above. The practitioner should apply a socio-economic adjustment before comparing the two exposures. Two adjustments are involved, and they should be kept apart. The first scales the replacement value of the asset, which is the role GDP-based country maximum damage values play in global damage databases (Huizinga et al., 2017). The second scales the damage ratio itself, using GDP per capita or an equivalent proxy for construction quality. The first adjustment lowers the absolute loss in the poorer region. The second raises the share of value destroyed. Merging the two into a single GDP multiplier can make the second warehouse look safer than the first.

Data for socio-economic variables

Demographic and macroeconomic indicators used as socio-economic proxies are generally sourced from multilateral and national statistical organizations: the International Monetary Fund, the World Bank, and the United Nations for cross-country series, and national statistical institutes or regional institutes such as Eurostat for the European Union for sub-national detail (IMF, World Bank, UN DESA, Eurostat…). These sources are the most widely available basis for the population, income, and poverty indicators discussed above.

VARIABLE	TYPICAL SOURCE	SCALE	COVERAGE
GDP, GDP per capita, macroeconomic projections	IMF World Economic Outlook database; World Bank World Development Indicators	National	Global
Population and age structure	UN DESA World Population Prospects	National	Global
Human development (health, education, income)	UNDP Human Development Index	National	Global
Subnational economic output	DOSE dataset (Wenz et al., 2023); EDHEC-CLIRMAP (Schneider, 2025)	First administrative level	1,661 regions in 83 countries (DOSE); 3,660+ regions (CLIRMAP)
Regional income, population, employment	Eurostat regional statistics	NUTS 2 and NUTS 3	European Union
Household income and poverty	INSEE	Municipality and 200 m grid	France
Composite disaster risk and coping capacity	INFORM Risk Index	National, plus 23 subnational models	Global
Social vulnerability	CDC/ATSDR SVI	Census tract	United States
Combined socio-economic, physical risk and housing data	SEPHER (Tedesco et al., 2022)	Census tract	United States
Insurance protection gap	EIOPA and ECB assessment	National	European Union

TABLE 49. Examples of socio-economic data variables, scale and coverage

Provider Box – MSCI: Climate VaR, Macroeconomic data and partnerships with other solutions

MSCI's Climate VaR is a single forward-looking measure of physical risk. It expresses the discounted future physical costs of a company relative to its enterprise value. The guide warns that the measure can be misread as a forecast of expected returns (Wang et al., 2026). On the vulnerability side, MSCI acknowledges that two assets with identical flood exposure can face different financial outcomes depending on building standards, sector sensitivity and local adaptation. It falls back on regional proxies when location data is missing, and its loss metrics do not evaluate insurance coverage (Wang et al., 2026).

MSCI adds a macroeconomic layer to the asset-level layer. The Macroeconomic Physical Risk Climate VaR uses the NGFS damage functions, supplemented with enhanced functions for non-linearities and tipping

points. Economy-wide effects reach company costs through supply chains, labour force participation, migration, the cost of capital, and capital expenditure for adaptation and reconstruction. Under current policies, the model estimates a loss of 19.8% of the market capitalisation of the MSCI ACWI IMI. In its worst-case scenarios, losses for global equity portfolios could be up to nine times greater than direct damage alone (Faigle, 2026).

Starting in 2027, Climate VaR physical risk will be powered by First Street hazard models, thanks to acquisition of First Street by MSCI that was finalized in August 2026.

On the data side, MSCI announced a partnership with Swiss Re in February 2025. It combines Swiss Re's physical risk analysis with MSCI's risk models and covers more than 20 hazard types aligned with the EU Taxonomy (MSCI Inc., 2025).

3.2.3 ADAPTATION AND RESILIENCE HYPOTHESIS

About adaptation

In 2019, the Global Commission on Adaptation, hosted by the World Resources Institute, estimated that investments in adaptation return between 2 and 10 dollars in benefits per dollar invested (Global Commission on Adaptation, 2019). Agricultural research and development shows even higher ratios, between 2:1 and 17:1 (Global Commission on Adaptation, 2019). These are hypothetical gains, just as most of the hazards and damages modelled in this report are hypothetical under future climate scenarios.



FIGURE 84. Nature-based solutions to climate hazards across five landscapes, from mountains to coasts
Source: WRI, 2019

The last design choice for exposure and vulnerability is what to assume about adaptation. The question is what has already been done, at the level of the individual asset or of the surrounding jurisdiction, to anticipate or reduce the impact of the modelled hazard. Across a portfolio the answer varies enormously. Some assets include substantial protection, such as reinforced concrete construction, elevated foundations, waterproofed facilities or seismic retrofitting. Others have little or none. A vulnerability function that ignores this variation implicitly assumes an

average level of protection, which may be wrong in either direction for any given asset. Adaptation is represented in the vulnerability component of the modelling chain, through the choice of a damage curve with lower damage ratios, and ignoring asset-level detail can lead to large errors when identifying the right adaptation response (CISL, 2019; Bressan et al., 2024).

Structural and operational adaptation measures

Adaptation measures fall into two broad categories. **Structural measures** physically change how well an asset resists a hazard. Consider for example reinforced construction, flood barriers and flood-proofing, elevated foundations, fire-resistant materials with defensible space around a structure, and seismic retrofitting. For flooding, these measures either prevent water from entering a building or limit damage once it has entered, and roof upgrades play the same role for tropical cyclones (CISL, 2019). Their cost-effectiveness depends on the asset. Elevating a home pays off mainly for high-value homes with low first floors in the most exposed surge zones, and cheaper options such as wet floodproofing or raising utilities suit the other cases (Montgomery & Kunreuther, 2018).

Operational and non-structural measures leave the asset unchanged and reduce the financial consequence of a given level of physical damage. Examples are early warning systems that allow inventory or operations to be moved before an event, business continuity plans, and diversification of the supply chain. Firms do adjust their supplier networks. Customers are 7% more likely to end a supplier relationship when the supplier's realized climate exposure exceeds prior expectations (Pankratz & Schiller, 2024). Insurance belongs to this second category, with one difference. It transfers the loss to another party and leaves the economic damage unchanged, so it reduces the loss borne by the asset owner and its lender without reducing the loss to the economy (Von Peter et al., 2024).

Both categories reduce expected loss through different mechanisms. A structural measure changes the damage ratio. An operational measure changes the conversion of physical damage into financial loss. A vulnerability function should state which of the two it captures, so that the same measure is never counted twice, once in the damage curve and once in the financial layer.

The adaptation data gap and the working-hypothesis problem

The practical difficulty is that asset-specific adaptation data is rarely available at the granularity of the hazard and exposure data it is meant to complement. Confirming that a given factory has been waterproofed, or that a given building has been seismically retrofitted, usually requires a site-level engineering survey or disclosure by the asset owner. Both are costly to collect at portfolio scale. Self-reported data may also be biased, since an owner has a direct financial incentive to overstate resilience measures in order to obtain cheaper insurance or a lower internal risk rating. Even reference frameworks for lenders work with assumed uptake rates for protection measures, because portfolio data rarely records adaptation attributes for individual assets (CISL, 2019).

Without asset-specific data, practitioners must adopt a working hypothesis about adaptation. The most common proxy is the vintage and jurisdiction of the applicable building code. Construction completed under a more stringent code can reasonably be assumed to include a higher baseline level of resilience than older stock. Examples are the seismic provisions of Eurocode 8 and the Florida Building Code adopted after Hurricane Andrew (CEN, 2004; Simmons et al., 2018).

The evidence supports this proxy. The 2001 Florida Building Code reduced windstorm losses by as much as 72%, with about six dollars of avoided loss for each dollar of added construction cost (Simmons et al., 2018). In the United States flood insurance programme, claims relative to home value are 42 to 45% higher for buildings constructed before the first flood insurance rate maps, which were built to lower standards (Kousky & Michel-Kerjan, 2017).

Adaptation also changes over time. Investment in resilience by an institution or a jurisdiction typically rises after a major loss event, as the Florida code shows. A vulnerability function calibrated on historical damage from before

such an event may therefore overstate future losses for the same nominal hazard intensity once later adaptation is considered.

VARIABLE	TYPICAL SOURCE	RESOLUTION	COVERAGE
Building age, type and identity	BDNB, Base de données nationale des bâtiments (CSTB)	Building	France, about 32 million buildings
Year built, foundation type and height	National Structure Inventory (USACE)	Structure point	United States
Wind mitigation features (roof covering, roof deck attachment, roof-to-wall connection, opening protection)	Florida Uniform Mitigation Verification Inspection Form OIR-B1-1802	Building, on inspection	Florida
Flood protection standards (design, policy and modelled layers)	FLOPROS (Scussolini et al., 2016)	Protection unit	Global
Roof type and material	Aerial and satellite imagery classified by machine learning (see box)	Building	Depends on provider
Building code vintage	National and regional building regulations	Jurisdiction	Varies

TABLE 50. Data for adaptation and resilience

Example: how much credit should adaptation receive?

Consider two single-family homes in the same coastal flood zone, insured and financed by the same bank. The first was built in 1970, before the local flood maps existed. The second was built in 2005 under current floodplain rules, and the owner declares that it is elevated. Should both receive the same flood vulnerability function?

They should not, and each adjustment should be documented. The vintage adjustment rests on observed evidence, since claims on pre-map buildings run 42 to 45% higher relative to value than on other buildings (Kousky & Michel-Kerjan, 2017). The elevation adjustment rests on a declaration. Homes known to be elevated show claims 16 to 18% lower (Kousky & Michel-Kerjan, 2017), but this credit should be applied only once the elevation has been verified through an inspection record, a certificate or imagery. Applying both adjustments to the 2005 home without checking would partly double count, because recent construction rules already require elevation in many flood zones. The practitioner should therefore apply the vintage modifier by default and add the elevation credit only when evidence exists.

Reading vulnerabilities from the roof: Addactis

Imagery now fills part of the adaptation data gap. Roof type drives much of a building's wind vulnerability, yet public databases rarely record it. A convolutional neural network trained on satellite images of individual buildings classifies roof types with an F1 score of 0.96. Applied to 161,772 houses in New Hanover County and Miami-Dade County, it maps roof types down to the census tract (Meng et al., 2023). Open datasets now extend the approach to

roofing materials. One of them gathers 51,503 labelled image tiles from 112 countries across 14 material classes, for underwriting and disaster preparedness (Law & Miura, 2025).

In France, ADDACTIS has brought this approach into insurance practice. In May 2025 the actuarial firm acquired the insurance business of namR, a provider of building-level climate data. namR's six data scientists joined ADDACTIS, along with building data and analytics covering 36 million French buildings. The stated aim is to compute risk as a combination of climate hazard and building vulnerability.

NamR's approach applies deep learning to aerial photographs at a resolution of about 20 cm per pixel to classify roof types and to detect solar panels and vegetated areas (Liora, 2026).

In June 2026 the acquisition became addactis Property Risk Intelligence. The platform is a digital twin of more than 36 million French buildings, with 200 to 300 attributes per building. Geocoding links each building to its surroundings, including roof shape and materials, distance to watercourses and soil type. Machine learning models trained on observed data turn these attributes into continuously updated vulnerability scores for flood, storm, hail, fire, theft and drought. Insurers use them to target prevention on the most vulnerable buildings, to detect risk concentrations for reinsurance, and to project how scores may evolve across home insurance portfolios.

Elsewhere, Mitigrate applies a similar vulnerability-first logic to flood risk in the Nordics and the United Kingdom. The platform recommends costed protection measures for each property and verifies installed protection with photographs of water entry points.

SOURCE	REPORTED RATIO	SCOPE	NATURE OF THE ESTIMATE
Global Commission on Adaptation (2019)	2:1 to 10:1 across domains; USD 7.1 trillion in net benefits for USD 1.8 trillion invested (net benefits to cost, about 4:1)	Five domains (early warning systems, resilient infrastructure, dryland agriculture, mangroves, water), world, 2020–2030	Synthesis of heterogeneous models and studies
World Bank, Lifelines (2019)	About 4:1 (USD 4.2 trillion in net benefits)	Resilient power, water, transport and telecommunication infrastructure in low- and middle-income countries	Model-based, combined with case studies and empirical analyses
National Institute of Building Sciences (2019)	4:1 for private-sector building retrofits, 6:1 for federal grants, 11:1 for adopting current building codes	Natural hazard mitigation, United States	Model-based on US data
Watkiss (2021)	2:1 to 10:1, some measures barely above 1	United Kingdom	Review of adaptation economics evidence for the third UK Climate Change Risk Assessment

SOURCE	REPORTED RATIO	SCOPE	NATURE OF THE ESTIMATE
WRI, Brandon et al. (2025)	About 10.5:1 over ten years	320 investments approved by development banks and climate finance institutions, 12 countries, 2014–2024	Ex ante, from project documents, including benefits that occur without disasters

TABLE 51. Reported benefit–cost ratios of climate adaptation investments

Broad programs cluster between 2:1 and 4:1. Higher ratios come either from specific measures, such as the adoption of current building codes in the United States (11:1), or from aggregates whose scope extends beyond avoided losses, such as the WRI study. The 8:1 found in the IFC and Scientific Climate Ratings case (see the provider box above) sits in this upper range

How much is adaptation worth? Reading benefit-cost ratios

Figures on the return of adaptation circulate widely, and they vary by an order of magnitude. The table below gathers the most cited estimates.

Three quantities are often merged under the phrase “value protected”. The first is the avoided expected loss, which forms the numerator of a benefit-cost ratio:

$$BCR = \frac{\sum_{t=0}^{T} B_t/(1+r)^t}{\sum_{t=0}^{T} C_t/(1+r)^t}, \qquad B_t = \mathbb{E}\left[L_t^0\right] - \mathbb{E}\left[L_t^1\right]$$

- $\boldsymbol{BCR}$ is the benefit-cost ratio.
- $\boldsymbol{B_t}$ is the benefit in year t, the avoided expected loss.
- $\boldsymbol{C_t}$ is the cost of the measure in year t
- $\boldsymbol{r}$ is the discount rate and T the horizon in years.
- $\boldsymbol{L_t^0}$ and L_t^1 are the losses in year t without and with adaptation.

A BCR of 8 corresponds to a return on investment of 7, since the ROI equals $\mathrm{BCR} - 1$. The second quantity is the reduction in value at risk at a confidence level α.

- It measures an avoided extreme loss at one date and does not add up across assets. The third is the change in the value of an asset or portfolio, which depends on the valuation model. For a bank or an insurer, the reduction in value at risk can be turned into a true financial return through the capital it frees:

$$B_t \approx k \cdot \Delta VaR_{\alpha,t}$$

- $\boldsymbol{k}$ is the cost of capital.
- $\boldsymbol{\Delta VaR_{\alpha,t}}$ is the reduction in value at risk at confidence level α in year t

In this framework, the freed capital earns its cost of capital, and it requires asset-level losses with an explicit baseline scenario.

Four features of a published ratio deserve checking before use:

The first is the **scope of benefits**. In the WRI study, avoided losses, induced economic benefits and social and environmental benefits each weigh about a third, so avoided losses account for roughly a third of the headline ratio, and part of the third dividend is carbon mitigation.

The second is the **discount rate**, which the GCA report does not state and WRI sets at 5% by default.

The third is **selection**. Projects approved on their projected benefits show systematic gaps ex post even when forecasts are unbiased (Eliasson & Fosgerau, 2013), and ex post evaluations remain rare (Josephson et al., 2024).

The fourth is the **difference between average and marginal returns**. An average ratio says nothing about the last euro invested, whose ratio tends towards 1 at the optimum.

The Hallegatte approach

Stéphane Hallegatte, lead economist at the World Bank, offers a different way to frame the return of adaptation. His work questions the metric before the figure. In *Unbreakable*, he and his co-authors measure disasters by their effect on well-being, expressed as an equivalent loss of consumption. The usual measure, the value of assets destroyed, counts poor households' losses as small simply because their assets are worth little. In all of the 117 countries studied, the effect on well-being, measured in terms of lost consumption, is found to be larger than asset losses. The distribution is striking. Globally, poor people are disproportionately affected by these losses: the poorest 20% experience only 11% of total asset losses, but 47% of wellbeing losses annually. The logic rests on the diminishing marginal utility of consumption:

$$\Delta W \approx \Delta C \cdot \left(\frac{\bar{c}}{c}\right)^{\eta}$$

- ΔW is the well-being loss, expressed as an equivalent loss of average consumption.
- ΔC is the monetary loss of the household.
- c is the consumption of the household and $\bar{c}$ the average consumption.
- η is the elasticity of the marginal utility of consumption.

The same dollar of damage therefore weighs more for a poor household than for a rich one. A ratio built on asset losses steers investment towards wealthy areas, while a ratio built on well-being gives weight to measures that protect the vulnerable. This is why *Unbreakable* also assesses resilience policies that reduce no physical damage, such as social protection, financial inclusion and insurance.

Hallegatte's second contribution concerns decisions under deep uncertainty. Shah et al. (2012) review four approaches: cost-benefit analysis under uncertainty, the same with real options, robust decision making, and climate-informed decision analysis. Their conclusion is that no method dominates, and that the choice depends on the lifetime, cost and vulnerability of the project. Robust decision making replaces the expected value with a regret criterion:

$$\text{Regret}(a, s) = \max_{a'} V(a', s) - V(a, s), \qquad a^* = \arg\min_{a} \max_{s} \text{Regret}(a, s)$$

- a is an adaptation option and a' any alternative option.
- s is a climate or socio-economic scenario.
- $V(a, s)$ is the value of option a under scenario s
- a^* is the option that minimises the worst-case regret across scenarios.

The robust option is not the one with the best expected ratio. It is the one that performs acceptably whatever the scenario, a sound criterion for long-lived assets whose climate is unknown. Hallegatte also signed the headline

figure of *Lifelines* (Hallegatte et al., 2019), about 4 dollars of benefit per dollar invested in resilient infrastructure in low- and middle-income countries, a model-based estimate like the others. Several sources of the Global Commission on Adaptation report draw on his work, so the agreement between its 2:1 to 10:1 range and the World Bank's 4:1 partly reflects shared studies, and the two do not count as independent evidence.

For a financial institution, the well-being metric serves public decision-making more than portfolio risk, which still needs asset-level losses linked to value at risk. The regret criterion transfers directly to the choice of adaptation measures for long-lived infrastructure.

3.3 FINANCIAL PROXY CHOICES

3.3.1 ACCOUNTING VARIABLES

Accounting variables are the first of three families of financial proxies covered in this section. The other two are market variables (Section 3.3.2) and actuarial variables (Section 3.3.3). Accounting variables translate a physical damage estimate into the financial metrics an institution already reports and manages, namely income statement items, balance sheet items and the return ratios derived from them. They are the natural proxy for company-level analysis, and for asset-level analysis whenever accounts can be traced to individual sites or subsidiaries.

Direct and indirect losses

Many, many, many accounting variables can carry a physical risk estimate, such as the book value of impaired assets or provisions for physical risk events. Whatever the variable chosen, the loss breaks down into two parts.

The replacement cost covers the repair or replacement of damaged capital, which is the direct damage. Business interruption covers indirect damage, through lost margin and extra operating costs, but also because of . Structural credit risk models of physical risk follow this split. Site-level hazard shocks become capital destruction and business interruption, which then reduce the cash flows of the firm (Mandel et al., 2025). For a site hit by a hazard of intensity I the loss can be written as:

$$L = \underbrace{\sum_{k} V_k(I) \cdot IV_k}_{\text{replacement cost}} + \underbrace{\frac{GM}{365}\left[d(I) + \sum_{t} s_t(I)\right] + \Delta C}_{\text{business interruption}}$$

- L is the total loss of the site for the event.
- k indexes the components of the site: building, machinery and equipment, contents, stock.
- $V_k(I)$ is the damage ratio of component k, between 0 and 1, given by its vulnerability function.
- IV_k is the value of component k, ideally its replacement value as new.
- GM is the annual gross margin of the site, measured as revenue minus variable costs.
- $d(I)$ is the number of days of full shutdown.
- $s_t(I)$ is the share of output lost on day t of reduced activity after restart, or during a hazard that slows activity without damaging the site, such as heat.
- ΔC is the extra cost incurred to limit the interruption, such as temporary premises, overtime or emergency logistics.

Each term can absorb many effects, and the formula is a structure for documenting them. The components react differently to the same intensity. A building loses value progressively with water depth, electrical equipment fails abruptly once water reaches it, and sensitive stock is lost almost entirely on contact.

Each term can absorb many effects, and the formula is a structure for documenting them. The components react differently to the same intensity. A building loses value progressively with water depth, electrical equipment fails abruptly once water reaches it, and sensitive stock is lost almost entirely on contact. Downtime $d(I)$ grows with the damage ratio, and it depends on cleaning and drying, the replacement lead time of critical equipment (often the dominant factor) and access to the site and its utilities. Business interruption does not stop at the site gate. A one standard deviation increase in hot days lowers firm revenues by 0.6% and operating income by 1.8% relative to average quarterly figures (Pankratz et al., 2023). Customers are 7% more likely to end a supplier relationship when the supplier's realized climate exposure exceeds prior expectations, so part of the loss outlasts the event (Pankratz & Schiller, 2024).

Damage to surrounding infrastructure adds further disruption. After the 2020 Labor Day wildfires on the US West Coast, repairs to impacted roads cost about 127,783 dollars per kilometre (Christiansen et al., 2024). Road agencies bear that cost, while the firm bears the loss of access. A severe or repeated event can also damage the firm's reputation and future revenue.

The inputs are easier to obtain than they seem. Firms declare replacement values and gross margin to their insurers each year, usually per site. Their definitions should be stated, since a depreciated book value in place of a replacement value, or an industrial margin in place of revenue minus variable costs, carries straight through to the result. The loss borne by the firm is finally net of insurance:

$$L^{net} = L - \min(\max(L - D, 0), Lim)$$

- L^{net} is the loss retained by the firm.
- D the deductible and Lim the policy limit.

Replacement costs are relatively easy to estimate, and exposure databases record them as the replacement value of each asset (CISL, 2019). Business interruption is much harder, and even within a single sector there is no standard method (Guzmán et al., 2023).

Income statement, balance sheet and return metrics

The loss above can be measured on several accounting lines, and the choice depends on where the user expects the impact to show.

French accounts break a company's profit down step by step, through what are called the soldes intermédiaires de gestion. The first step is what the company produces. The second is the value it adds once its suppliers are paid. The third is its operating profit before depreciation. Following a hazard through these steps shows where it hurts.

A flood that destroys stock lowers production, while a heatwave that raises cooling bills reduces the value added. International standards are moving the same way. From 2027, IFRS 18 requires companies to sort their income and expenses into three groups (operating, investing and financing) and to report a standard operating profit line, which makes companies easier to compare (IASB, 2024).

On the balance sheet, physical damage to a fixed asset usually appears as an impairment charge or as faster depreciation. Climate-related matters can indicate that an asset is impaired under IAS 36, and they can shorten the useful life or lower the residual value of an asset under IAS 16 (IFRS Foundation, 2023). Compensation from insurers for damaged assets is recognized in profit or loss when it becomes receivable (IAS 16, paragraph 65). Other balance sheet lines carry part of the loss: net property, plant and equipment, inventories written down after a flood, provisions for restoration costs, and capital expenditure for reconstruction in the cash flow statement.

For lenders, the relevant variable sits on their own balance sheet, as a change in loan loss provisions or in the probability of default of the borrower. IFRS 9 requires expected credit losses to reflect all reasonable and supportable information, including floods or wildfires that weaken a borrower's ability to repay or reduce the value

of its collateral (IFRS Foundation, 2023). Rising temperatures lower a firm's asset value and raise its probability of default in a Merton framework (Blanc-Blocquel et al., 2024). The effect also reaches lending volumes. Typhoons in China reduced bank lending by an amount equal to 2.8% of total bank assets (Avril et al., 2023). Credit metrics derived from the accounts translate the loss into covenant terms: net debt to EBITDA, interest coverage (EBIT over interest expense) and operating cash flow.

Return ratios (return on investment, return on equity and return on assets) express the impact relative to the capital base. They are the natural metric for comparing exposed assets or business units, or for benchmarking a hazard's impact against the cost of capital. Financial leverage amplifies direct losses for shareholders (Mandel et al., 2025), so the same damage weighs more heavily on return on equity than on return on assets, and more so the more leveraged the firm.

All these metrics can be sourced from the providers listed in the table below. Revenue, EBITDA, net income, net PP&E, inventories and impairment are standard tagged items in US filings (using for example XBRL) and in the French annual accounts filed with the commercial courts. Commercial databases harmonize them across countries. Site-level values remain the gap, since published accounts stop at the entity, and replacement values and gross margin per site come from insurance declarations or from the firm itself.

Choosing among accounting metrics

The metric to target follows from the question asked, and the figure that happens to be most readily available should not decide it. A shareholder or equity analyst looks first at EBITDA, net income and return on equity. An asset owner or insurer looks at the carrying value of the physical asset and its impairment. A lender looks at loan loss provisions and the borrower's probability of default. The same physical damage estimate can feed all three. The right accounting variable, and the balance sheet on which to measure it, depend on whose financial exposure the analysis is meant to inform.

Consider the financial impact of a flood on a single manufacturing subsidiary. Should the model target EBITDA, asset impairment or the probability of default on a loan?

- **Shareholder value at risk.** EBITDA and net income are the natural metrics. The replacement cost and the business interruption loss both pass through them.
- **Overvalued asset on the balance sheet.** Impairment and depreciation are the right metrics, net of any insurance receivable.
- **Riskier loan secured on the asset.** The relevant variable is the change in loan loss provisions on the lender's balance sheet. The subsidiary's own accounts only feed that calculation.

SOURCE	CONTENT	COVERAGE	ACCESS
SEC EDGAR filings and Financial Statement Data Sets	10-K and 10-Q filings, with numeric XBRL data from the primary statements, updated quarterly since 2009	US listed companies	Free
Company annual reports and national regulatory filings	Audited income statement, balance sheet and notes	Listed companies outside the US	Free
S&P Capital IQ	Standardized financials, estimates and ownership	Global, listed and large private companies	Commercial

SOURCE	CONTENT	COVERAGE	ACCESS
Moody's Orbis (Bureau van Dijk), including Orbis for France, formerly DIANE	Standardized financials with a focus on private companies	Global; millions of French companies	Commercial
INPI open data	Non-confidential annual accounts filed with commercial courts, including balance sheet and income statement	France, about 5.9 million active companies	Free

TABLE 52. Data for accounting metrics

3.3.2 MARKET VARIABLES

Market variables are the second of the three families of financial proxies covered in this section, after accounting variables (Section 3.3.1) and before actuarial variables (Section 3.3.3). Accounting figures are reported periodically and audited. Market variables capture, almost in real time, how investors price physical risk into the value of a listed asset. They are the natural proxy when the question concerns market value at risk, and less suited to questions about reported operating performance.

Stock price and volatility as market-based proxies

The most direct market variable is the reaction of a company's stock price around the date of a hazard event. It is usually measured with an event study (MacKinlay, 1997), which compares the observed return with the return the stock would have earned without the event:

$$AR_{i,t} = R_{i,t} - \left(\widehat{\alpha}_i + \widehat{\beta}_i R_{m,t}\right), \qquad CAR_i(t_1, t_2) = \sum_{t=t_1}^{t_2} A\,R_{i,t}$$

- $\boldsymbol{AR_{i,t}}$ is the abnormal return of firm i on day t
- $\boldsymbol{R_{i,t}}$ is the observed return of firm i on day t
- $\boldsymbol{R_{m,t}}$ is the return of the market index on day t
- $\widehat{\boldsymbol{\alpha}}_{\boldsymbol{i}}$ and $\hat{\beta}_i$ are the parameters of the market model, estimated on a window before the event.
- $\boldsymbol{CAR_i(t_1, t_2)}$ is the cumulative abnormal return over the event window from day t_1 to day t_2.

Insurance stocks show negative abnormal returns for up to 28 days after major wildfires (Tavor, 2024), and European airlines lost value after the 2010 Eyjafjallajökull ash cloud (Mazzocchi et al., 2010).

Volatility complements the price reaction. Realized volatility is computed from past returns. Implied volatility is extracted from option prices and reflects the uncertainty that investors expect. Both capture how a hazard, or its anticipation, raises uncertainty about a firm's future cash flows, whatever the direction of the price move. Options on firms with establishments in forecast and realized hurricane landfall regions exhibit large increases in implied volatility, and comparing this expected volatility with the volatility actually realized afterwards shows that investors significantly underestimate extreme weather uncertainty (Kruttli et al., 2025).

Market variables can also test whether prices anticipate slow-onset risks. In the agricultural sector, Hong et al. (2019) rank 31 countries each year by their long-term trend towards drought, measured by the Palmer Drought Severity Index.

They then regress the future returns of each country's food industry on this ranking:

$$r_{c,t+1} = a + b \cdot \text{DroughtTrend}_{c,t} + \gamma' X_{c,t} + \varepsilon_{c,t+1}$$

- $\boldsymbol{a}$ is the constant.
- $\boldsymbol{r_{c,t+1}}$ is the return of the food industry of country c in year $t+1$, net of the market.
- $\text{DroughtTrend}_{\boldsymbol{c,t}}$ is the country's long-term drought trend, publicly known in year t
- $\boldsymbol{X_{c,t}}$ is a vector of controls and γ is the vector of coefficients on the controls, and γ' its transpose, so in other words $X_{c,t}$ is their weighted sum.
- $\boldsymbol{\varepsilon_{c,t+1}}$ is the error term.
- $\boldsymbol{b}$ is the coefficient of interest. In an efficient market, public information should not predict future returns, so b should be zero.

A poor drought trend forecasts both lower profit growth and lower stock returns for the food companies of that country, which shows that food stock prices underreact to climate risk (Hong et al., 2019).

Market variables have structural limits that accounting variables do not share. They reflect investors' perception of physical risk, filtered through diversification, hedging and risk premia, and not the physical damage itself. They are also observable only for listed companies. Event studies are vulnerable to confounding events, since an abnormal return around a hazard date may reflect unrelated news released in the same window. A credible market-based estimate therefore needs a carefully specified counterfactual, and the raw price change is not enough.

Ownership as a control variable

Ownership structure is not itself a measure of financial impact. It covers the split between institutional and retail investors, insider ownership and the geographic concentration of shareholders. It is commonly used as a control variable in market-based studies of physical risk, because it shapes how a physical shock reaches the price. Higher institutional ownership is associated with more informationally efficient prices (Boehmer & Kelley, 2009), so new risk information is incorporated faster. Institutional investors also say that they consider climate risks, physical risks included, to have financial implications for their portfolios (Krueger et al., 2020). Investors' own location matters as well. Fund managers located near a climatic disaster overreact to it and underweight stocks in the affected region (Alok et al., 2020). Studies that estimate the market pricing of physical risk should therefore control for ownership alongside standard firm characteristics such as size and leverage, so that an ownership effect is not attributed to physical risk.

Example: how to choose?

Consider testing whether investors price flood risk into a portfolio of listed real estate companies. Should the model use raw stock returns, abnormal returns from an event study, or volatility?

Raw returns are unsuitable on their own, since broad market movements unrelated to flood risk dominate them. Abnormal returns are the right choice if the question is whether a specific flood triggered a repricing. Volatility is the right choice if the question is whether the market's ongoing uncertainty about flood exposure has increased, for instance after a change in regional flood risk disclosure, independently of any single event. A predictive regression fits a third question, whether prices already reflect a known long-term trend in the hazard (Hong et al., 2019).

DATA	TYPICAL SOURCE	COVERAGE	ACCESS
Daily stock prices, returns and volumes	CRSP (academic, US equities); Bloomberg; LSEG Workspace (formerly Refinitiv)	US (CRSP); global (Bloomberg, LSEG)	Academic licence; commercial
Option prices and implied volatility	OptionMetrics IvyDB (academic); Bloomberg; LSEG	US and major international markets	Academic licence; commercial
Market indices and factor returns	MSCI indices; Kenneth French Data Library	Global	Commercial; free
Institutional ownership	Form 13F filings via SEC EDGAR; LSEG and FactSet ownership modules	US (13F); global (commercial modules)	Free; commercial
Fund holdings	Morningstar; CRSP Mutual Fund database	Global; US	Commercial; academic licence

TABLE 53. Data for market metrics

3.3.3 ACTUARIAL VARIABLES

Actuarial variables are the third family of financial proxies. They express physical risk directly as a loss distribution, without going through accounting or market metrics. Their construction from a stochastic event set is presented in Part 2. This section covers the choice of metric and its data.

Loss exceedance metrics

The occurrence exceedance probability (OEP) gives the probability that a single event exceeds a given loss threshold in a year. The aggregate exceedance probability (AEP) gives the probability that the total of all losses in a year exceeds that threshold. OEP prices per-occurrence reinsurance, whose attachment point is triggered by one event. AEP prices annual aggregate covers and underlies internal models of solvency capital. Using the wrong curve for the structure being priced is a common and costly error.

Damage functions

A damage function converts hazard intensity into a loss ratio, the share of asset value destroyed (see Section 3.2). The choice between industry-standard functions (Hazus, FLEMOps), proprietary vendor functions or portfolio-specific functions calibrated on claims can weigh as much as the hazard data in the loss estimate. In a sensitivity analysis of an earthquake catastrophe model, construction type and catalogue length drive most of the uncertainty (Born et al., 2023). The function must match the asset class and construction type, and a generic function misprices non-standard assets.

Actuarial metrics live in the tail, where data are scarcest. Extreme value theory justifies extrapolating beyond the observed record. Above a high enough threshold, the losses of a wide class of distributions behave approximately like a generalized Pareto distribution (Balkema & de Haan, 1974; Pickands, 1975), whose shape parameter ξ sets how slowly the tail decays (see Part 2).

Three practical points follow for model users.

First, distributions with the same mean and variance can give very different 1-in-200-year losses, so variance alone does not describe catastrophe risk.

Second, the threshold above which the generalized Pareto applies has no objective, automatic selection rule, and in practice it is often set by business conventions such as a reinsurance attachment point, at the cost of some model bias.

Third, the tail index of a pooled portfolio equals the worst tail index of its components, so mixing a heavy-tailed risk with lighter ones penalises the whole portfolio. When ξ reaches 0.5, the variance of losses becomes infinite, and when it reaches 1, the mean does too. Estimates of ξ rising over time for flood or wildfire losses are therefore a warning sign for insurability. Tracking such a trend requires long records. On shorter series, the usual practice is to stress the model at the upper confidence bound of ξ.

What each institution must get right

An insurer needs OEP and AEP curves at the return periods that match its reinsurance attachment points and its capital requirement, which is 1-in-200 years under Solvency II. A bank needs actuarial losses translated into expected credit loss, the product of probability of default, loss given default and exposure at default. Damage to collateral raises loss given default, which feeds both provisions and regulatory capital. An asset manager needs expected annual loss and a value-at-risk over its holding period, and has little use for a multi-century tail metric it holds no capital against. Applying one institution's metric to another's book misstates the relevant risk and misdirects capital.

Example: how to choose?

A reinsurer is pricing a per-occurrence excess-of-loss cover on coastal windstorm, attaching at USD 500 million. Which curve? OEP, since the trigger is a single event exceeding a threshold, not the annual total. An annual aggregate stop-loss on the same portfolio would call for AEP.

An asset manager assessing the same coastline needs expected annual loss, not the 1-in-250-year OEP point the reinsurer cares about.

The CCR case: from climate scenarios to loss quantiles

France offers an unusual setting for actuarial climate analysis. The Caisse Centrale de Réassurance (CCR) reinsures the national natural catastrophe scheme and centralises the exposures of insurers. For the ACPR climate stress test of 2023–2024, CCR combined Météo-France climate simulations consistent with IPCC pathways with its vulnerability functions and these exposures, and computed physical losses for each insurer. Natural catastrophe losses in France would rise between 2022 and 2050 by 105% in the adverse scenario and by 42% in the reference scenario (ACPR, 2024). The supervisor's main lesson concerns the tail as much as the mean. As presented by the ACPR to the Cookbook working group, the annual mean loss rises from about EUR 2 billion to EUR 5 billion by 2050, while the 1-in-50-year loss rises from about EUR 16 billion to EUR 42 billion, inflation included. A move of the mean by a few billion is manageable, and a move of the tail by tens of billions changes the capital an insurer needs. Few countries can reproduce this exercise, since it relies on centralised exposure data and national climate modelling.

VARIABLE	DESCRIPTION	DATA SOURCE EXAMPLE
OEP curve	Probability that a single event exceeds a loss threshold in a year	Catastrophe model output (Moody's RMS, Verisk, Cotality); open models on the Oasis Loss Modelling Framework

VARIABLE	DESCRIPTION	DATA SOURCE EXAMPLE
AEP curve	Probability that annual aggregate loss exceeds a threshold	Catastrophe model output; reinsurance broker analytics
Damage function	Mapping from hazard intensity to loss ratio	Hazus (FEMA); FLEMOps; JRC global depth-damage functions (Huizinga et al., 2017); insurer claims-calibrated functions
Historical losses	Realised loss experience by peril and region	Munich Re NatCatSERVICE; Swiss Re sigma; PCS (Verisk) and PERILS industry loss indices; insurer and reinsurer claims databases

TABLE 54. Data for actuarial metrics

3.3.4 TEMPORAL CHOICES FOR FINANCIAL METRICS

Section 3.1.4 covered the temporal reference of the hazard itself, historical, real-time or projected. This section addresses how the financial metric is timed once a hazard estimate exists. Four choices matter: the time regime of the estimate, the scenario and its probability, the discount rate, and the reporting cadence.

Live feedback, projections and storylines

A financial metric can follow three time regimes. A live metric follows the event as it unfolds and supports monitoring and early loss estimates. Kayrros, for instance, tracks active fires near monitored assets and delivers damage maps within days of containment. A projected metric applies a climate scenario to a future horizon, as Schneider Electric's ClimVaR does with discounted cash flows to 2050 or CLIMACRED-PHYS with ISIMIP3b forcing. A storyline describes one plausible sequence of events in detail, without a probability. JBA keeps climate change and tipping points outside its stochastic catalogue and treats them through storylines, and the ACPR built its short-term stress test on a narrative of two drought years followed by severe floods (ACPR, 2024). The three regimes answer different questions, and a single model rarely serves all of them.

Choosing a scenario and assigning probabilities

Most scenario sets used in finance carry no probability. NGFS scenarios describe what could happen, not how likely it is, so a user who runs all of them obtains a range without weights. Two approaches fill the gap. The first selects one scenario that matches the purpose of the analysis: an adverse scenario for a stress test, a current-policy scenario for a central estimate. The second assigns probabilities to scenarios. Estimates of this kind place the median warming at about 2.7 °C by 2100, with a very low probability of limiting warming to 1.5 °C (Rebonato et al., 2025). Scientific Climate Ratings applies this logic to the NGFS set extended with two hotter pathways. It weights each scenario by its probability and uses the probability-weighted outcome as its default case (Scientific Climate Ratings, 2025b). The choice depends on the user. A supervisor running a stress test wants a severe but plausible scenario. An investor allocating capital wants an expected value across scenarios.

Point-in-time and through-the-cycle estimates

A bank that folds physical risk into its credit parameters faces two regimes that apply simultaneously. Under IFRS 9, expected credit losses must be point-in-time and forward-looking, reflecting current conditions and reasonable forecasts (IFRS Foundation, 2023). Under the Basel internal ratings-based approach, capital rests on long-run average probabilities of default and on loss given default estimated for an economic downturn (BCBS, 2023). A long-

run average estimated from history smooths a transient physical shock and, more importantly, misses a hazard trend that historical defaults have not yet revealed. The same physical scenario can therefore move provisions quickly and capital slowly.

Discounting future losses

A loss projected decades ahead is not comparable to a loss today without a discount rate. Insurers use the risk-free curves published by EIOPA for their technical provisions, asset managers apply their required return, and IAS 36 asks for a pre-tax rate reflecting the risks specific to the asset. The rate can reverse a decision, since a large loss in 2050 discounted at a high rate looks immaterial today. Whoever sets the rate implicitly sets the risk appetite, and the rate should be reported with the result.

Reporting frequency

Accounting variables (Section 3.3.1) update annually or quarterly, market variables (Section 3.3.2) continuously, and actuarial variables (Section 3.3.3) after each event or at each model release. Combining them requires a common cadence, and that choice decides whether a slow chronic trend or a sudden acute event dominates the result.

Example: how to choose?

A bank wants a single physical risk score per loan. How often should it be refreshed? The cadence should match the decision. Annual for provisioning, which follows the accounting cycle. Daily for a trading book, which follows the market. After each event for collateral revaluation, which follows the actuarial cycle. One model can give three legitimate answers.

VARIABLE	DESCRIPTION	DATA SOURCE EXAMPLE	COMPLIANCE FACTOR
Discount rate	Rate converting a future loss into present value	EIOPA risk-free rate term structures; central bank yield curves; internal WACC	Prescribed for insurers' technical provisions under Solvency II; effective interest rate for IFRS 9 expected credit losses; asset-specific pre-tax rate under IAS 36
Credit parameter regime	Point-in-time or long-run average PD and LGD	Internal rating model documentation; IFRS 9 impairment models	Long-run average PD and downturn LGD for Basel IRB capital; point-in-time, forward-looking estimates for IFRS 9 provisions
Scenario and probability	Climate pathway and the weight given to it	NGFS scenarios; IPCC SSP-RCP pathways; probability-weighted sets (Scientific Climate Ratings)	Supervisory stress tests prescribe their scenarios; climate scenario analysis expected in insurers' ORSA (EIOPA, 2021)
Reporting cadence	Update frequency of the financial metric	Company filings (annual, quarterly); market data feeds (daily); claims databases and catastrophe model releases (event-based, annual)	Annual for financial statements and sustainability reporting; at least annual for ORSA; no regulatory cadence for live monitoring

VARIABLE	DESCRIPTION	DATA SOURCE EXAMPLE	COMPLIANCE FACTOR
Live feed	Near-real-time hazard and damage observation	Satellite services (Kayrros); national flood forecasting services	No specific requirement; supports claims handling and post-event reporting

TABLE 55. Data for Temporal choices

3.3.5 COMPLIANCE ISSUES

The financial metric chosen for physical risk rarely serves analysis alone. It also feeds disclosures, supervisory exercises and internal governance, each with its own requirements. Three families of obligations shape the design choices covered in this Part.

Disclosure

Sustainability reporting standards ask companies to report the financial effects of physical risk. IFRS S2 requires disclosure of the current and anticipated financial effects of climate-related risks, supported by scenario analysis proportionate to the company's circumstances (ISSB, 2023). In the European Union, the climate standard of the CSRD (ESRS E1) asks for the anticipated financial effects of material physical risks, with phase-in provisions, and its scope has been revised under the 2025 simplification package. The EU Taxonomy adds a climate risk and vulnerability assessment for activities that claim to contribute to adaptation or to do no significant harm (European Commission, 2021). These frameworks rarely prescribe a model, but they require the user to state its scenario, horizon and main assumptions, which are precisely the choices discussed in this Part.

Prudential supervision

Supervisors expect banks and insurers to integrate physical risk into their risk management. For banks, the ECB guide on climate-related and environmental risks sets supervisory expectations, and the EBA guidelines on the management of ESG risks extend them across the European Union (ECB, 2020; EBA, 2025). For insurers, EIOPA expects climate change scenarios in the ORSA, including a long-term horizon (EIOPA, 2021). Supervisors also run their own climate stress tests, such as the ACPR exercises on insurers (ACPR, 2024). These exercises prescribe the scenarios and horizons, which constrains the temporal choices of Section 3.3.4.

Model risk and governance

A purchased model inside a capital or provisioning figure falls under the institution's model risk framework, which requires documented methodology, independent validation and version control (Federal Reserve & OCC, 2011) (see Section 3.4.5.).

A vendor that cannot explain a change between two releases exposes its client to a model risk finding. The practical consequence is that an institution should record, for each physical risk metric, the regulatory use it serves. A metric built for a voluntary disclosure may not meet the validation standard of a metric that enters provisions or capital.

3.4 AGGREGATION CHOICES

3.4.1 FROM ASSET TO PORTFOLIO

Every choice described so far produces a number attached to a single asset, whether it concerns the hazard, exposure and vulnerability, or the financial proxy. Banks, insurers and asset managers decide at the level of a portfolio. Consolidating asset-level results into a portfolio metric is a design choice of its own, and a naive consolidation can mislead more than any single upstream choice.

This section leaves aside the scale at which each model family operates. Part 2 already organizes model families by the financial object they address, and a building-level metric answers a different question from a portfolio metric. The question here is narrower. When a provider reports a portfolio figure, how was that figure built from the asset-level results underneath it?

Means add up, tails do not

The loss of a portfolio in a given year is the sum of the losses of its assets in that same year:

$$L_P = \sum_{i=1}^{N} L_i$$

- L_P is the portfolio loss over one year
- L_i is the loss of asset i over the same year
- N is the number of assets in the portfolio

The expected portfolio loss is always the sum of the expected asset losses, whatever the dependence between assets:

$$E[L_P] = \sum_{i=1}^{N} E[L_i]$$

- Where $E[L_P]$ is the expected value, here the average annual loss (AAL)

The tail behaves differently:

$$PML_T\,(L_P) \neq \sum_{i=1}^{N} PML_T(L_i)$$

- $PML_T\,(L_P)$ is the probable maximum loss at return period T, the loss exceeded with an annual probability of $\frac{1}{T}$
- the two sides are equal only when all assets lose together, in the same years

The tail of the portfolio therefore depends on how losses co-move across assets. A portfolio concentrated in one coastal region carries far more tail risk than a geographically spread portfolio with the same expected loss, because one event can hit every asset at once.

Consider this example, a portfolio holds 200 factories along the same monsoon flood system. Each factory has a 2% annual probability of flooding and loses €10 million when flooded.

AGGREGATION ASSUMPTION	EXPECTED ANNUAL PORTFOLIO LOSS	1-IN-100-YEAR PORTFOLIO LOSS
Assets flood independently	€40m	€90m
Sum of site-level 1-in-100 losses	€40m	€2,000m
Actual dependence, partial	€40m	Between €90m and €2,000m

TABLE 56. Illustrative figures

Under independence, the number of flooded sites follows a binomial distribution with 200 trials and probability 2%. Its mean is 4 sites, hence an expected loss of €40m, and its 99th percentile is 9 sites, hence €90m.

The expected loss is identical in every row. The two shortcuts err in opposite directions on the tail. Treating assets as independent understates it, because in the tail these factories fail together. Summing site-level PMLs implicitly assumes that every site floods in the same year, and overstates it. The true figure lies between the two, and only an explicit dependence model or a stochastic event set can locate it.

Five ways of building a portfolio figure

The market uses five aggregation methods, from the simplest to the most demanding. They differ in the assumption they make about dependence between assets, and therefore in what the portfolio figure can be trusted for.

METHOD	DEPENDENCE ASSUMPTION	PORTFOLIO OUTPUT	WHAT IT MISSES
1. No aggregation	None, results stay at asset level	Asset-level table, consolidated by the user	Everything above the asset, and the user often sums by default
2. Summation or weighted average	Independence for losses, none for scores	Expected loss, average hazard score	Concentration and the tail
3. Parametric dependence	Correlation matrix, copula or common regional factor	Loss distribution, VaR, credit VaR	The physical footprint of events, and dependence parameters calibrated on scarce joint data
4. Stochastic event set	Joint footprint of each simulated event	Full portfolio loss distribution (AEP, OEP), PML, TVaR	Indirect losses outside the assets held, and sampling error when few years are simulated
5. Network propagation	Explicit links between assets, firms, sectors or institutions	Indirect losses, network multiplier, systemic VaR	Asset-level precision, and the method requires heavy data on the links

TABLE 57. Aggregation methods for physical climate risk

Notice that Methods 3 and 4 relate to actuarial models and credit risks (sections 2.1.2 to 2.1.3 and 2.2.4), and method 5 to Network Based Models (Part 2.4).

Method 2 deserves particular attention, because it underlies many screening products. An exposure-weighted average of hazard scores tells a user where the portfolio sits on average. It says nothing about how much the portfolio

can lose in a single year, since averaging removes the very concentration that drives the tail. Such scores remain useful for screening and disclosure. They become misleading once a capital, provisioning or pricing decision rests on them.

The methods also combine. A provider can simulate event sets for its main hazard (method 4), fall back on summation for secondary hazards (method 2), and add supply chain scenarios on top (method 5). The portfolio figure then mixes several dependence assumptions, and the reader can only interpret it if each one is stated.

What each institution needs

An insurer needs the full portfolio loss distribution, because correlation drives the capital-relevant tail under Solvency II. Methods 4, and to a lesser extent 3, provide it.

A bank needs expected loss adjusted for concentration by exposure class, because risk-weighted capital and concentration risk depend on how losses cluster across borrowers. Method 3, through credit portfolio models, is the usual route (Part 2.2.4).

An asset manager needs the marginal contribution of each position to total portfolio risk, to decide whether adding or trimming it changes that risk. The Euler allocation gives this contribution:

$$RC_i = w_i \frac{\partial \rho(L_P)}{\partial w_i}$$

Where:

- RC_i is the risk contribution of position i
- w_i is the weight of position i in the portfolio
- $\rho(L_P)$ is the chosen portfolio risk measure, such as VaR or TVaR
- the contributions RC sum to $\rho(L_P)$

A contribution of this kind only carries meaning when $\rho(L_P)$ comes from a dependence-aware method. Under summation, each position contributes in proportion to its standalone risk, and the diversification signal disappears.

An aggregation field on the model card

The model cards in the present cookbook therefore document how each provider moves from asset to portfolio. The field describes the method without grading it. A “method 2” product can be the right choice for screening, and a “method 4” product can be the wrong choice for a user without the data to feed it.

VARIABLE	DESCRIPTION	DATA SOURCE EXAMPLE	ACCESS
Dependence structure	Co-movement of losses across assets in the same event	Event footprints in catastrophe models; studies of the spatial dependence of floods (Jongman et al., 2014; Kerkhofs et al., 2025)	Commercial; published research
Regional dependence parameters	Correlation between regions used to aggregate natural catastrophe risk	Solvency II standard formula: risk factors and correlation matrices by CRESTA zone (EIOPA)	Free

VARIABLE	DESCRIPTION	DATA SOURCE EXAMPLE	ACCESS
Stochastic event set	Simulated events with their joint spatial footprint over many years	Moody's RMS and Verisk (formerly AIR Worldwide) catalogues; open models on the Oasis Loss Modelling Framework; synthetic tropical cyclone tracks from STORM (Bloemendaal et al., 2020); CLIMADA event sets (Aznar-Siguan & Bresch, 2019); global risk model for national asset stocks (Cardona et al., 2013)	Commercial; open source
Portfolio loss distribution (AEP, OEP)	Full distribution of annual portfolio loss, beyond its mean	Portfolio output of catastrophe models; Oasis LMF; CLIMADA	Commercial; open source
Concentration metric	Share of exposure in one region, sector or event footprint	Internal portfolio geocoding; exposure accumulation by CRESTA zone; Herfindahl index of exposures	Internal; free
Network links	Supplier, customer, transport and financial links between entities	Firm-level supply chain data (FactSet Revere, used by Pankratz & Schiller, 2024); multi-regional input-output tables such as EXIOBASE (Stadler et al., 2018) and WIOD (Timmer et al., 2015) (see Part 2.4)	Commercial; free

TABLE 58. Aggregation field of the model card

Provider Box – From the ground up to the portfolio – The case of Altitude

AXA Climate is the climate subsidiary of the AXA group, and Altitude is its software business. The solution is built mainly for unlisted investors, and private equity and infrastructure funds represent about 80% of its clients. Industrial companies form a second group of clients, with fewer sites but more detailed information on each.

Altitude covers 18 physical climate hazards, and further hazards such as hail and coastal erosion are in development. Each hazard has its own model, developed by a scientist specialized in that hazard. The approach is physics-first, and machine learning and Monte Carlo methods are confined to post-processing. Flood is the only fully external component. For each horizon, scenario and return period, maps from Fathom give a water depth per 30 m pixel for river, rain and coastal flooding. Models are calibrated on ERA5, and future projections follow climate models under SSP scenarios, at horizons such as 2025, 2030 and 2050.

The investor describes the company with four items, which are the sector, the asset type, the address and the revenue. For a company with several sites, the value and revenue of each site may be added. The revenue dependency between sites is declared explicitly. When a product passes through three successive factories, all of the revenue depends on each factory, whereas three factories that produce independently each carry a third of it.

Two layers translate hazard into impact. The first is a risk score. A vulnerability matrix for each asset type selects the relevant indicator, and thresholds convert the change in that indicator into a risk level. The second layer quantifies impact through damage curves drawn from published and technical literature, such as the US Army Corps of Engineers and European Union reports. These curves turn a hazard intensity, such as a water depth, into a percentage of building value, a percentage of lost production or a number of days of interruption. The curves differ by region and asset type, and they do not separate stock, roof and structure. Results are expressed in currency when the client provides site values and revenues, and as percentages otherwise.

Aggregation is based on the reliability of averages. The damage curves are deliberately generic and represent broad means. A metre of water may correspond on average to 20% damage on a factory, while individual factories record 10% or 30%. Such an average is reliable across a large portfolio, where errors offset each other, and much less reliable for a single asset. This is the asset-level counterpart of the property set out above, that means add up while tails do not. Work under way on the cost and benefit of adaptation, such as the return on raising backup generators, points toward more asset-specific curves for industrial clients.

Provider Box – From the ground up to the portfolio – The case of Correntics

Correntics is a Zurich-based provider founded in 2021. Its clients are mainly large industrial corporates with long-lived fixed assets. The company sells its software-as-a-service solution directly to corporates, and also through consulting firms and other software providers serving asset managers and banks. Its team comes partly from reinsurance catastrophe modelling. Results are produced at the asset level and can be aggregated to company and portfolio level.

The software offers several aggregation methods to move from site-level assessments to a portfolio view, and the right one depends on the hazard. For its major acute hazards, such as tropical cyclones and floods, Correntics uses a stochastic approach to assess the likelihood of damage and business interruption. This approach addresses the aggregation problem described above. A model that holds one hazard value per site, such as the 100-year flood depth, can only combine site results after the fact. A stochastic event set keeps the footprint of each simulated event. One cyclone or one flood is applied to every site it reaches in the same simulated year, so the portfolio loss for that year already contains the joint failure of nearby assets, and the 1-in-100-year portfolio loss read from these simulated years differs from the sum of the 1-in-100-year losses of each site. This logic comes from insurance and reinsurance catastrophe modelling, and the global risk model prepared for the UN Global Assessment Report applied it to national asset stocks (Cardona et al., 2013). Chronic hazards, such as rising temperatures or water stress, unfold as gradual trends and call for other aggregation methods. Across methods, value-at-risk metrics describe the tail of the risk distribution and point risk managers to the main areas of concern, even in a large portfolio.

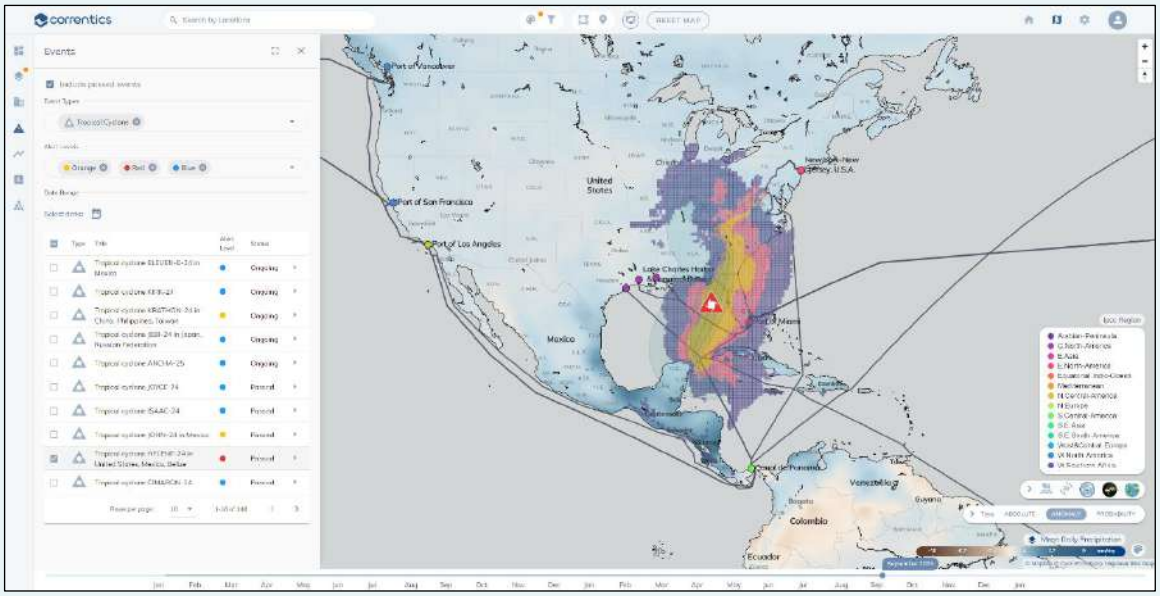


FIGURE 85. Correntics platform interface showing the footprint of Hurricane Helene (September 2024) across a logistics network of ports and shipping routes

Correntics extends the aggregation beyond the client's own sites. The platform assesses risk at nodes of the supply chain, within regions and along transport routes, and runs disruption scenarios on the value chain. A factory with low exposure can still lose output when a supplier or a port on its route is hit. Business interruption of industrial operations is the provider's main focus, and since downtime is the loss that spreads along these links, aggregation along the value chain matters as much as aggregation across owned assets. The same platform also monitors ongoing events. It overlays the forecast footprint of an active tropical cyclone on the client's sites, suppliers and shipping routes, with alert levels, before and during landfall.

At site level, the platform reports a risk matrix of likelihood against impact, with separate scores for damage and for outages. It shows the vulnerability curve used for each site, which links hazard intensity to damage ratio and can be adjusted by the user.

In the terms of the aggregation field, Correntics combines several methods. It uses stochastic event sets (method 4) for its major acute hazards, other aggregation methods for chronic hazards, and value chain scenarios (method 5) for indirect losses. The number of simulated years in its event sets is not disclosed, although the stability of the portfolio tail depends on it.

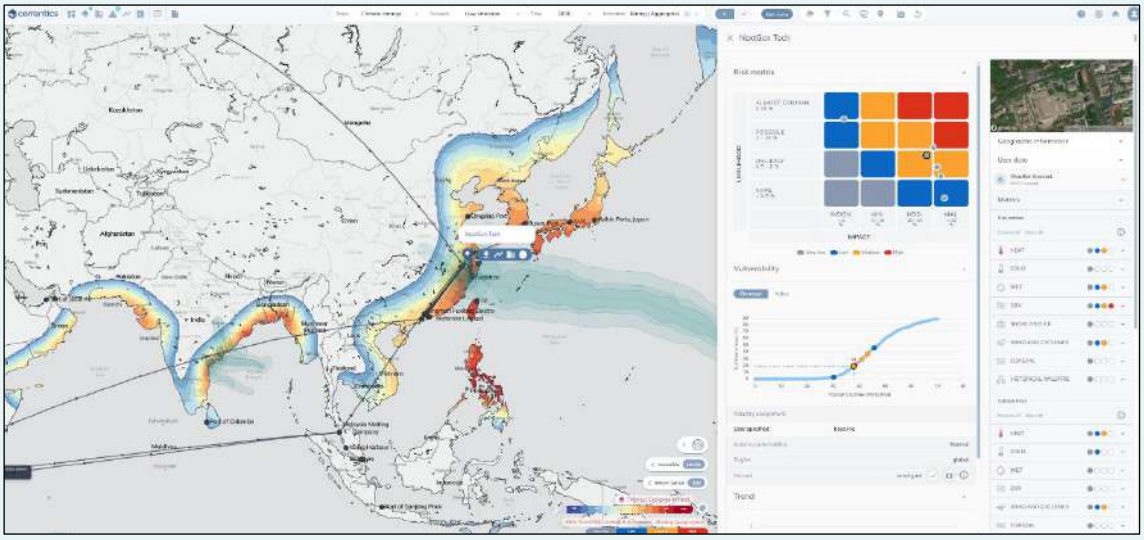


FIGURE 86. Correntics platform interface showing asset-level tropical cyclone risk for a site in East Asia under a low-emission scenario in 2030

3.4.2 THE UNCERTAINTY ISSUE

We have written extensively about the differences between providers. But what they have in common can be compared and therefore assessed. Uncertainty in the model's results can also be considered from one provider to the other.

PROVIDER	SCENARIOS	HAZARD	EXPOSURE	VULNERABILITY	FINANCIAL	ADAPTATION
Aon CRM	4 SSPs, CMIP6 median	In-house + Fathom + STORM	Client (site, occupancy)	Expert classes	% damage ranges, BI days	Advisory only
AXA Climate	3 SSPs	In-house (SpaSM) + Fathom	Client (type, address, revenue)	Literature curves	% loss, revenue-weighted	User-selected measures
Fathom	Any SSP / NGFS	In-house, flood only	Client	Engineering curves	Ground-up, gross and net loss (Oasis LMF)	If issuer supplies
RiskThinking.ai	All SSPs blended	In-house	Own database, 7.5M assets	Hybrid, in-house	DR, BI, CVaR, PD, LGD	Customizable
refinq	RCP 2.6–8.5, SSPs	Public datasets, own downscaling	Client (location, NACE)	Per hazard × NACE × place	EAL, OpEx, CapEx, NVaR	Logged actions
CCR (public, 2016 from Tinard et al.)	RCP 4.5	In-house (Météo-France)	CatNat portfolio	Claims-calibrated	Insured AAL	None

TABLE 59. Comparison of a few physical climate risk model providers for flood assessment

But such a comparison requires common basis of comparability, for example the two models compared must at least belong to the same model family (i.e. the output needs to be somewhat similar). A good example of an approach to compare between models has been recently written (Desaivres, 2026). Here is a breakdown of how the QUALYPSO R package is used to assess uncertainty between models.

Desaivres' model: using QUALYPSO as a way to decompose risks

Hazard, exposure and vulnerability. Each pillar draws on a different discipline, with its own data, spatial scale and level of maturity. Published metrics diverge widely, owing to inconsistent climate model outputs, limited access to reliable asset geolocation and a weakly characterized treatment of vulnerability (Hain et al., 2022).

For an asset a, the loss is the composition of a scenario, a sequence of climate and hazard models, an exposure specification and a vulnerability function (Equation 1). Each function has several defensible implementations, so the output is a distribution over modelling chains.

COMPOSITION OF THE MODELLING CHAIN

$$D_a = f_{\mathrm{DM}}\left(f_{\mathrm{HM}}\left(f_{\mathrm{RCM}}\left(f_{\mathrm{GCM}}(S)\right)\right), X_a\right) \quad (1)$$

- S is the climate scenario,
- X_a the exposure specification of asset a, and f_{GCM}, f_{RCM}, f_{HM} and f_{DM} the global climate, regional climate, hydrological and damage models.

Each f admits several implementations.

A multi-model ensemble (MME) widens the set of choices at every link and makes the dispersion of outputs observable. Measuring how much each link contributes to this dispersion shows where methodological effort raises predictive value fastest, and it separates what can be improved from what cannot.

Taxonomy of uncertainty

Uncertainty enters at several stages of mathematical modelling (Kennedy & O'Hagan, 2002). The table below lists the sources that matter for a modelling chain, with an illustration drawn from a flood chain.

SOURCE	DEFINITION	ILLUSTRATION IN A FLOOD CHAIN
Measurement error	Errors and approximations affecting observational data	Stage and discharge records behind rating curves
Parameter uncertainty	Limited confidence in the numerical values assigned to model inputs	Calibrated parameters of a hydrological model, unit prices in a damage function
Parametric variability	Input variables vary in ways the model structure does not describe	Flood duration, velocity or sediment load fixed at default values
Model ambiguity	Difficulty distinguishing the best model among rival candidates	Choice among four hydrological models or four damage models
Model misspecification	Suitability of the framework and assumptions of the model	Capped depth-damage curves that cannot extrapolate beyond observed depths
Residual variability (RV)	Stochastic fluctuation or irreducible noise preventing exact replication	Annual maxima of daily discharge
Numerical uncertainty	Interpolation, extrapolation or approximation during numerical resolution	Extension of rating curves beyond observed stages

TABLE 60. Sources of uncertainty along a modelling chain

These sources fall into two classes (Der Kiureghian & Ditlevsen, 2009). Aleatory uncertainty reflects the intrinsic randomness of the system and cannot be reduced. Epistemic uncertainty arises from imperfect knowledge or incomplete use of available information and can be reduced. In climate projections the aleatory component is called internal climate variability (Hawkins & Sutton, 2009, 2011).

How uncertainty is quantified

Variance-based methods measure the contribution of input factors to the statistical distribution of model outputs, Monte Carlo sampling estimates confidence intervals around target indicators, and one-at-a-time designs perturb a single input around a baseline (Douglas-Smith et al., 2019). In climate impact work, a first family of studies partitions the variance of an ensemble into scenario, model and internal variability with an additive ANOVA, which requires a balanced and complete ensemble (Hawkins & Sutton, 2009, 2011). The table below compares the main approaches.

APPROACH	PRINCIPLE	MAIN LIMITATION	REFERENCE
One-at-a-time	Perturb one input at a time around a baseline, as in the unsequa module of CLIMADA	Interactions between inputs remain invisible	Pianosi et al., 2016; Kropf et al., 2022
Monte Carlo	Repeated random sampling of inputs to estimate confidence intervals	Requires a probability distribution for each uncertain input	Le Guenedal et al., 2021
Additive ANOVA	Split variance into scenario, model and internal variability	Needs a balanced and complete ensemble	Hawkins & Sutton, 2009
Global sensitivity analysis	Variance-based indices on a spatialised flood damage model	Shares depend on the scale of evaluation	Saint-Geours et al., 2011
QE-ANOVA with data augmentation	Bayesian ANOVA on smoothed responses, with missing chains imputed	Additive structure with scenario-independent model effects	Evin et al., 2019

TABLE 61. Approaches to uncertainty quantification

Evidence from end-to-end flood studies

Uncertainty decomposition is well established for climate projections (Kundzewicz et al., 2008; Hawkins & Sutton, 2009; Deser et al., 2012) and for hydrological projections (Hawkins & Sutton, 2011; Lafaysse et al., 2014; Vetter et al., 2017; Lemaitre-Basset et al., 2021; Aitken et al., 2023; Evin et al., 2026).

Financial studies of flood risk that cover the full chain remain confined to single cities or catchments, and their results conflict (Apel et al., 2004; Saint-Geours et al., 2011; Eleutério, 2012). National studies usually stop at the hazard and do not propagate uncertainty to monetary damage.

Studies of damage models isolate that layer (Merz & Thieken, 2009; Jongman et al., 2012; Prahl et al., 2016), so they miss the rest of the chain and the full range of available models. Standard variance analysis also cannot process the incomplete ensembles that real model portfolios form (Desaivres, 2026).

STUDY	SETTING	FINDING
Apel et al., 2004	Germany, Monte Carlo split of aleatory and epistemic uncertainty	The choice of flood loss model weighs more than hazard modelling
Eleutério, 2012	GIS tool for sensitivity analysis on parallel models	Hazard modelling dominates, especially for frequent floods, with a non-negligible role for protection measures and estimation scale
Saint-Geours et al., 2011	France, variance-based global sensitivity analysis	The share carried by each layer depends on the scale of evaluation, asset or macro level
Jongman et al., 2012	Comparative assessment of damage models	Damage functions are sensitive to vulnerability parameterisation and to exposure land-use data
Merz & Thieken, 2009	Germany, flood risk curves with uncertainty bounds	Epistemic uncertainty, notably in loss estimation, produces wide bounds around flood risk curves
Prahl et al., 2016	Unified damage functions for several climate-related hazards	Damage function uncertainty is large and propagates directly into loss estimates

TABLE 62. Earlier studies on uncertainty in flood risk chains

The Explore2 ensemble

The hazard component rests on the Explore2 hydrological projections for France (Tramblay et al., 2024), distributed through DRIAS-Eau. The high-flow indicator is QJXA, the annual maximum of daily discharge, extracted for every hydrometric station, model combination and year from 1976 to 2100.

Three RCPs constrain 17 GCM-RCM pairs drawn from EURO-CORDEX at 12 km resolution (Coppola et al., 2021). Projections are bias-corrected with ADAMONT, a choice held fixed (Verfaillie et al., 2017). The ensemble is incomplete and unbalanced, since a GCM cannot be paired with every RCM and not every pair exists under each RCP.

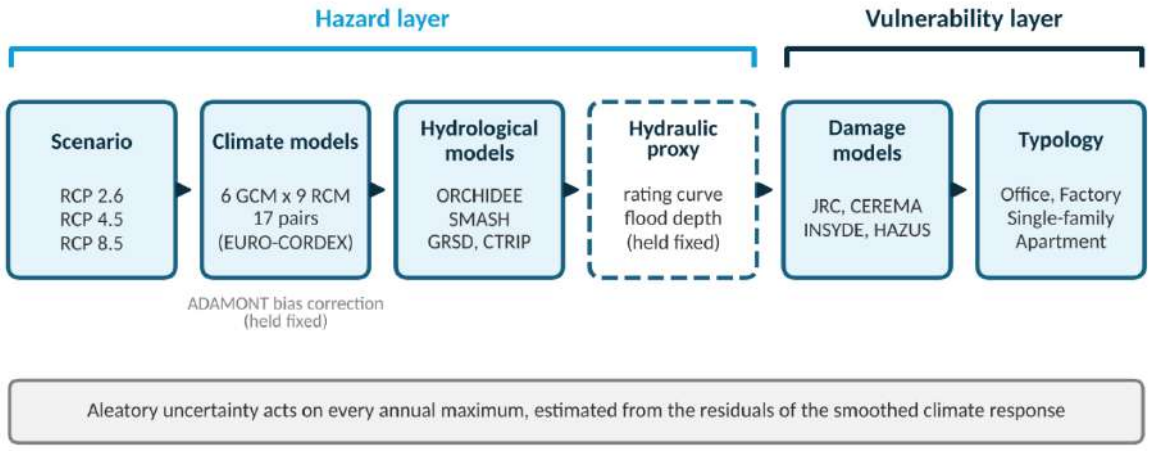


FIGURE 87. Modelling chain and ensemble layers of the case

***NB**: Dashed outlines mark layers held constant across simulations. Original schematic based on Desaivres (2026).*

LAYER	ALTERNATIVES	TREATMENT
Scenario	RCP 2.6, RCP 4.5, RCP 8.5	Varied
GCM	CNRM-CM5, EC-EARTH, IPSL-CM5A-MR, HadGEM2-ES, MPI-ESM-LR, NorESM1-M	Varied
RCM	ALADIN63, HadREM3-GA7-05, RACMO22E, RCA4, HIRHAM5, CCLM4-8-17, RegCM4-6, REMO2009, REMO2015, WRF381P	Varied
Hydrological model	ORCHIDEE, SMASH, GRSD, CTRIP	Varied
Bias correction	ADAMONT (CDF-t also available)	Fixed
Flood depth proxy	Rating curves, empirical curves, nearest-station curves	Fixed
Damage model	JRC, CEREMA, INSYDE, HAZUS	Varied
Building typology	Office building, factory, single-family home, apartment complex	Varied
Scope of losses	Direct tangible costs to buildings	Fixed

TABLE 63. Layers of the ensemble

Across 2,157 catchments, 2,048 modelling chains are available at each catchment. Data augmentation completes the design to 11,520 chains.

Hydrological models

MODEL	ORGANISATION AND TYPE	SPATIAL STRUCTURE	REFERENCE
ORCHIDEE	IPSL, physical continental surface model	Fully distributed, 8 km grid, about 3,600 simulation points	Krinner et al., 2005
SMASH	INRAE, conceptual	Fully distributed, 1 km, used operationally in the Vigicrues Flash warning system	Jay-Allemand et al., 2020
GRSD	INRAE, conceptual (GR4J with CemaNeige snow module)	Semi-distributed, sub-basins centred on stations	De Lavenne et al., 2016

MODEL	ORGANISATION AND TYPE	SPATIAL STRUCTURE	REFERENCE
CTRIP	Météo-France, physical routing model coupled with ISBA	Fully distributed, 1/12 degree, about 2,000 simulation points	Decharme et al., 2019

TABLE 64. Hydrological models of Explore2 used in the ensemble

Each model is calibrated and evaluated on its own set of stations (Sauquet et al., 2022). The number of simulation points ranges from about 2,000 for CTRIP to between 3,600 and 3,850 for the three other models.

Discharge and flood depth

Discharge is converted into stage with the rating curve of each station when one exists. A stage-discharge table reconstructed from archived observations and smoothed with LOWESS replaces it when no curve is available, and the curve of the nearest station serves as a last resort. Beyond observed stages, a power law of the form in Equation 2 is fitted when possible, with a log-log regression and then a linear regression as fallbacks (Herschy, 1998).

THE RATING CURVE

$$Q = a(h - h_0)^b \quad (2)$$

- Q is discharge,
- h is the stage,
- h_0 the stage at zero flow, and a and b are station-specific parameters.

Annual mean stage is taken as the non-overflowing water level. Flood depth is the difference between the annual maximum stage and this mean (Equation 3). The processing is identical across simulations, so it adds no dispersion to the ensemble.

FLOOD DEPTH PROXY

$$d_{c,m}(t) = h_{c,m}^{\max}(t) - h_{c,m}^{\text{mean}}(t) \quad (3)$$

- c indexes the station,
- m the modelling chain,
- t is the year.

The superscripts denote the annual maximum and the annual mean stage.

Vulnerability and exposure

Flood impacts are classified into direct, business interruption, indirect, intangible and risk mitigation costs (Meyer et al., 2013). The ensemble covers direct tangible costs to buildings only, which keeps the four damage models structurally comparable.

Damage models convert flood depth into monetary loss through depth-damage functions. They are either empirical, calibrated on reconstruction costs or insurance claims, or synthetic, built from expert judgement and engineering knowledge. Both types are tied to a cost level and a reference year, and the functions are capped, so extrapolation beyond observed depths is generally infeasible (Cammerer et al., 2013). Every model is converted to 2025 euros with a price correction factor.

MODEL	ORIGIN	CALIBRATION	FORM	INPUTS AND SCOPE
JRC	European Commission Joint Research Centre, 2007	Empirical, literature survey and international construction cost surveys	Relative to a maximum damage per m^2	Residential, commercial and industrial classes (Huizinga, 2007)
CEREMA	France, CEPRI 2014 and CEREMA 2024	Synthetic, Batiprix costs, SIMUDOM and Floodam	Absolute, per m^2	Residential curves by typology, basement and flood duration, commercial curves by NAF code, property damage only, 50% depreciation, floods shorter than 48 hours (Magnier et al., 2018; Bremond et al., 2024)
INSYDE	Italy, 2015	Synthetic, component by component, Italian prices of 2013	Absolute, sum of repair costs	Depth, velocity, sediment load, duration and water quality with a detailed building description (Dottori et al., 2016)
HAZUS	United States, FEMA, 2022	Empirical, diverse US sources	Relative to replacement cost, depth in feet	33 curves over seven occupancy classes, structural damage only (FEMA, 2024)

TABLE 65. Damage models of the ensemble

DAMAGE MODELS AS IMPLEMENTED

$$D_a^{\mathrm{JRC}} = \delta^{\mathrm{JRC}}(c,i) \times \mathrm{MaxV}_c \times S_a \times \pi_{2007,\mathrm{EU}} \quad (4)$$

$$D_a^{\mathrm{CEREMA}}(k,i,t) = \theta_{k,i,t} \times S_a \times \pi_{2016,\mathrm{FR}} \quad (5)$$

$$D_a^{\mathrm{INSYDE}} = \sum_{k=1}^{n} C_k\left(P_e, B_p, U_p\right) \times \pi_{2013,\mathrm{IT}} \quad (6)$$

$$D_a^{\mathrm{HAZUS}}(k,i) = \delta_{k,i} \times s_a \times r^{2022}_{\mathrm{USD/EUR}} \times \pi_{2022,\mathrm{FR}} \quad (7)$$

- δ is a relative damage ratio.
- $MaxVc$ the maximum damage per m^2 of typology class c.
- S_a and s_a the floor area in m^2 and ft^2, θ a damage per m^2 for typology or activity k, depth i and duration t
- C_k the repair cost of component k,
- P_e the flood event parameters,
- B_p the building characteristics,
- U_p the unit prices, r the USD to EUR exchange rate,

And π a correction factor to 2025 euros specific to each base year and country.

Exposure portfolio

A synthetic cluster of four buildings is placed at every station and exposed to the same flood depth, with each asset located below the water level. This design removes the effect of asset location, so that the ensemble isolates the convergence of the models.

ASSET	FLOOR AREA (M^2)	FLOORS	BASEMENT	BUILT	JRC CLASS	CEREMA CURVE	HAZUS CURVE
Office building	500	10	Yes, 250 m^2	2020	Commercial	Companies, NAF 6820A	217
Factory	1,000	0	No	2020	Industrial	Companies, NAF 3530Z	545
Single-family home	100	1	Yes, 50 m^2	1990	Residential	Residential	121
Apartment complex	200	4	No	1990	Residential	Residential	110

TABLE 66. Experimental portfolio

Floor numbers as reported in the source. The remaining INSYDE inputs take the default values proposed by the authors.

Variance partitioning with QUALYPSO

QUALYPSO couples a Bayesian analysis of variance with a time-series approach, which makes it applicable to incomplete and unbalanced ensembles (Evin et al., 2019). The R package implements three steps (Evin, 2026).

1. **Smoothing.** A cubic smoothing spline isolates the long-term climate response of each chain from interannual noise.
2. **Data augmentation.** Missing combinations are treated as latent variables, and the joint posterior of the ANOVA parameters and of the unobserved projections is sampled.
3. **QE-ANOVA.** The mean response and the main effect of each layer are estimated under an additive structure.

The climate change response is the smoothed response minus its value over the reference period. Aleatory variance is estimated from the residuals of the smoothing and averaged across chains (Equations 8 to 10).

CLIMATE CHANGE RESPONSE AND ALEATORY VARIANCE

$$\mathrm{CCR}_i(t) = \mathrm{CR}_i(t) - \mathrm{CR}_i(t_{\mathrm{ref}}) \quad (8)$$

$$\varepsilon_i(t) = Y_i(t) - \mathrm{CR}_i(t) \quad (9)$$

$$V_{\mathrm{AU}} = \frac{1}{N}\sum_{i=1}^{N} \mathrm{Var}\left(\varepsilon_i(t)\right) \quad (10)$$

- $Yi(t)$ is the raw damage projection of chain i,
- $\mathrm{CR}_i(t)$ is the smoothed climate response,
- t_{ref} the reference period 1975 to 2005,
- N the number of chains.

The response is an absolute change, so projections are centred and can be negative when future damage falls below the reference level.

Under additivity, the response of chain i is the sum of a grand mean, one main effect per layer and a residual (Equation 11). The residual variability RV collects interactions among models and scenarios, unmodelled sources and numerical approximations. **Epistemic variance is the most important indicator here. It is the sum of the variances of the main effects and of RV, and each uncertainty fraction is a ratio of variances** (Equations 12 to 14).

ADDITIVE DECOMPOSITION AND UNCERTAINTY FRACTIONS

$$\mathrm{CCR}_i(t) = M(t) + \sum_{x \in X} E_x^{(i)}(t) + \eta_i(t) \quad (11)$$

$$V_{\mathrm{TOT}}(t) = \sum_{x \in X} \mathrm{Var}\left(E_x(t)\right) + \mathrm{RV}(t) + V_{\mathrm{AU}} \quad (12)$$

$$\mathrm{UF}_x(t) = \frac{\mathrm{Var}\left(E_x(t)\right)}{V_{\mathrm{TOT}}(t)} \quad (13)$$

$$\mathrm{UF}_x^{\mathrm{EU}}(t) = \frac{\mathrm{Var}\left(E_x(t)\right)}{V_{\mathrm{EU}}(t)} \quad (14)$$

- M is the grand mean.
- $E_x^{(i)}$ is the main effect, in chain i, of the model chosen for layer x in $X = \{CS, GCM, RCM, HM, DM, Typo\}$, which stands for climate scenario, GCM, RCM, hydrological model, damage model and typology.
- The term η_i is the residual, $RV(t)$ its variance across chains, and $V_{EU}(t) = V_{TOT}(t) - V_{AU}$

Please note that three assumptions underpin the method.

- **Smooth climate response.** Short-term fluctuations around the climate response are attributed to internal variability.
- **Quasi-ergodicity.** Temporal variance estimated on one run reflects the variability across space and ensemble members.
- **Additivity.** Interactions are not represented explicitly and flow into the residual term.

PARAMETER	VALUE
Trend model	Cubic smoothing spline, smoothing parameter 10,000

PARAMETER	VALUE
Rolling window	30 years
Reference period	1975 to 2005
Future period	2006 to 2100
Change variable	Absolute
Short, mid and long term	2005 to 2040, 2040 to 2070, 2070 to 2100

TABLE 67. QUALYPSO configuration and horizons

Four diagnostics complete the partition.

The overall mean signal is the temporal mean of the response over the study period. The relative standard deviation compares dispersion to this mean (Equation 15). Ensemble quantiles at 5, 50 and 95% summarize the spread for RCP 8.5 and the factory typology. Individual model impact measures the share of a layer's main effect carried by each of its models.

A last indicator splits epistemic variance between a hazard set $H = \{CS, GCM, RCM, HM\}$ and a vulnerability set composed of DM and Typo (Equation 16).

RELATIVE DISPERSION AND HAZARD SHARE

$$\mathrm{RSD} = \frac{\sigma_{\mathrm{pred}}}{|\mu_{\mathrm{CCR}}|} \quad (15)$$

$$S_H(t) = \frac{\sum_{x \in H} \mathrm{Var}(E_x(t))}{\sum_{x \in X} \mathrm{Var}(E_x(t))} \quad (16)$$

- σ_{pred} is the standard deviation across predictions and μ_{CCR} is the overall mean signal.
- The vulnerability share is $S_V = 1 - S_H$, computed on the main effects without RV and AU.

Results

1. Dispersion of the damage signal

The standard deviation exceeds the overall mean signal at every catchment, with a relative standard deviation above 1. It rises above 5 at 10% of catchments and above 10 at 5% of them, which marks a strong lack of convergence among predictions (Desaivres, 2026). Dispersion is highest for the office and factory typologies, where extreme chains at both ends of the distribution produce pronounced outliers, while the single-family home and apartment typologies remain stable across chains.

2. Aleatory and epistemic uncertainty

Aleatory uncertainty dominates at every catchment and horizon. Its share of total variance falls from about 94% in the short term to about 85% in the long term, and its spatial dispersion widens at long horizons. Epistemic uncertainty carries the remainder, and within it residual variability is the largest component at both horizons.

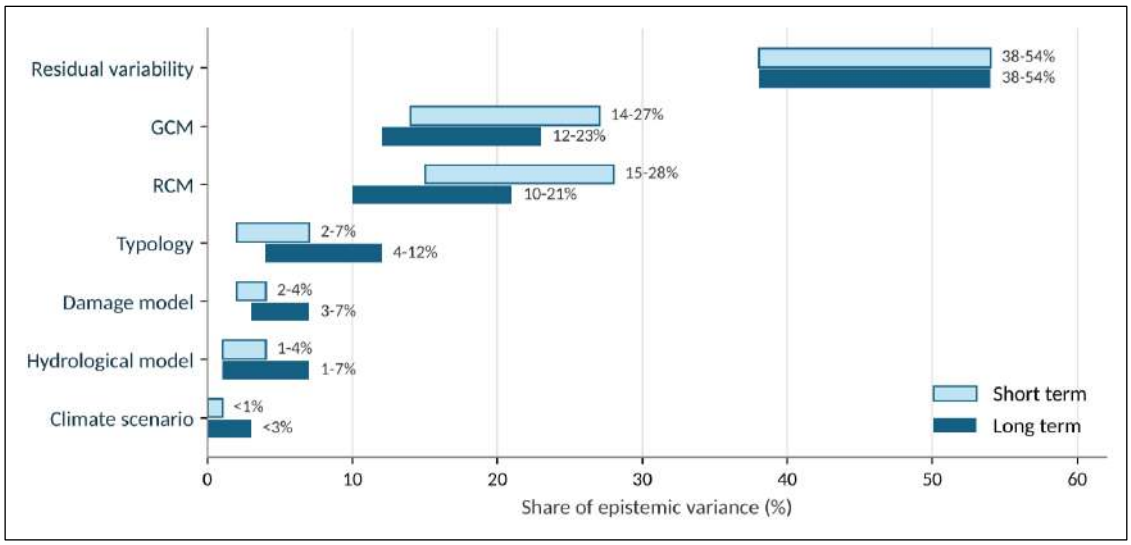


FIGURE 88. Range across catchments of each component's share of epistemic variance
Adapted from Desaivres (2026).

COMPONENT	SHORT TERM	LONG TERM	DIRECTION
Residual variability	38 to 54	38 to 54	Stable
GCM	14 to 27	12 to 23	Falls
RCM	15 to 28	10 to 21	Falls
Typology	2 to 7	4 to 12	Rises
Damage model	2 to 4	3 to 7	Rises
Hydrological model	1 to 4	1 to 7	Rises
Climate scenario	Below 1	Below 3	Rises slightly

TABLE 68. Share of epistemic variance by component (percent, range across catchments)

Once aleatory uncertainty and residual variability are set aside, climate models (GCM and RCM) carry 29 to 55% of the remaining epistemic variance in the short term, against 4 to 11% for the damage layer (damage model and typology) and 1 to 4% for hydrological models.

At long horizons the climate models fall to 22 to 44%, while the damage layer rises to 7 to 19% and hydrological models reach 1 to 7%. The hazard layer therefore constrains short-term assessments, and the vulnerability layer becomes a primary source of uncertainty at long horizons. The climate scenario stays below 3% throughout, which agrees with recent national findings on high flows (Evin et al., 2019, 2024, 2026).

Individual models

The main effect of an individual model measures how far it pulls its layer away from the ensemble mean.

LAYER	HIGHEST CONTRIBUTION	LOWEST CONTRIBUTION	REMARK
Climate scenario	RCP 8.5, 40% to 60% at long term	RCP 4.5, 5% to 28% at long term	All scenarios near 33% in the short term
GCM	IPSL-CM5A-MR, 28% (9 to 44%)	ICHEC-EC-EARTH, 8% (2 to 11%)	Average individual effects between 16 and 30%
RCM	WRF381P, 18% (9 to 25%)	HadREM3-GA7-05, 6% (2 to 8%)	High spatial dispersion for ALADIN63, REMO2015, REMO2009 and WRF381P
Hydrological model	CTRIP, 19 to 62%	GRSD, 6 to 27%	ORCHIDEE 14 to 57%, SMASH close to GRSD with more outliers
Damage model	CEREMA, 47 to 62%	JRC, 1 to 3%	HAZUS 14 to 43%, INSYDE 3 to 24%
Typology	Factory, 54 to 62%	Office building	Time plays a negligible role

TABLE 69.Individual model effects within each layer

Large contributions often come from a small subset of models (Evin et al., 2021). IPSL-CM5A-MR and NorESM1-M carry a wet signature, meaning a systematic tendency toward higher precipitation, soil moisture or streamflow (Evin et al., 2026; McSweeney et al., 2015).

Among RCMs, WRF381P is wet, while HadREM3-GA7-05, RACMO22E and RCA4 lean dry (Evin et al., 2021, 2026). ORCHIDEE retains more water in catchments and projects higher streamflow than other hydrological models under identical forcing, whereas CTRIP exerts the strongest opposing influence.

SMASH and GRSD are conceptual bucket models with comparable intermediate effects, CTRIP and ORCHIDEE are physically based models without automatic local recalibration, and models of one family share process descriptions that cluster their uncertainty (Sauquet et al., 2026).

Damage models and typologies

CEREMA yields the lowest damage with very low dispersion across catchments, HAZUS the highest with numerous outliers, and JRC and INSYDE sit between them with fewer outliers. The response is larger for the office and the factory than for residential assets, and the factory combined with HAZUS shows the strongest dispersion overall. A single-storey building with a large footprint and a long external perimeter concentrates value at ground level, which makes its loss estimates very sensitive to flood depth (Desaivres, 2026). The gap between models reflects different financial baselines.

MODEL	FINANCIAL BASELINE	CONSEQUENCE
JRC	Relative functions applied to maximum asset values derived from national economic indicators	Intermediate values
CEREMA	French cost-benefit framework for public investment, with depreciated values and capped flat-rate envelopes in place of market replacement cost	Lowest values, and commercial curves designed to be combined with equipment, stock and employee curves (Bremond et al., 2024)
INSYDE	Component-by-component replacement to like-new, simulating physical restoration	Intermediate to high values
HAZUS	Percentage of total replacement cost, calibrated on US construction data, insurance claims and engineering benchmarks	Highest values

TABLE 70. Financial baselines of the damage models

Extreme chains

Ranking chains by long-term mean signal for a fixed scenario and typology isolates the five highest and five lowest chains. High signals concentrate in three quadruplets of scenario, GCM, RCM and damage model that each appear four times, whereas low signals spread over ten quadruplets with four occurrences each (Desaivres, 2026).

LAYER	HIGH-SIGNAL CHAINS	LOW-SIGNAL CHAINS
GCM	IPSL-CM5A-MR	CNRM-CM5
RCM	ALADIN63, HIRHAM5, RCA4	ALADIN63, HIRHAM5, RCA4
Hydrological model	ORCHIDEE, GRSD	CTRIP, SMASH
Damage model	HAZUS, INSYDE	HAZUS, INSYDE

TABLE 71. Models recurring in extreme chains

Some pairs recur systematically. ALADIN63 always appears with CNRM-CM5 in extreme chains, and HadGEM2-ES with RegCM4-6 or HadREM3-GA7-05. These patterns point to multiplicative interactions that compound individual biases, and HAZUS and INSYDE act as signal multipliers at both ends of the distribution.

Geographical patterns

Median damage is projected to rise in the north and fall in the south, and the models disagree on the sign of the change across the territory. For the factory under RCP 8.5, changes reach about -200 euros in the south and +200 euros in the north. The table below summarizes the spatial structure of the partition.

INDICATOR	SPATIAL PATTERN
Relative standard deviation above 1,000%	Pyrénées-Atlantiques, Hautes-Pyrénées, Provence-Alpes-Côte d'Azur, north of Lozère, Savoie, Doubs, Meuse
Aleatory share	Above 60% everywhere in the short term, lower epistemic share in the south-east at long term, higher epistemic share in Finistère, Calvados, Orne, Sarthe, Gers and Hautes-Pyrénées
Climate scenario	Below 10% everywhere, higher in the south-east of Provence-Alpes-Côte d'Azur at long term
GCM	High in Bretagne and the Mediterranean basin in the short term, above 40% in Hérault, Aude and Gard at long term, lowest in Hauts-de-France and Morbihan
RCM	Highest in the south-west at about 40%, lowest in the north-east below 10%
Hydrological model	Below 10% in the west, above 20% in Meuse and Haute-Marne, above 30% in Doubs and Jura
Damage model and typology	Between 0 and 30%, rising over time in the north, with typology at 20 to 30% in Vaucluse, Bouches-du-Rhône and Drôme
Residual variability	Largest at long term in the north-east, Provence-Alpes-Côte d'Azur, south of Nouvelle-Aquitaine and eastern Bretagne

TABLE 72. Spatial structure of uncertainty in France

Individual models carry spatial signatures. CNRM-CM5 projects strong positive effects in the south and negative effects in the north, the exact opposite of NorESM1-M. ALADIN63 and HadREM3-GA7-05 project negative effects in the south and positive effects in the north, while REMO2015 and WRF381P show the reverse. ORCHIDEE gives strong positive effects across the territory, and CTRIP gives strong negative effects in the east, especially in Doubs and Jura.

Interpretation

The dominance of aleatory uncertainty reflects the chaotic behaviour of annual high-flow indicators around the underlying damage signal. It originates in internal climate variability (Deser et al., 2012) and in high-flow dynamics (Bonnet et al., 2020), and it cascades through the chain. Calibrating adaptation on the mean response alone omits a large envelope of plausible futures, including unusual or critical years.

The size of the residual term signals strong higher-order interactions. Candidates are non-additive interactions between GCM and RCM, between scenarios and hydrological models, and between damage models and building typologies, since the performance of a damage model depends on asset characteristics. Unmodelled sources and the numerical effects of data augmentation may contribute as well. The additive QUALYPSO framework cannot separate these terms, and ad hoc approaches (Evin et al., 2021) or Bayesian random-effect models (Northrop & Chandler, 2014) offer routes forward.

The spatial scatter of damage model and typology shares is two to three times smaller than that of climate and hydrological models. This follows from the design, since damage functions carry no spatial parameter and the portfolio is uniform. Spatial variations in damage shares are emergent properties of the decomposition and reflect the convergence of hazard drivers. A portfolio with real regional architecture would add geographic dependencies to the vulnerability layer (Desaivres, 2026).

Key Limitations

- **Ensemble dependency.** Consider that the variance shares depend on the composition of the ensemble. Bias correction and hydraulic processing were held fixed, and alternatives could change the shares.
- **Synthetic exposure.** A uniform portfolio isolates model convergence and omits real geographic variation in exposure.
- **Methodological constraints.** The incomplete RCP, GCM and RCM matrix prevents explicit estimation of interaction terms. QUALYPSO assumes scenario-independent GCM and RCM effects and stationary aleatory uncertainty.
- **Risk modifiers.** The framework evaluates raw physical damage and omits flood defences, emergency planning and insurance indemnification.

Generalisation to other hazards: A generic chain

The protocol applies to any metric built as a composition of L layers (Equation 17). The partition then extends to interactions among layers (Equation 18). In the additive model, RV stands in for the interaction terms, together with unmodelled sources and numerical effects.

GENERIC CHAIN AND PARTITION

$$D = (f_L \circ f_{L-1} \circ \cdots \circ f_1)(S, X) \qquad (17)$$

$$V_{\mathrm{TOT}} = \sum_{j=1}^{L} \mathrm{Var}(E_j) + \sum_{\mathrm{j<k}} \mathrm{Var}(I_{\mathrm{jk}}) + V_{\mathrm{AU}} \qquad (18)$$

Layers f_1 to f_L cover the scenario, the climate and hazard models, the exposure and the vulnerability functions. E_j is the main effect of layer j, I_{jk} the interaction between layers j and k, and V_{AU} the aleatory variance.

HAZARD	CLIMATE LAYER	HAZARD LAYER	VULNERABILITY LAYER	POINT OF ATTENTION
Fluvial flood	RCP, GCM, RCM	Hydrological model, rating curve, flood depth	Depth-damage function	Tested in the source study. Annual maxima, rating curve extrapolation, capped functions
Coastal flood	RCP, GCM, regional sea level	Surge and wave models, inundation depth	Depth-damage function	The vulnerability layer is shared with fluvial flood, so damage model divergence carries over
Wildfire	RCP, GCM, RCM	Fire weather index, ignition and spread model	Burn exposure, damage ratio	Several hazard indicators exist, so comparability across models has to be checked
Tropical cyclone and windstorm	RCP, GCM	Track and wind field model	Wind damage function	Few extreme events per sample, so a large aleatory share is expected
Heat stress	RCP, GCM, RCM	Wet-bulb or heat index	Productivity or energy demand function	Smoother indicators than annual maxima, so the aleatory share may be lower
Drought	RCP, GCM, RCM	Soil moisture or drought index, hydrological model	Yield or operating loss function	Persistent events, so stationarity of aleatory variance has to be tested

TABLE 73. Mapping of the chain to other hazards

This is an extension proposed for the Cookbook from the structure of the source study. The expectations in the last column are hypotheses that remain to be tested.

Conditions for transfer

1. **Several implementations at every layer.** A layer with a single model contributes no separable variance, and its uncertainty remains unquantified.
2. **Structurally comparable outputs.** Models must share a metric, a scope of losses and a price base year. The source study restricts losses to direct building costs and converts every damage function to 2025 euros.
3. **Explicit list of fixed layers.** Layers held constant, such as bias correction or hydraulic processing, are declared, because the partition is conditional on them.
4. **A design that is complete or can be augmented.** Data augmentation fills missing combinations, and its numerical error feeds the residual term.

5. **A check on additivity.** A large residual share calls for a dedicated interaction analysis before model selection.

Reading the results for decisions

Where the aleatory share is large, a metric should report the spread of outcomes in addition to a central estimate, since volatile and critical events remain plausible whatever the model choice. Where epistemic uncertainty is shared between hazard and vulnerability layers, effort should go to reducing model divergence and auditing dependencies among models. Extreme chains arise from specific pairs of models, so selection should operate at the level of the full chain.

Climate narratives offer a practical alternative to purely statistical ensembles. They select coherent combinations of scenario, climate model and hydrological model that describe distinct futures, in line with the storyline approach (Shepherd et al., 2018). Explore2 provides four contrasting narratives, selected on mean changes in temperature and precipitation (Calmel et al., 2025; Héraut et al., 2025).

GCM	RCM	RCP	STORYLINE
EC-EARTH	HadREM3-GA7-05	8.5	Strong warming and significant drying in summer and on an annual basis
CNRM-CM5	ALADIN63	8.5	Relatively minor future changes
HadGEM2-ES	CCLM4-8-17	8.5	Strong warming and marked seasonal variations in precipitation
HadGEM2-ES	ALADIN63	8.5	Significant warming and increased precipitation

TABLE 74. Climate narratives available in Explore2

Choosing the single best model in an ensemble is a problem of model ambiguity. Expected utility theory treats model uncertainty as risk with known probabilities, which assumes a precision of knowledge that is unavailable here (Knight, 1921/1965; Berger & Marinacci, 2020). Deep uncertainty calls for decision criteria that secure robust actions, meaning actions that perform acceptably across a plausible set of models (Brock & Hansen, 2018; Mukerji, 2009). The worst-case criterion of Equation 19 is the simplest example.

ROBUST CHOICE ACROSS MODELS

$$a^* = \mathbf{argmin}_{a \in A} \mathbf{max}_{m \in M} L(a, m) \qquad (19)$$

- A is the set of actions,
- M the set of plausible models,
- $L(a, m)$ the loss of action a under model m.

Protocol summarized

1. **Define the chain.** List every layer from scenario to vulnerability and record the alternatives available at each.
2. **Declare fixed layers.** State which layers are held constant and why.
3. **Build a comparable ensemble.** Harmonise units, loss scope and price base year across implementations.
4. **Fit the decomposition.** Smooth the responses, augment missing chains and estimate main effects, with QUALYPSO or an alternative that allows interactions.
5. **Report the partition.** Report the 7 indicators of the table below for each horizon.
6. **Inspect extreme chains.** Rank chains by mean signal and review recurrent model pairs.
7. **Read the geography.** Map shares by station, region or asset type to detect local trade-offs.
8. **Decide.** Select models, narratives or a robust criterion in light of the partition.

INDICATOR	DEFINITION	READING RULE
Aleatory share	V_{AU} divided by V_{TOT}	A large share calls for reporting the full envelope of outcomes
Relative standard deviation	Standard deviation divided by the absolute mean signal	A value above 1 means that dispersion exceeds the mean signal
Residual share of epistemic variance	RV divided by V_{EU}	A large share points to interactions among layers
Hazard and vulnerability shares	S_H and S_V of Equation 16	Compare across horizons to see where effort pays
Leading models per layer	Share of each layer's main effect carried by each model	Identify models that pull the ensemble
Extreme chains	Combinations recurring among the five highest and lowest chains	Review recurrent pairs for interactions
Spatial maps	Station-level shares of each component	Select models by region or asset type

TABLE 75. Reporting template

The Borowiak approach to model discrimination

The variance partition of the previous section shows how much each layer of a modelling chain contributes to the dispersion of outputs. It leaves open a second question, which of the rival models the observed data support. The table below names this source model ambiguity. QUALYPSO measures its size across an ensemble of projections, yet it never confronts a chain with observed losses and therefore cannot resolve it.

Model discrimination addresses that question. A discrimination procedure compares rival models against the same observations and returns a ranking, a selection or a formal test (Borowiak, 1989). This section presents the framework and its non-parametric extension by resampling (Kallel et al., 2002). It then proposes a way to apply both to physical risk models, as a tool to compare hazard maps, loss estimates and model performance. The approach has yet to be tested, and the application below is a research proposal.

Fit and stability: three questions

Two properties determine the usefulness of a model. Fit is the closeness between predicted and observed responses. Stability is the dependence of the fitted model on the particular sample that was observed. A model that fits one sample closely and shifts markedly on another offers weak ground for a financial decision, so any selection procedure has to assess both properties jointly (Borowiak, 1989).

The choice among rival models splits into three questions that differ in what is assumed before the data are examined (Atkinson, 1969). The table below lists them with the statistics usually associated with each.

QUESTION	ASSUMPTION AND QUESTION	USUAL STATISTIC	STRUCTURE
A1	A particular model is true. Is there evidence that another model fits the data better?	F test for nested models, Cox test for non-nested models (Cox, 1961, 1962)	Asymmetric, pairwise
A2	The models fit the data equally well. Is there evidence to the contrary?	Likelihood ratio test for non-nested models (Vuong, 1989)	Symmetric, pairwise
A3	One model of the candidate set is true. Is there enough evidence to select it?	Probabilities of correct and incorrect selection, AIC-type criteria	Global over the candidate set

TABLE 76. The three questions of model discrimination

The relation between candidate models decides which question applies. Econometric specifications of one family, obtained by adding or removing explanatory variables, are nested, and A1 or A3 apply with the F test as a natural statistic. A physically based damage model and an econometric or machine learning model fitted to the same hazard and exposure are non-nested. Neither reduces to the other and neither carries prior theoretical authority, so A2 is the question that rests on firm ground for such cross-family comparisons (Atkinson, 1969; Vuong, 1989).

Discrimination functions

Any statistic computed from observed responses and predicted values and used to select, eliminate or rank models is a discrimination function (Borowiak, 1989). The residual of each model at each observation is the starting point (Equation 20). The quadratic discrimination function weights the vector of residuals and adds a constant specific to the model (Equation 21).

RESIDUALS AND QUADRATIC DISCRIMINATION FUNCTION

$$R_j(i) = Y_j - f_j(i) \qquad (20)$$

$$Q(i) = R(i)^{\top} S(i)\, R(i) + c(i) \qquad (21)$$

- $R_j(i)$ is the residual of model i at observation j.
- Y_j is the observed response at observation j, for example a recorded loss.
- $f_j(i)$ is the value predicted by model i at observation j.
- $R(i)$ is the vector of the n residuals of model i, and the superscript T denotes its transpose.
- $S(i)$ is an n by n positive semi-definite weighting matrix.
- $c(i)$ is a constant specific to model i, used for instance to penalise the number of parameters.

The least quadratic method selects the model with the smallest Q(i).

Setting $S(i)$ to the identity matrix and $c(i)$ to zero gives the residual sum of squares, and selecting the model with the smallest value is the least squares rule. This rule is fair only when rival models carry the same number of parameters. Otherwise $S(i)$ and $c(i)$ must offset the difference, in the same spirit as the complexity penalty of the Akaike information criterion (Borowiak, 1989).

Two cases follow. When the models are completely determined, with no parameter estimated from the data, exact probabilities of correct and incorrect selection can be computed and formal tests of A1 and A3 constructed. When parameters are estimated from the data, the analysis calls for criterion functions, linearisation and diagnostics of parameter stability, and resampling methods such as the jackknife and the bootstrap replace normal theory when the error distribution is in doubt (Borowiak, 1989).

The Sørensen-Dice score as a discrimination function

Hazard maps produced by rival models are commonly compared with overlap scores such as the Sørensen-Dice coefficient. With maps coded 0 or 1, the residual equals plus or minus one at a disagreement and zero elsewhere, so the sum of squared residuals counts false positives and false negatives. The complement of the coefficient is then a quadratic discrimination function in the sense of Equation 21 (Equation 22).

SØRENSEN-DICE COMPLEMENT AS A QUADRATIC FORM

$$1 - \mathrm{Dice}(i) = \frac{FP(i)+FN(i)}{P+\hat{P}(i)} = R(i)^{\top} S(i)\, R(i), \quad S(i) = \frac{I}{P+\hat{P}(i)}, \quad c(i) = 0 \qquad (22)$$

- $FP(i)$ is the number of cells predicted exposed by model i and observed unexposed.
- $FN(i)$ is the number of cells observed exposed and missed by model i.
- P is the number of cells observed exposed.
- $\hat{P}(i)$ is the number of cells predicted exposed by model i.
- I is the identity matrix.

The weighting divides by the combined extent of observed and predicted exposure. It therefore corrects for the footprint of each model, a correction distinct from the parameter penalty of information criteria. The denominator grows with the predicted extent, which dampens the penalty on a model that over-predicts the hazard area, a bias to

document whenever maps of very different extent are compared. Used alone, the score is descriptive, with no significance level or selection probability attached. Reading it as a discrimination function opens the way to the resampling tools below, which would attach a measure of stability to each map comparison. A continuous intensity such as flood depth calls for a discrimination function built on the continuous residual.

Why parametric tools fail for complex models

Many models reviewed so far (Merz, Kreibich & Lall, 2013; Pilkington & Mahmoud, 2016; Wagenaar, de Jong & Bouwer, 2017) are nonlinear, calibrated by numerical optimisation, and in some cases built as neural networks.

A multilayer perceptron with p inputs and H hidden units carries $H(p+1)+H+1$ parameters that interact nonlinearly (Kallel et al., 2002). It has no closed-form likelihood, no guaranteed global minimum and no unique parameterisation, so the parametric machinery of discrimination functions becomes intractable. Training error favours over-fitted models, and a single held-out test set gives one noisy estimate with no measure of stability. Fit and stability remain the relevant criteria, and resampling offers a way to measure them without distributional assumptions.

Bootstrap discrimination

The bootstrap builds new samples by drawing observations with replacement from the observed data (Efron, 1979). Applied to model selection, it proceeds in three steps (Kallel et al., 2002).

1. **Resample.** Draw B bootstrap samples of n observations with replacement from the original data set, with B typically between 20 and 200.
2. **Refit.** Estimate the parameters of each candidate model on each bootstrap sample, which yields B parameter vectors per model.
3. **Evaluate.** Compute the test mean squared error of each refitted model on the original data set, then summarize the B values by their mean and standard deviation (Equations 23 to 25).

BOOTSTRAP MEASURES OF FIT AND STABILITY

$$\mathrm{TMSE}(b) = \frac{1}{n}\sum_{j=1}^{n}\left(y_j - y\left(x_j; \hat{\theta}^{*b}\right)\right)^2 \qquad (23)$$

$$\mu_{\mathrm{boot}} = \frac{1}{B}\sum_{b=1}^{B}\mathrm{TMSE}\,(b) \qquad (24)$$

$$\sigma_{\mathrm{boot}} = \left[\frac{1}{B-1}\sum_{b=1}^{B}(\mathrm{TMSE}(b) - \mu_{\mathrm{boot}})^2\right]^{1/2} \qquad (25)$$

- $\mathrm{TMSE}(b)$ is the test mean squared error of the model refitted on bootstrap sample b.
- n is the number of observations in the original data set.
- y_j is the observed response at observation j.
- x_j is the vector of explanatory variables at observation j.
- $\hat{\theta}^{*b}$ is the parameter vector estimated on bootstrap sample b.
- $y\left(x_j; \hat{\theta}^{*b}\right)$ is the prediction of the refitted model at observation j.
- B is the number of bootstrap samples.
- μ_{boot} is the mean error across samples and measures fit.
- σ_{boot} is the standard deviation of the error across samples and measures stability.

Adapted from Kallel et al. (2002).

The two statistics restore the two pillars of the framework without any parametric assumption. The preferred model offers the best compromise between them, ideally the lowest value of both. Each refit uses n observations, as the final model would, which spares the data a split into training and test sets (Kallel et al., 2002). Borowiak Table 2 summarises the evidence reported by the authors, who compare the bootstrap with leave-one-out resampling at an equal number of 50 replicates.

EXPERIMENT	CANDIDATES	BOOTSTRAP	LEAVE-ONE-OUT
Linear model, 500 simulated observations	True model with two regressors, against one and three regressors	Selects the true model, with a standard deviation of 0.016 against 0.43 to 0.60 for its rivals	No significant difference between models, with standard deviations one to two orders of magnitude larger
Simulated perceptron, true model with four hidden units	Two, four and six hidden units	Selects two hidden units, the most parsimonious adequate architecture, with near-equal means and the lowest standard deviation	Eliminates the true model and fails to separate the other two
Power peak control in nuclear reactor cores, 22 inputs	Two hidden layers, with 30, 35 or 40 units in the second	Selects 30 units, with the lowest mean and a standard deviation close to the smallest	Confirms the choice, with standard deviations about ten times larger

TABLE 77. Bootstrap and leave-one-out discrimination in three experiments

Source: Kallel et al. (2002)

Two cautions apply when reading this evidence. Each leave-one-out replicate is scored on a single held-out observation, whereas each bootstrap replicate averages n squared residuals, so part of the gap between the two standard deviations follows from the design of the comparison. In addition, a bootstrap sample contains on average about 63% of the distinct original observations, so scoring each replicate on the full original set mixes in-sample and out-of-sample points and makes the mean error optimistic. Scoring each replicate on the observations absent from its sample, the out-of-bag set, removes this bias (Efron & Tibshirani, 1993). The authors also point to the computing cost when n or p is large and to the risk that repeated extreme observations slow convergence (Kallel et al., 2002).

A proposed application to physical risk models

The framework could transfer to physical risk models wherever model outputs can be set against observed outcomes. Three uses appear promising. Hazard maps could be scored against observed footprints of past events with Equation 22. Loss estimates could be scored against insurance claims, reconstruction costs or default records with Equation 21. Model performance could then be summarised by the fit, stability and selection frequency of each candidate. Borowiak Table 3 sets the proposal against the variance partition of the previous section.

DIMENSION	VARIANCE PARTITION (QUALYPSO)	BOOTSTRAP DISCRIMINATION
Question answered	Which layer of the chain drives the dispersion of projections	Which model the observed data support, and how stable that verdict is
Data required	An ensemble of projections, without observations	Observed outcomes for a common sample of assets or events
Period	Future horizons	Historical calibration period
Output	Variance shares by layer and by model	Mean error, standard deviation and selection frequency by model
Main limitation	Additivity, with interactions pooled in the residual term	A good historical fit gives no guarantee under a changed climate

TABLE 78. Variance partition and bootstrap discrimination compared

Two configurations would arise in practice. Vendor models and published damage functions used as delivered would be completely determined in the sense of Borowiak (1989), since none of their parameters is re-estimated on the data of the user. Resampling would then leave the models unchanged, and the standard deviation would measure how sensitive the ranking is to the sample of observed losses. Models recalibrated on the data of the user, such as a maximum damage per square metre fitted to claims or an econometric loss model, would be data-fitted, and the standard deviation would then also capture parameter instability. Results would need to state which configuration applies to each model.

Repeating the selection on each replicate would yield a further indicator. The selection frequency of a model is the share of replicates in which it achieves the lowest error (Equation 26). It approximates the probability of correct selection that Borowiak uses to qualify a choice, and it requires all candidates to be scored on the same bootstrap samples.

BOOTSTRAP SELECTION FREQUENCY

$$\hat{\pi}_i = \frac{1}{B}\sum_{b=1}^{B} \mathbf{1}\left\{\mathrm{TMSE}_i(b) = \min_{k \in S} \mathrm{TMSE}_k(b)\right\} \qquad (26)$$

- $\hat{\pi}_i$ is the estimated selection frequency of model i.
- $\mathrm{TMSE}_i(b)$ is the test error of model i on bootstrap sample b, computed as in Equation 23.
- S is the set of candidate models, indexed by k.
- $\mathbf{1}\{\cdot\}$ is the indicator function, equal to one when the condition holds and zero otherwise.
- B is the number of bootstrap samples.

A frequency close to one supports the selection, and frequencies spread across several models signal persistent model ambiguity.

Conditions for a first test

1. **A common observed outcome.** Candidates must predict the same quantity as the observations, with the same loss scope and price base year, as for the variance partition.
2. **Coverage of the tail.** Claims are sparse for extreme events, and a sample dominated by frequent small losses rewards models that fit the body of the distribution. Weighting S(i) toward large losses or stratifying the resampling by event size keeps the tail in view.
3. **Paired resampling.** All candidates are refitted and scored on the same B samples, so that differences in error reflect the models alone.
4. **Out-of-bag scoring.** Errors are computed on the observations absent from each bootstrap sample.
5. **Preserved dependence.** Losses caused by one event are spatially correlated, so resampling whole events or blocks of years preserves a dependence that resampling individual assets would break.
6. **A limited claim on the future.** A fit over the historical period gives no certainty under projected climates, so a model selected in this way would enter the robust choice of Equation 19 as one plausible model among others.

Proposed protocol

1. **Define the candidate set.** List the rival models and identify nested and non-nested pairs.
2. **Choose the question.** Use A1 or A3 within a family of nested models and A2 across families.
3. **Assemble the observations.** Harmonise metric, loss scope and price base year between predictions and observed outcomes.
4. **Run the bootstrap.** Draw between 50 and 200 paired samples, resampling events when losses are correlated, and score each replicate out of bag.
5. **Report.** Give the mean error, standard deviation and selection frequency of each model, and state whether each model was refitted.
6. **Decide.** Read the results together with the variance partition, which shows where model choice matters, and with the robust criterion of Equation 19.

The Compound Risk issues

Section 3.1.1 addressed compound events at the hazard level, where two perils strike one asset together. The same problem reappears at aggregation. A portfolio can compound even when each asset faces a single hazard, if correlated hazards hit many assets within the same period and strain recovery capacity across the system.

A flood model and a windstorm model are often run separately and their loss distributions combined as if the two perils were independent. In Europe both perils tend to peak in the same wet and windy winters (Hillier et al., 2015). Independence then fails for the same statistical reason as in Section 3.1.1. Dependence is strongest in the tail, where the losses that matter for capital sit. Adding the 1-in-200-year loss of each peril errs the other way, because it assumes both extremes occur in the same year. The true joint loss lies between the two shortcuts.

An insurer needs a multi-peril aggregate exceedance curve built from jointly simulated years, in place of peril curves combined after the fact. A bank needs stress tests that layer physical perils on macroeconomic or credit shocks, since a flood during a recession is also a compound event. An asset manager rarely has the data for full joint modelling at portfolio scale. A scenario overlay applied across the book serves it better than a bespoke copula per asset.

Example: how to choose?

A reinsurer models flood and windstorm losses separately for the same European winter and combines them as if they were independent. This understates the joint 1-in-200-year loss, because both perils tend to peak in the same season. The joint extreme is more frequent than the independence assumption implies.

VARIABLE	DESCRIPTION	DATA SOURCE EXAMPLE
Multi-peril exceedance curve	Joint annual loss distribution across two or more perils	Catastrophe model multi-peril module, such as Moody's RMS or Verisk (Mitchell-Wallace et al., 2017)
Peril correlation matrix	Empirical or modelled correlation between peril losses at portfolio level	Reinsurance broker analytics; compound-event studies (Hillier et al., 2015; Zscheischler et al., 2018)
Combined stress scenario	Physical hazard shock layered on a macro or credit shock	NGFS scenarios with macroeconomic and physical damage components (NGFS, 2024)

TABLE 79. Data for Compound risk aggregation

3.4.3 NON-LINEARITIES

Most damage functions and financial proxies reviewed map hazard intensity to loss through a linear or gently curved function. Physical thresholds break that assumption. Past a threshold, a small increase in hazard intensity produces a disproportionate jump in loss, and the accumulation of warming beyond certain thresholds can lead to irreversible impacts (Bolton et al., 2020). The same logic applies at the scale of the Earth system, where several climate tipping points may be triggered within the 1.5 to 2°C range of global warming (Armstrong McKay et al., 2022).

Where jumps appear

A flood defence has a design standard, usually expressed as the return period of the flood it is built to withstand (Scussolini et al., 2016). Loss stays close to zero below that standard and becomes large as soon as the defence is overtopped. A protective ecosystem behaves in a similar way. A mangrove belt in front of a dike lowers the wave height reaching the structure and reduces the freeboard and revetment it requires (Verhagen, 2018), so the loss of the belt raises the hazard faced by every asset behind it at once. Heat produces the same pattern in agriculture. Yields increase with temperature up to about 29°C for corn, 30°C for soybeans and 32°C for cotton, and fall sharply above these thresholds (Schlenker & Roberts, 2009). The financial counterpart is visible in corporate earnings, which respond significantly to extreme temperatures in over 40% of US industries (Addoum et al., 2023).

Compound events add a further source of non-linearity. Two or more hazards that hit the same region within the same period can produce impacts greater than the sum of their individual effects (Loyer et al., 2026). All of these cases are structural breaks in the damage function itself, distinct from ordinary tail events drawn from a stable distribution.

Modelling implication

A linear or log-linear damage function fitted to historical, sub-threshold data cannot extrapolate across the threshold. It systematically understates loss just past the tipping point, and averaging exposure over time or space further dilutes the true non-linear response (Schlenker & Roberts, 2009). Where the physical threshold is known, as with a defence design standard or an ecological collapse point, the damage function should be specified piecewise with an explicit jump term.

$$D(h) = f(h) + J \cdot \mathbb{1}(h > h*)$$

In this expression, $D(h)$ is the loss, expressed as a share of asset value, at hazard intensity h.

The function $f(h)$ is the smooth damage function calibrated on sub-threshold observations.

The parameter $h*$ is the physical threshold, for instance the water level corresponding to the design standard of the defence. The parameter J is the additional loss incurred once the threshold is exceeded.

The indicator $\mathbb{1}(h > h*)$ equals 1 when the hazard intensity exceeds the threshold and 0 otherwise.

Where the threshold is unknown, stress testing across a range of hazard intensities is the more defensible fallback, in place of a single central estimate. Multi-hazard storylines serve the same purpose and allow an explicit examination of cascading effects, critical thresholds and indirect impacts (Loyer et al., 2026). Thresholds also interact at portfolio level, since a single event can push many assets past their breakpoints simultaneously.

Assumptions on the spatial correlation between events can then lead to large under- or overestimations of portfolio tail risk (Kerkhofs et al., 2025). The level of methodological complexity should remain proportionate to the decision the assessment supports (Loyer et al., 2026).

Example: how to choose?

A city's flood defence is rated to a 1-in-100-year storm surge. A linear damage function predicts a gentle rise in loss between a 1-in-90-year and a 1-in-110-year event. In reality, loss jumps as soon as the surge overtops the defence, so two events of similar probability produce radically different losses. The model must encode the threshold through the jump term J, and averaging across it would hide the loss that matters most.

VARIABLE	DESCRIPTION	DATA SOURCE EXAMPLE
Physical threshold / design standard	Known engineering or ecological breakpoint	FLOPROS global database of flood protection standards (Scussolini et al., 2016); climate tipping element thresholds (Armstrong McKay et al., 2022); crop temperature thresholds (Schlenker & Roberts, 2009)
Piecewise / jump damage function	Damage function with an explicit discontinuity at the threshold	Engineering studies of defended assets and protective forelands (Verhagen, 2018); threshold temperature responses of yields and earnings (Schlenker & Roberts, 2009; Addoum et al., 2023)
Multi-scenario stress set	Range of hazard intensities bracketing an uncertain threshold	CMIP6 climate model ensemble (Eyring et al., 2016); NGFS Phase V scenarios (NGFS, 2024); multi-hazard storylines (Loyer et al., 2026); correlated portfolio stress tests (Kerkhofs et al., 2025)

TABLE 80. Data for non-linearities

3.4.4 TYPE I, II AND III ERRORS

Choosing the wrong model has consequences that differ in kind as well as in size. Three error types matter for physical risk models, namely false positives, false negatives, and the correct answer to the wrong question.

Type I: false positive

The model flags an asset or region as high risk when it is not. The consequence is misallocated capital, through unnecessary de-risking or the refusal of a sound loan or insurance policy. Where the flagged decision affects a vulnerable population, it also creates reputational and regulatory exposure.

Type II: false negative

The model fails to flag a genuinely high-risk asset. The consequence is under-reserved capital and losses that materialize without warning. For an insurer or a bank this is the more dangerous error, since it threatens solvency directly, whereas a Type I error misallocates capital.

Both rates can be measured against realized losses during model validation.

$$FPR = FP / (FP + TN) \qquad FNR = FN / (FN + TP)$$

In these expressions, FPR is the false positive rate, that is the Type I error rate, and FNR is the false negative rate, that is the Type II error rate. FP is the number of assets flagged as high risk that incurred no material loss.

TN is the number of assets correctly classified as low risk. FN is the number of assets classified as low risk that incurred a material loss. TP is the number of assets correctly flagged as high risk.

Type III: the right answer to the wrong question

The model correctly answers a mis-specified question. The error of the third kind is defined as having solved the wrong problem when one should have solved the right one, and reducing Type I and Type II errors at its expense amounts to misplaced precision (Mitroff & Featheringham, 1974).

A model may precisely quantify a 1-in-200-year flood loss while the real threat to the portfolio is a chronic, slow-onset drought that was never modelled. The error has a strategic root. When calculation and the control of the present displace the definition of ends, an organisation keeps optimising its means against a problem it has stopped questioning, a condition described as the strategic void (Baumard, 2012). This error is the least visible and the most institutionally dangerous. The output looks precise and can pass every internal validation test, while the underlying risk goes unaddressed.

Climate service guidance treats this risk as a problem of fitness for purpose, since a product used by the wrong person or in the wrong context can lead to decisions its providers never intended (Villwock, 2023). The safeguard lies in the scoping phase, which defines the system boundaries, the objectives, the time horizons and the decision context before any modelling starts, and in a hazard screening that starts from the most comprehensive list of potential hazards (Loyer et al., 2026).

An insurer must weight Type II errors most heavily in model validation, since under-reserved tail risk is existential. A bank must track its Type I rate as closely as its Type II rate, since unjustified loan declines carry direct reputational and regulatory cost. An asset manager is most exposed to Type III error. A sophisticated model built for the wrong hazard family, or covering physical risk alone in a sector where transition risk dominates, provides false confidence in place of insight. Materiality assessments should therefore consider physical and transition risks together, under several scenarios and timescales (Loyer et al., 2026).

Backtesting detects Type I and Type II errors. Type III errors only surface through a review of the conceptual soundness of the model and its fit with the intended use, carried out by objective and informed parties able to challenge the model choice itself (Board of Governors of the Federal Reserve System et al., 2026).

Example: how to choose?

A bank builds a highly validated flood-only credit risk model for its Sahel agricultural loan book. The model passes every backtest. The dominant physical driver of the book, however, is drought. The model is right about the wrong risk, which is a Type III error.

VARIABLE	DESCRIPTION	DATA SOURCE EXAMPLE
Model validation / backtest results	False positive and false negative rates against realized losses	Internal model validation reports; outcomes analysis and back-testing under model risk guidance (Board of Governors of the Federal Reserve System et al., 2026)
Portfolio hazard scope review	Confirmation that the modelled hazard matches the dominant driver of the portfolio	Hazard screening by sector, site typology and geography, starting from a comprehensive hazard list (Loyer et al., 2026)
Independent model risk sign-off	Second-line review of model choice, in addition to calibration	Internal model risk management function; effective challenge requirements of SR 26-2, which supersedes SR 11-7 (Board of Governors of the Federal Reserve System et al., 2026)

TABLE 81. Data for type 1 and 2 errors

3.4.5 TRANSPARENCY AND REPLICABILITY

Transparency and replicability go together. Replicability is necessary for auditability, but it is insufficient on its own. A fully documented process that is wrong remains auditably wrong and is of no use for decision-making. Transparency and correctness are therefore required together.

Why replicability matters

A vendor hazard or damage-function output that cannot be traced to its inputs and methodology cannot be defended to a regulator, a rating agency or an internal model risk committee. If a supplier changes its methodology between two annual runs without explanation, the institution cannot tell whether the resulting change in reported risk is real or an artefact of the update.

This kind of behaviour has long motivated independent reviews of physical risk models. Commercial solutions often rest on proprietary models whose assumptions, downscaling and bias correction are insufficiently documented or validated. Such black-box assessments limit the ability of regulators and stakeholders to evaluate the credibility of the results (Loyer et al., 2026).

Current guidance therefore asks climate data and tool providers to disclose their evidence sources, modelling approaches and assumptions, to let independent experts assess the scientific quality of their outputs, and to apply standard software quality practices such as version control, testing, code review and clear documentation (Loyer et al., 2026).

Users find it difficult to judge whether data are fit for a given application, which places a particular responsibility on the provider (Villwock, 2023). Model evaluation depends on this transparency. Suppliers should document the methodology, data sources, key assumptions and known limitations, and share the validation tests they performed, including the parts of the model that could not be validated (Norris, 2026).

Risk maps offer a concrete test. A catastrophe model produces hazard maps for key return periods as an intermediate output, and these maps should be compared with scientific datasets, peer-reviewed studies, official maps and alternative models (Norris, 2026). Material differences should be explained and documented. The alternative sources serve as benchmarks, since they carry their own differences in methodology, data vintage and resolution (Norris, 2026).

What does an institution need? An insurer or a bank generally operates under a regulatory model risk framework that requires demonstrable methodology and version control for any purchased model feeding into capital calculations. An asset manager faces less direct regulatory pressure, but the reputational risk of an unexplained swing in a disclosed climate risk metric is just as real.

Example: how to choose ?

A catastrophe model vendor updates its wildfire module, and the reported 1-in-100-year loss for an unchanged portfolio jumps 40%. If the vendor cannot walk through what changed and why, the institution has no way to know whether the new number is more accurate or simply different, and cannot defend either figure to a regulator. For comparison, model evaluation guidance treats a gap of around 30% between modelled and recorded losses as a warning flag that calls for further analysis (Norris, 2026). A vendor that can show the hazard maps before and after the update, together with the change log behind them, gives the institution the means to decide.

VARIABLE	DESCRIPTION	DATA SOURCE EXAMPLE
Methodology documentation	Vendor's published model assumptions and version history. At a minimum, each vendor should be able to describe each model in the format of the model cards.	Vendor technical documentation; minimum transparency standards for vendors (Loyer et al., 2026); supplier documentation checklist (Norris, 2026)
Version control / change log	Record of what changed between model releases	Vendor release notes; version control and provenance metadata as software quality practice (Loyer et al., 2026)
Independent validation	Third-party or internal replication of vendor output on a known case	Internal model validation function; comparison of hazard maps and historical losses against alternative sources (Norris, 2026); model discrimination on fit and stability (Borowiak, 1989)

TABLE 82. Data for Transparency and replicability

4 MODEL CLUSTERING

Parts 2 and 3 describe the recipes and their ingredients one at a time. A practitioner who opens the literature review behind this Cookbook faces 195 coded models that link a physical hazard to a financial outcome, each built for one hazard, one scale and one end-user. Part 4 answers a question that the model-by-model reading leaves open. **Which models can be compared with which, and which of them fits a given need?**

The answer fits in one proposition. The literature contains many physical risk models, but few can be compared directly or validated independently. Mapping their methodological characteristics shows which models serve similar purposes, where gaps persist in the corpus, and how a user can identify relevant alternatives.

A direct comparison of predictive performance is out of reach. Source codes, calibrated parameters and a common input-output design are almost never published, so two models cannot be run on the same benchmark. The comparison therefore relies on what every paper does report, namely the hazard it takes as input, the resolution of that input, the financial indicator it produces, the method linking the two and the user it serves.

Clustering turns these descriptions into a map. Models sharing a functional signature fall into the same region, isolated research communities appear as separated regions, and empty regions point to hazard and indicator combinations that the corpus does not cover, within the limits of the bibliographic search. **The opacity of the model landscape, the second hidden risk introduced in Part 1, is a measurable structure** (Vigneron, 2025).

BOX 4.1. MODELS AND SIGNATURES

$$M_i : \mathcal{X}_{climate} \rightarrow \mathcal{Y}_{risk} \quad (4.1)$$

$$q_i = \Phi(M_i) \in \mathcal{Q} \quad (4.2)$$

$$\Psi : \mathcal{Q} \rightarrow \mathbb{R}^2 \quad (4.3)$$

With

- $\boldsymbol{M_i}$ is the financial impact model reported in article i
- $\boldsymbol{\mathcal{X}_{climate}}$ is the space of hazard and exposure inputs (flood depth, wind speed, drought index, asset location)
- $\boldsymbol{\mathcal{Y}_{risk}}$ is the space of financial outputs (losses, probability of default, returns, Climate VaR)
- $\boldsymbol{q_i}$ is the qualitative functional signature of model i
- $\boldsymbol{\Phi}$ is the coding function defined by the framework variables
- $\boldsymbol{\mathcal{Q}}$ is the coding space, which collects every admissible combination of modalities
- $\boldsymbol{\Psi}$ is the projection learned by the clustering algorithm, which places each signature on a plane

The position $\Psi(q_i)$ reflects methodological proximity between models. It says nothing about which model predicts better, a question that only a shared reference design could settle. Each cluster can then be summarised by one real model, its medoid, which remains open to inspection, unlike an abstract average.

BOX 4.2. PROTOTYPE OF A CLUSTER

$$M_C = arg\,min_{Mi \in C} \Sigma_{Mj \in C}\, d(q_i, q_j) \quad (4.4)$$

With

- C is a cluster of models
- M_C is the prototype of cluster C, the model whose signature is closest on average to all other members
- M_i, M_j is two models belonging to C
- $d(q_i, q_j)$ is the distance between the signatures of models i and j

How to cluster physical climate risks financial impacts models

The procedure has three steps. Encoding describes each model with the framework variables. Transformation converts these categorical descriptions into a numerical matrix whose geometry respects the χ^2 distance of correspondence analysis (Benzécri, 1973). Projection places the models on a plane, and three projection regimes are used in this part.

REGIME	NUMBER OF GROUPS	METHODS	OUTPUT FOR THE MODEL USER	SECTION
Closed clusters	Fixed before training	Kohonen self-organizing map	A partition of models into grid cells with known neighbours	4.1
Open clusters	Emerges from the data	MCA, VAE with t-SNE	A continuous map where borders are left to interpretation	4.2
User-guided filtering	Set by the user need	Labels, filters, ranking	A shortlist of comparable models for one use case	4.3

TABLE 83. Three regimes of model clustering

Each article of the literature review is coded in a framework along ten groups of variables with constrained admissible values. An article enters the corpus only if it reports at least one proxy of physical hazard and at least one proxy of financial impact, which guarantees that every coded model actually links the two domains. The maps below use a harmonized version of this framework, with 21 active variables and 93 modalities, shared by all 195 coded models.

GROUP	VARIABLES	ADMISSIBLE VALUES (ABRIDGED)	ROLE IN CLUSTERING
1. Bibliographic identification	Link, paper name, paper type	Empirical, literature review, methodological, other	Identification only
2. Geographic scope	Market scope, coverage type, sector	Local, national, international; 10 sectors from agriculture to multi-sectoral	Active
3. Temporal dimension	Baseline, projection timeframe, timeframe type	Historical, near present, future, historical with projection, event-based, simulation	Active
4. Financial impact	Precision, type, proxies, source, sample size	Asset to national level; specific, general, hybrid; 9 proxy families from Climate VaR to accounting financials	Active
5. Physical risk	Resolution, spatial resolution, type, proxies, source	Asset to national level; water, wind, fire, earth, compound, other	Active
6. Modelling	Method, method type, control variables	OLS, panel, event study, DiD, Monte Carlo, Bayesian, ML, physics-based	Active (controls as free text)
7. Scenarios	Projection (y/n), projection scenario	RCP 2.6, 4.5, 8.5, SSP equivalents, NGFS	Active
8. Geographic bias	Developed or developing	Developed, developing, unspecified	Active
9. Specific approaches	Damage function, transition risk	Yes, no	Active
10. Model quality	Accuracy	R^2, RMSE, MAE, AIC, BIC, log-likelihood, Christoffersen tests, other	Supplementary (free text)

TABLE 84. Framework coding groups

Every modality of an active variable becomes one binary column of a complete disjunctive table. The Burt table then counts how often two modalities appear in the same model, and a normalization by the margins prevents frequent modalities from dominating the geometry (Box 4.3). Two geometries follow. The normalized Burt table describes how modalities associate, so it supports a map of modalities. A distance between models needs a second step, which rescales each row of the disjunctive table so that the Euclidean distance between two models equals their χ^2 distance. Each map below states which of the two matrices it uses.

BOX 4.3. FROM CODED ARTICLES TO A GEOMETRY OF MODALITIES AND A GEOMETRY OF MODELS

$$Z = [z_{kj}] \in \{0,1\}^{n \times p},\ z_{kj} = 1\ if\ model\ k\ displays\ modality\ j \quad (4.5)$$

$$B = Z^{\top} Z \quad (4.6)$$

$$\tilde{B}_{ij} = B_{ij} / \sqrt{(B_{i+} B_{+j})} \quad (4.7)$$

$$x_{kj} = \frac{z_{kj}}{Q\sqrt{c_j}},\ c_j = \frac{1}{nQ}\sum_k z_{kj} \quad (4.8)$$

$$d^2(k,l) = \sum_j \left(x_{kj} - x_{lj}\right)^2 \quad (4.9)$$

With

- $\boldsymbol{Z}$ is the complete disjunctive table of the corpus
- $\boldsymbol{z_{kj}}$ is equal to 1 when model k takes modality j and 0 otherwise
- $\boldsymbol{n}$ is the number of coded models
- $\boldsymbol{p}$ is the total number of modalities across all active variables
- $\boldsymbol{B}$ is the Burt table, a symmetric p × p matrix of modality co-occurrences
- $\boldsymbol{B_{ij}}$ is the number of models displaying both modality i and modality j
- $\boldsymbol{B_{i+}, B_{+j}}$ is the row and column totals of B
- $\tilde{\boldsymbol{B}}_{\boldsymbol{ij}}$ is the normalised co-occurrence of modalities i and j. The Euclidean distance between two rows of $\tilde{B}$ approximates the χ^2 distance between modalities, so this matrix places modalities
- $\boldsymbol{x_{kj}}$ is the coordinate of model k on modality j, its row of the disjunctive table rescaled by the modality frequency
- $\boldsymbol{Q}$ is the number of active variables
- $\boldsymbol{c_j}$ is the relative frequency of modality j
- $\boldsymbol{d(k,l)}$ is the distance between models k and l, equal to their χ^2 distance in multiple correspondence analysis

Equation 4.7 gives a geometry of modalities and Equation 4.9 a geometry of models. The two are linked, since a model sits at the barycentre of its modalities in correspondence analysis, but a map built on one cannot be read directly as a map of the other.

4.1 CLOSED CLUSTERS

A clustering is closed when the number of groups is chosen before training and every model ends up in exactly one group. The Kohonen self-organizing map adds a second property, since its groups sit on a grid and adjacent cells hold similar models (Kohonen, 1982). A closed map can therefore be read like a table of contents, each cell holding a model family and each neighbour a close substitute.

Each neuron of the grid carries a prototype, a vector living in the same space as the coded models. Training presents the models one at a time, finds the closest prototype and pulls it, together with its grid neighbours, toward the presented model (Box 4.4).

BOX 4.4. KOHONEN TRAINING RULE

$$c(t) = arg\,min_{u \in U} \| x(t) - w_u(t) \| \qquad (4.10)$$

$$w_u(t+1) = w_u(t) + \eta(t)\, h_{cu}(t) [\, x(t) - w_u(t) \,] \qquad (4.11)$$

$$h_{cu}(t) = exp(- \| r_c - r_u \|^2 / 2\sigma(t)^2) \qquad (4.12)$$

$$\eta(t) = \eta_0\, exp(-t/\tau_1), \quad \sigma(t) = \sigma_0\, exp(-t/\tau_2) \qquad (4.13)$$

With

- t is the training iteration
- $x(t)$ is the input presented at iteration t, a modality profile (Equation 4.7) or a model profile (Equation 4.8)
- U is the set of neurons of the grid
- $w_u(t)$ is the prototype vector of neuron u at iteration t
- $c(t)$ is the best matching unit, the neuron whose prototype is closest to x(t)
- $\eta(t)$ is the learning rate, which sets the size of each update
- $h_{cu}(t)$ is the neighbourhood function, equal to 1 for the best matching unit and decreasing with grid distance
- r_c, r_u is the grid coordinates of neurons c and u
- $\sigma(t)$ is the neighbourhood radius
- η_0, σ_0 is the initial learning rate and radius
- τ_1, τ_2 is the time constants of the two decays

At convergence, neighbouring cells encode similar signatures, so distance on the grid becomes a proxy for methodological distance. Cluster borders are read from the U-matrix, which displays the distance between the prototypes of adjacent neurons, light zones marking homogeneous regions and dark zones marking boundaries (Ultsch & Siemon, 1990). Kohonen maps already serve to classify firms in corporate finance, where they stabilise bankruptcy forecasts over longer horizons (du Jardin & Séverin, 2011).

Two cautions apply to categorical data. The Euclidean distance suits continuous inputs, hence the normalisations of Equations 4.7 and 4.8. The rows of the Burt table describe modalities, so a map trained on $\tilde{B}$ places modalities. Placing the models themselves requires training the map on the model profiles of Equation 4.8, or a dedicated algorithm that classifies individuals and modalities jointly (Cottrell & Letrémy, 2005). The two maps answer different questions, and Section 4.1.3 shows that they lead to different conclusions on banking models.

4.1.1 A FIRST CLOSED MAP OF THE CORPUS

The first closed clustering of the literature review was presented at the Institut Louis Bachelier Risk Forum. It trained a 3 × 3 Kohonen grid on the raw Burt table of the coded corpus and labelled each cell after its representative articles.

	COLUMN 1	COLUMN 2	COLUMN 3
Row 1	n = 12 Wildfire budget risk planning	n = 7 Seismic non-structural loss dynamics	n = 11 Asset-level climate tail risk
Row 2	n = 10 Compound flood impact pathways	n = 5 Pluvial flood data science modelling	n = 7 Disasters and financial regulation
Row 3	n = 13 Coastal flooding and agricultural losses	n = 9 Climate justice and housing risk data	n = 12 Drought and water policy economics

TABLE 85. First closed clustering of the corpus on a 3 × 3 Kohonen grid
Adapted from Vigneron (2025).

The densest cells sit in the corners, with 11 to 13 units each, and carry hazard-specific traditions such as public wildfire budgets, drought and water policy economics, coastal flooding with agricultural losses, and asset-level climate tail risk. The central cell holds 5 units only and marks a sparse zone between these traditions. Flood modelling splits across two cells of the middle row, one for compound flood pathways and one for pluvial flood data science, next to a cell dedicated to disaster regulation.

The map established three results. The corpus is structured, since modalities do not spread evenly across the grid. Hazard specificity drives the structure, since most cells are named after one peril. Regulation and tail risk models form their own cells, apart from the damage-oriented traditions. The same exercise also revealed five limits that shaped the next iteration.

- **Resolution.** Nine cells are too few for a corpus of 154 models, which merges distinct families, as the cell joining coastal flooding and agricultural losses shows.
- **Labels.** Cell names describe topics drawn from representative articles, so a user looking for a model family still has to open each cell.
- **Metric.** The raw Burt table lets frequent modalities dominate the Euclidean distance, and rare hazards such as heat, about 3% of the corpus, are absorbed into larger cells. To correct this, a possibility is adding more cells (e.g. a 100*100 grid instead of a 3*3 grid can solve this problem).

- **Object mapped.** A map trained on the Burt table places modalities, and the occupation counts of t sum to 86. This total matches a count of modalities and falls far below the size of the corpus, so each cell holds a group of co-occurring modalities labelled after the articles that display them.
- **Borders and stability.** Without a U-matrix and with a single training run, the boundaries between cells and their robustness to initialisation remain unknown.

4.1.2 A REFINED CLOSED MAP

The second iteration addresses the resolution, metric and border limits. It trains a 5 × 5 grid on the normalized Burt matrix of Equation 4.7 and reads the borders from the U-matrix. Like the first map, it places modalities, and each region is read from the modalities it gathers and the articles that display them. Retrained with the published training rule on the coded corpus of 195 models, it yields four regions.

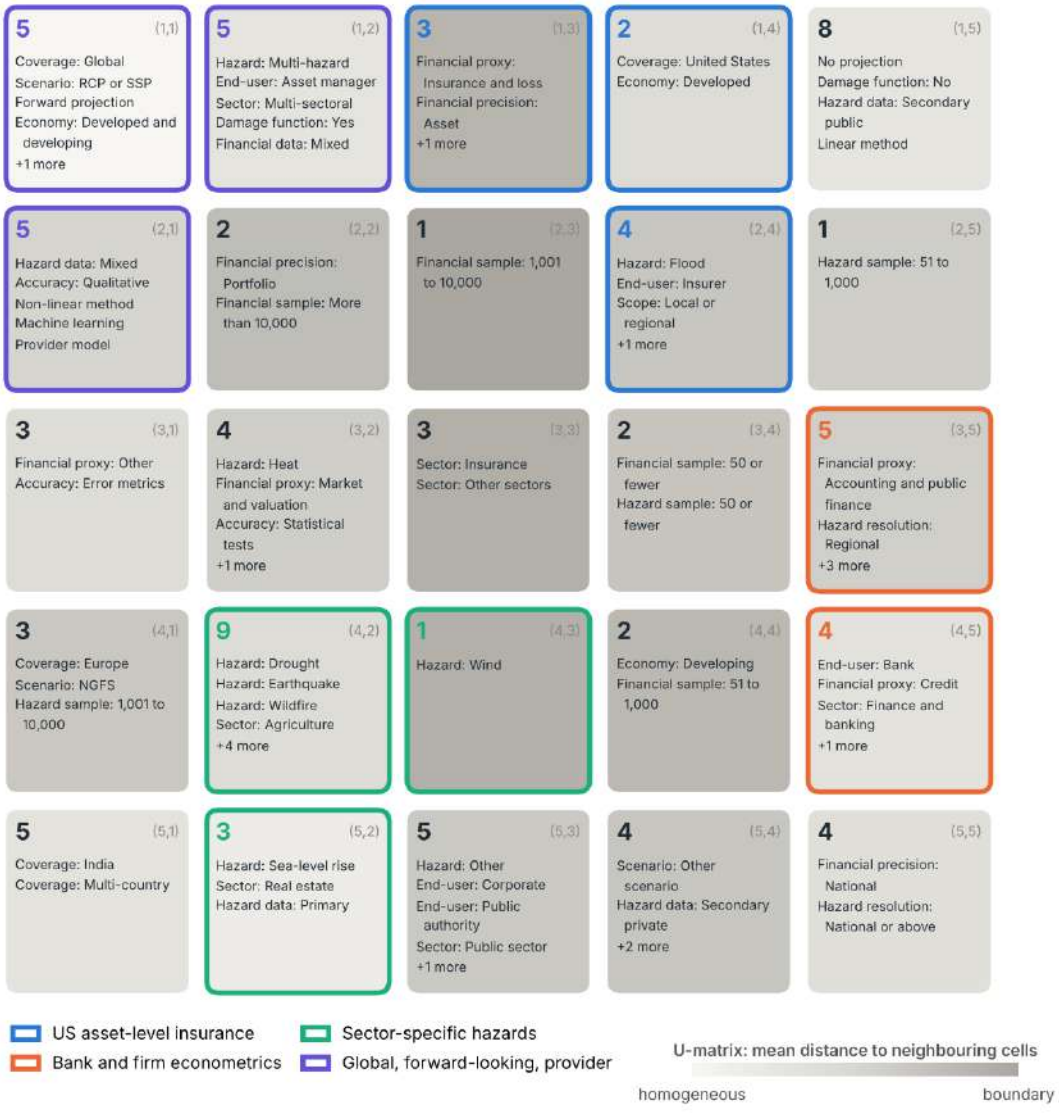


FIGURE 89. Kohonen map of the 93 harmonized modalities of the 195 coded models
Source: authors, training rule of Vigneron and Vigneron (2026).

REGION	CELLS (ROW, COLUMN)	SHARED SIGNATURE	MAIN END-USERS
US asset-level insurance	(1,3), (1,4), (2,4), top right	Insurance and loss proxies, asset-level financial precision, US coverage of developed economies; (2,4) adds local scope and flood hazard	Insurers, reinsurers
Bank and firm econometrics	(4,5), (3,5), right edge	Bank end-users, finance and banking sector, company-level precision, credit proxies; (3,5) adds accounting and public-finance proxies, private financial data, regional hazard resolution and R^2	Banks, supervisors
Sector-specific hazards	(4,2), (4,3), (5,2), bottom left	Drought, earthquake, wildfire and wind models of agriculture, energy and infrastructure, often on primary data with no reported accuracy; sea-level rise and real estate in (5,2)	Corporates, public authorities
Global, forward-looking, provider	(1,1), (1,2), (2,1), top left	Global coverage, forward projections under RCP or SSP, multi-hazard proxies and damage functions; provider, non-linear and machine-learning models in (2,1)	Asset managers, investors

TABLE 86. Regions of the 5 × 5 Kohonen map

A fifth strand gathers corporate and public-authority end-users, Chinese coverage and the public sector in (5,3), next to developing economies in (4,4). The banking result revises the first reading. On the raw 3 × 3 grid, banking themes seemed scattered across regulation and tail risk cells, whereas the normalized 5 × 5 grid gathers the banking modalities (bank end-users, finance and banking sector, credit proxies, company-level precision) in one cell at the edge of the grid, next to the accounting and public-finance econometrics of (3,5) (Vigneron, 2025; Vigneron & Vigneron, 2026). Two settings changed at once, the metric and the grid size, so the revision cannot be attributed to the metric alone. Conclusions drawn from a closed map therefore depend on how the map is built, and Section 4.1.3 tests which of them survive.

4.1.3 WHAT SURVIVES A CHANGE OF OBJECT, GRID AND SEED

Three checks answer the open questions of the first two maps. They use the coded corpus of 195 models and 21 harmonized active variables, the training rule of the published map (200 epochs, initial learning rate 0.1, initial radius of half the grid size) and 20 random initialisations per configuration. Agreement between two runs is measured with the adjusted Rand index, which equals 1 when two runs give the same partition and is close to 0 when their agreement is no better than chance (Hubert & Arabie, 1985).

Object mapped. A 5 × 5 map trained on the normalized Burt table places the four banking modalities in one or two cells in all 20 runs. A 5 × 5 map trained on the model profiles of Equation 4.8 does not isolate banking models. The cell holding most bank models gathers about a fifth of them (21% on average), and that cell with its four neighbours about 39%. The banking modalities co-occur tightly, while the banking models spread over several cells, because they share hazard, geography and data choices with other families.

Grid and metric. On model profiles, replacing the raw disjunctive table by the χ^2 rescaling of Equation 4.8 changes the results little, and the grid size drives most of the difference between configurations. The 3 × 3 versus 5 × 5 contrast of Section 4.1.2 therefore mixes two effects, and the grid size weighs more.

Stability. Partitions agree moderately to well across initialisations, with a mean adjusted Rand index between 0.66 and 0.72 depending on the configuration, a narrower range than on the earlier corpus of 154 models (0.55 to 0.76). The 5 × 5 grid on χ^2 profiles remains the least stable, at 0.66. Cell borders are zones, and the cell of a given model is one plausible assignment among several.

Coding redundancy. The earlier version of the dataset contained redundant variables. The sector variable appeared twice under two names, the end-user variable duplicated the model family variable for 152 of the 154 models, and the damage function appeared in two variables with inconsistent capitalisation, so banking was encoded in four columns. The harmonized schema keeps one variable per concept, and the banking result above holds on it.

The banking signal is therefore real but weaker than a single isolated cell suggests. A closed map of modalities shows which characteristics travel together, and a closed map of models shows which models resemble each other. Both readings are useful, provided the map states which object it places.

Reading a closed map as a model user

A closed map works as an index to the Cookbook. The user locates the cell matching the need, defined by hazard, scale and financial output, reads the prototype of that cell (Equation 4.4), then compares the models inside it. Sharing a signature makes models candidates for the same use, while their reported fit statistics remain hard to compare, since samples, target variables and evaluation protocols differ (Section 4.3.2). Neighbouring cells provide substitutes when the prototype misses one requirement, for instance a model with an explicit damage function one cell away from a damage-function-free valuation model.

Closedness also has a cost. The number of cells is fixed in advance, every model is forced into a single cell, and hybrid models suffer most, although multi-hazard studies represent 38% of the coded corpus. Models sitting between two traditions become invisible on a discrete grid. Section 4.2 lifts this constraint with open clusters, where the borders between model families are left to the data.

4.2 OPEN CLUSTERS

A clustering is open when neither the number of groups nor their borders is fixed in advance. Each model receives continuous coordinates, model families appear as denser regions, and the analyst reads their limits from the map. Open clusters keep hybrid models visible between two traditions, where a closed grid would force them into one cell. Two methods are used here. Multiple correspondence analysis gives a linear benchmark, and a variational autoencoder gives a non-linear projection of the same coded corpus.

4.2.1 MULTIPLE CORRESPONDENCE ANALYSIS TO DISCRIMINATE MODELS

Multiple correspondence analysis decomposes the complete disjunctive table of Equation 4.5 into orthogonal axes ranked by the share of total inertia they carry (Benzécri, 1973; Lebart et al., 2000). Models and modalities are projected on the same plane. Two models lie close when they share most of their modalities, and a modality sits at the barycentre of the models that display it (Box 4.5).

BOX 4.5. MULTIPLE CORRESPONDENCE ANALYSIS OF THE CODED CORPUS

$$P = Z / (nQ) \quad \textbf{(4.14)}$$

$$S = D_r^{-1/2} (P - r\, c^\top)\, D_c^{-1/2} = U\, \Lambda^{1/2}\, V^\top \quad \textbf{(4.15)}$$

$$F = D_r^{-1/2}\, U\, \Lambda^{1/2} \quad \textbf{(4.16)}$$

$$\tau_k = \lambda_k / \Sigma_l\, \lambda_l \quad (4.17)$$

With

- $\boldsymbol{P}$ is the correspondence matrix, the disjunctive table rescaled to sum to one
- $\boldsymbol{Z}$ is the complete disjunctive table of Equation 4.5
- $\boldsymbol{n}$ is the number of coded models
- $\boldsymbol{Q}$ is the number of active variables
- $\boldsymbol{r}$ is the vector of row masses, equal to 1/n for every model
- $\boldsymbol{c}$ is the vector of column masses, the frequency of each modality divided by nQ
- $\boldsymbol{D_r, D_c}$ is the diagonal matrices of row and column masses
- $\boldsymbol{S}$ is the matrix of standardised residuals
- $\boldsymbol{U, V}$ is the left and right singular vectors of S
- $\boldsymbol{\Lambda}$ is the diagonal matrix of eigenvalues λ_k
- $\boldsymbol{F}$ is the matrix of model coordinates on the factorial axes
- $\boldsymbol{\tau_k}$ is the share of total inertia carried by axis k

The shares τ_k stay low in any MCA built on many modalities, so a first plane carrying less than 15% of inertia can still describe a clear structure (Lebart et al., 2000). The reading relies on the most contributive modalities and on the position of the models around them.

A first exploration on a sample of 30 articles projected title and keyword terms on a plane carrying 9.1% and 8.6% of inertia (Vigneron, 2025). This first plane described the vocabulary of the field, and the next iteration moved to the coded modalities of the framework.

On the coded corpus, the first two axes carry 8.20% and 6.80% of inertia, or 41.1% and 26.0% after the Benzécri correction, which discounts the inertia created by the coding itself (Benzécri, 1979; Vigneron & Vigneron, 2026). Axis 1 opposes historical, linear econometric models to forward-looking models with scenario projections, provider tools and machine learning. Axis 2 opposes local, asset-level insurance loss models to company-level bank and firm econometrics, often on China and private financial data. A permutation test confirms that this structure is real. When each variable is shuffled independently, which destroys every association between variables, the first axis carries 3.3% of inertia on average and never more than 4.0% over 30 draws.

REGION OF THE PLANE	MOST CONTRIBUTIVE MODALITIES	MODEL FAMILY
Upper left	Insurer end-users, local or regional scope, insurance and loss proxies, asset-level precision	Insurance and real-estate loss models, mostly on the United States
Right	Forward projection, RCP or SSP scenarios, provider models, machine learning, non-linear methods, global coverage	Climate stress-testing and provider models
Lower left	Finance and banking sector, company-level financial precision, private financial data, credit proxies, Chinese coverage	Bank and firm credit econometrics, largely from the recent literature
Lower center	Multi-hazard proxies, international scope, developed and developing economies	Transversal multi-hazard studies

TABLE 87. Reading the first MCA plane of the coded corpus

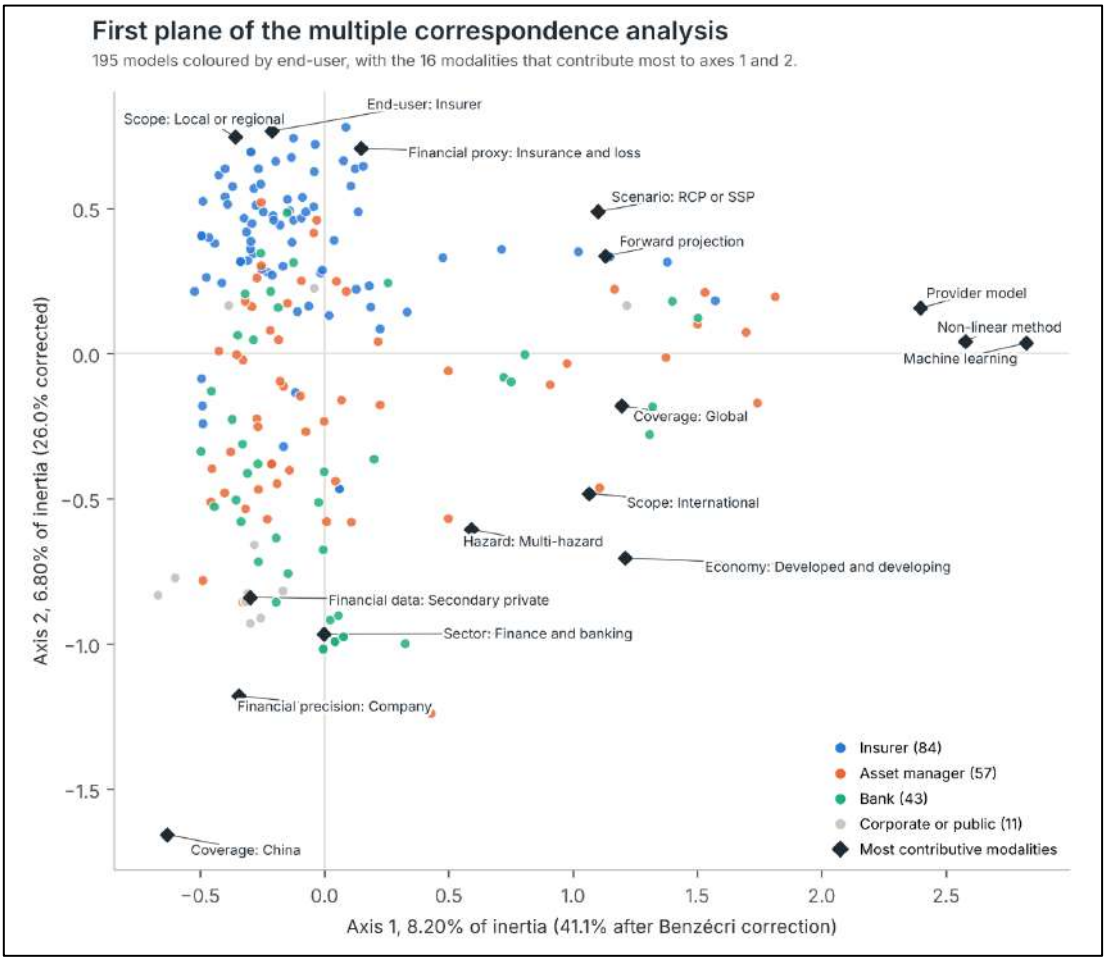


FIGURE 90. First plane of the multiple correspondence analysis of the 195 coded models

4.2.2 VAE LATENT SPACE PROJECTIONS TO DISCRIMINATE MODELS

A variational autoencoder learns a non-linear compression of the coded models (Kingma & Welling, 2013). The encoder maps each model to a Gaussian distribution in a latent space of low dimension, the decoder reconstructs the coded profile from a sample of that distribution, and training balances reconstruction accuracy against the regularity of the latent space (Box 4.6). The posterior means are then projected on a plane with t-SNE, which preserves local neighbourhoods (perplexity of 25; van der Maaten & Hinton, 2008).

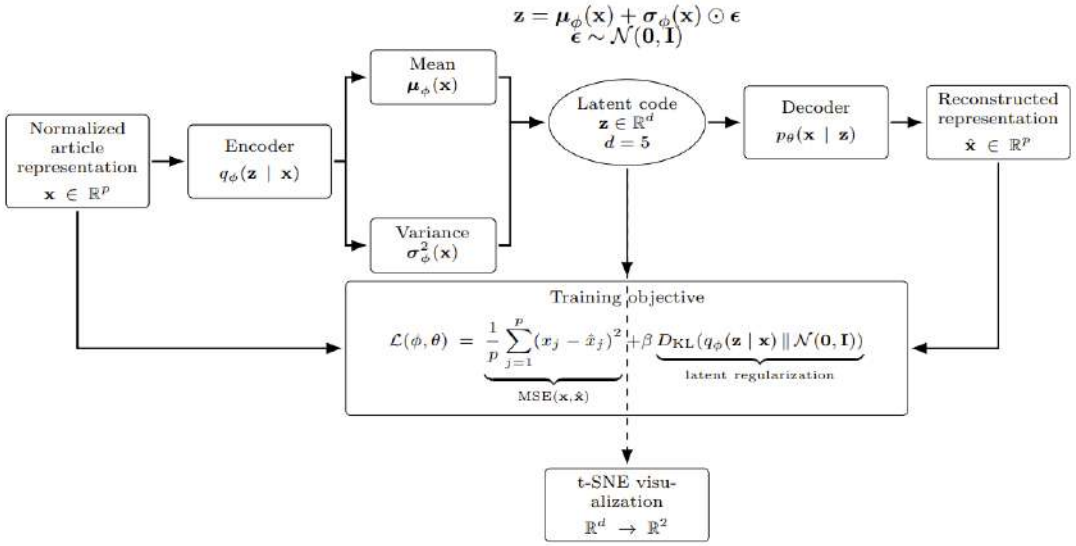


FIGURE 91. Architecture of the variational autoencoder
Source: Vigneron and Vigneron (2026).

BOX 4.6. VARIATIONAL AUTOENCODER ON THE CODED CORPUS

$$x_{kj} = z_{kj} / (Q\sqrt{c_j}),\ c_j = (1/(nQ))\,\Sigma_k\, z_{kj} \quad (4.18)$$

$$q_\varphi(\zeta \mid x) = \mathcal{N}(\mu_\varphi(x), diag\ \sigma^2_\varphi(x)) \quad (4.19)$$

$$\zeta = \mu_\varphi(x) + \sigma_\varphi(x) \odot \varepsilon,\ \varepsilon \sim \mathcal{N}(0, I) \quad (4.20)$$

$$\mathcal{L}(\varphi, \theta) = (1/p)\,\Sigma_j\,(x_j - \hat{x}_j)^2 + \beta\, D_{KL}(\, q_\varphi(\zeta \mid x) \,\|\, \mathcal{N}(0, I)\,) \quad (4.21)$$

With

- x_{kj} is the input of model k on modality j, rescaled to the MCA metric
- z_{kj} is the entry of the disjunctive table of Equation 4.5
- c_j is the relative frequency of modality j
- $q_\varphi(\zeta \mid x)$ is the encoder, a Gaussian distribution over the latent code given the coded model x
- ζ is the latent code of the model, of dimension d = 5
- $\mu_\varphi(x), \sigma_\varphi(x)$ is the mean and standard deviation produced by the encoder
- φ, θ is the parameters of the encoder and of the decoder

BOX 4.6. VARIATIONAL AUTOENCODER ON THE CODED CORPUS

- $\boldsymbol{\varepsilon}$ is a standard Gaussian draw, which makes sampling differentiable
- $\hat{\boldsymbol{x}}_{\boldsymbol{j}}$ is the reconstruction of modality j by the decoder
- $\boldsymbol{p}$ is the number of modalities
- $\boldsymbol{\beta}$ is the weight of the regularisation term, set to 1 in the standard formulation and to 10^{-4} here, reached after a linear warm-up over 300 epochs
- $\boldsymbol{D_{KL}}$ is the Kullback–Leibler divergence between the encoder distribution and the standard Gaussian prior

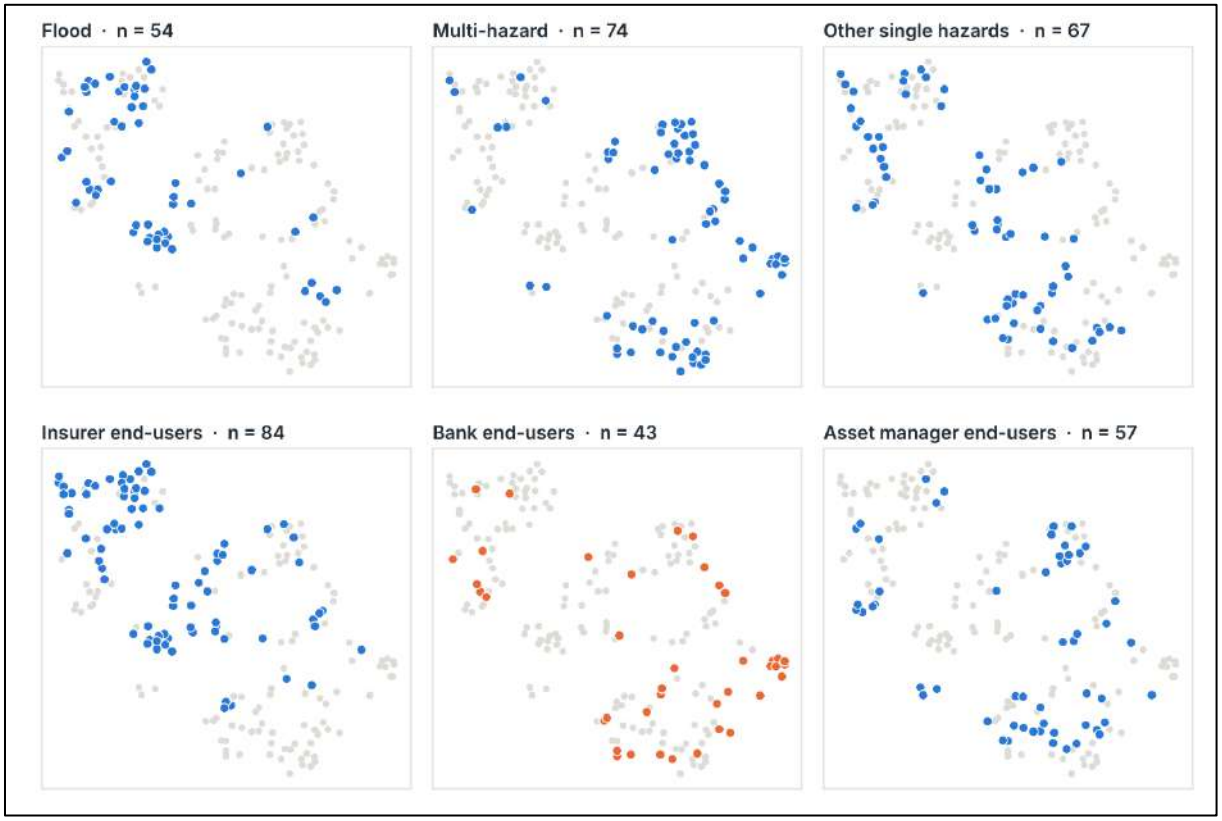


FIGURE 92. T-SNE projection of the VAE latent space, one panel per group of models

The latent code is written ζ here to keep z for the disjunctive table. A first attempt on the Risk Forum corpus produced a latent representation too diffuse to yield interpretable groups, a result attributed to the small sample and to the binary nature of the inputs. The current iteration, trained on the rescaled input of Equation 4.18 for the 195 models of the coded corpus, is read one group at a time.

- Flood models and insurer end-users share the upper-left region, in compact pockets of near-identical profiles, consistent with the homogeneous hedonic and damage-function tradition found in the Kohonen map.
- Other single-hazard models (drought, wildfire, wind, earthquake, heat, sea-level rise) run along the left and center of the map without a group of their own.
- Multi-hazard models fill the right half of the map, with their densest pockets in the lower right.

- Bank models spread over the lower right, which matches the model-level Kohonen result of Section 4.1.3.

A neighbourhood count makes this reading precise. Among the five nearest neighbours of a model in the latent space, 58% share its hazard (24% expected by chance) and 58% its end-user (32% by chance). The first five MCA axes give almost the same neighbourhoods (53% and 60%), so the non-linear compression adds little to the linear benchmark at this sample size. Its validation error, about 1.8 times its training error, also points to partial memorisation. The open map also shows its own fragility. Its layout changes with the random seed, the weight β and the latent dimension, proximity reflects shared coding choices that may hide different theoretical assumptions, and small sub-literatures tend to be absorbed by larger groups because separating a few models barely reduces the average reconstruction error. t-SNE adds its own variability, since it keeps local neighbours together but distorts the distances between groups, so the size of a group and the gap between two groups on the plane should not be interpreted. A latent map is one plausible realisation of the corpus, to be read across several runs.

4.2.3 COMPARING MODELS FROM THE LITERATURE

The three maps agree on one diagnosis. The corpus is structured, as the permutation test of Section 4.2.1 shows, but it does not split into sealed families. A k-means partition of the first five MCA axes reaches a best silhouette of 0.39, below the 95th percentile of the shuffled data (0.49).

Model families are therefore regions along gradients, which closed maps cut into cells and open maps display as continua. The three projections share the same coding and the same χ^2 geometry, so their agreement is a consistency check, short of an independent test of stability. A model family that appears in the closed grid, on the MCA plane and in the latent space is consistent across methods. A family that appears in one map only depends on the method and calls for caution. Stability proper rests on the checks of Section 4.1.3 and on the coding audit of Section 4.3.2.

MODEL FAMILY	KOHONEN 5 × 5	MCA PLANE	VAE LATENT SPACE	OBSERVATION
Actuarial and insurance	(1,3), (1,4), (2,4), top right	Upper left, insurance and loss	Upper left, with flood models	Consistent across the three methods
Banking and credit risk	(4,5), next to (3,5), on the map of modalities	Lower left, company-level credit models	Lower right, spread	Consistent on maps of modalities, spread on maps of models
Forward-looking and corporate	(1,1), (1,2), (2,1)	Right side of axis 1	Not resolved	Consistent in two of the three maps
Compound and multi-hazard	(1,2), with asset managers	Lower center	Right half, dispersed	Consistent as a transversal zone
Developing-country models	(4,4) and (5,3), with China and public authorities	Lower left, with Chinese coverage	Not resolved	Visible on two maps

TABLE 88. Consistency of model families across the three projections

Once a family is located, its members can be compared on common ground. The table lists four families placed by the maps, each with one representative model from the literature review. Models inside a family share hazard, financial indicator and method, so their reported coefficients and fit statistics can be read side by side, with care, since samples, target variables and evaluation protocols still differ. Across families, the comparison stays at the level of the signature.

MODEL FAMILY	HAZARD	FINANCIAL INDICATOR	METHOD	REPRESENTATIVE MODEL
Public finance and wildfires	Wildfire	Annual suppression expenditure	Time series and linear regression with cost scenarios	State wildfire budgets in the western US (McIver et al., 2021)
Agricultural exposure to drought	Drought	Farm income, margin, profit per hectare	Panel regression	Australian farm businesses (Kingwell & Xayavong, 2017)
Abnormal returns and natural catastrophes	Earthquake	Stock return, volatility, market index	Event study	Market reaction to the 2011 Tohoku earthquake (Tao et al., 2019)
Portfolio exposure to physical risk	Compound, multi-hazard	Portfolio and credit losses	Network and scenario model	Global financial risk mapping (Mandel et al., 2025)

TABLE 89. Model families located by the maps, with one representative each

4.3 USER-GUIDED FILTERING

The maps of Sections 4.1 and 4.2 are unsupervised and describe the literature as it stands. A model user starts from a need, defined by a hazard, an exposure, a financial output and a regulatory frame. This need sets labels or constraints on the coded corpus. No model is trained on these labels, so the step is a guided selection on top of unsupervised maps, distinct from supervised learning in the statistical sense. Three operations follow, labelling, filtering and ranking, and together they turn a map of the field into a shortlist of comparable models.

4.3.1 MAPS AND FILTERS

Labelling colours an unsupervised map with a variable chosen by the user, as the latent space of Section 4.2.2 was read by hazard and end-user. The map is learned without the label, so a label that falls into compact regions reveals a variable that structures the field, and a scattered label reveals one that cuts across it. Filtering keeps the models whose signature matches the constraints of the need, and ranking orders the remaining models by their distance to a target profile (Box 4.7).

Since the corpus forms gradients, these operations work best locally. Filtering first and then training a small map on the filtered subset orders the remaining models so that grid neighbours are close substitutes, and the neighbourhood counts of Section 4.2.2 show that nearest neighbours share hazard and end-user far more often than chance. The grid must shrink with the subset, for instance 3 × 3 for 25 to 50 models, and the filter variable becomes constant, so a local map shows variation within one need and never across needs. The figure below applies this to the banking models.

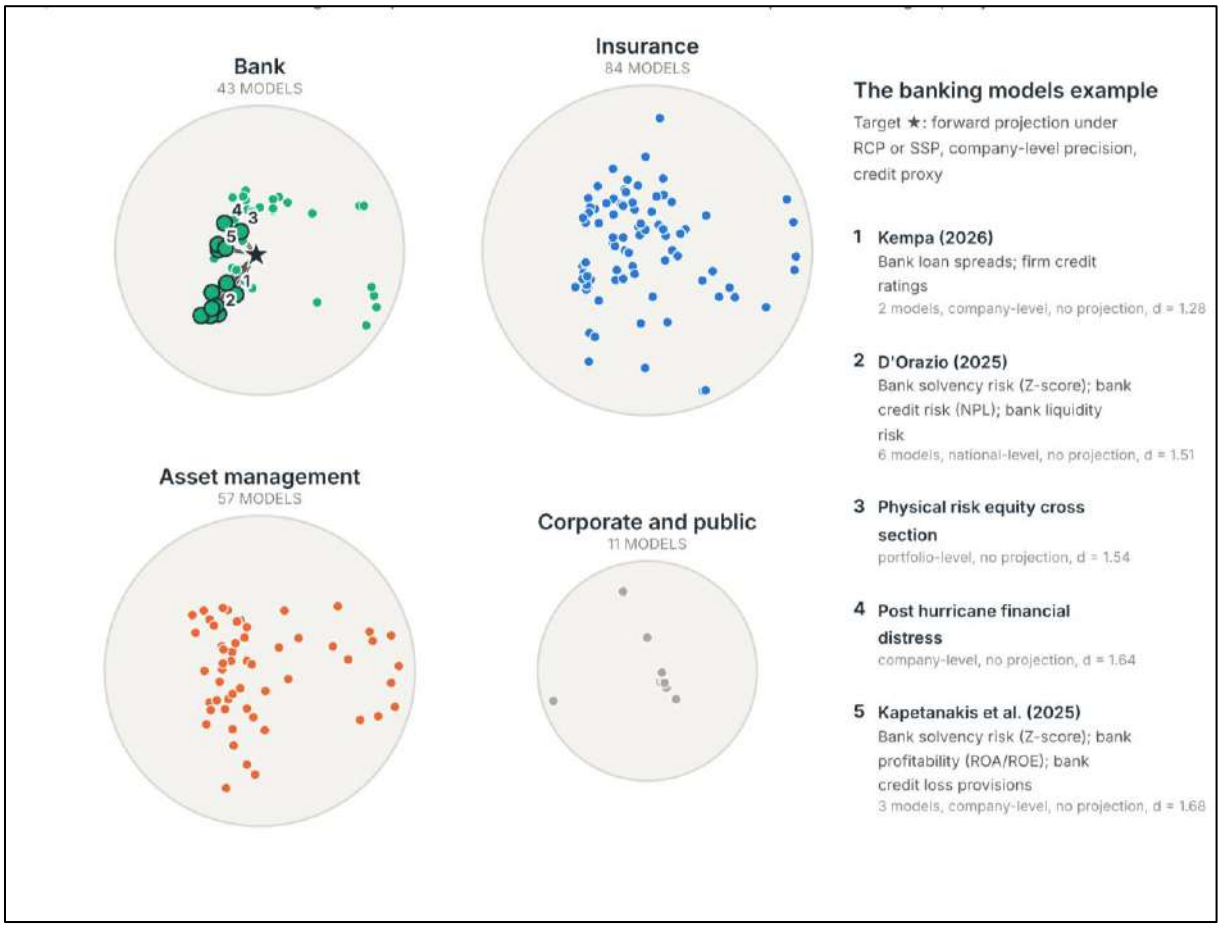


FIGURE 93. Filtered map of the coded corpus

The filtered map of the coded corpus shows models that are filtered by end-user and placed within each group by classical multidimensional scaling of their χ^2 distances (Equation 4.9).

The star marks a target profile for a bank, and the numbers give the ranking of Equation 4.23, computed on the full distance. Models from the same article are listed together.

BOX 4.7. FILTERING AND RANKING THE CODED CORPUS

$$\mathcal{F}(\gamma) = \{ i : q_{ij} = \gamma_j \ for\ every\ j \in J \} \quad (4.22)$$

$$s_i = d(q_i, q^*), \ i \in \mathcal{F}(\gamma) \quad (4.23)$$

With

- $\mathcal{F}(\gamma)$ is the filtered set of models that satisfy the constraints γ
- i is the index of a coded model
- J is the set of variables constrained by the user
- q_{ij} is the modality taken by model i on variable j
- γ_j is the modality required by the user on variable j
- s_i is the ranking score of model i, lower scores marking closer models
- q^* is the target profile, the signature of the ideal model for the need
- d is the distance between models of Equation 4.9, computed on the remaining free variables

Filtering the coded corpus shows how fast it narrows.

Of the 195 coded models, 51 belong to the financial markets and banking sector. Among them, 26 measure the financial impact at asset or company level, and only 3 add a forward projection. One model prices heat stress in US assets for bank end-users, one maps multi-hazard losses on assets worldwide, and one projects physical risk on the default risk of banks worldwide (Abinzano et al., 2026). A bank looking for a forward-looking, borrower-level model therefore starts from a handful of candidates, which shows that forward-looking hazard-to-credit models remain scarce in the corpus.

FILTER STEP	BANKING USE CASE	INSURANCE USE CASE	CONSTRAINT APPLIED
Coded corpus	195	195	None
Sector	51	29	Financial markets and banking, or insurance
Granularity or method	26	5	Asset or company-level financial precision, or explicit damage function

FILTER STEP	BANKING USE CASE	INSURANCE USE CASE	CONSTRAINT APPLIED
Time horizon	3	Not applied	Forward projection under a climate scenario

TABLE 90. Two filters applied to the coded corpus

4.3.2 TOWARD A SYSTEMATIC COMPARISON OF MODELS

A shortlist identifies comparable models without saying which one performs best. Proximity on a map reflects shared design choices, and predictive performance can only be compared on a shared reference design (Vigneron & Vigneron, 2026). The accuracy information reported in the corpus shows how far the literature stands from such a design. About one model in ten reports no accuracy measure, one in seven reports a qualitative validation only, and the accuracy of another one in eight could not be established. R^2 is the most common measure (42%). For the 41 most recently coded models, where each validation type is recorded separately, only 3, all from one article, report an out-of-sample test, and none reports a backtest.

This is the main finding of Part 4 for the model user. The literature offers many ways to quantify the financial impact of physical hazards, but the information needed to judge their predictive performance remains thin. More than a third of the coded models (36%) report no quantitative accuracy measure, and out-of-sample validation or backtesting remains exceptional. This explains why models remain hard to compare and why a standardized characterization, such as our model cards is needed .

MAIN ACCURACY MEASURE REPORTED	MODELS	SHARE OF 195
R^2 or adjusted R^2	81	42%
Qualitative validation only	27	14%
Accuracy unknown	24	12%
No accuracy measure reported	19	10%
Error metrics (RMSE, MAE)	12	6%
Statistical tests	12	6%
Information criteria (AIC, BIC, log-likelihood)	10	5%
Other quantitative indicator or convergence check	5	3%
Not yet coded	5	3%

A systematic comparison therefore requires four extensions of the framework. The first is a common reference design, a set of inputs on which every model of a family is run, so that the distance between two models becomes a

distance between their outputs (Box 4.8). The second is a reporting standard that makes the accuracy and validation fields mandatory, and for insurance models records which Solvency II channel the model feeds (Vigneron & Vigneron, 2026). The third is a double-coding protocol, where two reviewers code a random subsample and their agreement is measured, together with a cleaning step that removes duplicated variables and harmonises the spelling of modalities (Section 4.1.3). The fourth extends the stability checks of Section 4.1.3 by perturbing a share of the coded modalities and retraining the maps.

BOX 4.8. COMPARING MODELS ON A SHARED DESIGN AND AUDITING THE CODING

$$d(M_i, M_j) = (1/N)\, \Sigma_k \parallel M_i(x_k) - M_j(x_k) \parallel \quad (4.24)$$

$$\kappa = (p_o - p_e) / (1 - p_e) \quad (4.25)$$

With

- $d(M_i, M_j)$ is the output distance between models i and j on the shared design
- $M_i(x_k)$ is the financial output of model i for input k
- x_k is the k-th input of the reference design, a hazard and exposure scenario
- N is the number of inputs in the reference design
- κ is the inter-coder agreement corrected for chance (Cohen, 1960)
- p_o is the observed share of modalities coded identically by the two reviewers
- p_e is the share of agreement expected by chance from each reviewer's marginal frequencies

Updatability completes the framework. The coding step can be partly automated with a pipeline in which a bibliographic API retrieves new articles, a language model extracts the framework variables, and the resulting table feeds the Kohonen, MCA and VAE algorithms. A human verification step stays mandatory before any result is treated as final. The authors estimate that processing 100 articles takes about one day of computation and one month of verification, against several months to a year for a thematic or systematic review (Vigneron & Vigneron, 2026).

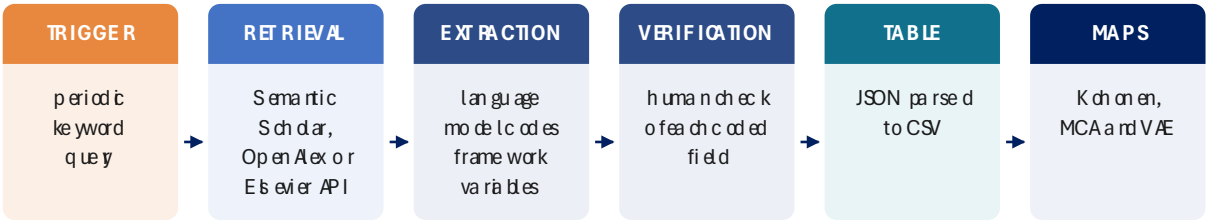


FIGURE 94. Updatable pipeline of the self-organizing literature review.
Source: Adapted from Vigneron and Vigneron (2026).

Part 4 thus moves from description to selection. Closed clusters give an index of model families, open clusters show the transitions between them, and user-guided filtering turns both into a shortlist for one use case. The maps describe a structured field organized along gradients, with no sealed families, so they serve best as navigation tools. Across the three, one result dominates. Many models are documented and few are validated out of sample, so a user finds comparable candidates more easily than evidence on which of them performs best. Part 5 draws the insights that follow from this structure and offers a concrete application for considering this workflow.

5 THE CLIMATE FINANCIAL RISK HUB, REFERENCES AND APPENDIX

The Cookbook is the result of a sequence of steps. Each step answered a question that the previous one had left open. This section describes that sequence and the reasons behind it, since the order of the steps shaped the content of the volume.

5.1 THE CLIMATE-RELATED FINANCIAL RISK HUB: AN ONLINE PLATFORM TO REVIEW MODELS

Models gathered in the Climate Risk Models Hub

Every model in the Cookbook, whether it comes from the literature or the market, is recorded in the same format, the model card. The card is the basic unit of the research protocol. It is used to code each model, to compare models with one another and to display them in the Climate risk models hub.

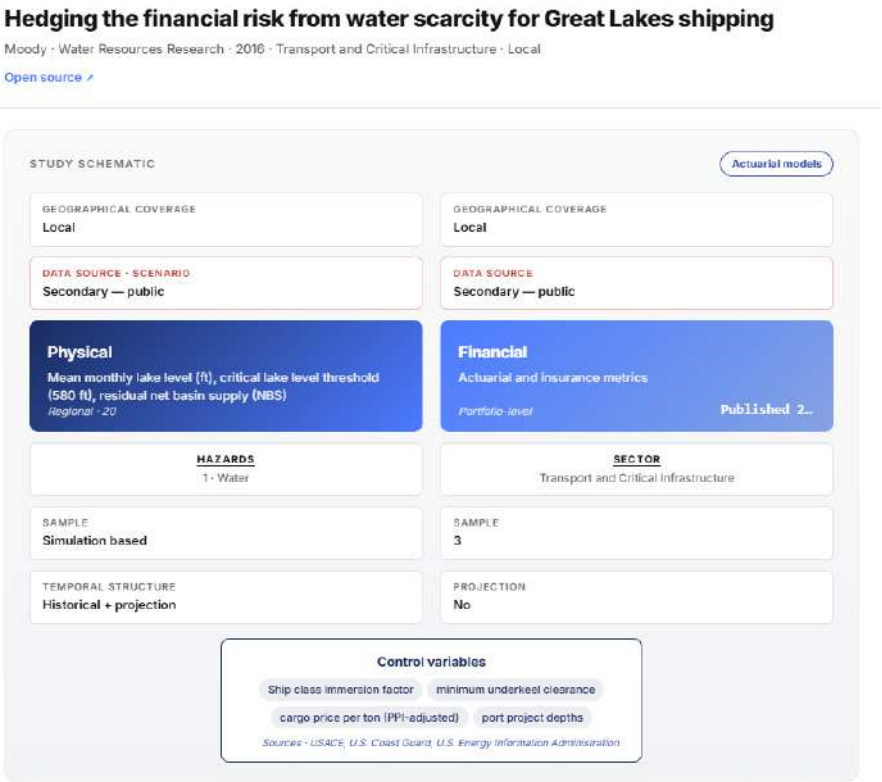


FIGURE 95. The model card as the unit of record

The card opens with a header. The header gives the title of the study or the name of the solution, the author or provider, the outlet, the year, the sector, the geographic scope and a link to the source when it is openly accessible. A tag shows the model family to which the model belongs, such as actuarial models in the example shown.

The body of the card is a study schematic in two columns. The left column describes the physical side of the model and the right column describes the financial side. Each side gives the geographic coverage and the type of data source, with a distinction between primary data, public secondary data and private secondary data.

At the center of the card, two blocks summarize what the model takes in and what it produces. The physical block names the physical proxy used, for example a mean monthly lake level compared with a critical navigation threshold, together with its spatial scale. The financial block names the category of financial metric produced, such as actuarial and insurance metrics, together with the level at which it is computed, from asset to portfolio.

Below the schematic, the card records the hazards covered, the sector, the size and nature of the sample, the temporal structure (historical, projected or both) and whether the financial results are projected forward. A box of control variables lists the parameters that condition the result, with their data sources. The card closes with a short summary of findings.

This format serves three purposes. A reader can grasp the logic of a model in one page, without reading the full paper or methodology note. Literature models and commercial solutions are described with the same fields, so they can be compared directly. Finally, each field corresponds to a coded variable, which allows the cards to be counted, filtered and clustered.

The dashboard of the Climate risk models hub

The model cards are gathered in the Climate risk models hub. Its dashboard gives an overview of the whole corpus.

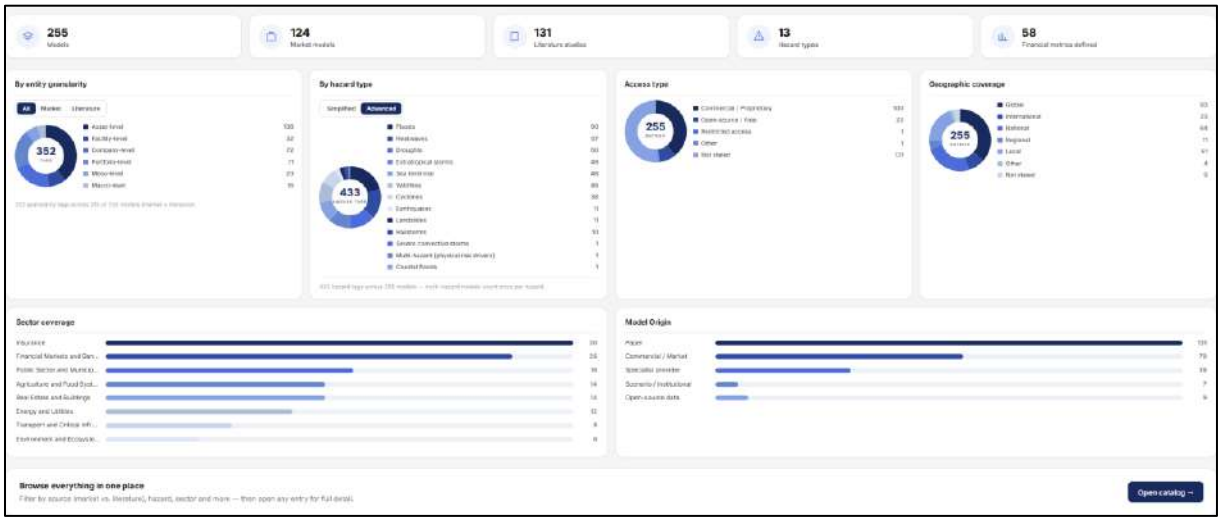


FIGURE 96. Climate risk models hub dashboard

At the time of writing, the Hub records 255 models. Of these, 131 are studies from the literature and 124 are market models. They cover 13 hazard types and 58 defined financial metrics. The dashboard breaks the corpus down along six dimensions.

- **Entity granularity** shows the level at which the financial impact is computed. Asset-level models dominate (138 tags), followed by company-level (72) and portfolio-level (71) models. Facility, meso and macro levels are less frequent. One model can carry several granularity tags.
- **Hazard type.** Floods are the most covered hazard (90 models), followed by heatwaves (67), droughts (60), extratropical storms (49), sea level rise (48), wildfires (46) and cyclones (38). Earthquakes, landslides and hailstorms appear far less often. Multi-hazard models count once per hazard.
- **Access type** separates commercial and proprietary solutions (100) from open-source or free solutions (22).
- **Geographic coverage.** Global models (93) are the largest group, followed by national (64) and local (51) models.

- **Sector coverage.** Insurance (30) and financial markets and banking (26) lead, followed by the public sector, agriculture, real estate and energy.
- **Model origin.** This separates models from academic papers (131), from commercial market solutions (70), specialist providers (38), scenario and institutional frameworks (7) and open-source data products (9).

Each view can be filtered by source, so the literature and the market can be compared on any dimension. A catalogue gives access to the full model card of each entry, with filters by source, hazard, sector and other criteria.

The dashboard already shows several features of the field. The market concentrates on asset-level, global and multi-hazard solutions. The literature is more often local or regional and focused on a single hazard. Some hazards, such as hail and convective storms, remain poorly covered by both.

An open and evolving catalogue

The Hub extends the Cookbook beyond this current document. Providers can add their own models through the same fields as the model card, and update them when their methodology changes. Users can search the catalogue by hazard, financial metric, scale, sector or institution type, and identify the models that match their question. Information supplied by providers is kept separate from information validated by the research team. The Hub follows the same editorial line as the Cookbook, which is to describe models and not to rank them.

The Cookbook gives a picture of the field at a given date. The market for physical risk models moves quickly. New providers enter, others are acquired by financial institutions or data vendors, and methodologies are updated as climate data and scientific knowledge improve.

SOLR for Financial Impact Physical Climate Risks Models

The literature on the financial impacts of physical climate risks grows quickly, and it spans actuarial science, banking, finance, economics, hydrology and climatology. A thematic review explores the field but covers only part of it. A systematic review aims at exhaustivity, but it takes about a year for a hundred articles and becomes outdated soon after completion. Neither approach fits a field that expands every month.

The Self-Organizing Literature Review (SOLR) was developed to address this problem (Vigneron & Vigneron, 2025). Its purpose is to classify a large and growing body of models according to their content, and to produce a map of the field that can be updated. Generation of the map for a hundred articles takes about a day. Human verification takes between one week and one month, depending on the document length and resources available to verify.



FIGURE 97. A simplified view of our model classification pipeline

The protocol was built in two stages.

In the first stage, the corpus was assembled and coded by hand. A Scopus search on the terms “physical risk” and “financial impacts” in titles, abstracts and keywords returned more than 150 peer-reviewed articles. Each article was coded on more than 25 variables. These include the hazard type, the physical proxy and its scale, the financial proxy and its level, the method, the sector, the geographic scope, the data sources, the scenarios and the intended end user. The qualitative variables were converted into binary indicators and combined into a single table of co-occurrences. A self-organizing map and a variational autoencoder were then trained on this table to group the articles into families. The method and its results are presented in Part 4.

In the second stage, the collection and coding steps were automated. The pipeline shown in the figure runs as follows. A trigger, either periodic or keyword-based, launches a query to the Semantic Scholar API. The articles

retrieved form a new sample. Each article is passed to a large language model, which extracts the coding variables into a structured JSON file. The JSON is validated and parsed into a table, and the table feeds the self-organizing map, the multiple correspondence analysis and the variational autoencoder in Python. The workflow is orchestrated with Node-RED.

The model card, the Hub and SOLR are three parts of one system. The model card defines what is recorded for each model. The Hub is the interface through which these records are consulted. SOLR is the engine that collects, codes and organizes them.

For the literature, SOLR aims to feed the Hub directly. New articles are detected, coded with the same variables as the model card, verified by the research team and added to the catalogue. For the market, the same role is played by the providers themselves, who fill in the model card fields through the Hub. Both sources therefore enter the catalogue in the same format. The clustering can then be rerun on the whole corpus, so the families described in Part 2 are revised as the field evolves instead of being fixed at the date of publication.

This design turns the Cookbook from a one-off review into a living cartography of the field. Three benefits follow. The classification is reproducible, since the same query and the same coding produce the same map. It is transparent, since every model can be traced back to its card and every card to its source. It also scales, since the cost of adding a hundred new articles falls from months of reading to days of verification.

Starting from the literature

The project began with a review of the academic and institutional literature. The aim was to establish a common grammar before meeting the market. The literature sets out the transmission chain from hazard to financial impact, the main model families and the assumptions on which each family rests. Starting with providers would have meant adopting the vocabulary of each solution and comparing products on their own terms. The literature review produced a corpus of more than 200 models and a first taxonomy, against which commercial solutions could then be read.

Meeting the providers

The literature describes methods, while practice is shaped by providers and many of the choices that determine a final figure appear in no academic paper. These include the source and resolution of hazard data, the selection of damage curves, the treatment of missing asset information and the aggregation from site to company or portfolio. As public documentation on these points is often partial but many providers are making efforts to increase this aspect of their solution, more than 40 interviews were conducted with providers, in order to understand how models are built, calibrated and used in practice.

Standardizing through a survey

The interviews produced rich but heterogeneous material. Each conversation followed the logic and vocabulary of the solution discussed, and the answers could not be compared directly. A standardized questionnaire was therefore sent to providers after each interview. It covered scope, hazards, the modelling chain, the nature of the damage function, the level of analysis, uncertainty, backtesting, transparency, data sources and use cases. The questionnaire turned qualitative exchanges into comparable fields. It also gave each provider the opportunity to confirm or correct the description of its own solution.

Turning to end users

At this stage, the supply side was mapped. The demand side remained to be understood. Banks, insurers and asset managers do not ask the same questions of a model, and the same output can serve one decision well and another poorly. Discussions with end users and supervisors identified the decisions that models are expected to inform, such

as underwriting, credit assessment, portfolio allocation, stress testing and regulatory reporting. They also identified the constraints these institutions face in data, budget, transparency and auditability. These exchanges revealed a recurring gap. Models are often selected on the reputation of a provider, and less often on the fit between the model and the question asked.

Clustering last

The clustering was carried out last, once literature models and market solutions had been coded on the same grid. Only at that point could both corpora be placed on a single map, with no a priori separation between research and practice. The clustering confronted the taxonomy built from the literature with the structure that emerged from the data. Where the two matched, the taxonomy was confirmed. Where they diverged, the model families were revised. The method and its results were presented in Part 4.

5.2 A PANORAMIC VIEW OF THE FINANCIAL RISK MODELLING ECOSYSTEM

The physical risks such as 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.

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 (Higuera Roa et al., 2025), 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.

We established that the ecosystem of physical climate risk modelling has expanded rapidly over the last decade. Once a relatively specialized domain largely confined to catastrophe modelling and insurance analytics physical risk modeling has progressively evolved into a broad and heterogeneous market involving financial institutions, climate analytics providers, consulting firms, data vendors, academic laboratories, public institutions, regulators, and technology companies.

The importance of climate-related financial risks within prudential supervision, investment management, insurance underwriting, and strategic corporate decision-making is linked to this expansion.

For non-specialists, this diversity may appear difficult to navigate.

Physical risk methodologies contain technical vocabularies combining climate science, geospatial analytics, engineering, economics, actuarial science, and quantitative finance. For a large number of proprietary platforms, their models operate as "*black boxes*" (Colman, 2021), climate-adjusted financial indicators without always making explicit the assumptions, proxies, or simplifications embedded within the modelling chain.

This opaque approach generates important concerns for financial institutions and regulators alike. Physical climate risk models increasingly influence investment decisions, prudential exercises, underwriting strategies, portfolio allocation, stress testing procedures, and disclosure frameworks. Yet different models may produce substantially different results even when analyzing similar exposures under identical climate scenarios. Divergence between outputs may emerge from differences in hazard datasets, vulnerability assumptions, downscaling methodologies, financial proxies, temporal horizons, or aggregation mechanisms.

These divergences generate what may broadly be described as model risk (***See Section 1.4.3***). Model risk here refers to the possibility that methodological assumptions, incomplete datasets, inappropriate proxies, calibration biases, or misunderstood outputs lead to inaccurate estimations of physical climate financial impacts. A model may appear scientifically sophisticated while remaining poorly adapted to the operational realities or prudential requirements of financial institutions. High predictive performance under laboratory conditions does not automatically guarantee regulatory usability, operational robustness, explainability, or decision relevance.

This distinction between theoretical performance and practical usability is central to the current evolution of the market. Certain models prioritize scientific precision and physical realism. Others prioritize scalability,

interpretability, regulatory alignment, computational efficiency, or integration within existing financial workflows. Financial institutions therefore face a complex trade-off between methodological sophistication, transparency, explainability, operational integration, and prudential compatibility.

Regulators increasingly monitor these issues closely. Prudential authorities do not merely seek accurate climate metrics; they also seek methodologies that remain understandable, auditable, coherent across institutions, and compatible with supervisory expectations. Excessive fragmentation of modelling approaches may create inconsistencies in climate disclosures, prudential exercises, and risk assessments across the financial sector. In extreme cases, overreliance on poorly understood models could generate false confidence regarding the true scale of physical climate vulnerabilities.

For this reason, one of the objectives of this cookbook is to provide a structured framework capable of organizing this heterogeneous ecosystem into coherent analytical categories. Rather than treating all models as interchangeable, the cookbook seeks to identify distinct modelling families according to their methodological architecture, underlying assumptions, intended users, hazard focus, financial translation mechanisms, and operational objectives. Such categorization helps clarify what different models are actually designed to do, which questions they attempt to answer, and where their limitations may lie.

Another dimension is about the relationship between institutional context and model selection. Different financial institutions face different regulatory obligations, portfolio structures, temporal horizons, and operational constraints. Insurance companies frequently prioritize high-resolution catastrophe modelling and probabilistic loss estimation. Banks often focus on credit deterioration, collateral valuation, and stress testing compatibility. Asset managers may prioritize portfolio aggregation tools and long-term valuation analysis. Public institutions and central banks may instead focus on macroeconomic resilience and systemic stability. **The "best" model therefore depends heavily on the institutional use case for which it is deployed.**

This cookbook does not attempt to review every possible climatic hazard or modelling approach currently available on the market. In order to maintain analytical coherence and regulatory relevance, the analysis focuses primarily on the physical hazards identified within the European Taxonomy framework. This includes both acute hazards such as floods, storms, wildfires, heatwaves, and droughts, and chronic hazards such as sea-level rise, long-term temperature shifts, changing precipitation patterns, and water stress. Anchoring the analysis within the European Taxonomy ensures direct relevance for European financial institutions operating under evolving prudential and disclosure requirements.

The following sections present the methodological protocol used to construct this panorama of physical climate risk modelling ecosystems. Attention is devoted to data collection procedures, literature review methodologies, market interviews, provider analysis, end-user perspectives, and exploratory classification techniques used to organize the rapidly growing modelling landscape.

5.2.1 A COLLECTION AND TREATMENT

The physical climate risk modelling ecosystem evolves quickly and spans several disciplines, industries and institutional settings. No single source covers the range of methods currently in use. This Cookbook therefore combines scientific literature, exchanges with the academic community, interviews with model providers and interviews with end-users in financial institutions.

The aim was to identify existing models and to document how they operate, which assumptions they rely on, which financial objects they target, and where they sit within the climate analytics ecosystem. Academic and commercial

models were coded with a common set of variables, so that they can be compared on the same basis (Vigneron & Vigneron, 2026). The following subsections describe each source, and Section 1.7.7 presents the classification methods applied to the coded data.

5.2.2 AUTOMATED SOURCES OF COLLECTION

The academic corpus was built through a systematic search on Scopus, OpenAlex and Semantic Scholar, with keywords such as “physical risk” and “financial impacts” in titles, abstracts and keywords, completed by other sources. The search covered physical climate risk modelling, catastrophe modelling, climate finance, geospatial vulnerability analysis, macroeconomic climate impacts and financial transmission mechanisms.

To be retained, a study had to report **at least one proxy of physical hazard**, such as flood depth, wildfire extent, drought or heatwave intensity, **and at least one proxy of financial impact**, such as losses, returns, solvency or credit indicators. This rule kept our corpus focused on the link between physical hazards and financial outcomes, which remains the least standardized part of climate risk analytics (Vigneron & Vigneron, 2026).

Each retained study was coded along two axes, physical risk and financial impact. The variables cover geographic scope and sector, baseline and projection horizon, the type and precision of the financial impact, the hazard type and its spatial resolution, the modelling method, the scenarios used and the presence of a damage function. The coded corpus brings out the recurring modelling architectures, the dominant hazards, the most common financial proxies and the main methodological limits of the field (see sections 4, 4.3.1 and 4.3.2 for more information).

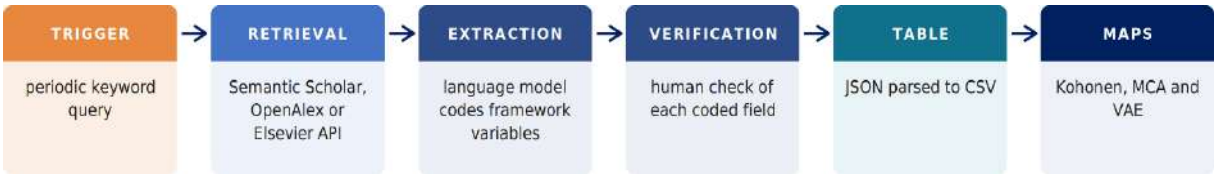


FIGURE 98. The updatable pipeline of the self-organizing literature review. Adapted from Vigneron and Vigneron (2026).

5.2.3 INSIGHTS FROM THE FIELD

Scientific literature alone does not fully capture the operational realities of the market. Many commercially deployed methodologies remain partially proprietary and evolve more rapidly than academic publication cycles. Therefore, we also incorporated insights from the field. Many targeted interviews were made on the climate physical risks and their financial impacts in the insurance, asset management and banking sector.

For this reason, the research process incorporated these extensive field insights collected through discussions and various documents from institutions, analytics providers, consultants, regulators, and industry practitioners. These exchanges provided valuable information regarding operational constraints, model usability, client expectations, data limitations, integration challenges, and regulatory concerns. Field insights also revealed substantial differences between theoretical modelling capabilities and the actual tools deployed within institutional decision-making processes.

5.2.4 ACADEMIC INSIGHT

Our research process sought to first identify patterns underlying the academic model's ecosystem. This included examining how different disciplines conceptualize physical climate risks, how uncertainty is treated across modelling traditions, and how hazard, exposure, vulnerability, and financial translation mechanisms are operationalized.

To that end, we relied on a working group within the PARC's Experts and Scientific committees. This working group gathered over several times, every two months to discuss views on the research questions, the current state of the art and solutions on the market. The Financial Risks Forum, an annual scientific gathering at the Institut Louis Bachelier's initiative held its 19th edition in 2026. This was the occasion for us to present our work on model classification.

The French Actuary congress held by the French Actuary Institute or Institut des actuaires was also the occasion to present our work in front of insurance companies and researchers in actuarial science.

During these sessions, we gathered the feedback. The analysis highlighted recurring methodological tensions between physical realism, computational tractability, interpretability, and scalability of the models used. Certain approaches prioritize detailed engineering realism at local scale, while others prioritize broad economic coverage at global scale. These trade-offs strongly influence model outputs and applicability. Given the broad view that all those actors and institutions provided, only one type of actor remained to be interviewed.

5.2.5 PROVIDERS INSIGHT

A substantial component of the research relied on direct exchanges with climate analytics providers and executives operating within the market. These discussions helped clarify how providers position their methodologies commercially, which hazards and financial metrics they prioritize, how they differentiate themselves technologically, and how they perceive evolving regulatory expectations.

Provider discussions also revealed important strategic trends shaping the market, including the growing importance of geospatial analytics, the integration of artificial intelligence techniques, increasing demand for explainability, and the tension between proprietary competitive advantages and transparency requirements. You will find several use cases all throughout parts 2 and 3. These are the results of those extensive discussions.

We conducted over 30 interviews with them.

5.2.6 END-USER INSIGHTS

Understanding the ecosystem also requires understanding how models are actually used by financial institutions. End-user discussions therefore focused on operational integration, interpretability, reporting requirements, governance structures, and decision-making constraints.

Many institutions emphasized that model adoption depends not only on predictive sophistication, but also on explainability, auditability, compatibility with internal systems, regulatory defensibility, and integration within existing risk management frameworks.

End-users frequently expressed concerns regarding excessive methodological opacity, inconsistent outputs across providers, and the difficulty of comparing models operating at different scales or relying on incompatible

assumptions. We conducted over a dozen interviews with them and regular group discussions within the working group and beyond.

5.2.7 PREDICTIVE MODEL CLASSIFICATIONS

We decided, given the diversity of methodologies identified during the research process, that classification techniques were required to structure the ecosystem into coherent analytical categories. The objective of these classifications is to identify recurring methodological families, modelling clusters into a limited number of dominant analytical patterns.

The ecosystem contains a very large number of possible combinations involving hazards, physical proxies, financial metrics, spatial resolutions, temporal horizons, sectors, and modelling techniques. Traditional linear classifications rapidly become insufficient to represent such complexity.

For this reason, exploratory projection and clustering techniques may provide useful complementary perspectives. Self-Organizing Maps based on Kohonen architectures allow high-dimensional methodological ecosystems to be projected into interpretable topological structures revealing similarities between models. Variational Autoencoder latent space projections also help identify latent structures and hidden relationships across modelling approaches, providers, hazards, or financial outputs.



FIGURE 99. Another form of pipeline for self-organizing literature reviews.

These exploratory approaches do not replace expert judgement or detailed methodological review. They instead provide additional analytical tools capable of organizing a rapidly expanding ecosystem whose complexity increasingly exceeds traditional manual classification approaches.

5.3 REFERENCES

Acemoglu, D., Ozdaglar, A., & Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review, 105*(2), 564–608. https://doi.org/10.1257/aer.20130456

Acharya, V. V., Johnson, T., Sundaresan, S., & Tomunen, T. (2022). *Is physical climate risk priced? Evidence from regional variation in exposure to heat stress* (NBER Working Paper No. 30445). National Bureau of Economic Research. https://doi.org/10.3386/w30445

ACPR. (2021). *Les principaux résultats de l'exercice pilote climatique 2020* (Analyses et synthèses No. 122). Autorité de contrôle prudentiel et de résolution, Banque de France. https://acpr.banque-france.fr/system/files/import/acpr/medias/documents/20210602_as_exercice_pilote.pdf

Actuaries Climate Index. (2018). Actuaries Climate Index: Development and design. American Academy of Actuaries, Canadian Institute of Actuaries, Casualty Actuarial Society, & Society of Actuaries. https://actuariesclimateindex.org

Addoum, J. M., Ng, D. T., & Ortiz-Bobea, A. (2023). Temperature shocks and industry earnings news. *Journal of Financial Economics, 150*(1), 1–45. https://doi.org/10.1016/j.jfineco.2023.07.002

Aerts, J. C. J. H., & Botzen, W. J. W. (2011). Climate change impacts on pricing long-term flood insurance: A comprehensive study for the Netherlands. *Global Environmental Change, 21*(3), 1045–1060. https://doi.org/10.1016/j.gloenvcha.2011.04.005

Agliardi, E., & Agliardi, R. (2021). Pricing climate-related risks in the bond market. Journal of Financial Stability, 54, 100868. https://doi.org/10.1016/j.jfs.2021.100868

Albertazzi, U., Djekic, D., & Ponte Marques, A. (2026). *Physical climate risk, credit risk and lending activity* (Working Paper Series No. 3224). European Central Bank. https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp3224~8735166667.en.pdf

Albrecher, H., & Beirlant, J. (2025). Statistics of extremes for the insurance industry (to appear in Handbook of Statistics of Extremes, Chapman and Hall). arXiv:2511.22272.

Alfieri, L., Burek, P., Dutra, E., Krzeminski, B., Muraro, D., Thielen, J., & Pappenberger, F. (2013). GloFAS: Global ensemble streamflow forecasting and flood early warning. *Hydrology and Earth System Sciences, 17*(3), 1161–1175. https://doi.org/10.5194/hess-17-1161-2013

Alfieri, L., Burek, P., Feyen, L., Forzieri, G. (2015). Global warming increases the frequency of river floods in Europe. Hydrology and Earth System Sciences, 19, 2247-2260. https://doi.org/10.5194/hess-19-2247-2015

Allen, T. I., & Wald, D. J. (2007). *Topographic slope as a proxy for seismic site-conditions (VS30) and amplification around the globe* (Open-File Report 2007-1357). U.S. Geological Survey. https://doi.org/10.3133/ofr20071357

Almufti, I., & Willford, M. (2013). REDi™ rating system: Resilience-based earthquake design initiative for the next generation of buildings. Arup.

Alok, S., Kumar, N., & Wermers, R. (2020). Do fund managers misestimate climatic disaster risk? *The Review of Financial Studies, 33*(3), 1146–1183.

American Society of Civil Engineers. (2022). Minimum design loads and associated criteria for buildings and other structures (ASCE/SEI 7-22). ASCE.

Apel, H., Thieken, A. H., Merz, B., & Blöschl, G. (2004). Flood risk assessment and associated uncertainty. Natural Hazards and Earth System Sciences, 4(2), 295–308. https://doi.org/10.5194/nhess-4-295-2004

APRA. (2021). *Climate Vulnerability Assessment* [Information paper]. Australian Prudential Regulation Authority. https://www.apra.gov.au/news-and-publications/climate-vulnerability-assessment

Aprea, I. L., Perla, F., Petrella, R. C., Pietroluongo, M., & Scognamiglio, S. (2025). The role of climate risk in shaping longevity dynamics: Extending stochastic mortality models. *Socio-Economic Planning*

Sciences. Advance online publication. https://doi.org/10.1016/j.seps.2025.102353

Association of British Insurers (ABI). (2009). *The financial risks of climate change*. Association of British Insurers.

Atkinson, A. C. (1969). A test for discriminating between models. Biometrika, 56(2), 337–347. https://doi.org/10.1093/biomet/56.2.337

Atkinson, A. C. (1970). A method for discriminating between models. *Journal of the Royal Statistical Society: Series B (Methodological), 32*(3), 323–345. https://doi.org/10.1111/j.2517-6161.1970.tb00845.x

Autorité de Contrôle Prudentiel et de Résolution. (2021). A first assessment of financial risks stemming from climate change: The main results of the 2020 climate pilot exercise. Banque de France.

Avril, P., Levieuge, G., & Turcu, C. (2023). *Do bankers want their umbrellas back when it rains? Evidence from typhoons in China* (INFER Working Paper No. 2023.08). International Network for Economic Research. https://ideas.repec.org/p/inf/wpaper/2023.08.html

Aznar-Siguan, G., & Bresch, D. N. (2019). CLIMADA v1: a global weather and climate risk assessment platform. Geoscientific Model Development, 12(7), 3085–3097. https://doi.org/10.5194/gmd-12-3085-2019

Azzone, M., Ghesini, M., Stocco, D., & Viola, L. (2025). Climate physical risk assessment in asset management. arXiv:2504.19307.

Bahrami, O., Farahani, S. , Muraleetharan, K. K., (2025). A high-resolution methodology for quantifying financial losses from natural hazards with application to pricing multi-peril catastrophe bonds. *International Journal of Disaster Risk Reduction*. https://www.sciencedirect.com/science/article/abs/pii/S2212420925003930

Baldauf, M., Garlappi, L., & Yannelis, C. (2020). Does climate change affect real estate prices? Only if you believe in it. *The Review of Financial Studies, 33*(3), 1256–1295. https://doi.org/10.1093/rfs/hhz073

Balkema, A. A., & de Haan, L. (1974). Residual life time at great age. *The Annals of Probability, 2*(5), 792–804.

Barbaglia, L., Fatica, S., & Rho, C. (2023). Flooded credit markets: Physical climate risk and small business lending (JRC Working Papers in Economics and Finance No. 2023-14). Joint Research Centre, European Commission.

Bardoscia, M., Battiston, S., Caccioli, F., & Caldarelli, G. (2015). DebtRank: A microscopic foundation for shock propagation. PLoS ONE, 10(6), e0130406. https://doi.org/10.1371/journal.pone.0130406

Barratt, J., Kath, J., Mushtaq, S., Collins, B., Chenu, K., Christopher, J., & An-Vo, D.-A. (2025). Strategic use of index-based frost insurance to reduce financial risk and improve income stability for wheat producers in Australia. Agricultural Systems, 226, 104306. https://doi.org/10.1016/j.agsy.2025.104306

Basel Committee on Banking Supervision. (2022). Principles for the effective management and supervision of climate-related financial risks. Bank for International Settlements.

Battaglioli, F., Taszarek, M., Groenemeijer, P., Púčik, T., & Rädler, A. T. (2026). Contrasting trends in very large hail events and related economic losses across the globe. Nature Geoscience, 19, 52–58. https://doi.org/10.1038/s41561-025-01868-0

Battiston, S., Farmer, J. D., Flache, A., Garlaschelli, D., Haldane, A. G., Heesterbeek, H., Hommes, C., Jaeger, C., May, R., & Scheffer, M. (2016). Complexity theory and financial regulation. *Science, 351*(6275), 818–819. https://doi.org/10.1126/science.aad0299

Battiston, S., Mandel, A., Monasterolo, I., Schütze, F., & Visentin, G. (2017). A climate stress-test of the financial system. *Nature Climate Change, 7*(4), 283–288. https://doi.org/10.1038/nclimate3255

Battiston, S., Puliga, M., Kaushik, R., Tasca, P., & Caldarelli, G. (2012). DebtRank: Too central to fail? Financial networks, the FED and systemic risk. Scientific Reports, 2, 541. https://doi.org/10.1038/srep00541

BCG. (2024). *How banks can transform physical climate risk into an opportunity* (A. Castoldi, G. Lucini,

B. Micale, A. Benayad, & M. Coppola). Boston Consulting Group. https://media-publications.bcg.com/How-Banks-Can-Transform-Physical-Climate-Risk-into-an-Opportunity.pdf

Benzécri, J.-P. (1973). L'analyse des données. Tome II: L'analyse des correspondances. Dunod.

Berger, L., & Marinacci, M. (2020). Model uncertainty in climate change economics: A review and proposed framework for future research. Environmental and Resource Economics, 77(3), 475–501.

Bernstein, A., Gustafson, M. T., & Lewis, R. (2019). Disaster on the horizon: The price effect of sea level rise. Journal of Financial Economics, 134(2), 253–272. https://doi.org/10.1016/j.jfineco.2019.03.013

Berry-Stölzle, T. R., & Hao, Y. (2026). Wildfire risk and municipal bond yields. *Journal of Risk and Insurance, 93*(1), 118–162. https://doi.org/10.1111/jori.70021

Bevacqua, E., Maraun, D., Vousdoukas, M. I., Voukouvalas, E., Vrac, M., Mentaschi, L., & Widmann, M. (2019). Higher probability of compound flooding from precipitation and storm surge in Europe under anthropogenic climate change. *Science Advances, 5*(9),eaaw5531. https://doi.org/10.1126/sciadv.aaw5531

Bhattacharya-Craven, A., Golnaraghi, M., Thomson, M., & Caplan, T. (2024). *Climate change: What does the future hold for health and life insurers?* The Geneva Association.

Blanc-Blocquel, A., Ortiz-Gracia, L., & Sanfelici, S. (2024). Climate-related default probabilities. *Risks, 12*(11), Article 181. https://doi.org/10.3390/risks12110181

Blickle, K. S., & Santos, J. A. C. (2022). Unintended consequences of “mandatory” flood insurance (Federal Reserve Bank of New York Staff Reports No. 1012). Federal Reserve Bank of New York.

Bloemendaal, N., Haigh, I. D., de Moel, H., Muis, S., Haarsma, R. J., & Aerts, J. C. J. H. (2020). Generation of a global synthetic tropical cyclone hazard dataset using STORM. *Scientific Data, 7*, 40.

Bloomfield, R., Nelson, M. W., & Soltes, E. (2016). Gathering data for archival, field, survey, and experimental accounting research. *Journal of Accounting Research, 54*(2), 341–395. https://doi.org/10.1111/1475-679X.12104

Boehmer, E., & Kelley, E. K. (2009). Institutional investors and the informational efficiency of prices. *The Review of Financial Studies, 22*(9), 3563–3594.

Bonnet, R., Boé, J., & Habets, F. (2020). Influence of multidecadal variability on high and low flows: The case of the Seine basin. Hydrology and Earth System Sciences, 24(4), 1611–1631. https://doi.org/10.5194/hess-24-1611-2020

Bordeman, A., & Demerjian, P. (2022). Do borrowers intentionally avoid covenant violations? A reexamination of the debt covenant hypothesis. Journal of Accounting Research, 60(5), 1741–1774.

Born, P., Dumm, R., & Johnson, M. E. (2023). Epistemic uncertainty in catastrophe models: A base level examination. Risk Management and Insurance Review, 26(2), 247–269. https://doi.org/10.1111/rmir.12246

Borowiak, D. S. (1989). Model discrimination for nonlinear regression models (Statistics: Textbooks and Monographs, Vol. 101). Marcel Dekker.

Boumezoued, A., Elfassihi, A., Germain, V., & Titon, E.-E. (2022). Modelling the impact of climate risks on mortality: Detailed case studies on exposure to high temperatures, alternative modelling for vector-borne diseases [White paper]. Milliman.

Bourget, M., Boudreault, M., Carozza, D. A., Boudreault, J., & Raymond, S. (2024). A data science approach to climate change risk assessment applied to pluvial flood occurrences for the United States and Canada. *ASTIN Bulletin, 54*(3), 495–517. https://doi.org/10.1017/asb.2024.19

Bove, M. C., O'Brien, J. J., Elsner, J. B., Landsea, C. W., & Niu, X. (1998). Effect of El Niño on U.S. landfalling hurricanes, revisited. Bulletin of the American Meteorological Society, 79(11), 2477–2482. https://doi.org/10.1175/1520-0477(1998)079<2477:EOENOU>2.0.CO;2

Bremond, P., Tarrit, R., Grelot, F., & Richert, C. (2024). Dommages des inondations aux activités économiques (technical report). Ministère chargé du développement durable; Inrae; Cerema. https://hal.inrae.fr/hal-04665053

Bresch, D. N., & Aznar-Siguan, G. (2021). CLIMADA v1.4.1: Towards a globally consistent adaptation options appraisal tool. *Geoscientific Model Development, 14*(1), 351–363. https://doi.org/10.5194/gmd-14-351-2021

Bressan, G., Đuranović, A., Monasterolo, I., & Battiston, S. (2024). Asset-level assessment of climate physical risk matters for adaptation finance. Nature Communications, 15, Article 5371. https://doi.org/10.1038/s41467-024-48820-1

Brock, W. A., & Hansen, L. P. (2018). Wrestling with uncertainty in climate economic models (Working Paper No. 2019-71). University of Chicago, Becker Friedman Institute for Economics. https://ssrn.com/abstract=3008833

Brownlees, C., & Engle, R. F. (2017). SRISK: A conditional capital shortfall measure of systemic risk. The Review of Financial Studies, 30(1), 48–79. https://doi.org/10.1093/rfs/hhw060

Brunnermeier, M. K., & Pedersen, L. H. (2009). Market liquidity and funding liquidity. The Review of Financial Studies, 22(6), 2201–2238. https://doi.org/10.1093/rfs/hhn098

Burke, M., Hsiang, S. M., & Miguel, E. (2015). Global non-linear effect of temperature on economic production. Nature, 527(7577), 235-239.

Burnecki, K., Kukla, G., & Weron, R. (2000). Property insurance loss distributions. Physica A: Statistical Mechanics and its Applications, 287(1 to 2), 269 to 278. https://doi.org/10.1016/S0378-4371(00)00453-2

Busby, J. W., Baker, K., Bazilian, M. D., Gilbert, A. Q., Grubert, E., Rai, V., Rhodes, J. D., Shidore, S., Smith, C. A., & Webber, M. E. (2021). Cascading risks: Understanding the 2021 winter blackout in Texas. Energy Research & Social Science, 77, 102106. https://doi.org/10.1016/j.erss.2021.102106

Calabrese, R, Dombrowski, T, Mandel, A, Pace, RK & Zanin, L 2024, 'Impacts of extreme weather events on mortgage risks and their evolution under climate change: A case study on Florida', European Journal of Operational Research, vol. 314, no. 1, pp. 377-392. https://doi.org/10.1016/j.ejor.2023.11.022

Calmel, B., Vidal, J.-P., Héraut, L., Sauquet, E., & Torremocha, A. (2025). Définition de narratifs hydrologiques par niveau de réchauffement défini par la TRACC (narraTRACC). https://doi.org/10.57745/FE356V

Caloia, F. G., & Jansen, D.-J. (2021). Flood risk and financial stability: Evidence from a stress test for the Netherlands (Working Paper No. 730). De Nederlandsche Bank.

Cammerer, H., Thieken, A. H., & Lammel, J. (2013). Adaptability and transferability of flood loss functions in residential areas. Natural Hazards and Earth System Sciences, 13(11), 3063–3081. https://doi.org/10.5194/nhess-13-3063-2013

Cannon, A. J., Sobie, S. R., & Murdock, T. Q. (2015). Bias correction of GCM precipitation by quantile mapping: How well do methods preserve changes in quantiles and extremes? Journal of Climate, 28(17), 6938–6959. https://doi.org/10.1175/JCLI-D-14-00754.1

Cardona, O. D., Ordaz, M. G., Yamín, L. E., Barbat, A. H., et al. (2013). Probabilistic modelling of natural risks at the global level: Global risk model. Background paper for the Global Assessment Report on Disaster Risk Reduction 2013. UNISDR and CIMNE.

Castoldi, A., Lucini, G., Micale, B., Benayad, A., & Coppola, M. (2024). *How banks can transform physical climate risk into an opportunity*. Boston Consulting Group. https://media-publications.bcg.com/How-Banks-Can-Transform-Physical-Climate-Risk-into-an-Opportunity.pdf

Cavallo, E., Galiani, S., Noy, I., & Pantano, J. (2013). Catastrophic natural disasters and economic growth. *Review of Economics and Statistics, 95*(5), 1549–1561. https://doi.org/10.1162/REST_a_00413

Centers for Disease Control and Prevention/Agency for Toxic Substances and Disease Registry. (2024).

CDC/ATSDR Social Vulnerability Index 2022 database, US [Data set]. https://www.atsdr.cdc.gov/place-health/php/svi/index.html

Chatelain, P. (2023). *Tarification à l'adresse en assurance habitation individuelle* [Doctoral dissertation, Université Claude Bernard Lyon 1].

Chen, Y.-C., Hung, M., & Wang, Y. (2018). The effect of mandatory CSR disclosure on firm profitability and social externalities: Evidence from China. Journal of Accounting and Economics, 65(1), 169–190. https://doi.org/10.1016/j.jacceco.2017.11.009

Chen, Y., Ding, Z., Barbaglia, L., Calabrese, R., & Fatica, S. (2026). A climate stress testing exercise on loans to European small and medium enterprises. Energy Economics, 155, 109177. https://doi.org/10.1016/j.eneco.2026.109177

Christiansen, K., Al Assi, A., Mostafiz, R. B., & Rohli, R. V. (2024). Charting the scorched trails: A comparative analysis of roadway damage from historical megafires to the unprecedented 2020 Labor Day wildfires on the USA West Coast. *Emergency Management Science and Technology, 4*(1), 1–10. https://doi.org/10.48130/emst-0024-0016

Christiansen, K., Mostafiz, R. B., Al Assi, A., Rohli, R. V., & Friedland, C. J. (2024). Financial impacts of 2020 Labor Day wildfires to Oregon highways. *PLOS Climate, 3*(10), Article e0000489. https://doi.org/10.1371/journal.pclm.0000489

CISL. (2019). *Physical risk framework: Understanding the impact of climate change on real estate lending and investment portfolios*. Cambridge Institute for Sustainability Leadership. https://www.cisl.cam.ac.uk/resources/sustainable-finance-publications/physical-risk-framework-understanding-the-impact-of-climate-change-on-real-estate-lending-and-investment-portfolios

Coppola, E., Nogherotto, R., Ciarlò, J. M., Giorgi, F., van Meijgaard, E., Kadygrov, N., … Wulfmeyer, V. (2021). Assessment of the European climate projections as simulated by the large EURO-CORDEX regional and global climate model ensemble. Journal of Geophysical Research: Atmospheres, 126(4), e2019JD032356. https://doi.org/10.1029/2019JD032356

Correa, R., He, A., Herpfer, C., & Lel, U. (2022). The rising tide lifts some interest rates: Climate change, natural disasters, and loan pricing (International Finance Discussion Papers No. 1345). Board of Governors of the Federal Reserve System. https://doi.org/10.17016/IFDP.2022.1345

Cottrell, M., & Letrémy, P. (2005). How to use the Kohonen algorithm to simultaneously analyze individuals and modalities in a survey. Neurocomputing, 63, 193–207. https://doi.org/10.1016/j.neucom.2004.04.011

Cox, D. R. (1961). Tests of separate families of hypotheses. In J. Neyman (Ed.), Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability (Vol. 1, pp. 105–123). University of California Press.

Cox, D. R. (1962). Further results on tests of separate families of hypotheses. Journal of the Royal Statistical Society: Series B, 24(2), 406–424.

Czajkowski, J., Kunreuther, H., & Michel-Kerjan, E. (2011). A methodological approach for pricing flood insurance and evaluating loss reduction measures: Application to Texas. Risk Management and Decision Processes Center, The Wharton School, University of Pennsylvania.

Damodaran, A. (2026). Operating and net margins by sector (US) [Data set]. NYU Stern School of Business. https://pages.stern.nyu.edu/~adamodar/New_Home_Page/datafile/margin.html

De Lavenne, A., Thirel, G., Andréassian, V., Perrin, C., & Ramos, M.-H. (2016). Spatial variability of the parameters of a semi-distributed hydrological model. Proceedings of the International Association of Hydrological Sciences, 373, 87.

De Maximy, F., Pouderoux, V., & Le Guenedal, T. (2025). *Modeling direct and indirect climate-related physical risks* [Conference paper]. GRASFI 2025 Annual Conference. https://doi.org/10.2139/ssrn.4792467

De Paula Salgado, I., Guenther, E., & Mechtcherine, V. (2025). Integrated sustainability and resilience assessments of concrete infrastructures subjected to hazards: A systematic literature review. Sustainable and Resilient Infrastructure, 10(5), 450–

471.
https://doi.org/10.1080/23789689.2025.2471119

Decharme, B., Delire, C., Minvielle, M., Colin, J., Vergnes, J.-P., Alias, A., … Voldoire, A. (2019). Recent changes in the ISBA-CTRIP land surface system for use in the CNRM-CM6 climate model and in global off-line hydrological applications. Journal of Advances in Modeling Earth Systems. https://doi.org/10.1029/2018MS001545

DeRosa, S. E., Kimura, Y., Stadtherr, M. A., McGaughey, G., McDonald-Buller, E., & Allen, D. T. (2019). Network modeling of the U.S. petrochemical industry under raw material and Hurricane Harvey disruptions. *Industrial & Engineering Chemistry Research, 58*(28), 12801–12815. https://doi.org/10.1021/acs.iecr.9b01035

Desaivres, S. (2026). Disentangling epistemic uncertainty in fluvial flood risk assessment for financial decision-making. Working paper, Centre d'économie de la Sorbonne, Université Panthéon-Sorbonne (Paris 1).

Deser, C., Phillips, A., Bourdette, V., & Teng, H. (2012). Uncertainty in climate change projections: The role of internal variability. Climate Dynamics, 38(3), 527–546.

Dey, A. K., & Das, K. P. (2016). Modeling extreme hurricane damage using the generalized Pareto distribution. American Journal of Mathematical and Management Sciences, 35(1), 55-66. https://doi.org/10.1080/01966324.2015.1075926

Di Lascio, F. M. L., Noy, I., & Perazzini, S. (2025). Modeling spatial correlation between earthquake insured losses in New Zealand: A mixed-effects analysis. Risk Analysis, 45(2), 387–398. https://doi.org/10.1111/risa.16638

Dice, J., Hossain, M., & Rodziewicz, D. (2024). *Flood risk exposures and mortgage-backed security asset performance and risk sharing* (Research Working Paper No. 24-05). Federal Reserve Bank of Kansas City. https://doi.org/10.18651/RWP2024-05

Dietzenbacher, E. (1997). In vindication of the Ghosh model: A reinterpretation as a price model. Journal of Regional Science, 37(4), 629–651. https://doi.org/10.1111/0022-4146.00073

Döhrmann, D., Gürtler, M., & Hibbeln, M. (2017). Insured loss inflation: How natural catastrophes affect reconstruction costs. *Journal of Risk and Insurance, 84*(3), 851–879. https://doi.org/10.1111/jori.12134

Dottori, F., Figueiredo, R., Martina, M. L. V., Molinari, D., & Scorzini, A. R. (2016). INSYDE: A synthetic, probabilistic flood damage model based on explicit cost analysis. Natural Hazards and Earth System Sciences, 16(12), 2577–2591. https://doi.org/10.5194/nhess-16-2577-2016

Dottori, F., Salamon, P., Bianchi, A., Alfieri, L., Hirpa, F. A., & Feyen, L. (2016). Development and evaluation of a framework for global flood hazard mapping. *Advances in Water Resources, 94*, 87–102. https://doi.org/10.1016/j.advwatres.2016.05.002

Dotzek, N., Groenemeijer, P., Feuerstein, B., & Holzer, A. M. (2009). Overview of ESSL's severe convective storms research using the European Severe Weather Database ESWD. *Atmospheric Research, 93*(1–3), 575–586. https://doi.org/10.1016/j.atmosres.2008.10.020

Douglas-Smith, D., Iwanaga, T., Croke, B., & Jakeman, A. (2019). Certain trends in uncertainty and sensitivity analysis: An overview of software tools and techniques. Environmental Modelling and Software, 124. https://doi.org/10.1016/j.envsoft.2019.104588

Dowrick, D. J., & Rhoades, D. A. (1995). Damage ratios for plant, equipment and stock in the 1987 Edgecumbe, New Zealand earthquake. Bulletin of the New Zealand Society for Earthquake Engineering, 28(4), 265–278. doi:10.5459/bnzsee.28.4.265-278.

Drusch, M., Del Bello, U., Carlier, S., Colin, O., Fernandez, V., Gascon, F., Hoersch, B., Isola, C., Laberinti, P., Martimort, P., Meygret, A., Spoto, F., Sy, O., Marchese, F., & Bargellini, P. (2012). Sentinel-2: ESA's optical high-resolution mission for GMES operational services. *Remote Sensing of Environment, 120*, 25–36. https://doi.org/10.1016/j.rse.2011.11.026

Du Jardin, P., & Séverin, E. (2011). Predicting corporate bankruptcy using a self-organizing map: An empirical study to improve the forecasting horizon of a financial failure model. Decision Support Systems, 51(3), 701–711.

Dupuis, D. J., & Jones, B. L. (2006). Multivariate extreme value theory and its usefulness in understanding risk. *North American Actuarial Journal, 10*(4), 1–27. https://doi.org/10.1080/10920277.2006.10597411

Dutfoy, A., Parey, S., & Roche, N. (2014). Multivariate extreme value theory: A tutorial with applications to hydrology and meteorology. *Dependence Modeling*, *2*(1), 1–19. https://doi.org/10.2478/demo-2014-0003

Eberenz, S., Stocker, D., Röösli, T., & Bresch, D. N. (2020). Asset exposure data for global physical risk assessment. *Earth System Science Data, 12*(2), 817–833. https://doi.org/10.5194/essd-12-817-2020

ECB/ESRB Project Team on Climate Risk Monitoring. (2021). Climate-related risk and financial stability. European Central Bank/European Systemic Risk Board.

Economou, T., Stephenson, D. B., Pinto, J. G., Shaffrey, L. C., & Zappa, G. (2015). Serial clustering of extratropical cyclones in a multi-model ensemble of historical and future simulations. Quarterly Journal of the Royal Meteorological Society, 141(693), 3076–3087. https://doi.org/10.1002/qj.2591

Edgeley, C. M., Colavito, M. M., vonHedemann, N., & Burnett, J. T. (2024). Voluntary uptake of flood insurance in post-fire landscapes: Insights from household surveys after two wildfires in northern Arizona. *Environmental Research Communications, 6*(6), Article 061006. https://doi.org/10.1088/2515-7620/ad5647

EDHEC Climate Institute. (2025, June 20). *EDHEC launches Scientific Climate Ratings to quantify the financial impact of climate risk*. https://climateinstitute.edhec.edu/news/edhec-launches-scientific-climate-ratings-quantify-financial-impact-climate-risk

Efron, B. (1979). Bootstrap methods: Another look at the jackknife. The Annals of Statistics, 7(1), 1–26. https://doi.org/10.1214/aos/1176344552

Efron, B., & Tibshirani, R. J. (1993). An introduction to the bootstrap. Chapman & Hall.

Eid, M. S., & El-Adaway, I. H. (2015). Optimizing disaster recovery strategies using agent-based simulation. In *Computing in Civil Engineering 2015: Proceedings of the 2015 ASCE International Workshop on Computing in Civil Engineering* (pp. 379–386). American Society of Civil Engineers. https://doi.org/10.1061/9780784479247.047

EIOPA & ECB. (2023, April 24). *Policy options to reduce the climate insurance protection gap* [Staff paper]. European Insurance and Occupational Pensions Authority & European Central Bank. https://www.eiopa.europa.eu/publications/staff-paper-policy-options-reduce-climate-insurance-protection-gap_en

EIOPA. (2021). *Opinion on the supervision of the use of climate change risk scenarios in ORSA* (EIOPA-BoS-21-127). European Insurance and Occupational Pensions Authority.

Eleutério, J. (2012). Flood risk analysis: Impact of uncertainty in hazard modelling and vulnerability assessments on damage estimations [Doctoral thesis, Université de Strasbourg]. https://theses.hal.science/tel-00821011

Emanuel, K. (2005). Increasing destructiveness of tropical cyclones over the past 30 years. Nature, 436, 686-688.

Emanuel, K. (2011). Global warming effects on U.S. hurricane damage. Weather, Climate, and Society, 3(4), 261–268. https://doi.org/10.1175/WCAS-D-11-00007.1

Embrechts, P., Kluppelberg, C., & Mikosch, T. (2013). Modelling extremal events for insurance and finance. Springer.

Embrechts, P., McNeil, A., & Straumann, D. (2002). Correlation and dependence in risk management: Properties and pitfalls. In M. Dempster (Ed.), Risk management: Value at risk and beyond (pp. 176–223). Cambridge University Press.

Epelbaum, Thomas, ClimVaR: Discounted Cash Flow Approach on Climate Risks (January 09, 2026). Available at SSRN: https://ssrn.com/abstract=6047297 or http://dx.doi.org/10.2139/ssrn.6047297

European Commission Joint Research Centre. (2026). *INFORM Risk Index mid 2026* [Data set]. Disaster Risk Management Knowledge Centre. https://drmkc.jrc.ec.europa.eu/inform-index/INFORM-Risk/Results-and-data

European Commission. (2021). 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. *Official Journal of the European Union, L 442*, 1–349. http://data.europa.eu/eli/reg_del/2021/2139/oj

European Committee for Standardization. (2005). EN 1991-1-4: Eurocode 1: Actions on structures – Part 1-4: General actions – Wind actions. CEN.

European Parliament & Council of the European Union. (2009). Directive 2009/138/EC of 25 November 2009 on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II). *Official Journal of the European Union, L 335*, 1–155.

Evin, G. (2026). QUALYPSO: Partitioning uncertainty components of an incomplete ensemble of climate projections (R package version 3.1). https://doi.org/10.32614/CRAN.package.QUALYPSO

Evin, G., Hingray, B., Blanchet, J., Eckert, N., Morin, S., & Verfaillie, D. (2019). Partitioning uncertainty components of an incomplete ensemble of climate projections using data augmentation. Journal of Climate, 32(8), 2423–2440. https://doi.org/10.1175/JCLI-D-18-0606.1

Evin, G., Hingray, B., Reverdy, A., Ducharne, A., & Sauquet, E. (2024). Ensemble de projections Explore2: Changements moyens et incertitudes associées. https://doi.org/10.57745/KWH320

Evin, G., Hingray, B., Thirel, G., Ducharne, A., Strohmenger, L., Corre, L., … Sauquet, E. (2026). Uncertainty sources in a large ensemble of hydrological projections: Regional climate models and internal variability matter. Hydrology and Earth System Sciences, 30(4), 1023–1051. https://doi.org/10.5194/hess-30-1023-2026

Evin, G., Somot, S., & Hingray, B. (2021). Balanced estimate and uncertainty assessment of European climate change using the large EURO-CORDEX regional climate model ensemble. Earth System Dynamics, 12(4), 1543–1569. https://doi.org/10.5194/esd-12-1543-2021

Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., & Taylor, K. E. (2016). Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization. *Geoscientific Model Development, 9*(5), 1937–1958. https://doi.org/10.5194/gmd-9-1937-2016

Faigle, N. (2026, February 9). *The macroeconomic side of physical risk*. MSCI. https://www.msci.com/research-and-insights/blog-post/the-macroeconomic-side-of-physical-risk

Federal Emergency Management Agency (FEMA). (2018). *Seismic performance assessment of buildings, volume 1 – Methodology* (FEMA P-58-1, 2nd ed.). Applied Technology Council.

Federal Emergency Management Agency. (2022, January 18). National Flood Insurance Program Risk Rating 2.0: Methodology and data sources. U.S. Department of Homeland Security. https://www.fema.gov/sites/default/files/documents/FEMA_Risk-Rating-2.0_Methodology-and-Data-Appendix__01-22.pdf

FEMA. (2022). *Hazus flood model technical manual* (Hazus 5.1). Federal Emergency Management Agency.

FEMA. (2024). Hazus flood model technical manual 6.1. FEMA.

Financial Stability Board. (2025). Assessment of climate-related vulnerabilities: Analytical framework and toolkit. https://www.fsb.org/2025/01/assessment-of-climate-related-vulnerabilities-analytical-framework-and-toolkit/

Flanagan, B. E., Gregory, E. W., Hallisey, E. J., Heitgerd, J. L., & Lewis, B. (2011). A social

vulnerability index for disaster management. *Journal of Homeland Security and Emergency Management, 8*(1), Article 3. https://doi.org/10.2202/1547-7355.1792

Froot, K. A. (2001). The market for catastrophe risk: A clinical examination. *Journal of Financial Economics, 60*(2–3), 529–571. https://doi.org/10.1016/S0304-405X(01)00052-6

FSB & NGFS. (2022). *Climate scenario analysis by jurisdictions: Initial findings and lessons*. Financial Stability Board and Network for Greening the Financial System. https://www.fsb.org/2022/11/climate-scenario-analysis-by-jurisdictions-initial-findings-and-lessons/

Fuerst, F., & Warren-Myers, G. (2021). Pricing climate risk: Are flooding and sea level rise risk capitalised in Australian residential property? Climate Risk Management, 34, Article 100361. https://doi.org/10.1016/j.crm.2021.100361

Gai, P., & Kapadia, S. (2010). Contagion in financial networks. *Proceedings of the Royal Society A, 466*(2120), 2401–2423. https://doi.org/10.1098/rspa.2009.0410

Gao, L., & Shi, P. (2025). Risk modeling of property insurance claims from weather events. ASTIN Bulletin: The Journal of the IAA, 55(2), 242 to 262. https://doi.org/10.1017/asb.2025.7

Garcia-Jorcano, L., & Sanchis-Marco, L. (2024). Forecasting the effect of extreme sea-level rise on financial market risk. *International Review of Economics & Finance*, *93*(Part B), 1–27. https://doi.org/10.1016/j.iref.2024.03.079

Garnier, J., Gaudemet, J.-B., & Gruz, A. The Climate Extended Risk Model (CERM). Working paper. https://josselin-garnier.org/wp-content/uploads/2022/04/cerm.pdf

GARP. (2025). *A risk professional's guide to physical risk assessments: A benchmarking study of climate risk vendors* (J. Paisley & M. Nelson). Global Association of Risk Professionals. https://www.garp.org/white-paper/risk-professionals-guide-251023

Gaupp, F., Hall, J., Hochrainer-Stigler, S., & Dadson, S. (2020). Changing risks of simultaneous global breadbasket failure. Nature Climate Change, 10(1), 54–57. https://doi.org/10.1038/s41558-019-0600-z

Gesualdo, G. C., Benso, M. R., Sass, K. S., & Mendiondo, E. M. (2024). Index-based insurance to mitigate current and future extreme events financial losses for water utilities. International Journal of Disaster Risk Reduction, 100, 104218. https://doi.org/10.1016/j.ijdrr.2023.104218

Giglio, L., Boschetti, L., Roy, D. P., Humber, M. L., & Justice, C. O. (2018). The Collection 6 MODIS burned area mapping algorithm and product. *Remote Sensing of Environment, 217*, 72–85. https://doi.org/10.1016/j.rse.2018.08.005

Global Commission on Adaptation. (2019). Adapt now: A global call for leadership on climate resilience. World Resources Institute. https://gca.org/reports/adapt-now-a-global-call-for-leadership-on-climate-resilience/

Golnaraghi, M., & Vichev, Z. (2025). *Safeguarding home insurance: Reducing exposure and vulnerability to extreme weather*. The Geneva Association. https://www.genevaassociation.org/sites/default/files/2025-06/safeguarding_home_insurance_040625.pdf

Gould, I. J., Wright, I., Collison, M., et al. (2020). The impact of coastal flooding on agriculture: A case-study of Lincolnshire, United Kingdom. *Land Degradation & Development, 31*(12), 1545–1559. https://doi.org/10.1002/ldr.3551

Greenwald, D. L. (2019). Firm debt covenants and the macroeconomy: The interest coverage channel [Working paper]. MIT Sloan School of Management.

Grossi, P., & Kunreuther, H. (Eds.). (2005). Catastrophe modeling: A new approach to managing risk. Springer.

Guzmán, D. A., Mohor, G. S., & Mendiondo, E. M. (2023). Multi-driver ensemble to evaluate the water utility business interruption cost induced by hydrological drought. *Urban Water Journal, 20*(10), 1517–1531. https://doi.org/10.1080/1573062X.2022.2058564

Hain, L. I., Kölbel, J. F., & Leippold, M. (2022). Let's get physical: Comparing metrics of physical climate risk. Finance Research Letters, 46, 102406. https://doi.org/10.1016/j.frl.2021.102406

Haldane, A. G., & May, R. M. (2011). Systemic risk in banking ecosystems. *Nature, 469*(7330), 351–355. https://doi.org/10.1038/nature09659

Hansen, J., Lacis, A., Ruedy, R., & Sato, M. (1992). Potential climate impact of Mount Pinatubo eruption. *Geophysical Research Letters, 19*(2), 215–218. https://doi.org/10.1029/91GL02788

Haraguchi, M., & Lall, U. (2015). Flood risks and impacts: A case study of Thailand's floods in 2011 and research questions for supply chain decision making. International Journal of Disaster Risk Reduction, 14, 256-272 https://doi.org/10.1016/j.ijdrr.2014.09.005

Hawkins, E., & Sutton, R. (2009). The potential to narrow uncertainty in regional climate predictions. Bulletin of the American Meteorological Society, 90(8), 1095–1108.

Hawkins, E., & Sutton, R. (2011). The potential to narrow uncertainty in projections of regional precipitation change. Climate Dynamics, 37(1), 407–418.

Heranval, A., Lopez, O., Ngatcha, D., & Nkameni, D. (2026). A Wasserstein GAN-based climate scenario generator for risk management and insurance: The case of soil subsidence (arXiv:2605.06678). arXiv. https://doi.org/10.48550/arXiv.2605.06678

Héraut, L., Sauquet, E., Calmel, B., & Vidal, J.-P. (2025). Sélections de narratifs hydrologiques (narraTRACC) par niveau de réchauffement sur chaque secteur hydrographique. https://doi.org/10.57745/KAHIWJ

Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., . . . Thépaut, J.-N. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society, 146*(730), 1999–2049. https://doi.org/10.1002/qj.3803

Herschy, R. W. (1998). Hydrometry: Principles and practice. John Wiley & Sons.

Hervouet, J.-M. (2007). *Hydrodynamics of free surface flows: Modelling with the finite element method*. Wiley. https://doi.org/10.1002/9780470319628

Hillier, J. K., Bloomfield, H. C., Manning, C., & Garry, F. (2025). Increasingly seasonal jet stream raises risk of co-occurring flooding and extreme wind in Great Britain. *International Journal of Climatology*.

Higuera Roa, O., Bachmann, M., Mechler, R., Šakić Trogrlić, R., Reimann, L., Mazzoleni, M., Aerts, J. C. J. H., Buskop, F. E., Pirani, A., & Mysiak, J. (2025). Challenges and opportunities in climate risk assessment: Future directions for assessing complex climate risks. *Environmental Research Letters, 20*(5), 053003. https://doi.org/10.1088/1748-9326/adc756

Hillier, J. K., Macdonald, N., Leckebusch, G. C., & Stavrinides, A. (2015). Interactions between apparently 'primary' weather-driven hazards and their cost. *Environmental Research Letters, 10*(10), Article 104003. https://doi.org/10.1088/1748-9326/10/10/104003

Hochrainer-Stigler, S., Balkovic, J., Silm, K., & Timonina-Farkas, A. (2019). Large scale extreme risk assessment using copulas: An application to drought events under climate change for Austria. Computational Management Science, 16(4), 651 to 669. https://doi.org/10.1007/s10287-018-0339-4

Hodges, K., Cobb, A., & Vidale, P. L. (2017). How well are tropical cyclones represented in reanalysis datasets? *Journal of Climate, 30*(14), 5243–5264. https://doi.org/10.1175/JCLI-D-16-0557.1

Holgate, S. J., Matthews, A., Woodworth, P. L., Rickards, L. J., Tamisiea, M. E., Bradshaw, E., Foden, P. R., Gordon, K. M., Jevrejeva, S., & Pugh, J. (2013). New data systems and products at the Permanent Service for Mean Sea Level. *Journal of Coastal Research, 29*(3), 493–504. https://doi.org/10.2112/JCOASTRES-D-12-00175.1

Hong, H., Li, F. W., & Xu, J. (2019). Climate risks and market efficiency. *Journal of Econometrics, 208*(1), 265–281. https://doi.org/10.1016/j.jeconom.2018.09.015

Hsiang, S. M., & Jina, A. S. (2014). *The causal effect of environmental catastrophe on long-run economic growth: Evidence from 6,700 cyclones* (NBER Working Paper No. 20352). National Bureau of Economic Research. https://doi.org/10.3386/w20352

Huizinga, J. (2007). Flood damage functions for EU member states (technical report).

Huizinga, J., de Moel, H., & Szewczyk, W. (2017). *Global flood depth-damage functions: Methodology and the database with guidelines* (EUR 28552 EN). Publications Office of the European Union. https://doi.org/10.2760/16510

Humphreys, N. (2021). Exceedance probability in catastrophe modeling. Casualty Actuarial Society E-Forum, Winter 2021, 1–61. https://www.casact.org/sites/default/files/2021-03/02_humphreys.pdf

Huynh, T. D., Nguyen, T. H., & Truong, C. (2020). Climate risk: The price of drought. *Journal of Corporate Finance, 65*, 101750. https://doi.org/10.1016/j.jcorpfin.2020.101750

IASB. (2024). IFRS 18 Presentation and disclosure in financial statements. IFRS Foundation.

IFRS Foundation. (2023). *Effects of climate-related matters on financial statements* [Educational material]. https://www.ifrs.org/content/dam/ifrs/supporting-implementation/documents/effects-of-climate-related-matters-on-financial-statements.pdf

Integrated Research on Disaster Risk. (2014). *Peril classification and hazard glossary* (IRDR DATA Publication No. 1). IRDR. https://www.preventionweb.net/publications/view/36979

International Monetary Fund (IMF), Bank for International Settlements (BIS), & Financial Stability Board (FSB). (2009). *Guidance to assess the systemic importance of financial institutions, markets and instruments: Initial considerations* (Report to the G-20 Finance Ministers and Central Bank Governors). https://www.bis.org/publ/othp07.htm

International Organization for Standardization. (2021). *Adaptation to climate change: Guidelines on vulnerability, impacts and risk assessment* (ISO Standard No. 14091:2021). https://www.iso.org/standard/68508.html

IPCC - Fox-Kemper, B., Hewitt, H. T., Xiao, C., Aðalgeirsdóttir, G., Drijfhout, S. S., Edwards, T. L., Golledge, N. R., Hemer, M., Kopp, R. E., Krinner, G., Mix, A., Notz, D., Nowicki, S., Nurhati, I. S., Ruiz, L., Sallée, J.-B., Slangen, A. B. A., & Yu, Y. (2021). Ocean, cryosphere and sea level change. In V. Masson-Delmotte et al. (Eds.), *Climate change 2021: The physical science basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (pp. 1211–1362). Cambridge University Press. https://doi.org/10.1017/9781009157896.011

IPCC - Seneviratne, S. I., Zhang, X., Adnan, M., Badi, W., Dereczynski, C., Di Luca, A., Ghosh, S., Iskandar, I., Kossin, J., Lewis, S., Otto, F., Pinto, I., Satoh, M., Vicente-Serrano, S. M., Wehner, M., & Zhou, B. (2021). Weather and climate extreme events in a changing climate. In V. Masson-Delmotte et al. (Eds.), Climate change 2021: The physical science basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change (pp. 1513–1766). Cambridge University Press. https://doi.org/10.1017/9781009157896.013

Jay-Allemand, M., Javelle, P., Gejadze, I., Arnaud, P., Malaterre, P.-O., Fine, J.-A., & Organde, D. (2020). On the potential of variational calibration for a fully distributed hydrological model: Application on a Mediterranean catchment. Hydrology and Earth System Sciences, 24, 5519.

Jewson, S., Maynard, T., & Dottori, F. (2023). A service to help insurers understand the financial impacts of changing flood risk in Europe, based on PESETA IV. *Climate Services, 30*, Article 100395. https://doi.org/10.1016/j.cliser.2023.100395

Joe, H. (1997). Multivariate models and dependence concepts. Chapman & Hall.

Johnson, C. W., Fu, Y., & Bürgmann, R. (2017). Seasonal water storage, stress modulation, and California seismicity. *Science, 356*(6343), 1161–1164. https://doi.org/10.1126/science.aak9547

Jongman, B., Hochrainer-Stigler, S., Feyen, L., Aerts, J. C. J. H., Mechler, R., Botzen, W. J. W., Bouwer, L. M.,

Pflug, G., Rojas, R., & Ward, P. J. (2014). Increasing stress on disaster-risk finance due to large floods. *Nature Climate Change, 4*(4), 264–268.

Jongman, B., Kreibich, H., Apel, H., Barredo, J., Bates, P., Feyen, L., ... Ward, P. (2012). Comparative flood damage model assessment: Towards a European approach. Natural Hazards and Earth System Sciences, 12, 3733–3752. https://doi.org/10.5194/nhess-12-3733-2012

Joseph, C., Wang, Q. J., Mushtaq, S., Li, Y., Barratt, J., & Barratt, T. (2026). Advancing weather index insurance for climate risk management: A review of modelling techniques and implementation strategies. Environmental Impact Assessment Review, 119, 108376. https://doi.org/10.1016/j.eiar.2026.108376

Kallel, R., Cottrell, M., & Vigneron, V. (2002). Bootstrap for neural model selection. Neurocomputing, 48(1–4), 175–183. https://doi.org/10.1016/S0925-2312(01)00650-6

Karimi, N., Shokrollahi, F. (2026). Pricing Excess-of-Loss Reinsurance and CAT Bonds under Climate Uncertainty: A Cox Process Framework with Temperature-Dependent Stochastic Intensity. arXiv:2606.14830

Kennedy, M. C., & O'Hagan, A. (2002). Bayesian calibration of computer models. Journal of the Royal Statistical Society Series B: Statistical Methodology, 63(3), 425–464. https://doi.org/10.1111/1467-9868.00294

Kerkhofs, R., Bernhofen, M., Borsuk, M., Baer, M., Ranger, N., Schoutens, W., & Shrimali, G. (2025). An asset-level analysis of financial tail risks under extreme weather events. *Environmental Research: Climate, 4*(2). https://doi.org/10.1088/2752-5295/addf6f

Keys, B. J., & Mulder, P. (2024). *Property insurance and disaster risk: New evidence from mortgage escrow data* (NBER Working Paper No. 32579). National Bureau of Economic Research. https://doi.org/10.3386/w32579

Kiureghian, A. D., & Ditlevsen, O. (2009). Aleatory or epistemic? Does it matter? Structural Safety, 31(2), 105–112. https://doi.org/10.1016/j.strusafe.2008.06.020

Kling, G., Volz, U., Murinde, V., & Ayas, S. (2021). The impact of climate vulnerability on firms' cost of capital and access to finance. *World Development, 137*, 105131. https://doi.org/10.1016/j.worlddev.2020.105131

Klinker, Frederick L., Glenn G. Meyers, and David A. Lalonde. 2003. "The Aggregation and Correlation of Reinsurance Exposure." CAS Forum 2003 (May).

Klotzbach, P. J., Bell, M. M., Bowen, S. G., Gibney, E. J., Knapp, K. R., & Schreck, C. J., III. (2020). Surface pressure a more skillful predictor of normalized hurricane damage than maximum sustained wind. Bulletin of the American Meteorological Society, 101(6), E830–E846. https://doi.org/10.1175/BAMS-D-19-0062.1

Klotzbach, P. J., Bowen, S. G., Pielke, R., Jr., & Bell, M. (2018). Continental U.S. hurricane landfall frequency and associated damage: Observations and future risks. Bulletin of the American Meteorological Society, 99(7), 1359–1376. https://doi.org/10.1175/BAMS-D-17-0184.1

Knapp, K. R., Kruk, M. C., Levinson, D. H., Diamond, H. J., & Neumann, C. J. (2010). The International Best Track Archive for Climate Stewardship (IBTrACS): Unifying tropical cyclone best track data. *Bulletin of the American Meteorological Society, 91*(3), 363–376. https://doi.org/10.1175/2009BAMS2755.1

Knight, F. (1965). Risk, uncertainty and profit. Harper & Row.

Knutson, T., Camargo, S. J., Chan, J. C. L., Emanuel, K., Ho, C.-H., Kossin, J., Mohapatra, M., Satoh, M., Sugi, M., Walsh, K., & Wu, L. (2020). Tropical cyclones and climate change assessment: Part II: Projected response to anthropogenic warming. Bulletin of the American Meteorological Society, 101(3), E303 to E322. https://doi.org/10.1175/BAMS-D-18-0194.1

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. Biological Cybernetics, 43(1), 59–69. https://doi.org/10.1007/BF00337288

Kopp, R. E., Horton, R. M., Little, C. M., Mitrovica, J. X., Oppenheimer, M., Rasmussen, D. J., Strauss, B. H., & Tebaldi, C. (2014). Probabilistic 21st and 22nd century sea-level projections at a global network of tide-gauge sites. *Earth's Future, 2*(8), 383–406. https://doi.org/10.1002/2014EF000239

Krinner, G., Viovy, N., de Noblet-Ducoudré, N., Ogée, J., Polcher, J., Friedlingstein, P., … Prentice, I. C. (2005). A dynamic global vegetation model for studies of the coupled atmosphere-biosphere system. Global Biogeochemical Cycles. https://doi.org/10.1029/2003GB002199

Kropf, C. M., Ciullo, A., Otth, L., Meiler, S., Rana, A., Schmid, E., McCaughey, J. W., & Bresch, D. N. (2022). Uncertainty and sensitivity analysis for probabilistic weather and climate-risk modelling: An implementation in CLIMADA v.3.1.0. *Geoscientific Model Development, 15*(18), 7177–7201. https://doi.org/10.5194/gmd-15-7177-2022

Krueger, P., Sautner, Z., & Starks, L. T. (2020). The importance of climate risks for institutional investors. *The Review of Financial Studies, 33*(3), 1067–1111.

Kruttli, M. S., Roth Tran, B., & Watugala, S. W. (2025). Pricing Poseidon: Extreme weather uncertainty and firm return dynamics. *The Journal of Finance, 80*(2), 783–832. https://doi.org/10.1111/jofi.13416

Kummu, M., Taka, M., & Guillaume, J. H. A. (2018). Gridded global datasets for Gross Domestic Product and Human Development Index over 1990–2015. *Scientific Data, 5*, 180004. https://doi.org/10.1038/sdata.2018.4

Laidlaw, J., Lord, R., MacFarland, M., & Singh, K. (2023, November 20). Quantifying the financial costs of climate change physical risks for companies. S&P Global Sustainable1. https://www.spglobal.com/sustainable1/en/insights/special-editorial/quantifying-the-financial-costs-of-climate-change-physical-risks

Law, N. T., & Miura, Y. (2025). *RoofNet: A global multimodal dataset for roof material classification* (arXiv:2505.19358). arXiv. https://arxiv.org/abs/2505.19358

Le Guenedal, T., & Tankov, P. (2025). Corporate debt value under transition scenario uncertainty. Mathematical Finance, 35(1), 40–73.

Le Guenedal, T., Drobinski, P., & Tankov, P. (2021). Measuring and pricing cyclone-related physical risk under changing climate. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3850673

Li, J., Cai, Z., Liu, Y., & Ling, C. (2023). Extremal analysis of flooding risk and its catastrophe bond pricing. *Mathematics, 11*(1), Article 114. https://doi.org/10.3390/math11010114

Li, X. (2025). Physical climate change exposure and firms' adaptation strategy. Strategic Management Journal, 46(3), 750–789. https://doi.org/10.1002/smj.3674

Ling, C., Li, J., Liu, Y., & Cai, Z. (2021). Extremal analysis of flooding risk and management. arXiv:2112.00562.

Liora. (2026, January 29). *L'IA et l'open data au service de la rénovation énergétique des bâtiments*. https://liora.io/lopen-data-au-service-de-la-renovation-energetique-des-batiments

MacKinlay, A. C. (1997). Event studies in economics and finance. *Journal of Economic Literature, 35*(1), 13–39.

Maes, M. J. A., Gonzales-Hishinuma, A., Haščič, I., Hoffmann, C., Banquet, A., Veneri, P., Bizeul, A., Risquez Martin, A., & Quadrelli, R. (2022). *Monitoring exposure to climate-related hazards: Indicator methodology and key results* (OECD Environment Working Papers No. 201). OECD Publishing. https://doi.org/10.1787/da074cb6-en

Magnier, C., et al. (2018). Analyse multicritère des projets de prévention des inondations: Guide méthodologique.

Mahalingam, A., Coburn, A., Jung, C. J., Yeo, J. Z., Cooper, G., & Evan, T. (2018). Impacts of severe natural catastrophes on financial markets. Cambridge Centre for Risk Studies, University of Cambridge Judge Business School.

Mandel, A., Battiston, S., & Monasterolo, I. (2025). Mapping global financial risks under climate change.

Nature Climate Change, 15(3), 329–334. https://doi.org/10.1038/s41558-025-02244-x

Mandel, A., Tiggeloven, T., Lincke, D., Koks, E., Ward, P., & Hinkel, J. (2021). Risks on global financial stability induced by climate change: The case of flood risks. *Climatic Change, 166*(1–2), 4. https://doi.org/10.1007/s10584-021-03092-2

Maneerat, P., Rungskunroch, P., & Persaud, P. (2024). Seismic hazard analysis and financial impact assessment of railway infrastructure in the US West Coast: A machine learning approach. *PLOS ONE, 19*(8), Article e0308255. https://doi.org/10.1371/journal.pone.0308255

Martius, O., Pfahl, S., & Chevalier, C. (2016). A global quantification of compound precipitation and wind extremes. *Geophysical Research Letters, 43*(14), 7709–7717. https://doi.org/10.1002/2016GL070017

Mazzocchi, M., Hansstein, F., & Ragona, M. (2010). The 2010 volcanic ash cloud and its financial impact on the European airline industry. CESifo Forum, 11(2), 92–100.

McGuire, B. (2010). Potential for a hazardous geospheric response to projected future climate changes. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 368*(1919), 2317–2345. https://doi.org/10.1098/rsta.2010.0080

McKee, T. B., Doesken, N. J., & Kleist, J. (1993). The relationship of drought frequency and duration to time scales. In Proceedings of the 8th Conference on Applied Climatology (pp. 179–184). American Meteorological Society.

McSweeney, C. F., Jones, R. G., Lee, R. W., & Rowell, D. P. (2015). Selecting CMIP5 GCMs for downscaling over multiple regions. Climate Dynamics, 44(3-4), 907–928. https://doi.org/10.1007/s00382-014-2418-8

Meng, S., Soleimani-Babakamali, M. H., & Taciroglu, E. (2023). *Automatic roof type classification through machine learning for regional wind risk assessment* (arXiv:2305.17315). arXiv. https://arxiv.org/abs/2305.17315

Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. The Journal of Finance, 29(2), 449–470. https://doi.org/10.1111/j.1540-6261.1974.tb03058.x

Merz, B., Kreibich, H., & Lall, U. (2013). Multi-variate flood damage assessment: A tree-based data-mining approach. *Natural Hazards and Earth System Sciences, 13*(1), 53–64. https://doi.org/10.5194/nhess-13-53-2013

Météo-France. (2020). Météo-France dans le dispositif CATNAT sécheresse. https://meteofrance.fr/sites/meteofrance.fr/files/files/editorial/Catnat03022023.pdf

Meyer, V., Becker, N., Markantonis, V., Schwarze, R., van den Bergh, J. C. J. M., Bouwer, L. M., … Viavattene, C. (2013). Review article: Assessing the costs of natural hazards, state of the art and knowledge gaps. Natural Hazards and Earth System Sciences, 13(5), 1351–1373. https://doi.org/10.5194/nhess-13-1351-2013

Mitchell-Wallace, K., Jones, M., Hillier, J., & Foote, M. (2017). *Natural catastrophe risk management and modelling: A practitioner's guide*. Wiley-Blackwell. https://doi.org/10.1002/9781118906057

Moftakhari, H. R., Salvadori, G., AghaKouchak, A., Sanders, B. F., & Matthew, R. A. (2017). Compounding effects of sea level rise and fluvial flooding. Proceedings of the National Academy of Sciences, 114(37), 9785 to 9790. https://doi.org/10.1073/pnas.1620325114

Moncoulon, D., Veysseire, M., Naulin, J.-P., Wang, Z., Tinard, P., Desarthe, J., Hajji, C., Onfroy, T., Regimbeau, F., & Déqué, M. (2016). Modelling the evolution of the financial impacts of flood and storm surge between 2015 and 2050 in France. International Journal of Safety and Security Engineering, 6(2), 141–149. https://doi.org/10.2495/SAFE-V6-N2-141-149

Montgomery, M., & Kunreuther, H. (2018). Pricing storm surge risks in Florida: Implications for determining flood insurance premiums and evaluating mitigation measures. Risk Analysis, 38(11), 2275–2299. https://doi.org/10.1111/risa.13127

Moody's KMV. Crosbie, P., & Bohn, J. (2003). *Modeling default risk*.

MSCI Inc. (2025, April 22). *Quarterly update 1Q25*. https://ir.msci.com/static-files/e16bd774-2386-42fa-bea8-ab65ce9d40af

Mühlhofer, E., Kropf, C. M., Riedel, L., Bresch, D. N., & Koks, E. E. (2024). OpenStreetMap for multi-faceted climate risk assessments. Environmental Research Communications, 6(1), 015005. https://doi.org/10.1088/2515-7620/ad15ab

Mukerji, S. (2009). Foundations of ambiguity and economic modelling. Economics and Philosophy, 25(3), 297–302.

Munich Re. (2026). Natural disasters in 2025 [Fact sheet]. NatCatSERVICE.

Murfin, Justin, & Matthew Spiegel & Jose Scheinkman, 2020. "Is the Risk of Sea Level Rise Capitalized in Residential Real Estate?," The Review of Financial Studies, Society for Financial Studies, vol. 33(3), pages 1217-1255.

Nelsen, R. B. (2006). An introduction to copulas (2nd ed.). Springer.

Network for Greening the Financial System. (2024). *NGFS climate scenarios technical documentation* (Version 5.0). NGFS. https://www.ngfs.net/system/files/2025-01/NGFS%20Climate%20Scenarios%20Technical%20Documentation.pdf

Network of Central Banks and Supervisors for Greening the Financial System (NGFS), (2024). NGFS climate scenarios: Technical documentation (Phase V). Network for Greening the Financial System.

Network of Central Banks and Supervisors for Greening the Financial System (NGFS), (2024). Note on damage functions, NGFS scenarios, and the economic commitment of climate change. Network for Greening the Financial System.

Network of Central Banks and Supervisors for Greening the Financial System (NGFS), (2025, December 3). Statement regarding physical risk estimates in Phase V of NGFS long-term scenarios

Network of Central Banks and Supervisors for Greening the Financial System (NGFS), (2025). Guide to climate scenario analysis for central banks and supervisors: Update. Network for Greening the Financial System.

Network of Central Banks and Supervisors for Greening the Financial System (NGFS), (2025). NGFS short-term climate scenarios for central banks and supervisors. Network for Greening the Financial System.

Network of Central Banks and Supervisors for Greening the Financial System (NGFS), 2025, Leveraging physical climate risk data

Nguyen, D. D., Ongena, S., Qi, S., & Sila, V. (2022). Climate change risk and the cost of mortgage credit. Review of Finance, 26(6), 1509–1549. https://doi.org/10.1093/rof/rfac013

Northrop, P., & Chandler, R. (2014). Quantifying sources of uncertainty in projections of future climate. Journal of Climate, 27, 8793–8808. https://doi.org/10.1175/JCLI-D-14-00265.1

Norton, E. C., Dowd, B. E., & Maciejewski, M. L. (2019). Marginal effects, quantifying the effect of changes in risk factors in logistic regression models. JAMA, 321(13), 1304 to 1305. https://doi.org/10.1001/jama.2019.1954

Noth, F., & Schüwer, U. (2023). Natural disasters and bank stability: Evidence from the U.S. financial system. *Journal of Environmental Economics and Management, 119*, 102792. https://doi.org/10.1016/j.jeem.2023.102792

Oosterhaven, J. (1988). On the plausibility of the supply-driven input-output model. Journal of Regional Science, 28(2), 203–217. https://doi.org/10.1111/j.1467-9787.1988.tb01208.x

Ortega, F., & Taşpınar, S. (2018). Rising sea levels and sinking property values: Hurricane Sandy and New York's housing market. *Journal of Urban Economics, 106*, 81–100.

Pagani, M., Garcia-Pelaez, J., Gee, R., Johnson, K., Poggi, V., Silva, V., Simionato, M., Styron, R., Viganò, D., Danciu, L., Monelli, D., & Weatherill, G. (2020).

The 2018 version of the Global Earthquake Model: Hazard component. *Earthquake Spectra, 36*(1_suppl), 226–251. https://doi.org/10.1177/8755293020931866

Palmer, W. C. (1965). Meteorological drought (Research Paper No. 45). U.S. Department of Commerce, Weather Bureau.

Panjer, H. H. (1981). Recursive evaluation of a family of compound distributions. *ASTIN Bulletin*, *12*(1), 22–26. https://doi.org/10.1017/S0515036100006796

Pankratz, N. M. C., & Schiller, C. M. (2024). Climate change and adaptation in global supply-chain networks. *The Review of Financial Studies, 37*(6), 1729–1777. https://doi.org/10.1093/rfs/hhad093

PARC. (2024). The alignment cookbook 2: A technical panorama of the alignment methodologies and metrics used by and applied to the financial sector, with a view to inform consolidated alignment assessments. Fondation PARC – Paris Agreement Research Commons. https://institutlouisbachelier.org/en/the-alignment-cookbook-2-2

Parker, M. (2018). The impact of disasters on inflation. *Economics of Disasters and Climate Change, 2*(1), 21–48. https://doi.org/10.1007/s41885-017-0017-y

Perkins, S. E., & Alexander, L. V. (2013). On the measurement of heat waves. *Journal of Climate, 26*(13), 4500–4517. https://doi.org/10.1175/JCLI-D-12-00383.1

Pianosi, F., Beven, K., Freer, J., Hall, J. W., Rougier, J., Stephenson, D. B., & Wagener, T. (2016). Sensitivity analysis of environmental models: A systematic review with practical workflow. Environmental Modelling and Software, 79, 214–232. https://doi.org/10.1016/j.envsoft.2016.02.008

Pickands, J. (1975). Statistical inference using extreme order statistics. *The Annals of Statistics, 3*(1), 119–131.

Pilkington, S. F., & Mahmoud, H. N. (2016). Using artificial neural networks to forecast economic impact of multi-hazard hurricane-based events. *Sustainable and Resilient Infrastructure, 1*(1–2), 63–83. https://doi.org/10.1080/23789689.2016.1179529

Ponticelli, J., Xu, Q., & Zeume, S. (2023). *Temperature, adaptation, and local industry concentration* (NBER Working Paper No. 31533). National Bureau of Economic Research. https://doi.org/10.3386/w31533

Pozdyshev, V., Lobanov, A., & Ilinsky, K. (2025). Incorporating physical climate risks into banks' credit risk models, BIS Working Papers No. 1274

Prohaska, A. (2020). Still struggling: Intersectionality, vulnerability, and long-term recovery after the Tuscaloosa, Alabama USA tornado. *Critical Policy Studies, 14*(4), 466–487. https://doi.org/10.1080/19460171.2020.1724549

Ramanathan, K. (2017, April 24). Modeling fundamentals: Anatomy of a damage function. Verisk. https://www.verisk.com/blog/modeling-fundamentals-anatomy-of-a-damage-function/

Ranasinghe, R., Ruane, A. C., Vautard, R., Arnell, N., Coppola, E., Cruz, F. A., Dessai, S., Islam, A. S., Rahimi, M., Ruiz Carrascal, D., Sillmann, J., Sylla, M. B., Tebaldi, C., Wang, W., & Zaaboul, R. (2021). Climate change information for regional impact and for risk assessment. In V. Masson-Delmotte et al. (Eds.), *Climate change 2021: The physical science basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (pp. 1767–1926). Cambridge University Press. https://doi.org/10.1017/9781009157896.014

Rebonato, R., Melin, L., & Zhang, F. (2025). *How to assign probabilities to climate scenarios*. EDHEC Climate Institute. https://climateinstitute.edhec.edu/publications/how-assign-probabilities-climate-scenarios

Rijal, B., Raulier, F., Martell, D. L., & Gauthier, S. (2018). The economic impact of fire management on timber production in the boreal forest region of Quebec, Canada. *International Journal of Wildland Fire*, *27*(12), 831–844. https://doi.org/10.1071/WF18041

Robock, A. (2000). Volcanic eruptions and climate. *Reviews of Geophysics, 38*(2), 191–219. https://doi.org/10.1029/1998RG000054

Roudier, P. et al. (2016). Projections of future floods and hydrological droughts in Europe under a +2°C global warming. Climatic Change, 135(2), 341-355. https://doi.org/10.1007/s10584-015-1570-4

Saint-Geours, N., Lavergne, C., Bailly, J.-S., & Grelot, F. (2011). Analyse de sensibilité globale d'un modèle spatialisé pour l'évaluation économique du risque d'inondation. Journal de la Société Française de Statistique, 152(1), 49–71.

Sampson, C. C., Smith, A. M., Bates, P. D., Neal, J. C., Alfieri, L., & Freer, J. E. (2015). A high-resolution global flood hazard model. *Water Resources Research, 51*(9), 7358–7381. https://doi.org/10.1002/2015WR016954

San-Miguel-Ayanz, J., Schulte, E., Schmuck, G., Camia, A., Strobl, P., Liberta, G., Giovando, C., Boca, R., Sedano, F., Kempeneers, P., McInerney, D., Withmore, C., Santos de Oliveira, S., Rodrigues, M., Durrant, T., Corti, P., Oehler, F., Vilar, L., & Amatulli, G. (2012). Comprehensive monitoring of wildfires in Europe: The European Forest Fire Information System (EFFIS). In J. P. Tiefenbacher (Ed.), *Approaches to managing disaster: Assessing hazards, emergencies and disaster impacts* (pp. 87–108). InTech. https://doi.org/10.5772/28441

Sastry, P., Sen, I., & Tenekedjieva, A.-M. (2023). *When insurers exit: Climate losses, fragile insurers, and mortgage markets* (SSRN Working Paper No. 4674279). https://doi.org/10.2139/ssrn.4674279

Sauquet, E., Evin, G., Siauve, S., Aissat, R., Arnaud, P., Beniston, M., ... Vidal, J.-P. (2026). A large transient multi-scenario multi-model ensemble of future streamflow and groundwater projections in France. Hydrology and Earth System Sciences, 30(8), 2277–2300. https://doi.org/10.5194/hess-30-2277-2026

Sauquet, E., Vergnes, J.-P., Thirel, G., & Strohmenger, L. (2022). Identification de stations hydrométriques et piézomètres pour l'évaluation des modélisations hydrologiques et hydrogéologiques (technical report). INRAE; BRGM. https://doi.org/10.57745/DE9SPE

Schmid, T., Portmann, R., Villiger, L., Schröer, K., & Bresch, D. N. (2024). An open-source radar-based hail damage model for buildings and cars. Natural Hazards and Earth System Sciences, 24(3), 847–872. https://doi.org/10.5194/nhess-24-847-2024

Schneider, N. (2025). *EDHEC-CLIRMAP: EDHEC-CLimate-Induced Regional MAcroimpacts Projector. How-to guide*. EDHEC Climate Institute. https://climateinstitute.edhec.edu/sites/default/files/2025-08/REVISED_How-To-Guide_EDHEC-CLIRMAP_vSchneider2025.pdf

Schroeder, W., Oliva, P., Giglio, L., & Csiszar, I. A. (2014). The new VIIRS 375 m active fire detection data product: Algorithm description and initial assessment. *Remote Sensing of Environment, 143*, 85–96. https://doi.org/10.1016/j.rse.2013.12.008

Scientific Climate Ratings. (2025). Climate risk rating methodology (Version CRRmod_25.03.00). https://publishing.scientificratings.com/sitepublicmethodologies/climate_risk_rating_methodology.pdf

Scientific Climate Ratings. (2025). Climate scenario probabilities methodology (Version Scenarios_25.01.00). https://publishing.scientificratings.com/sitepublicmethodologies/climate_scenario_probabilities_methodology.pdf

Seklecka, M., Pantelous, A. A., & O'Hare, C. (2017). Mortality effects of temperature changes in the United Kingdom. *Journal of Forecasting, 36*(7), 824–841. https://doi.org/10.1002/for.2473

Shepherd, T. G., Boyd, E., Calel, R. A., Chapman, S. C., Dessai, S., Dima-West, I. M., ... Wade, R. (2018). Storylines: An alternative approach to representing uncertainty in physical aspects of climate change. Climatic Change, 151(3), 555–571. https://doi.org/10.1007/s10584-018-2317-9

Shleifer, A., & Vishny, R. W. (1992). Liquidation values and debt capacity: A market equilibrium approach. The Journal of Finance, 47(4), 1343–1366. https://doi.org/10.1111/j.1540-6261.1992.tb04661.x

Siddikee, M. N., & Rahman, M. M. (2017). Effect of catastrophic disaster in financial market contagion.

Cogent Economics & Finance, 5(1), Article 1288772. https://doi.org/10.1080/23322039.2017.1288772

Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. Publications de l'Institut de Statistique de l'Université de Paris, 8, 229–231.

Stadler, K., et al. (2018). EXIOBASE 3: Developing a time series of detailed environmentally extended multi-regional input-output tables. *Journal of Industrial Ecology, 22*(3), 502–515.

Stanley, T., and D.B. Kirschbaum. 2017. "A Heuristic Approach to Global Landslide Susceptibility Mapping." Natural Hazards 87, no. 1: 145-164; doi:10.1007/s11069-017-2757-y.

Swindles, G. T., Watson, E. J., Savov, I. P., Lawson, I. T., Schmidt, A., Hooper, A., Cooper, C. L., Connor, C. B., Gloor, M., & Carrivick, J. L. (2018). Climatic control on Icelandic volcanic activity during the mid-Holocene. *Geology, 46*(1), 47–50. https://doi.org/10.1130/G39633.1

Task Force on Climate-related Financial Disclosures. (2017). *Recommendations of the Task Force on Climate-related Financial Disclosures: Final report*. Financial Stability Board. https://www.fsb-tcfd.org/recommendations/

Tavor, T. (2024). Assessing the financial impacts of significant wildfires on US capital markets: Sectoral analysis. Empirical Economics, 67(3), 1115–1148. https://doi.org/10.1007/s00181-024-02574-3

Tedesco, M., Hultquist, C., & de Sherbinin, A. (2022). A new dataset integrating public socioeconomic, physical risk, and housing data for climate justice metrics: A test-case study in Miami. *Environmental Justice, 15*(3), 149–159. https://doi.org/10.1089/env.2021.0059

Thieken, A. H., Samprogna Mohor, G., Kreibich, H., & Müller, M. (2022). Compound inland flood events: Different pathways, different impacts and different coping options. Natural Hazards and Earth System Sciences, 22(1), 165–185. https://doi.org/10.5194/nhess-22-165-2022

Thompson, P. and Lynne, G.D. (1994), POLICY DROUGHT: THE CASE OF SOUTH FLORIDA. JAWRA Journal of the American Water Resources Association, 30: 19-26. https://doi.org/10.1111/j.1752-1688.1994.tb03269.x

Timmer, M. P., Dietzenbacher, E., Los, B., Stehrer, R., & de Vries, G. J. (2015). An illustrated user guide to the World Input–Output Database: The case of global automotive production. *Review of International Economics, 23*(3), 575–605.

Tinard, P., Wang, Z.-X., Moncoulon, D., Veysseire, M., Regimbeau, F., & Déqué, M. (2016). Coupling a global climatic model with insurance impact models for flood and drought: An estimation of the financial impact of climate change. *E3S Web of Conferences, 7*, Article 02003. https://doi.org/10.1051/e3sconf/20160702003

Tramblay, Yves; Sauquet, Éric; Arnaud, Patrick; Rousset, Fabienne; Soubeyroux, Jean-Michel; Hingray, Benoit; Jaouen, Tristan; Jeantet, Alexis; Munier, Simon; Vergnes, Jean-Pierre, 2024, "Scénarios d'extrêmes hydrologiques", https://doi.org/10.57745/2XDJ5H, Recherche Data Gouv, V4

Trigg, M. A., Birch, C. E., Neal, J. C., Bates, P. D., Smith, A., Sampson, C. C., Yamazaki, D., Hirabayashi, Y., Pappenberger, F., Dutra, E., Ward, P. J., Winsemius, H. C., Salamon, P., Dottori, F., Rudari, R., Kappes, M. S., Simpson, A. L., Hadzilacos, G., & Fewtrell, T. J. (2016). The credibility challenge for global fluvial flood risk analysis. *Environmental Research Letters, 11*(9), Article 094014. https://doi.org/10.1088/1748-9326/11/9/094014

Tuo, X., & Guo, W. (2011). Flood insurance bonds pricing based on the Monte Carlo simulation. *Systems Engineering Procedia*, *2*, 199–204. https://doi.org/10.1016/j.sepro.2011.10.023

Ultsch, A., & Siemon, H. P. (1990). Kohonen's self organizing feature maps for exploratory data analysis. In B. Widrow & B. Angeniol (Eds.), Proceedings of the International Neural Network Conference (INNC-90), Paris (Vol. 1, pp. 305–308). Kluwer.

UNEP Finance Initiative. (2024). *2024 climate risk landscape report*. United Nations Environment Programme Finance Initiative. https://www.unepfi.org/wordpress/wp-

content/uploads/2024/04/Climate-Risk-Landscape-2024.pdf

United Nations Development Programme. (2025). *Human development report 2025. A matter of choice: People and possibilities in the age of AI*. https://hdr.undp.org/content/human-development-report-2025

United Nations, Department of Economic and Social Affairs, Population Division. (2024). *World population prospects 2024* [Data set]. https://population.un.org/wpp/

Upper, C. (2011). Simulation methods to assess the danger of contagion in interbank markets. *Journal of Financial Stability, 7*(3), 111–125. https://doi.org/10.1016/j.jfs.2010.12.001

Van Wagner, C. E. (1987). *Development and structure of the Canadian Forest Fire Weather Index System* (Forestry Technical Report No. 35). Canadian Forestry Service.

Vasicek, O. (1987). *Probability of loss on loan portfolio*. KMV Corporation.

Vasicek, O. (2002). The distribution of loan portfolio value. *Risk, 15*(12), 160–162.

Verfaillie, D., Déqué, M., Morin, S., & Lafaysse, M. (2017). The method ADAMONT v1.0 for statistical adjustment of climate projections applicable to energy balance land surface models. Geoscientific Model Development, 10(11), 4257–4283. https://doi.org/10.5194/gmd-10-4257-2017

Vicente-Serrano, S. M., Beguería, S., & López-Moreno, J. I. (2010). A multiscalar drought index sensitive to global warming: The Standardized Precipitation Evapotranspiration Index. Journal of Climate, 23(7), 1696–1718. https://doi.org/10.1175/2009JCLI2909.1

Vigneron, P. (2026). Lifting the double hidden risk: Toward a self-organizing review of climate physical risks financial impact models [Working paper]. Institut Louis Bachelier Risk Forum.

Vigneron, P., & Vigneron, V. (2026). Toward a self organizing literature review: An application to the financial impacts of physical climate risks [Conference paper]. Paris Conference on AI and Digital Ethics (PCAIDE).

Von Peter, G., von Dahlen, S., & Saxena, S. C. (2024). *Unmitigated disasters? Risk-sharing and macroeconomic recovery in a large international panel* (BIS Working Papers No. 1175). Bank for International Settlements. https://www.bis.org/publ/work1175.htm

Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. Econometrica, 57(2), 307–333. https://doi.org/10.2307/1912557

Wagenaar, D., de Jong, J., & Bouwer, L. M. (2017). Multi-variable flood damage modelling with limited data using supervised learning approaches. *Natural Hazards and Earth System Sciences, 17*(9), 1683–1696. https://doi.org/10.5194/nhess-17-1683-2017

Wagner, S., Thiam, S., Dossoumou, N. I. P., Hagenlocher, M., Souvignet, M., & Rhyner, J. (2022). Recovering from financial implications of flood impacts. The role of risk transfer in the West African context. *Sustainability, 14*(14), Article 8433. https://doi.org/10.3390/su14148433

Wald, D. J., Worden, B. C., Quitoriano, V., & Pankow, K. L. (2005). *ShakeMap manual: Technical manual, user's guide, and software guide* (Techniques and Methods 12-A1). U.S. Geological Survey. https://doi.org/10.3133/tm12A1

Wang, X., Mammadov, E., Lee, M., Ruiz-Vergote, S., Towey, K., Saffioti, C., & Edwards, J. (2026). *MSCI physical risk guide: For asset owners and asset managers*. MSCI. https://www.msci.com/research-and-insights/paper/physical-risk-guide-for-asset-owners-and-asset-managers

Ward, P. J., Jongman, B., Salamon, P., Simpson, A., Bates, P., De Groeve, T., Muis, S., Coughlan de Perez, E., Rudari, R., Trigg, M. A., & Winsemius, H. C. (2015). Usefulness and limitations of global flood risk models. *Nature Climate Change, 5*(8), 712–715. https://doi.org/10.1038/nclimate2742

Wenz, L., Carr, R. D., Kögel, N., Kotz, M., & Kalkuhl, M. (2023). DOSE. Global data set of reported sub-national economic output. *Scientific Data, 10*, Article 425. https://doi.org/10.1038/s41597-023-02323-8

Wing, O. E. J., Bates, P. D., Quinn, N. D., Savage, J. T. S., Uhe, P. F., Cooper, A., et al. (2024). A 30 m global flood inundation model for any climate scenario. *Water Resources Research*, 60, e2023WR036460. https://doi.org/10.1029/2023WR036460

Wing, O. E. J., Pinter, N., Bates, P. D., & Kousky, C. (2020). New insights into US flood vulnerability revealed from flood insurance big data. Nature Communications, 11, 1444. https://doi.org/10.1038/s41467-020-15264-2

World Bank. (2012). Thai flood 2011: Rapid assessment for resilient recovery and reconstruction planning. World Bank.

Worthington, A. C., & Valadkhani, A. (2004). Measuring the impact of natural disasters on capital markets: An empirical application using intervention analysis. Applied Economics, 36(19), 2177–2186. https://doi.org/10.1080/0003684042000282489

XDI. (2024). 2024 XDI gross domestic climate risk report. The Cross Dependency Initiative.

Xiao, Z., 2026, Labor exposure to climate risk, productivity loss, and capital deepening. *Journal of Financial and Quantitative Analysis*. https://www.cambridge.org/core/journals/journal-of-financial-and-quantitative-analysis/article/labor-exposure-to-climate-risk-productivity-loss-and-capital-deepening/130DF9993D0589A07160509C8C272C25

Xiao, Zhanbing, Labor Exposure to Climate Risk, Productivity Loss, and Capital Deepening (August 1, 2021). *Journal of Financial and Quantitative Analysis, Forthcoming*, Available at SSRN: https://ssrn.com/abstract=4475467 or http://dx.doi.org/10.2139/ssrn.4475467

Yamazaki, D., Kanae, S., Kim, H., & Oki, T. (2011). A physically based description of floodplain inundation dynamics in a global river routing model. Water Resources Research, 47, W04501. https://doi.org/10.1029/2010WR009726

Yang, C.-Y. (2013). Measuring U.S. hurricane risk associated with natural climate cycle and global warming effects. Asia-Pacific Journal of Risk and Insurance, 7(2), 1–26.

Zhou, N., Vilar-Zanón, J. L., Garrido, J., & Heras Martínez, A. J. (2023). On the definition of an actuarial climate index for the Iberian Peninsula. Anales del Instituto de Actuarios Españoles, 29, 37–59. https://doi.org/10.26360/2023_3

Zscheischler, J., Martius, O., Westra, S., Bevacqua, E., Raymond, C., Horton, R. M., van den Hurk, B., AghaKouchak, A., Jézéquel, A., Mahecha, M. D., Maraun, D., Ramos, A. M., Ridder, N. N., Thiery, W., & Vignotto, E. (2020). A typology of compound weather and climate events. *Nature Reviews Earth & Environment, 1*(7), 333–347. https://doi.org/10.1038/s43017-020-0060-z

Zscheischler, J., Westra, S., van den Hurk, B. J. J. M., Seneviratne, S. I., Ward, P. J., Pitman, A., AghaKouchak, A., Bresch, D. N., Leonard, M., Wahl, T., & Zhang, X. (2018). Future climate risk from compound events. *Nature Climate Change, 8*(6), 469–477. https://doi.org/10.1038/s41558-018-0156-3

APPENDIX 1 – QUESTIONS TO ASK CLIMATE DATA AND MODEL PROVIDERS

A model can be outsourced, but the responsibility for understanding it cannot. The questions below turn the uncertainty disclosure card and the scale suitability field into a checklist that institutions can put to any provider, and that supervisors can expect to see documented within model risk management.

- "Can you show me the Risk Map?"
- Which hazards, scenarios and time horizons does the model cover, and which are left out?
- At what spatial and temporal resolution is the hazard observed, and at what granularity is the financial output produced?
- Which method function links the hazard to the loss (damage function, regression, structural model), and how was it calibrated and validated?
- How does the model treat hazards that fall outside the historical record used for calibration?
- Are adaptation and resilience assumptions included, and can the user modify them?
- Is the output a point estimate, an interval or a full distribution, and which confidence margin is disclosed?
- Which parameters were subject to sensitivity analysis, and with what results?
- How are several hazards combined, and how are double-counting and correlation effects handled?
- At which scale is the result directly estimated, and from which scale onward is it aggregated, extrapolated or propagated?
- How is uncertainty carried through when results are aggregated from asset to portfolio level?
- How often are the data and the model updated, and how are changes communicated to users?
- Which parts of the methodology are proprietary, and what can be shared with validators and supervisors?
- What independent validation exists, and can the institution reproduce or challenge the result within its own model risk management framework?

SURVEY FORM – MODEL VALIDATION

Climate Risk Cookbook — Physical Impact Models Survey

Organization: _____ **Date:** _____

1. Scope

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Hazards covered (e.g. floods, storms, earthquakes, landslides, soil erosion, extreme temperatures, hail...) (excluding model-based scenarios, covered in section 13)	☐	☐	

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Coverage period (baseline, live feed, historical...)	☐	☐	

2. Climate risk scoring & adaptation maturity

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Does your solution produce a gross risk score for issuers (physical risk exposure before any adaptation measures)?	☐	☐	
Does your solution assess adaptation maturity (degree to which an issuer has implemented measures to reduce its physical risk exposure)?	☐	☐	
Does your solution produce a net risk score (residual physical risk after accounting for an issuer's adaptation measures)?	☐	☐	
Does your solution produce a resilience score (capacity to absorb shocks and recover from physical risk events)?	☐	☐	

3. Uncertainty quantification

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Does your solution quantify and communicate uncertainty around its financial impact estimates?	☐	☐	
If yes: how is uncertainty communicated?	☐	☐	

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Select all that apply: ☐ Confidence intervals ☐ Probability distributions ☐ Scenario ranges ☐ Model ensembles ☐ Sensitivity analysis ☐ Other			
How would you characterise the overall confidence level you currently assign to your model outputs? Select one: ○ Very high (>90%) ○ High (80–90%) ○ Moderate (70–80%) ○ Indicative (60–70%) ○ Exploratory (<60%)	☐	☐	

4. Backtesting & historical validation

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Has your model been backtested against historical climate events?	☐	☐	
If yes: please describe the methodology and key performance metrics used in your backtesting approach (e.g. RMSE, Brier score, hit rate, return period calibration…)	☐	☐	

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Do you monitor the consistency between your damage predictions and the outcomes your clients subsequently observe (ongoing performance tracking)?	☐	☐	
If yes: what measure do you use for ongoing performance monitoring?	☐	☐	

5. Target users

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Primary target (insurers, banks, investment funds, other...)	☐	☐	

6. Competitive differentiation

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
What makes your solution particularly strong and distinguishes it from other platforms?	☐	☐	

7. Methodology

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Method type Select all that apply: ☐ Damage-based	☐	☐	

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
☐ PD-based (probability of default) ☐ Accounting-based ☐ Actuarial ☐ Other			
Use of machine learning (ML)	☐	☐	
Tipping points modelled (critical thresholds beyond which climate impacts become non-linear and potentially irreversible)	☐	☐	
Control variables used	☐	☐	
Is any part of your methodology published in open access? Is any part of your methodology shared with or accessible to your clients?	☐	☐	

8. Modelling chain transparency

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Which parts of your modelling chain are documented and accessible to clients? Select all that apply: ☐ Hazard module ☐ Vulnerability module ☐ Damage function	☐	☐	

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
☐ Financial translation module ☐ Scenario assumptions ☐ None / fully proprietary			

9. Damage function nature

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
How are physical hazards translated into damages in your solution? Select all that apply: ☐ Engineering / catastrophe model ☐ Empirical damage function (based on observed losses) ☐ Statistical relationship (regression-based) ☐ Machine learning ☐ Expert judgement ☐ Hybrid approach ☐ Other	☐	☐	

10. Primary level of analysis

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
What is the primary level of analysis at which your solution operates? Select all that apply: ☐ Asset / Infrastructure	☐	☐	

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
☐ Facility / Site ☐ Company / Issuer ☐ Portfolio ☐ Sector ☐ Region / Sub-national ☐ Sovereign / Country ☐ Other			

11. Regulatory & operational use cases

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Which use cases does your solution support? Select all that apply: ☐ Regulatory stress testing ☐ ORSA (Own Risk and Solvency Assessment) ☐ Capital allocation ☐ Investment decisions ☐ Underwriting ☐ Strategic planning ☐ Disclosure (TCFD / ISSB / CSRD) ☐ Other	☐	☐	

12. Model types

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Which of the following model categories does your solution rely on? Select all that apply and correct where necessary. See annex for model category descriptions Select all that apply: ☐ Econometric / regression-based (OLS, GLM, Logit, Probit...) ☐ EVT-based — extreme value theory (GEV, GPD, tail risk modelling...) ☐ Simulation-based (Monte Carlo, GARCH, Markov chain...) ☐ Machine learning (Random Forest, XGBoost, SVM, neural networks...) ☐ Time-series models (ARMA, ARIMA, LSTM...) ☐ Causal inference (SEM, DiD, DAG...) ☐ Hybrid / stacking approaches ☐ Other	☐	☐	

13. Model technical fields — physical dimension

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Model name	☐	☐	
Geographical coverage	☐	☐	
Time horizon considered	☐	☐	
Number of hazards	☐	☐	

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Physical proxy (e.g. wind speed, precipitation, temperature anomaly...)	☐	☐	
Physical risk data source	☐	☐	
Physical resolution	☐	☐	
Physical refresh rate	☐	☐	
Sample events (illustrative events used in the model)	☐	☐	
Scenario(s) considered (e.g. RCP 2.6, 4.5, 8.5 / SSP...)	☐	☐	

14. Model technical fields — financial dimension

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Financial granularity (asset level, portfolio level, sector level...)	☐	☐	
Financial proxy (e.g. revenue loss, asset value impairment, EAD...)	☐	☐	
Financial impact data source	☐	☐	

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Financial refresh rate	☐	☐	
Geographical coverage (financial)	☐	☐	
Sample units (e.g. companies, loans, real estate assets...)	☐	☐	
Sector(s) covered	☐	☐	
Personalized scoring (y/n). The solution offers a personalized scoring output, simplifying the results interpretation.	☐ Yes ☐ No	☐ Yes ☐ No	

15. Other data

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
External data sources — physical	☐	☐	
External data sources — financial	☐	☐	
Internal data sources (proprietary sensors, satellite networks, network-based...)	☐	☐	

Optional — Sample for publication

This section is not mandatory. Your response will help us illustrate the diversity of solutions covered in our research.

QUESTION / RECORDED ANSWER	OK	CORR.	CORRECTIONS / COMMENTS
Would you be open to us contacting you again to obtain a data sample that we could feature on our website (e.g. an illustrative output, anonymised extract, or visual)?	☐	☐	

16. Additional comments

Any other remarks, corrections, or information you would like to add or remove:

Thank you for taking the time to review this information. Please return this form at your earliest convenience.

APPENDIX 2 - MODEL CATEGORY DESCRIPTIONS

The following paragraphs describe each model category listed in section 12.

Econometric / regression-based (OLS, GLM, Logit, Probit...)

Statistical models estimating relationships between variables via regression techniques (OLS, GLM, etc.). Suitable for explaining damage as a function of hazard intensity and exposure characteristics.

EVT-based — extreme value theory (GEV, GPD, tail risk modelling...)

Models characterising the distribution of rare, high-impact events (100-year or 200-year return levels for floods, storms, heatwaves, etc.). Specifically designed for tail risk estimation.

Simulation-based (Monte Carlo, GARCH, Markov chain...)

Models generating large numbers of simulated scenarios to estimate outcome distributions (uncertainty propagation, non-linear dynamics, volatility clustering, etc.).

Machine learning (Random Forest, XGBoost, SVM, neural networks...)

Data-driven algorithms learning complex patterns from large datasets without explicit functional form assumptions (damage prediction, risk classification at scale, etc.).

Time-series models (ARMA, ARIMA, LSTM...)

Models capturing temporal dependencies in sequential data (hazard intensity trends, financial exposure forecasts over time, etc.).

Causal inference (SEM, DiD, DAG...)

Methods estimating causal effects rather than mere correlations (isolating the financial impact of a hazard event while controlling for confounders, etc.).

Hybrid / stacking approaches

Combinations of two or more model families (a physics-based hazard model feeding into a machine learning damage function, stacking multiple predictors, etc.).

Other

Any model type not covered by the categories above (please describe in the correction field).

TABLE OF CONTENTS

FOREWORD 3

AUTHORS AND CONTRIBUTORS 4

EXECUTIVE SUMMARY 5

KEY OBJECTIVES 7

KEY FINDINGS 10

READER'S GUIDE 17

THE MENU – UNDERSTANDING PHYSICAL CLIMATE RISK AND FINANCIAL IMPACT 19

1.1 WHAT IS A PHYSICAL CLIMATE RISK? 19

1.2 FROM PHYSICAL DAMAGE TO FINANCIAL IMPACT 28

1.3 FINANCIAL OBJECTS AT RISK 31

1.4 PHYSICAL RISK AND FINANCIAL IMPACTS: MODELLING ONTOLOGY 35

1.4.1 What is a model? 35

1.4.2 The specificities of physical climate risk models 35

1.4.3 Model Risk in Physical Climate Risk Assessment 35

1.4.4 What to expect from Physical Risk Models and how to use them responsibly? 36

1.5 THE PURPOSE OF THIS COOKBOOK AND ITS INTENDED USERS 38

1.6 RESEARCH QUESTIONS 40

1.6.1 How to predict financial impacts from physical RISKS? 40

1.6.2 How to provide an accurate financial assessment in case of a new physical risk? The Bayesian approach to physical risks 41

THE RECIPES – FAMILIES OF PHYSICAL RISK FINANCIAL IMPACT MODELS 44

AN OVERVIEW 44

2.1 DAMAGE-BASED MODELS 49

2.1.1 Damage functions and precise simulation modelling 49

2.1.2 Actuarial models 70

2.1.3 Loss and frequency curves 86

2.1.4 Real effects: An approach to physical risks TO GO BEYOND THE FINANCIAL IMPACTS 97

2.2 PROBABILISTIC AND DEFAULT BASED MODELS 101

2.2.1 Physical risks and probability of default 101

2.2.2 Probit, logit and survival models 110

2.2.3 Structural Credit Models with Physical Risk: the Merton-Climate Framework 124

2.2.4 Portfolio Credit Risk and Physical Climate Credit VaR 130

2.2.5 From One Risk to Another: From Credit Risk to Liquidity Risk 134

2.3 ACCOUNTING AND VALUATION IMPACT MODELS 136

2.3.1 Price Loss, Turnover loss, capex and opex 136

2.3.2 Enterprise value impact 146
2.3.3 Financial ratios 158
2.3.4 Micro and meso models 164
2.4 NETWORK AND SUPPLY-CHAIN IMPACT ASSESSMENT MODELS 169
2.4.1 Physical risks and macro-economic frameworks 169
2.4.2 Shock propagation 193
2.4.3 Supply chain impact 198
2.4.4 Systemic risk 204
2.4.5 Lagged effects 209
2.5 SCENARIO-BASED AND PROJECTIVE MODELS 215
2.5.1 SSP-RCP projections 215
2.5.2 Stress-test analysis 221
2.5.3 Cat Models 224
2.5.4 ML-Based models 230
2.5.5 Predictive Limits 232
THE INGREDIENTS – DESIGN CHOICES AND DATA 236
3.1 THE HAZARD CHOICES AND THEIR REPRESENTATION 236
3.1.1 Single vs Multiple Hazard 237
3.1.2 Hazard type 255
3.1.3 Spatial resolution 264
3.1.4 Historical vs Real Time vs Projection 278
3.2 EXPOSURE AND VULNERABILITY MODELLING CHOICES 282
3.2.1 Asset location 283
3.2.2 Socio-Economic Variables 286
3.2.3 Adaptation and resilience hypothesis 290
3.3 FINANCIAL PROXY CHOICES 296
3.3.1 Accounting variables 296
3.3.2 Market variables 299
3.3.3 Actuarial variables 301
3.3.4 Temporal choices for Financial Metrics 303
3.3.5 Compliance issues 305
3.4 AGGREGATION CHOICES 306
3.4.1 From asset to portfolio 306
3.4.2 The uncertainty issue 312
3.4.3 Non-linearities 336
3.4.4 Type I, II and III errorS 338

3.4.5 Transparency and replicability 339
MODEL CLUSTERING 341
4.1 CLOSED CLUSTERS 345
4.1.1 A first closed map of the corpus 346
4.1.2 A refined closed map 347
4.1.3 What survives a change of object, grid and seed 348
4.2 OPEN CLUSTERS 350
4.2.1 Multiple correspondence analysis to discriminate models 350
4.2.2 VAE latent space projections to discriminate models 352
4.2.3 Comparing models from the literature 354
4.3 USER-GUIDED FILTERING 356
4.3.1 Maps and filters 356
4.3.2 Toward a systematic comparison of models 358
THE CLIMATE FINANCIAL RISK HUB, REFERENCES AND APPENDIX 360
5.1 THE CLIMATE FINANCIAL RISK HUB: FROM A MODEL REVIEW TO AN ONLINE PLATFORM 360
5.2 A PANORAMIC VIEW OF THE FINANCIAL RISK MODELLING ECOSYSTEM 365
5.2.1 A collection and treatment 366
5.2.2 Automated sources of collection 367
5.2.3 Insights from the field 367
5.2.4 Academic insight 368
5.2.5 Providers insight 368
5.2.6 End-user insights 368
5.2.7 Predictive Model classifications 369
5.3 REFERENCES 370
APPENDIX 1 – QUESTIONS TO ASK CLIMATE DATA AND MODEL PROVIDERS 390
APPENDIX 2 - MODEL CATEGORY DESCRIPTIONS 401
TABLE OF CONTENTS 402
LIST OF FIGURES 404
LIST OF TABLES 408
LIST OF PROVIDER BOXES 411

LIST OF FIGURES

FIGURE 1. FROM CLIMATE SIGNAL TO FINANCIAL IMPACT: THE TRANSMISSION CHAIN OF PHYSICAL CLIMATE RISK 6
FIGURE 2. STRUCTURE OF THE MODEL CARD, FROM THE PHYSICAL SIGNAL TO THE FINANCIAL OUTPUT 9
FIGURE 3. PHYSICAL RISKS AND TRANSITION RISKS 19
FIGURE 4. STANDARD RISK MATRIX (ADAPTED FROM ISO 31000 AND ISO/IEC 31010) 21
FIGURE 5. REVERSED RISK MATRIX: A PHYSICAL CLIMATE RISK PERSPECTIVE 22
FIGURE 6. HAZARD, EXPOSURE AND VULNERABILITY, THE THREE COMPONENTS OF PHYSICAL RISK 22
FIGURE 7. TOTAL DAMAGE RELATED TO NATURAL CATASTROPHES FROM 1983 TO 2024 (IN $BN). 23
FIGURE 8. TRANSMISSION CHAIN FROM CLIMATE SIGNAL TO FINANCIAL IMPACT 28
FIGURE 9. THE LIBBY BOXES OF PHYSICAL RISK FINANCIAL IMPACT MODELS 29
FIGURE 10. GEOGRAPHICAL AND FINANCIAL SCALES OF ANALYSIS 32
FIGURE 11. FINANCIAL OBJECTS AT RISK, FROM GRANULAR PHYSICAL OBJECTS TO THE FINANCIAL SYSTEM 33
FIGURE 12. THREE-STEP APPROACH TO MODELLING: SELECTION, CALIBRATION, TESTING AND EXPLORATION 38
FIGURE 13. FRAMEWORK FOR THE ASSESSMENT OF CLIMATE-RELATED VULNERABILITIES 39
FIGURE 14. THE FAMILIES OF MODELS 44
FIGURE 15. MODEL CARD TEMPLATE USED TO DESCRIBE MODELS 48
FIGURE 16. THE MODELLING CHAIN OF PHYSICAL RISKS FINANCIAL IMPACT MODELS 49
FIGURE 17. HISTOGRAM OF DAMAGE STATE FREQUENCY FOR 1000 SIMULATION RUNS 57
FIGURE 18. AN EXAMPLE OF OUTPUT USING CLIMADA OPEN-SOURCE DAMAGE FUNCTIONS. 65
FIGURE 19. MODEL REPRESENTATION CARD – DAMAGE FUNCTION / VULNERABILITY MODEL (SALT FARMING IN INDONESIA) 66
FIGURE 20. MODEL REPRESENTATION CARD – DAMAGE FUNCTIONS (CAT MODEL) APPLIED TO FLOOD RISK IN FRANCE 67
FIGURE 21. POISSON DISTRIBUTION OF NORTH ATLANTIC HURRICANE FREQUENCY BASED ON 30-YEAR AVERAGE (1990 – 2020) 71
FIGURE 22. AN EXAMPLE OUTPUT FROM PCRM1.0 81
FIGURE 23. MODEL REPRESENTATION CARD – COMPOUND POISSON / ACTUARIAL CATASTROPHE MODEL (TEXAS, CZAJKOWSKI ET AL., 2011) 85
FIGURE 24. GALVESTON COUNTY 2009 NFIP PREMIUM DATA 86
FIGURE 25. AN OEP CURVE EXAMPLE 87
FIGURE 26. AN EP CURVE EXAMPLE 89
FIGURE 27. EXAMPLE OF AN AAL CURVE 90
FIGURE 28. AN EXAMPLE OF AN RP CURVE (LOG SCALE) 91
FIGURE 29. MODEL REPRESENTATION CARD – EP FRAMEWORK (LOG-LINEAR WARMING SCALING, JEWSON ET AL., 2023 96
FIGURE 30. FRENCH DEATH RATE ATTRIBUTABLE TO HIGH TEMPERATURE EXPOSURE (BOTH SEXES, FRANCE, FOR 100,000) 99
FIGURE 31. THE LINK BETWEEN DEFAULT PROBABILITIES AND TEMPERATURE ANOMALIES 102
FIGURE 32. DISTRIBUTION OF LOSSES ON A GLOBAL UNIVERSE OF BONDS AND EQUITIES ACROSS 1,000 CLIMATE REALIZATIONS ... 108
FIGURE 33. MODEL REPRESENTATION CARD – PD-BASED CREDIT MODEL (CLIMACRED-PHYS) 109
FIGURE 34. ODACC INTERFACE 116

FIGURE 35. ODACC EXPOSED ASSETS INTERFACE 117
FIGURE 36. PORTFOLIO SUMMARY 121
FIGURE 37. MODEL REPRESENTATION CARD – CLIMATE-ADJUSTED CREDIT MODEL (S&P CLIMATE CREDIT ANALYTICS)..... 122
FIGURE 38. MODEL REPRESENTATION CARD – REDUCED-FORM HAZARD MODEL (FLOOD RISK AND LOAN PRICING, BARBAGLIA ET AL., 2023)..... 123
FIGURE 39. COMPARISON OF % ADD-ON ON 5 YEAR VAR 95% 128
FIGURE 40. MODEL REPRESENTATION CARD – STRUCTURAL CREDIT MODEL WITH CLIMATE JUMPS (AZZONE ET AL., 2025) 129
FIGURE 41. ANALYTICAL DIMENSIONS AND STEPS INTEGRATED IN CRG RE (FIG.1) 132
FIGURE 42. MODEL REPRESENTATION CARD – SURVIVAL MODEL (EXTREME WEATHER AND MORTGAGE DEFAULT, CALABRESE ET AL., 2024 132
FIGURE 43. MODEL REPRESENTATION CARD – P&L BRIDGE MODEL (THE BRAZILIAN CASE) 143
FIGURE 44. SCHNEIDER ELECTRIC ADVISORY SERVICES' CLIMATE RISK DASHBOARD (FORMERLY ECLR, NOW RA+ CLIMATE RISK). 153
FIGURE 45. MODEL REPRESENTATION CARD – ENTERPRISE VALUE IMPACT (EXTREME HEAT AND FIRM PERFORMANCE, PANKRATZ ET AL., 2023)..... 155
FIGURE 46. MODEL REPRESENTATION CARD – ENTERPRISE VALUE MODEL (VALUATION UNDER CLIMATE UNCERTAINTY, WANG ET AL., 2025)..... 156
FIGURE 47. MODEL REPRESENTATION CARD – FINANCIAL RATIO MODEL (FLOODS AND BANK SOLVENCY, CALOIA ET AL., 2023) 163
FIGURE 48. FRAMEWORK FOR THE ASSESSMENT OF CLIMATE-RELATED VULNERABILITIES 171
FIGURE 49. NBFI BLOCK OF THE FSB FRAMEWORK AND NBFI RISK TRANSMISSION CHANNELS 172
FIGURE 50. BANK EXPOSURES TO PHYSICAL RISK ACROSS JURISDICTIONS 172
FIGURE 51. DIFFERENT USE CASES FOR PHYSICAL RISK DATA BEING CONSIDERED BY SURVEYED INSTITUTIONS 173
FIGURE 52. ECONOMIC AND FINANCIAL IMPACTS ASSESSED BY CENTRAL BANKS AND SUPERVISORS..... 174
FIGURE 53. A QUICK LOOK AT AN ASSET PORTFOLIO – CLIMATEMETRICS..... 187
FIGURE 54. SELECTED EXAMPLE OF A CLIMATE EXPOSURE RATING – HORIZON 2050 FOR A GAS-FIRED POWER GENERATION ASSET IN THE USA..... 188
FIGURE 55. SELECTED EXAMPLE: 2050 HORIZON SOVEREIGN RISK RATING FOR THE UNITED STATES OF AMERICA..... 190
FIGURE 56. CLIRMAP SSP2-RCP4.5 MODERATE WARMING SCENARIO (EDHEC CLIRMAP, 2025)..... 191
FIGURE 57. MODEL REPRESENTATION CARD – NGFS MACRO MODEL (NGFS PHYSICAL RISK FRAMEWORK, NGFS 2023–2025) . 192
FIGURE 58. MODEL REPRESENTATION CARD – SHOCK PROPAGATION MODEL (TROPICAL CYCLONES, DE MAXIMY ET AL., 2024).... 197
FIGURE 59. MODEL REPRESENTATION CARD – SUPPLY CHAIN BUSINESS INTERRUPTION (CLIMVAR, EPELBAUM, 2026)..... 202
FIGURE 60. THE CASE FOR CONSIDERING MACROECONOMIC VARIABLES – IMPACT OF EQ-LA SCENARIO ON KEY MACROECONOMIC VARIABLES. 208
FIGURE 61. MODEL REPRESENTATION CARD – FLOOD-INDUCED SYSTEMIC RISK THROUGH FINANCIAL NETWORKS (MANDEL ET AL., 2021)..... 208
FIGURE 62. FROM CLIMATE SCENARIO TO FINANCIAL LOSS: THE MODELLING CHAIN..... 217
FIGURE 63. MODEL REPRESENTATION CARD – PROBABILISTIC CLIMATE SCENARIOS..... 220
FIGURE 64. MODEL REPRESENTATION CARD – CLIMATE-ADJUSTED FLOOD CATASTROPHE MODEL (CONTIGUOUS US, WING ET AL., 2022)..... 229

FIGURE 65. A THREE-STEP EXAMPLE ON HOW DOWNSCALING OPERATES 241
FIGURE 66. EARTHSCAN PLATFORM INTERFACE 243
FIGURE 67. 200-YEAR FLUVIAL FLOOD DEPTH MAP IN PAKISTAN. 245
FIGURE 68. TAIL DEPENDENCE OF FOUR COPULAS WITH IDENTICAL KENDALL'S TAU 251
FIGURE 69. UNDERESTIMATION OF TAIL LOSSES BY SINGLE-HAZARD MODELS COMPARED WITH COMPOUND HAZARD EVENTS 252
FIGURE 70. AN ILLUSTRATION OF THE BREADTH VS. DEPTH DILEMMA – NUMBER OF HAZARDS BY PROVIDER 255
FIGURE 71. ILLUSTRATION OF THE VOLCANIC ACTIVITY AND GLACIER EXTENT. 257
FIGURE 72. SATELLITE-BASED WILDFIRE MONITORING WORKFLOW 264
FIGURE 73. (A) FLOOD HAZARD CLASSES ACROSS ITALY; (B) DOWNSCALING OF A EUROPEAN TEMPERATURE FIELD FROM 100 KM TO 5 KM RESOLUTION 267
FIGURE 74. WWF WATER RISK FILTER INTERFACE SHOWING BASIN IN NORTHERN CHINA AND THE KOREAN PENINSULA 268
FIGURE 75. STATIC FLOOD HAZARD MAP COMPARED WITH A STOCHASTIC EVENT CATALOGUE FOR A PORTFOLIO OF 30 PROPERTIES 270
FIGURE 76. FROM LOCALIZATION TO DIGITAL TWIN 271
FIGURE 77. WHY ASSET SPECIFIC BOUNDARIES MATTER 273
FIGURE 78. CLIMATE RISK HITS EQUITY BROADLY, DEBT NARROWLY 275
FIGURE 79. FOR DEBT, THE RISK IS CONCENTRATED IN A SMALL SET OF EXPOSED FIRMS 275
FIGURE 80. HISTORICAL GLOBAL TEMPERATURE CHANGE AND PROJECTIONS TO 2100 UNDER THREE CMIP6 SCENARIOS (SSP1-2.6, SSP2-4.5 AND SSP5-8.5) 278
FIGURE 81. GLOBAL WARMING PATHWAYS TO 2100 UNDER IPCC/CMIP6, NGFS AND IEA SCENARIOS 282
FIGURE 82. SIZE OF THE ASSET UNIVERSES MAPPED BY SELECTED VENDORS, AS DECLARED BY THE VENDORS (LOG SCALE) 286
FIGURE 83. ILLUSTRATION OF DIFFERENCES IN HAZARD IMPACT IN DIFFERENT AREAS 287
FIGURE 84. NATURE-BASED SOLUTIONS TO CLIMATE HAZARDS ACROSS FIVE LANDSCAPES, FROM MOUNTAINS TO COASTS 290
FIGURE 85. CORRENTICS PLATFORM INTERFACE SHOWING THE FOOTPRINT OF HURRICANE HELENE (SEPTEMBER 2024) ACROSS A LOGISTICS NETWORK OF PORTS AND SHIPPING ROUTES 311
FIGURE 86. CORRENTICS PLATFORM INTERFACE SHOWING ASSET-LEVEL TROPICAL CYCLONE RISK FOR A SITE IN EAST ASIA UNDER A LOW-EMISSION SCENARIO IN 2030 311
FIGURE 87. MODELLING CHAIN AND ENSEMBLE LAYERS OF THE CASE 315
FIGURE 88. RANGE ACROSS CATCHMENTS OF EACH COMPONENT'S SHARE OF EPISTEMIC VARIANCE 322
FIGURE 89. KOHONEN MAP OF THE 93 HARMONIZED MODALITIES OF THE 195 CODED MODELS 347
FIGURE 90. FIRST PLANE OF THE MULTIPLE CORRESPONDENCE ANALYSIS OF THE 195 CODED MODELS 351
FIGURE 91. ARCHITECTURE OF THE VARIATIONAL AUTOENCODER 352
FIGURE 92. T-SNE PROJECTION OF THE VAE LATENT SPACE, ONE PANEL PER GROUP OF MODELS 353
FIGURE 93. FILTERED MAP OF THE CODED CORPUS 356
FIGURE 94. UPDATABLE PIPELINE OF THE SELF-ORGANIZING LITERATURE REVIEW. 359
FIGURE 95. THE MODEL CARD AS THE UNIT OF RECORD 360
FIGURE 96. CLIMATE RISK MODELS HUB DASHBOARD 361
FIGURE 97. A SIMPLIFIED VIEW OF OUR MODEL CLASSIFICATION PIPELINE 362
FIGURE 98. THE UPDATABLE PIPELINE OF THE SELF-ORGANIZING LITERATURE REVIEW. 367

FIGURE 99. ANOTHER FORM OF PIPELINE FOR SELF-ORGANIZING LITERATURE REVIEWS..... 369

LIST OF TABLES

TABLE 1. MODEL FAMILIES REVIEWED IN PART 2, WITH THEIR OUTPUTS AND USES 8

TABLE 2. FINANCIAL OBJECTS AT RISK, FROM THE MOST GRANULAR TO THE BROADEST SCALE OF ANALYSIS 11

TABLE 3. SCALE SUITABILITY OF THE MODEL FAMILIES 12

TABLE 4. READER'S GUIDE: QUESTIONS ANSWERED BY EACH PART AND SECTION OF THE COOKBOOK 18

TABLE 5. DNSH'S ANNEX A TABLE LISTING THE MOST WIDESPREAD CLIMATE HAZARDS 24

TABLE 6. MAIN SUPERVISORY REQUIREMENTS ON CLIMATE RISKS FOR BANKS AND INSURERS 27

TABLE 7. TIME HORIZON, SPACE CONSIDERED AND DAMAGE TYPE BY TYPE OF FINANCIAL ACTOR 28

TABLE 8. FINANCIAL OBJECTS AT RISK 31

TABLE 9. MODEL FAMILIES, OUTPUTS AND USES 47

TABLE 10. EXAMPLES OF HAZARD-SPECIFIC DAMAGE FUNCTIONS 55

TABLE 11. DAMAGE STATE PROBABILITIES AND REPEATED RANDOM SAMPLING 57

TABLE 12. JRC DAMAGE FUNCTIONS LINKING FLOOD DEPTH TO DAMAGE BY REGION 59

TABLE 13. EXAMPLES OF OPEN-SOURCE TOOLS FOR DAMAGE-BASED MODELLING 61

TABLE 14. FLOOD DAMAGE FUNCTION APPLIED TO AN INDUSTRIAL PORTFOLIO 63

TABLE 15. HAZARDS AND RESOLUTIONS COVERED BY MORNINGSTAR'S SUSTAINALYTICS 82

TABLE 16. BASEL III RISK WEIGHTS FOR RESIDENTIAL REAL ESTATE EXPOSURES THAT DEPEND ON THE PROPERTY'S CASH FLOWS, BY LTV BAND (CRE20) 104

TABLE 17. BASEL III RISK WEIGHTS FOR RESIDENTIAL REAL ESTATE EXPOSURES THAT DO NOT DEPEND ON THE PROPERTY'S CASH FLOWS, BY LTV BAND (CRE20) 104

TABLE 18. EFFECT OF A €20,000 FLOOD SHOCK ON COLLATERAL VALUE, LTV, RISK WEIGHT AND RWA FOR A €240,000 LOAN ... 105

TABLE 19. THE RATIO IMPACT SUMMARY 162

TABLE 20. THE SCALE HIERARCHY IN PHYSICAL CLIMATE RISK MODELLING 165

TABLE 21. AGGREGATE PROXY METRICS 175

TABLE 22. FINANCIAL EXPOSURE TO HAZARD-PRONE LOCATIONS 175

TABLE 23. FINANCIAL SECTOR LOSS METRICS 176

TABLE 24. PHYSICAL AND FINANCIAL RISK UNDER THE NGFS SCENARIOS 178

TABLE 25. AVAILABLE DATABASES ON TRANSMISSION CHANNELS 183

TABLE 26. HAZARD TABLE – CLIMATEMETRICS 185

TABLE 27. IMPLEMENTATION AND USE CASES: NGFS MACRO MODEL 193

TABLE 28. IMPLEMENTATION AND USE CASES: THE SYSTEMIC RISK MODEL 209

TABLE 29. THE FOUR TEMPORAL PHASES 211

TABLE 30. THE DIFFERENT IPCC SCENARIOS AND THEIR NGFS EQUIVALENT 216

TABLE 31. THE MAJOR REGULATORY STRESS TEST FRAMEWORKS 222

TABLE 32. THE COMPONENTS AND THEIR ROLE 225

TABLE 33. FOUR EXAMPLES OF PERILS 239

TABLE 34. THE 28 MUNICHRE HAZARDS 241

TABLE 35. CLIMATE INDICES COMMONLY USED TO MEASURE PHYSICAL HAZARDS IN INSURANCE AND FINANCIAL RISK ASSESSMENT 248

TABLE 36. CLIMATE INDICES COMMONLY USED TO MEASURE PHYSICAL HAZARDS IN INSURANCE AND FINANCIAL RISK ASSESSMENT 249
TABLE 37. COMPARISON: SINGLE-HAZARD VS COMPOUND HAZARD DATA REQUIREMENTS 253
TABLE 38. CAUSAL LINKS BETWEEN HUMAN ACTIVITY, HAZARDS, CLIMATE TRENDS AND FINANCIAL IMPACT 258
TABLE 39. PHYSICS-BASED VARIABLES AND EXAMPLE DATA SOURCES BY HAZARD TYPE 259
TABLE 40. REQUIRED SPATIAL RESOLUTION VS. RESOLUTION OF REPRESENTATIVE GLOBAL DATASETS 265
TABLE 41. TABLE OF DATASETS USED IN THE WWF RISK FILTER SUITE 269
TABLE 42. THE 6 SCR MODELLING STEPS 273
TABLE 43. SPATIAL RESOLUTIONS USED BY SCR 277
TABLE 44. COMPARING SPATIAL UNIT AND RESOLUTION BY PROVIDER 277
TABLE 45. RISK HORIZONS AND DATA EMPHASIS OF PHYSICAL CLIMATE RISK ASSESSMENT BY TYPE OF FINANCIAL INSTITUTION 280
TABLE 46. MAIN CLIMATE SCENARIO FRAMEWORKS USED IN PHYSICAL RISK ASSESSMENT 281
TABLE 47. SOURCES FOR IDENTIFYING AND LOCATING ASSETS, FROM LEGAL ENTITY TO PHYSICAL SITE 284
TABLE 48. ASSET LOCATION METHODS OF SELECTED PROVIDERS, FROM CLIENT INPUT TO LOCATION OUTPUT 286
TABLE 49. EXAMPLES OF SOCIO-ECONOMIC DATA VARIABLES, SCALE AND COVERAGE 289
TABLE 50. DATA FOR ADAPTATION AND RESILIENCE 292
TABLE 51. REPORTED BENEFIT–COST RATIOS OF CLIMATE ADAPTATION INVESTMENTS 294
TABLE 52. DATA FOR ACCOUNTING METRICS 299
TABLE 53. DATA FOR MARKET METRICS 301
TABLE 54. DATA FOR ACTUARIAL METRICS 303
TABLE 55. DATA FOR TEMPORAL CHOICES 305
TABLE 56. ILLUSTRATIVE FIGURES 307
TABLE 57. AGGREGATION METHODS FOR PHYSICAL CLIMATE RISK 307
TABLE 58. AGGREGATION FIELD OF THE MODEL CARD 309
TABLE 59. COMPARISON OF A FEW PHYSICAL CLIMATE RISK MODEL PROVIDERS FOR FLOOD ASSESSMENT 312
TABLE 60. SOURCES OF UNCERTAINTY ALONG A MODELLING CHAIN 313
TABLE 61. APPROACHES TO UNCERTAINTY QUANTIFICATION 314
TABLE 62. EARLIER STUDIES ON UNCERTAINTY IN FLOOD RISK CHAINS 315
TABLE 63. LAYERS OF THE ENSEMBLE 316
TABLE 64. HYDROLOGICAL MODELS OF EXPLORE2 USED IN THE ENSEMBLE 317
TABLE 65. DAMAGE MODELS OF THE ENSEMBLE 318
TABLE 66. EXPERIMENTAL PORTFOLIO 319
TABLE 67. QUALYPSO CONFIGURATION AND HORIZONS 321
TABLE 68. SHARE OF EPISTEMIC VARIANCE BY COMPONENT (PERCENT, RANGE ACROSS CATCHMENTS) 322
TABLE 69. INDIVIDUAL MODEL EFFECTS WITHIN EACH LAYER 323
TABLE 70. FINANCIAL BASELINES OF THE DAMAGE MODELS 324
TABLE 71. MODELS RECURRING IN EXTREME CHAINS 324
TABLE 72. SPATIAL STRUCTURE OF UNCERTAINTY IN FRANCE 325

TABLE 73. Mapping of the chain to other hazards 327
TABLE 74. Climate narratives available in Explore2 328
TABLE 75. Reporting template 329
TABLE 76. The three questions of model discrimination 330
TABLE 77. Bootstrap and leave-one-out discrimination in three experiments 333
TABLE 78. Variance partition and bootstrap discrimination compared 334
TABLE 79. Data for Compound risk aggregation 336
TABLE 80. Data for non-linearities 337
TABLE 81. Data for type 1 and 2 errors 339
TABLE 82. Data for Transparency and replicability 340
TABLE 83. Three regimes of model clustering 342
TABLE 84. Framework coding groups 343
TABLE 85. First closed clustering of the corpus on a 3 × 3 Kohonen grid 346
TABLE 86. Regions of the 5 × 5 Kohonen map 348
TABLE 87. Reading the first MCA plane of the coded corpus 351
TABLE 88. Consistency of model families across the three projections 354
TABLE 89. Model families located by the maps, with one representative each 355
TABLE 90. Two filters applied to the coded corpus 358

LIST OF PROVIDER BOXES

CLIMADA AND CLIMADA TECH 63

RREFINQ 68

THE CASE OF ADAPTATION/S – USING INQUIRIES ABOUT INDUSTRIAL SITES TO FIT THE MODEL . 70

THE CASE OF RMS CLIMATE ON DEMAND 76

MORNINGSTAR'S SUSTAINALYTICS: A RAPID EVOLUTION IN PHYSICAL RISKS PORTFOLIO MANAGEMENT 81

SWISS RE TURNS ITS NATURAL CATASTROPHE EXPERTISE INTO A SET OF CLIMATE RISK SCORES 83

EQECAT / CORELOGIC / COTALITY 83

MUNICHRE LOSS FREQUENCY APPROACH TO RISK 95

MILLIMAN'S APPROACH TO MORTALITY 99

STRAIGHT OUT OF ACADEMIA TO THE INDUSTRY: CLIMACRED / CLIMAFIN AND A UNIFIED APPROACH TO PD-BASED MODELLING 107

INSIGHTS FROM THE FIELD – ODACC 116

THE LONG ROAD TO CREDIT ASSESSMENT – S&P GLOBAL SUSTAINABLE1 117

S&P PHYSICALRISK AND CLIMATE RISKGAUGE 131

QUANT FOUNDRY – COMBINING MODELS TO ASSESS A COMPANY'S SURVIVAL CAPACITY 134

XDI (THE CROSS DEPENDENCY INITIATIVE): FROM ENGINEERING DAMAGE TO THE P&L BRIDGE 144

PLANETVIEW 146

CLIMATE VAR: A FIGHT FOR THE RIGHT DEFINITIONS 152

PROBABILITIES FOR SCENARIOS AND QUANTIFYING ANY TYPE OF ASSETS – THE SCIENTIFIC CLIMATE RATINGS AMBITION 183

EY CLIMATE ANALYTICS PLATFORM (EY CAP) 218

RISKTHINKING.AI 230

MORE HAZARDS MORE PROBLEMS: THE MUNICH RE APPROACH TO CLIMATE PHYSICAL RISK .. 238

MITIGA SOLUTION – EARTHSCAN AND BAYESIAN INFERENCE 242

THE CASE FOR FATHOM FLOOD MODELS 243

HYDROCLIMAT, A HAZARD-FOCUSED PHYSICS-BASED APPROACH 245

THE CASE OF CLIMAFIN MULTI-HAZARD APPROACH 253

IMAGE RECOGNITION AND FIRES: THE KAYRROS APPROACH 263

HYPERMETEO: A PRACTICAL USE OF DOWNSCALING 266

INTENSEL – AN ASIA-FOCUSED MODEL 267

SCREENING AT BASIN SCALE: THE WWF RISK FILTER SUITE 268

FIVE METERS AND TEN THOUSAND YEARS: JBA FLOOD MAPS 269

A DIGITAL TWIN OF EVERY BUILDING: FIRST STREET (MSCI) 270

CLIMATE RISK SERVICES – STRONG MODELS FOR ASSETS LOCATED IN DEVELOPING COUNTRIES 271
GAIA VISION – A JAPANESE-FOCUSED MODEL 272
SCIENTIFIC CLIMATE RATINGS (SCR) – A POLYGON APPROACH 272
MSCI: CLIMATE VAR, MACROECONOMIC DATA AND PARTNERSHIPS WITH OTHER SOLUTIONS . 293
FROM THE GROUND UP TO THE PORTFOLIO – THE CASE OF ALTITUDE 309
FROM THE GROUND UP TO THE PORTFOLIO – THE CASE OF CORRENTICS 310





CONTACT

www.parc-research.org
parc@institutlouisbachelier.org

Institut Louis Bachelier et al. (2026). Climate Related Financial Risks Cookbook: A Practical Guide to Comparing and Selecting Models That Assess the Financial Impacts of Climate Hazards.

Artificial intelligence tools were used in the preparation of this publication. All content has been reviewed and validated by the authors.

© 2026 PARC Foundation – Institut Louis Bachelier. All rights reserved.