

BEMPS –

**Bozen Economics & Management
Paper Series**

NO 121/ 2026

**From Macro Climate Risk
to Site-Level Corporate Loss:
A Geoclimatic Probable Maximum Loss Framework**

**Valter Carturo, Daniele Cortolezzis, Pietro Domenichini,
Davide Ferrari, Francesco Ravazzolo**

From Macro Climate Risk to Site-Level Corporate Loss: A Geoclimatic Probable Maximum Loss Framework

Valter Carturo¹ Daniele Cortolezzis² Pietro Domenichini³ Davide Ferrari² Francesco Ravazzolo^{2,4}

¹Carturo & Partners

²Free University of Bozen-Bolzano

³ISR Institute

⁴BI Norwegian Business School

August 2026

Abstract.

Macro climate and climate-finance research quantify aggregate damages, valuation effects, and financial propagation, but they provide limited guidance on how a macro climate state becomes an exposure at a particular operating site and within a particular business process. We develop Geoclimatic Probable Maximum Loss (GCPML), a transparent macro-to-micro framework that evaluates physical hazards at verified facility coordinates, separates site-specific physical risk from ex ante hazard-importance weights, transmits hazards through damage mechanisms to eleven business processes, mediates exposure through audit-based organisational resilience, and scales the result by production value. The terminology denotes a conditional management assessment under existing audited protections and no additional adaptation; the pilot does not estimate a loss quantile. Macro-financial evidence on climate-policy uncertainty describes how climate risk is priced and propagated; GCPML supplies the complementary upstream physical and organisational exposure that can subsequently be linked to cash flows, financing conditions, and valuations. To assess feasibility, we apply the framework to a pilot study of 17 audited sites. The pilot identifies material address-geocoding errors and shows that coordinate-only screening retains much of the intensity ranking (Spearman correlation 0.917) but reproduces the dominant process area at only five sites; the audit relocates it at the other twelve. The higher correlation for absolute exposure is partly mechanical because screening and full assessment share the same site production values. The exercise evaluates implementation, sensitivity, and diagnostic content, not predictive accuracy or actuarial calibration.

JEL classification: C43, D22, E44, G32, Q54.

Keywords: physical climate risk; climate finance; enterprise risk management; facility location; composite indicators; probable maximum loss; organisational resilience.

1. Introduction

Climate change is simultaneously a macroeconomic process, a spatially heterogeneous source of physical hazards, and a managerial problem inside firms. The macroeconomic literature estimates how temperature and other climate variables affect aggregate output, investment, productivity, and welfare. Seminal country-panel evidence finds large and persistent effects in poorer economies (Dell et al., 2012); later work documents nonlinear global production effects (Burke et al., 2015), long-run cross-country damages (Kahn et al., 2021), and an active debate over whether global temperature variation captures channels that country-level temperature misses (Bilal and Känzig, 2026). Yet a manager does not protect “the country” or “the representative firm.” The decision object is a production site, warehouse, hospital, office, data centre, or logistics node with a precise location, a specific organisation, and a specific amount of economic value at risk.

The finance literature shows that these real effects are accompanied by valuation and funding channels. Climate risks can alter expected cash flows, discount rates, risk premia, and financial stability (Dietz et al., 2016; Battiston et al., 2017; Giglio et al., 2021a). Carbon-intensive firms display distinctive return and valuation dynamics, climate news can be hedged in traded portfolios, and green and brown assets respond asymmetrically to transition risk (Engle et al., 2020; Bolton and Kacperczyk, 2021, 2023; Pástor et al., 2022; Ardia et al., 2023). More recently, Guidolin and Ionta (2026) identify climate-policy uncertainty shocks using innovations in the realised volatility of U.S. carbon futures. They find tighter financial conditions, lower industrial production and equity price–dividend ratios, and more persistent contractions in carbon-intensive sectors. Although this evidence establishes a macro-financial amplification channel, it does not by itself identify the facility, physical hazard, or organisational process from which firm-level exposure originates.

This paper develops a bridge from macro-level climate risk to site-level corporate exposure. We define *geoclimatic risk* as the exposure of a specific facility to physical climate hazards and the vulnerability of its business processes to those hazards. To operationalise this concept, we propose the Geoclimatic Probable Maximum Loss (GCPML) framework, which follows the three-part climate-risk logic of hazard, exposure, and vulnerability (Portner et al., 2022) and adapts it to enterprise risk management. Climate and hydrological fields characterize the hazards affecting each facility; verified facility coordinates and production value determine physical and economic exposure; and an audit of eleven business processes captures organisational vulnerability and resilience. GCPML integrates these components to produce process-level and site-level measures of monetary exposure.

Throughout, “probable” recognises that physical hazards and protective effectiveness remain uncertain conditional on the climate pathway and site information. “Maximum” refers to the no-additional-action reference relative to a specified set of risk-reducing interventions, with existing audited protections retained. This is an operational management convention: the pilot estimates neither a calibrated loss probability nor an absolute maximum over physical events. Engineering practice likewise requires explicit scenario and protection

assumptions (Mbatha and Nkuna, 2021); the precise convention adopted here is stated in Section 3.

Spatially resolved physical-risk analysis is already well established. In particular, Bressan et al. (2024) develop an asset-level methodology for productive assets and show that using headquarters as proxies for asset locations, or omitting acute hazard tails, can substantially understate financial risk. Our contribution is therefore more specific. GCPML links site-level climate hazards to the business processes through which physical shocks affect a facility, incorporates organisational vulnerability and resilience through a process audit, and translates the resulting exposure into monetary terms. This structure allows us to examine how much information is retained when assessment relies only on facility location and physical hazards, and how much additional information is provided by the organisational audit. The distinction is operationally important: identifying which sites warrant attention requires less information than identifying which processes within a facility drive exposure and should be targeted for intervention. GCPML should therefore be interpreted as a site-level measure of physical climate exposure rather than as an asset-pricing model. Whether this exposure is reflected in equity values, loan spreads, insurance prices, or collateral haircuts is a separate empirical question requiring external validation.

The proposed theoretical framework has four elements. First, a macro climate state and scenario generate geospatial hazard fields, that is, spatial distributions that assign a hazard value to each geographic location s . Each field is evaluated at the coordinate of facility s , where s indexes facilities, rather than represented by a country or headquarters average. This clarifies the role of geocoding error: if the address-derived location of facility s differs from its verified location by a displacement vector u_s , the hazard error is locally proportional to the spatial gradient of the field. A given coordinate error may have little effect when the hazard varies gradually across space, but a large effect when the facility lies near a sharp hazard boundary, such as the edge of a floodplain.

Second, the framework separates measurement from transmission. The normalized score $h_{js\tau}$ records how intense hazard j is at site s in climate window τ . The weight λ_j records ex ante relative hazard importance, using information distinguishable from the local physical assessment, while the two transmission matrices allocate weighted hazard pressure first to damage mechanisms and then to business processes. Because the matrices are row stochastic, they conserve the pre-assigned weight λ_j : they answer where the pressure goes, but they cannot determine how important one hazard is relative to another. This conservation identity makes the distinction between hazard importance and its allocation explicit.

Third, organisational quality enters through a bounded multiplier that increases with adverse audit findings and decreases with protective findings. The multiplier is process specific, because a firm may be resilient in information technology but fragile in logistics, or strong in business continuity but weak in supplier concentration. Economic value and a sector-process correction then translate process pressure into a monetary quantity. The

full mode uses verified coordinates and the audit; a screening mode sets the organisational multiplier to one and is designed for rapid portfolio triage.

Fourth, the framework separates aggregation from reporting. Aggregation can conceal important differences across sites: an improvement at one site can offset deterioration at others. At the portfolio level, a large site that improves can therefore mask deterioration at most sites. Changing the aggregation rule does not eliminate this problem. It is therefore necessary to report the portfolio aggregate together with the median site change, the dominant component, and its value share.

The case study is a feasibility and diagnostic assessment of the proposed framework using 17 audited sites in 13 countries and five broad NACE sections, with total annual production value of EUR 1.128 billion (STAIRS Srl, 2026). It asks three bounded questions: whether the model can be implemented end to end; which conclusions depend on verified locations and process-level organisational information; and how the constructed outputs respond to declared modelling perturbations. It does not test predictive accuracy, causal effects, or actuarial calibration. Six of a 17-hazard operational catalogue are populated using river-discharge, precipitation, wind, snow, heat, and dry-spell data for current and 2050 windows. The pilot identifies six address-derived locations more than one kilometre from the verified facility, four of which are labelled “high precision” by the service. The current partial-coverage output is EUR 331.6 million, or 29.4 per cent of production value, but it is a model-based exposure measure rather than a loss quantile. Coordinate-only screening has a Spearman correlation of 0.917 with the full assessment for site intensity and 0.983 for absolute exposure; the latter is partly stabilised mechanically by the production-value scale shared by both modes. The dominant physical hazard is identical by construction and is therefore not a screening-performance statistic. The dominant process area agrees at only five sites, so the audit relocates it at twelve. Finally, the 2050 portfolio aggregate rises only 2.1 per cent although the median site rises 11.2 per cent and the worst site rises 35.1 per cent, because one site representing 39 per cent of production value improves by 9.3 per cent. The exercise therefore illustrates masking, not sign reversal, under a fixed 2026 portfolio.

These results should be interpreted conservatively. The sample demonstrates the feasibility and internal behaviour of the framework, rather than its predictive validity or population-level accuracy. Nine of the 17 sites belong to one industrial group. Only six hazards are populated, corresponding to 63 per cent of the operational hazard weights. One climate pathway is used. The reference population for normalization is the sample itself. No realised loss claims are available, so the paper does not estimate predictive validity, a return period, or an actuarial PML. The comparison between screening and full assessment is an internal ablation exercise, not a comparison with observed damages. Those limitations are central to the proposed research agenda rather than relegated to a final disclaimer.

The paper is organised as follows. Section 2 places GCPML in the macro-climate, climate-finance, firm, asset-level, production-network, and composite-indicator literatures. Section 3 develops the model. Section 4 describes the case-study design and reports the

results. Section 5 draws managerial and financial implications. Section 6 concludes. Proofs and additional implementation details appear in the Appendices.

2. Related literature and contribution

A large empirical literature studies the relationship between weather, climate, and aggregate economic activity. Dell et al. (2012) use within-country temperature fluctuations and report strong growth effects in poorer economies. Burke et al. (2015) estimate a nonlinear temperature-production relationship, while Kahn et al. (2021) analyse long-run macroeconomic effects in a cross-country setting. Reviews by Carleton and Hsiang (2016) and Auffhammer (2018) emphasize both the breadth of documented channels and the sensitivity of damage estimates to functional form, adaptation, spatial scale, and extrapolation.

The distinction between global and local climate signals is particularly relevant here. Bilal and Känzig (2026) argue that global temperature shocks predict a broader bundle of extremes than local temperature shocks and therefore generate much larger estimated output effects. Whatever the final magnitude of that debate, it reinforces a conceptual point: a macro climate state and a local facility hazard are not substitutes. Global climate conditions generate spatial fields and correlated extremes; local coordinates determine which part of those fields reaches an asset. A macro-to-micro corporate framework therefore requires both levels rather than a choice between them.

The GCPML framework does not estimate a macro damage function. It takes scenario-conditioned climate and hydrological fields as external inputs and asks how they reach a firm. This places it downstream of macro climate econometrics. Its purpose is decision translation: to turn external physical pressure into an auditable map of organisational exposure and a monetary management quantity.

The climate-finance literature asks how physical and transition risks enter asset values, risk premia, and systemic vulnerability. Barnett et al. (2020) develop an asset-pricing treatment of climate uncertainty. Dietz et al. (2016) quantify a climate value-at-risk concept for global financial assets, while Battiston et al. (2017) trace climate-sensitive exposures through the financial system. Giglio et al. (2021a) review the broader climate-finance agenda. At the asset level, sea-level exposure is capitalised into property prices (Bernstein et al., 2019), and long-dated real-estate transactions reveal how markets discount climate damages that materialise far in the future (Giglio et al., 2021b).

A closely related asset-pricing strand studies transition risk and investor beliefs. Bolton and Kacperczyk (2021) and Bolton and Kacperczyk (2023) document the pricing of corporate carbon emissions in U.S. and global equity markets. Engle et al. (2020) construct portfolios that hedge innovations in climate news, Pástor et al. (2022) decompose green returns, and Ardia et al. (2023) study the relative performance of green and brown stocks as climate concerns vary. Together these papers show that climate information affects both expected cash flows and discount rates, but their typical unit is a security, firm, sector, or financial network rather than a verified facility and its internal processes.

Climate-policy uncertainty provides the most direct bridge between macroeconomics and

finance. Newspaper-based measures were introduced by Gavriilidis (2021) and extended in a structural macroeconomic setting by Gavriilidis et al. (2026); related evidence links environmental-policy uncertainty to lower low-carbon investment and greater financial volatility (Noailly et al., 2022; Palikhe et al., 2024). Guidolin and Ionta (2026) complement textual indicators with a market-based proxy constructed from innovations in the realised volatility of ICE U.S. carbon futures. Their Bayesian Proxy-SVAR shows that transition uncertainty raises broad equity volatility, depresses output and price–dividend ratios, and produces a green–brown gradient whose severity increases with sectoral emission intensity. This contribution and GCPML address different risk families. Guidolin and Ionta identify transition-policy uncertainty and its macro-financial propagation; GCPML identifies physical hazards at a facility and their organisational transmission. A joint empirical design could therefore test whether physical exposure and transition uncertainty have distinct or interacting effects on cash flows, financing conditions, and valuations.

Firm-level evidence is more heterogeneous than aggregate evidence. Using U.S. establishment data, Addoum et al. (2020) find little robust effect of temperature exposure on sales or productivity. In a global sample, Pankratz et al. (2023) find that extreme heat reduces revenues and operating income and that investors do not fully anticipate those effects. These results need not conflict: firms differ by climate, technology, sector, adaptation, network structure, and the location of their productive assets.

A second strand measures firms’ perceived or disclosed climate exposure. Sautner et al. (2023) infer physical, regulatory, and opportunity exposure from earnings calls. Li et al. (2024) construct textual measures of physical and transition risk and document heterogeneous corporate responses. Li (2025) links forecast physical exposure to adaptation strategies and shows that organisational capabilities and time horizons shape whether firms act. This literature makes organisational heterogeneity central, but textual measures cannot by themselves identify the exact facility, physical hazard, or process through which a shock propagates.

GCPML is complementary. It starts from observed or projected physical fields rather than management language, and it makes organisational mediation an input rather than an inferred response. In future work, the two approaches can be combined: textual measures may provide scalable priors or external validation for audit-based resilience, while the site model can provide physically grounded outcome variables.

The closest methodological precursor is Bressan et al. (2024). Their asset-level framework combines geolocated productive assets, acute and chronic hazards, a macroeconomic model, and financial valuation. They show that using headquarters as a proxy for asset location can materially understate value at risk. Related work maps floods, storms, and wildfires at production facilities into globally held securities and highlights cross-border financial transmission (Mandel et al., 2025). The present paper therefore does not claim to be the first macro-to-asset physical-risk model. Its incremental contribution is the process-level organisational layer and the empirical error budget for address-to-building location in an audited industrial sample.

Spatial precision is not a cosmetic issue. Physical-risk vendors can produce divergent scores even when they purport to measure the same concept (Hain et al., 2022). A separate geocoding literature shows that street interpolation, parcel centroids, and address points have different positional accuracy and can create exposure misclassification (Cayo and Talbot, 2003; Zandbergen, 2008; Kinnee et al., 2020). GCPML brings this measurement problem into corporate climate risk. It distinguishes the address used to locate a counterparty from the verified coordinate of the facility and treats location confidence as part of the model's error budget.

The macro-to-micro link also connects to production-network research. Aggregate fluctuations can originate in and propagate through input-output structures (Acemoglu et al., 2012). Firm-specific shocks travel through supplier relationships when inputs are difficult to substitute (Barrot and Sauvagnat, 2016); natural disasters provide direct evidence of propagation through firm linkages (Boehm et al., 2019; Carvalho et al., 2021). GCPML uses an internal rather than interfirm network: climate hazards reach damage mechanisms, which reach business processes, whose vulnerability depends on organisational design. Supply-chain and logistics processes can then connect the internal map back to external production networks.

The immediate methodological precedent is the Geopolitical Probable Maximum Loss framework (Domenichini and Ravazzolo, 2026). GPML connects macro geopolitical indicators to eleven corporate process areas, an organisational audit, sectoral corrections, and production value. The climate adaptation replaces the country-level geopolitical pressure with site-resolved physical hazards while retaining the common process architecture. That common architecture is useful for enterprise risk management, but the risk families should not be mechanically summed: geopolitical, legal, and geoclimatic assessments have distinct risk perimeters, scenario assumptions and potentially overlapping economic exposures.

GCPML is also a composite indicator. The composite-indicator literature stresses that normalization, weighting, aggregation, and sensitivity analysis are substantive modelling choices rather than presentation details (OECD and European Commission Joint Research Centre, 2008; Greco et al., 2019). Linear and generalized-mean aggregation impose compensability, meaning that strength in one component can offset weakness in another. Noncompensatory approaches make a different normative choice and often increase complexity (Munda and Nardo, 2009).

The composite-indicator literature implies three requirements that are used in the model below. First, any cross-hazard consequence weights must be specified as external inputs and cannot be inferred from a row-normalized allocation matrix. Second, sensitivity analysis must vary the normalization rule, the consequence weights, and the organisational multiplier jointly, because each is a substantive modelling choice. Third, an aggregate must be reported together with the distribution and the component that dominates it. These are *ex ante* design requirements derived from the literature; the case study later evaluates their practical relevance rather than serving as their justification.

Table 1: Positioning the GCPML contribution

Literature strand	Main unit of analysis	Main outcome	Remaining interface addressed by GCPML
Macro climate econometrics	Country, region, or world economy	Output, growth, welfare	Does not identify the facility, process, or organisational intervention.
Climate finance and policy uncertainty	Security, firm, sector, or financial system	Returns, valuations, volatility, financing, and stress losses	Prices and propagates climate risk but rarely resolves verified facility coordinates and internal process resilience.
Firm exposure and disclosure	Listed firm or establishment	Performance, valuation, disclosed risk	Often lacks verified asset coordinates and process-level transmission.
Asset-level physical risk	Geolocated productive asset and portfolio	Damage, value at risk, financial valuation	Closest precursor; typically no audit-based map across internal business processes.
Production networks	Firm and supplier network	Propagation of shocks	Explains external linkages but not a facility's internal resilience architecture.
Composite indicators	Multidimensional index	Ranking and benchmarking	Provides the weighting, aggregation, and robustness discipline needed for GCPML.
GCPML	Verified facility, process, and portfolio	Model-based monetary exposure	Integrates site hazard, process transmission, organisational mediation, and economic value.

Notes: "Remaining interface" describes the paper's intended incremental contribution; it is not a claim that prior literatures contain no related elements.

3. A macro-to-micro model of geoclimatic risk

Let $s = 1, \dots, S$ index facilities, $p = 1, \dots, P$ business processes, $j = 1, \dots, J$ physical hazards, $c = 1, \dots, C$ damage mechanisms, and $\tau = 1, \dots, T$ index climate windows or scenario horizons. The application uses $P = 11$ and $J = 17$ physical hazards in the operational catalogue, of which six are populated with data. Let $\ell_s \in \mathbb{R}^2$ denote the verified geographic location of facility s and C_τ the macro climate state or scenario associated with window τ .

For hazard j , let $X_j(\ell, \tau; C_\tau)$ be a geospatial physical field. The site-level physical quantity is

$$x_{js\tau} = X_j(\ell_s, \tau; C_\tau). \quad (1)$$

The same field can be decomposed conceptually into a broad regional component and a local deviation:

$$x_{js\tau} = \bar{x}_{j,r(s),\tau} + \delta_{js\tau}, \quad (2)$$

where $r(s)$ is a country or region. A country-level assessment sets $\delta_{js\tau} = 0$ for all facilities in the region. GCPML retains the deviation. Equation (2) is a measurement decomposition, not a statistical identifying equation in the current case.

The term “probable maximum loss” has both probabilistic and scenario-based uses in catastrophe and industrial-risk practice (Grossi and Kunreuther, 2005). For engineering insurance, Mbatha and Nkuna (2021, pp. 7, 27, 29–30) describe a scenario assessment with explicit assumptions about protections, warn that protective effectiveness may be overestimated, and require a documented monetary-loss calculation for the selected adverse scenario. This provides a practical precedent for making the assessment assumptions explicit; it does not validate the current GCPML calculation or prescribe the intervention-based convention used here.

In GCPML, “probable” recognises that the physical-risk inputs describe uncertain conditions, conditional on the stated climate pathway and the information available at the facility. The scores may summarise observed frequencies, severity statistics, climate-model projections or expert assessments, depending on the underlying hazard variable; these ingredients and their transformations must be documented hazard by hazard. A reproducible calculation at fixed inputs does not make the physical outcome certain. Conversely, uncertainty in an input does not establish that its normalized score is an event probability or that the final monetary output is a loss quantile. The current six-band scores have the operational interpretation in (5), not a calibrated probability of loss. “Maximum” is defined relative to additional risk-reducing interventions in (17), with existing audited protections retained. Actuarial interpretation would additionally require a mapping to engineering damage and a loss distribution with a stated probability or return-period convention.

A public address produces an estimated coordinate $\tilde{\ell}_s = \ell_s + u_s$, where ℓ_s is the verified facility coordinate and $u_s = \tilde{\ell}_s - \ell_s$ is positional error. The measured physical quantity is $\tilde{x}_{js\tau} = X_j(\tilde{\ell}_s, \tau; C_\tau)$, compared with $x_{js\tau} = X_j(\ell_s, \tau; C_\tau)$ at the verified location. If the

field is differentiable near the facility, a first-order expansion gives

$$\tilde{x}_{jst} - x_{jst} \approx \nabla_{\ell} X_j(\ell_s, \tau; C_{\tau})' u_s. \quad (3)$$

Equation (3) is a local approximation: its usefulness requires a hazard surface that is sufficiently smooth over the displacement u_s and a displacement that is small relative to the field's scale of variation. It can fail when the two coordinates cross a floodplain or coastline boundary, an abrupt elevation change, a categorical-map boundary, or the edge of a coarse grid cell. The preferred measure of location-induced hazard error is therefore the direct two-coordinate difference,

$$e_{jst}^{\text{loc}} = \tilde{x}_{jst} - x_{jst} = X_j(\tilde{\ell}_s, \tau; C_{\tau}) - X_j(\ell_s, \tau; C_{\tau}). \quad (4)$$

The positional displacement u_s is common to all hazards at a facility, but e_{jst}^{loc} is hazard-specific because each field has its own spatial structure. A large coordinate error need not produce a large hazard or exposure error, while a small displacement may matter near a sharp boundary. The distance $\|u_s\|$ is thus a location-quality diagnostic, not by itself a measure of physical-risk error. The effect on hazard scores and monetary exposure must be assessed through the corresponding physical fields and the subsequent model transformations.

In the 17-site pilot, the positional displacements are observed, but every hazard has not yet been evaluated at both the address-derived and verified coordinates. The evidence therefore quantifies coordinate unreliability, not the resulting distribution of errors in h_{jst} or GCPML. Verified coordinates are used throughout the assessment; the full two-coordinate recomputation remains a required validation step.

Raw hazards differ in units, support, and meaning. Let $B_{j\Omega}(x)$ be a six-level banding function defined relative to a reference population Ω . The operational method extends a five-band z-score rule by adding level zero for verified physical absence of the hazard. Define

$$h_{jst} = \frac{B_{j\Omega}(x_{jst})}{sc} \in [0, 1]. \quad (5)$$

The division by sc is a simple rescaling that preserves the operational ordering.¹ The zero level matters. Without presence screening, a site at which a hazard cannot physically occur may receive the same low score as a site at which it occurs rarely. The distinction is especially important when normalized indicators are later multiplied by value. Presence screening also makes the missing-hazard problem explicit: unpopulated hazards are not treated as absent.

The operational normalization used in the pilot is deliberately simple but has two limitations. The six-level score in (5) is ordinal, yet is treated as cardinal when it enters arithmetic or generalized means. It discards within-band variation, introduces discontinuities at band thresholds, and implicitly treats adjacent score increments as

¹We use $sc = 5$ in the case-study.

comparable in aggregation. These features aid screening and communication but require sensitivity assessment when the score scales a monetary exposure. The second limitation is benchmark dependence: changing the population used to construct the bands can change an otherwise unchanged site's score. The reported movement of some levels by up to 22 per cent illustrates why the normalization benchmark must be declared.

For longitudinal and cross-portfolio assessment, a fixed external reference population is preferable to a benchmark re-estimated from the current portfolio. A broad climatological population or another pre-defined external universe can provide such a reference; the audited sample and a sector-country peer group may also be reported for relative comparison. For each hazard, the application should document the source and composition of Ω , any spatial weights, the physical variable and temporal statistic, units, calibration period, estimation rule and version. The variable definitions must be consistent across assessment windows. A fixed-benchmark comparison holds its fitted reference distribution constant across the current and future windows, rather than re-estimating it as the portfolio or climate window changes.

A natural extension is a percentile normalization based on a reference cumulative distribution function $F_{j,\Omega}$. To distinguish it from the banded score actually used in the pilot, write h_{jst}^{PIT} . When larger values of the physical variable imply greater hazard, define

$$h_{jst}^{\text{PIT}} = F_{j,\Omega}(x_{jst}); \quad (6)$$

when lower values imply greater hazard, define

$$h_{jst}^{\text{PIT}} = 1 - F_{j,\Omega}(x_{jst}). \quad (7)$$

With a continuous reference distribution, this provides a continuous percentile scale without the six-band discretization. More generally, the transformation is monotone and retains distinctions wherever the fitted distribution varies. An empirical distribution can still have steps and ties, and scores may saturate at zero or one outside its observed support. The estimation and interpolation rules, treatment of ties and tail behaviour must therefore be disclosed; information preservation is not automatic for every estimator.

Verified physical absence should retain a separate presence flag, with the score set to zero when absence is established. The flag distinguishes physical impossibility from a very low percentile of an otherwise feasible hazard, since a percentile of zero alone does not identify absence. An unobserved physical input remains missing: its omission from the partial-coverage calculation does not establish that the hazard is physically absent.

The percentile interpretation also needs an explicit boundary. A value $h_{jst}^{\text{PIT}} = 0.95$ places the raw physical measure at approximately the ninety-fifth percentile of the chosen reference distribution when larger values imply greater hazard, or approximately the fifth percentile when lower values imply greater hazard. It does not imply a 95 per cent probability of a damaging event or economic loss. A common percentile scale does not make equal percentile increments economically equivalent across hazards. Mapping physical risk to

engineering damage, interruption or a calibrated loss distribution remains a separate task.

The present case retains the six-level banded specification to reproduce the operational model, and every reported numerical result remains conditional on that implementation and the sample-based benchmark. The fixed-reference percentile specification is a preferred extension to compare with the pilot through recomputation and external validation; it has not been implemented in the reported case. The alternative-population outputs and normalization sensitivities still required are set out in Section 4.

Let $\lambda_j \geq 0$ denote the declared ex ante relative risk importance of hazard j , normalized so that $\sum_j \lambda_j = 1$. Let D_{jc} be the share of hazard j transmitted through damage mechanism c , and let I_{cp} be the share of mechanism c reaching process p . Assume both maps are row stochastic:

$$\sum_c D_{jc} = 1 \quad \text{for each } j, \quad \sum_p I_{cp} = 1 \quad \text{for each } c. \quad (8)$$

The hazard-process transmission matrix is

$$F_{jp} = \lambda_j \sum_{c=1}^C D_{jc} I_{cp}. \quad (9)$$

The decomposition separates two expert questions: what damage a hazard causes and which organisational processes that damage affects. It is more auditable than assigning every hazard-process cell directly.

The score h_{jst} captures the site- and scenario-specific physical expression of hazard j : where the hazard is present and its assessed intensity at facility s in climate window τ . The coefficient λ_j captures relative risk importance across hazard classes and may reflect conditional economic severity, general likelihood, historical losses, engineering evidence or structured expert judgement. The distinction concerns the information represented, not a division into a deterministic site component and a probabilistic scenario component. Both physical variability and imperfect knowledge may remain after location verification and auditing.

The matrices D and I contain conditional allocation shares: given weighted hazard pressure, they determine which damage mechanisms and business processes receive it. Row normalization removes absolute magnitude from these maps; hazard-class importance enters through λ_j , while site and scenario heterogeneity enter through h_{jst} . These objects meet in (13). Information used to construct λ_j must be distinguishable from the local physical information already represented in h_{jst} ; statistical correlation between the two is neither required nor sufficient to diagnose duplicated inputs.

Weight conservation. Under (8), the row sum of the transmission matrix equals the ex ante hazard weight:

$$\sum_{p=1}^P F_{jp} = \lambda_j. \quad (10)$$

Consequently, if an aggregate hazard weight is reconstructed from row sums,

$$\omega_j = \frac{\sum_p F_{jp}}{\sum_r \sum_p F_{rp}}, \quad (11)$$

then $\omega_j = \lambda_j$ whenever $\sum_r \lambda_r = 1$.

The derivation is retained in Appendix A. Economically, F reallocates the declared hazard weights across processes; it does not identify their relative importance independently. Expert or empirical hazard weights must be justified separately. In the application, a numerical check of the identity yields a residual of 2.8×10^{-17} (STAIRS Srl, 2026).

The empirical dimensions and evidence classes of D , I , and F are therefore documented in Section 4, together with the other case-study inputs, rather than in the theoretical definition.

For each process, normalize the incoming transmission weights:

$$\phi_{jp} = \frac{F_{jp}}{\sum_{r=1}^J F_{rp}}, \quad \sum_j \phi_{jp} = 1, \quad (12)$$

when the denominator is positive. Process-level physical pressure is a weighted generalized mean,

$$H_{sp\tau}^{(q)} = \left(\sum_{j=1}^J \phi_{jp} h_{j\tau}^q \right)^{1/q}, \quad q \geq 1. \quad (13)$$

The operational default is $q = 1$, the weighted arithmetic mean across hazards within a site–process cell. This supplies an interpretable composite of the hazards under declared weights instead of making the largest score the sole determinant. The exponent and the weights are separate choices: changing q changes compensability, whereas changing ϕ_{jp} changes relative contributions. Higher q emphasises large scores but does not eliminate compensability at finite values. The arithmetic specification is retained as a transparent operational baseline and compared with alternatives in the case study. Averaging may reduce the influence of an extreme component relative to a maximum operator, but it does not establish that the resulting euro amount is a realistic loss estimate or an upper-tail quantity. This aggregation choice is distinct from the no-additional-action interpretation of “maximum” in (17).

To avoid double counting, the elicitation protocol must identify the information entering each λ_j and any overlap with the construction of $h_{j\tau}$. Hazard-class likelihood or uncertainty can legitimately inform relative importance without duplicating a particular site’s flood intensity or other local physical input. The available material does not disclose the individual weights or their elicitation protocol, so the empirical results remain conditional on that external vector. The specific disclosures and computations still required are set out in the case-study discussion.

For site s and process p , let $a_{sp} \geq 0$ summarize audit findings that aggravate loss and

$b_{sp} \geq 0$ summarize protective findings. The organisational multiplier is

$$M_{sp} = \text{clip}_{[0,2]} \left(1 + \frac{a_{sp} - b_{sp}}{a_{sp} + b_{sp}} \right), \quad (14)$$

with $M_{sp} = 1$ by convention when $a_{sp} + b_{sp} = 0$. The multiplier is neutral at one, increases with aggravating evidence, and decreases with protective evidence. The bound prevents a single audit layer from producing negative loss or unbounded amplification.

The eleven process areas inherited from GPML are management responsibility, finance, supply chain, sales and marketing, anti-bribery, human resources, health and safety, environment, privacy, logistics, and information technology (Domenichini and Ravazzolo, 2026). Not every climate hazard materially reaches every process. That selectivity is expressed by F_{jp} , while M_{sp} describes how well the receiving process absorbs the pressure.

The audit multiplier is a structured expert measure. Its defensibility depends on documented questions, scoring anchors, inter-auditor reliability, and version control (Aven, 2016). Public management-system certifications are attractive as scalable proxies, but the case shows that their process-level variation is too small to reproduce the audit signal.

An audit establishes evidence about a protection, not certainty about its performance during a future event. For example, a flood wall may be observed while its effective capacity, maintenance condition or adequacy for a future flood remains uncertain. The audit can reduce this information uncertainty without eliminating it. Its treatment must follow the model boundary: if $h_{js\tau}$ measures physical conditions outside the protection, protective effectiveness can enter the vulnerability assessment; if $h_{js\tau}$ already measures residual hazard behind the protection, the same protective effect must not be credited again through M_{sp} . Distinct residual uncertainties may remain in both layers, but the same reduction in exposure must be counted only once.

Let V_s be annual production value at site s , $n(s)$ its NACE or ISIC sector, and $k_{n(s),p}$ a sector-process correction. With $P = 11$, organisational elasticity $\gamma > 0$, and non-negative economic coefficients $1 + k_{n(s),p} \geq 0$, define process-level GCPML as

$$\text{GCPML}_{sp\tau} = \frac{V_s}{P} H_{sp\tau}^{(q)} (1 + k_{n(s),p}) M_{sp}^\gamma. \quad (15)$$

The case sets $\gamma = 1$. Site and portfolio totals are

$$\text{GCPML}_{s\tau} = \sum_{p=1}^P \text{GCPML}_{sp\tau}, \quad \text{GCPML}_{\tau}^{\text{port}} = \sum_{s=1}^S \text{GCPML}_{s\tau}. \quad (16)$$

The no-additional-action reference retains the protections represented by the baseline audit; it does not remove existing defences or equate them with the neutral screening multiplier $M_{sp} = 1$. Let $\alpha = 0$ denote this reference, and let $\mathcal{A}_s$ contain 0 and only additional interventions represented through the organisational layer such that $M_{sp}(\alpha) \leq M_{sp}(0)$ for every process. Holding hazard scores, transmission weights, production value and sector

corrections fixed, the non-negative multiplicative structure and $\gamma > 0$ imply

$$\text{GCPML}_{s\tau}(\alpha) \leq \text{GCPML}_{s\tau}(0) = \max_{\beta \in \mathcal{A}_s} \text{GCPML}_{s\tau}(\beta), \quad \alpha \in \mathcal{A}_s. \quad (17)$$

This is the precise management interpretation of “maximum” used in this paper: an upper reference within a stated set of risk-reducing interventions, under fixed physical conditions. It is a conditional property of the model, not an estimated treatment effect or a maximum across possible physical-event realizations. The pilot does not quantify intervention effectiveness or identify a worst plausible event-loss scenario. Interventions that change residual physical hazard require re-evaluating that layer as well and are outside the fixed-hazard comparison in (17). The economic scale factor V_s requires an explicit accounting definition, because the term “production value” does not denote the same aggregate across reporting frameworks. Under Italian GAAP it corresponds to heading A) of the income statement prescribed by Article 2425 of the Civil Code – revenue from sales and services, changes in inventories of work in progress and finished goods, own work capitalised, and other income (Repubblica Italiana, 1942; Organismo Italiano di Contabilità, 2024). Under IFRS the closest construct is revenue from contracts with customers (IFRS Foundation, 2014), which excludes inventory movements and capitalised internal work and is therefore not equivalent. Because the case spans 13 countries and five NACE sections, cross-site comparability depends on this choice rather than on the climate layer. In this paper V_s is measured under local GAAP (Generally Accepted Accounting Principles) for the correspondent financial year. Moreover, the site values are reported gross of intra-group transactions. Then, the aggregate economic value at risk exceeds the value that would be recognised on consolidation, where intercompany revenue is eliminated.²

Uniform allocation of V_s/P replaces an earlier circular specification that allocated value using the model’s own damage shares. Uniform allocation is deliberately simple, and its consequence should be stated rather than merely acknowledged: because the value base is identical across processes by construction, the identity of the dominant process area is determined entirely by $H_{s\tau}^{(q)}$, $k_{n(s),p}$, and M_{sp} . The alternatives are not equivalent, and an application should treat them as an ordered hierarchy rather than a list. The preferred basis is the contribution margin attributable to the process plus the fixed costs that cannot be avoided over the expected interruption period, allocated by cost or responsibility centre: this is the quantity that business-interruption cover actually indemnifies, and it is the only candidate that aligns the denominator of the model with the economic loss it is meant to approximate. Where analytical accounting by cost centre is unavailable, operating-segment disclosure offers a coarser but externally verifiable substitute (IFRS Foundation, 2006). Replacement cost is appropriate for asset-destruction hazards but overstates exposure for processes whose loss is primarily an interruption of flow. Gross profit is the weakest of the


²The alternative consolidated values, net of intra-group transactions, would reduce the denominator of the intensity ratio and, to the extent that internal flows are concentrated in particular facilities, would also change relative site weights.

four, since it mixes avoidable and unavoidable cost structures. Whichever basis is adopted, the allocation must remain exogenous to the model that consumes it.

Equation (15) is multiplicative because none of the three layers alone is decision sufficient. A dangerous site with low economic value and strong organisation may have modest exposure. A moderate hazard at a high-value, fragile site may dominate. This is the enterprise analogue of the IPCC interaction among hazard, exposure, and vulnerability.

Macro-to-micro geoclimatic transmission architecture

Climate scenario information is spatially resolved, transmitted to corporate processes, mediated by organisation, and scaled by economic value.



Screening mode omits the audit multiplier. Full assessment combines verified coordinates, process audit, and production value.

Figure 1: Macro-to-micro transmission architecture. The equation implemented in the final block is (15). Climate fields are evaluated at a verified facility coordinate, normalized, transmitted to processes, mediated by the process audit, and scaled by production value.

The full assessment uses verified coordinates, M_{sp} from an on-site audit, and the sector-process correction. The screening mode uses the physical layer at the available coordinate and sets

$$M_{sp} = 1, \quad k_{n(s),p} = k_p^0, \quad (18)$$

where k_p^0 is a declared default. Screening is intended for large portfolios in which the first question is which sites require attention. Full assessment is intended for industrial management, where the question is which process should be changed and how much loss exposure the change may avoid.

The difference between the modes is an ablation of organisational information. It measures the value of that layer inside the model; it does not establish which mode is closer to realised losses. External validation requires observed damage or business-interruption outcomes.

Aggregation, weights, and reporting

The power mean in (13) is often used to reduce compensation by increasing q . This is useful but should not be misinterpreted.

Compensability of hazard aggregation. For positive components z_i and weights w_i

summing to one, write

$$A_q(z) = \left(\sum_i w_i z_i^q \right)^{1/q} \quad (19)$$

For finite $q \geq 1$, each component has a positive marginal effect,

$$\frac{\partial A_q}{\partial z_i} = w_i z_i^{q-1} A_q^{1-q} > 0, \quad (20)$$

and the marginal rate of substitution between positive components is finite. Changes can therefore compensate within any finite power mean. As q increases without bound, A_q approaches $\max_i z_i$ and retains only the largest component, changing the estimand from an aggregate to the worst component. Appendix A retains the differentiation and limiting argument.

The positive-component qualification matters because presence screening can assign a zero hazard score. When the remaining components generate positive aggregate pressure, the marginal derivative at a zero component is zero for $q > 1$; this statement does not cover the all-zero boundary or imply that a finite change has no effect. Physical absence must therefore be established from the hazard evidence rather than inferred from this derivative property, and remains distinct from a missing input.

The operational sweep from the default $q = 1$ up to $q = 20$ confirms the practical implication in a monodominant case: increasing the exponent changes emphasis but does not restore the desired ordering before the measure approaches the maximum (STAIRS Srl, 2026). That exercise is evidence for the case, not a universal calibration of q . Raising the exponent to $q = 2$ lifts the reported portfolio total by 12.9 per cent, from EUR 331.6 million to EUR 374.4 million, and aggregate intensity from 29.4 to 33.2 per cent, while leaving the substantive conclusions intact: the rank correlation between the two exponents is 0.998 for absolute exposure and 0.956 for intensity, and the dominant process area coincides at sixteen of seventeen sites. The level of the monetary output is therefore sensitive to the exponent; the ordering and the dominance structure are not.

Define site intensity $r_{s\tau} = \text{GCPML}_{s\tau}/V_s$ and value share $v_s = V_s/\sum_i V_i$. Portfolio intensity is

$$\bar{r}_\tau = \sum_{s=1}^S v_s r_{s\tau}. \quad (21)$$

This mean across site intensities is distinct from the mean across hazards in (13). It has a direct accounting interpretation: $\bar{r}_\tau = \text{GCPML}_\tau^{\text{port}}/\sum_s V_s$, because the monetary portfolio output is the sum of site outputs. The mean therefore supplies a clear value-normalized portfolio statistic without selecting the worst site as the portfolio's representative. It can nevertheless mask concentrated site risks, and should be accompanied by the median, dispersion and tail diagnostics below. Neither averaging site assessments nor summing their monetary outputs establishes a jointly realisable maximum portfolio loss; that would require an event specification and the dependence across sites. The statistical summary

and the management convention for “maximum” answer different questions. Comparisons across climate windows are normally expressed in relative terms. Let $g_s = \Delta r_s / r_s$ be the relative change at site s and $\bar{g} = \Delta \bar{r} / \bar{r}$ the relative change of the aggregate. Since $v_s r_s = \text{GCPML}_s / \sum_i V_i$, equation (21) implies

$$\bar{g} = \sum_{s=1}^S \theta_s g_s, \quad \theta_s = \frac{\text{GCPML}_s}{\sum_i \text{GCPML}_i}. \quad (22)$$

The weights that govern the aggregate relative change are exposure shares, not value shares. The two coincide only when site intensities are equal and are close whenever intensities lie in a narrow band, but the distinction should be stated because it determines which share a report must disclose.

Two diagnostics distinguish an aggregate below the typical site from an aggregate that reverses the direction of change when the typical site deteriorates.

Aggregate masking. Let g^{med} be the median site change. The aggregate understates the typical site, $\bar{g} < g^{\text{med}}$, if and only if

$$\sum_s \theta_s (g_s - g^{\text{med}}) < 0, \quad (23)$$

The masking gap $\Gamma = g^{\text{med}} - \bar{g}$ measures the difference directly. This is an identity following from the exposure-share representation; its derivation is retained in Appendix A.

Aggregate reversal. If one dominant site improves, $g_d < 0$, while every other site deteriorates, $g_i > 0$ for $i \neq d$, the aggregate improves, $\bar{g} < 0$, precisely when

$$\theta_d |g_d| > \sum_{i \neq d} \theta_i g_i. \quad (24)$$

Equivalently, $\bar{g} < 0$ if and only if $\Gamma > g^{\text{med}}$. In the setting considered here, where the median site change is positive, this makes reversal a stronger outcome than masking. Appendix A retains the full algebraic argument.

Both diagnostics depend on the joint configuration of exposure shares and site changes. Unequal shares can amplify the difference between the aggregate and the typical site, but masking can also occur with equal shares, and reversal does not require a particular concentration threshold: its condition is the contribution comparison in (24). Reporting the aggregate, median and masking gap makes these configurations visible. A similar dominance structure appears in generalized means. Differentiating the log of A_q gives

$$d \log A_q = \sum_i \pi_i(q) d \log z_i, \quad \pi_i(q) = \frac{w_i z_i^q}{\sum_r w_r z_r^q}. \quad (25)$$

The effective contribution of component i depends on both its declared weight and its level. Increasing q can intensify dominance rather than solve it.

The model therefore adopts a reporting rule: no aggregate indicator should be published alone. At minimum, a portfolio report should include

- (i) total monetary GCPML and intensity relative to production value;
- (ii) median and interquartile site intensity or change;
- (iii) the dominant site, hazard, and process, with their value or contribution shares;
- (iv) the tail of the site distribution, including the worst observed site;
- (v) the masking gap Γ between the median site change and the aggregate change;
- (vi) sensitivity to normalization population, hazard weights, and organisational multipliers.

This rule preserves the convenience of a single number while exposing the component that can reverse its interpretation. It is consistent with the composite-indicator literature's emphasis on transparency and robustness (OECD and European Commission Joint Research Centre, 2008; Greco et al., 2019).

One qualification applies to item (v). The exposure shares θ_s in (22) are not exogenous data: they are quotients of model output, and therefore inherit its sensitivity to the power-mean exponent, to the reference population used for normalization, and to the organisational multiplier. A perturbation that leaves the ordering of sites unchanged can still move the shares, and hence the gap. The masking gap should accordingly be reported as an interval over the same perturbation design used for the rank statistics rather than as a point estimate, since reporting it as a single number would attribute to the diagnostic a precision that the underlying weights do not possess.

4. Case-study exercise

Case-study design

The numerical exercise has three purposes. First, it tests whether the proposed multi-layer framework can be implemented end to end using observed site coordinates, physical hazard inputs, declared risk weights, transmission matrices, audit-based organisational information, and economic values. Second, it performs an internal ablation analysis by comparing coordinate-only screening with the full assessment, thereby quantifying how much of the site ranking is retained when organisational information is omitted and how often the identified dominant business process changes. Third, it evaluates the sensitivity of rankings, dominant components, and portfolio aggregates to declared modelling choices and perturbations. The exercise therefore assesses feasibility, internal information content, and model sensitivity; it does not establish predictive validity, causal effects, sampling uncertainty, population parameters, or actuarial calibration. Throughout this section, the reported euro amount is interpreted as the current framework's model-based monetary exposure, not as a return-period loss quantile. To make that perimeter auditable, the empirical objects are classified below as directly observed or externally sourced inputs,

expert or declared inputs, audit-derived inputs, and quantities generated mechanically by the model.

The case comprises 17 sites in 13 countries across four continents. Five broad NACE sections are represented: manufacturing, transport, financial services, health, and other services. Nine sites belong to one manufacturing group, two to a second group, and six are independent organisations. Total annual production value is EUR 1,128,450,000, with site values ranging from EUR 450,000 to EUR 450 million. Every site was visited by an auditor, and the organisational archive covers eleven process areas over twelve years (STAIRS Srl, 2026).

Site values originate in several currencies and are expressed in euro throughout. The conversion convention is material for a framework whose stated purpose is to isolate a climate channel: closing rates, average rates for the period, and constant rates fixed at a base date produce different portfolio totals and different relative site weights, none of which reflects physical hazard. Where a site operates in an economy meeting the hyperinflation criteria, restatement was required before translation (IFRS Foundation, 1989). The point is sharper in the 2050 exercise. The portfolio is held fixed so that changes in GCPML reflect changes in hazard rather than changes in the asset base, and consistency requires the monetary layer to be held fixed as well: forward values are therefore expressed at constant currency and constant prices, so that the reported change is not a composite of climate and monetary effects.

Table 2: Case-study sample

Dimension	Case-study coverage
Sites and geography	17 sites, 13 countries, four continents
NACE coverage	Five broad sections: manufacturing, transport, financial services, health, and other services
Ownership structure	Nine sites in one group, two in a second group, six independent
Production value	EUR 1.12845 billion; site range EUR 0.45 million to EUR 450 million
Organisational information	Eleven process areas; all sites visited; twelve-year audit archive
Main sampling limitation	Concentration in one industrial group; feasibility sample rather than an inferential design

The dimensional counts belong to this implementation layer. The hazard-to-mechanism matrix D has $17 \times 10 = 170$ entries; 28 of them (16 per cent) are reported as relying on expert judgement alone, while the remainder use measured, modelled, or proxy evidence. The mechanism-to-process matrix I has $10 \times 11 = 110$ positions, of which 92 are non-zero. The final hazard-process matrix F has $17 \times 11 = 187$ entries and is generated mechanically from D , I , and λ ; no element of F is directly assigned. These 187 entries are model coefficients, not empirical observations.

Table 3: Empirical status of the principal model objects

Object	Empirical content and source	Unit or transformation	Status in the case
$x_{js\tau}$	Physical climate or hydrological quantity evaluated at the verified site coordinate	Hazard-specific physical units	Externally sourced or modelled physical input
$h_{js\tau}$	Six-level score obtained from $x_{js\tau}$ and the reference population	Unitless, 0–1; band divided by five	Mechanically transformed input; ordinal bands treated as cardinal
λ_j	Declared relative consequence of one unit of normalized exposure to hazard j	Unitless; sums to one over 17 hazards	External or expert input; individual values and elicitation basis are not reported in the supplied material
D_{jc} and I_{cp}	Hazard-to-mechanism and mechanism-to-process allocation shares	Row-stochastic shares	Structured mapping inputs; evidence classes differ across cells
$H_{sp\tau}$	Process-level physical pressure combining h , λ , D , and I	Unitless generalized mean	Generated mechanically by the model
a_{sp}, b_{sp}, M_{sp}	Aggravating and protective audit findings and the resulting bounded multiplier	Audit scores; $M \in [0, 2]$	Findings are audit-derived; M is mechanically generated
V_s	Site production value translated to euro, gross of intra-group transactions	Euro under local GAAP	Accounting input
$GCPML_{sp\tau}$	Process-level model exposure	Euro	Generated model output, not observed loss

Notes: The status classification is epistemic: it distinguishes data, judgement, and model output. It does not imply that the 187 site–process cells are independent observations.

The unit of analysis is the physical site, not necessarily the legal entity. This matters for groups with multiple facilities and for enterprises whose registered office differs from the production plant. Site identities are anonymised in the paper.

Each site was first geocoded from its public address. The resulting point was then checked against satellite imagery and the auditor’s knowledge of the visited facility. The verification identifies the building rather than accepting the service’s reported confidence class.

The location exercise is not a random audit of global geocoders. It is a diagnostic for this portfolio. Nor does it imply that every address error biases a portfolio in the same direction. At the individual site, an error can cross a floodplain or elevation boundary. Across a large portfolio, positive and negative site errors may partly average out. The operational implication is targeted verification: verify sites for which coordinate disagreement, address class, or hazard gradients make misclassification costly.

The operational catalogue contains 17 geoclimatic hazards grouped by damage mechanism. The current case populates six, summarized in Table 4. River discharge is taken from GloFAS-related data (Alfieri et al., 2013); precipitation, wind, snow, heat, and dry-spell metrics are based on Copernicus and CMIP6-derived information (Eyring et al., 2016; Hersbach et al., 2020). The operational data pipeline uses Open-Meteo as an access layer for some climate-model series.

The six hazards represent 63 per cent of the catalogue’s declared economic weights. The reported GCPML is therefore a partial-coverage lower-bound metric in the accounting sense that omitted hazards contribute zero to the current calculation. It is not necessarily a statistical lower bound on realised loss because interactions and calibration are not estimated.

Table 4: Populated physical hazards in the case

Hazard	Operational variable	Principal source family
River flooding	River discharge / flow pressure	GloFAS
Extreme precipitation	High precipitation statistic	Copernicus / reanalysis and CMIP6-derived series
Wind	Extreme or high wind statistic	Copernicus / reanalysis and CMIP6-derived series
Snow	Snow occurrence or load-related statistic	Copernicus / reanalysis and CMIP6-derived series
Heat	Days above 32°C	Copernicus / CMIP6-derived series
Drought	Duration of dry spells	Copernicus / CMIP6-derived series

Table 5: Hazard-weight information available for the pilot

Catalogue group	Number of hazards	Total declared weight	Individual λ_j values
Populated hazards	6	0.63	Not reported in the supplied case material
Unpopulated hazards	11	0.37	Not reported in the supplied case material

Notes: The aggregate 63/37 split is reported by the operational synthesis. Individual hazard weights cannot be reconstructed from that split and are not inferred here.

The aggregate 0.63/0.37 split does not identify the monetary contribution of the eleven unpopulated hazards. For $q = 1$, and holding the full-catalogue process weights ϕ_{jp} , audit multipliers, sector corrections and production values fixed, the full-catalogue model output is the sum of the observed and omitted non-negative hazard contributions, $E_{\text{full}} = E_O + E_U$. The current $E_O = \text{EUR } 331.6$ million is therefore a lower bound within that fixed additive specification, not a lower bound on realised loss. The omitted contribution cannot be quantified from the aggregate catalogue weights alone. In particular, equal simple portfolio-average hazard intensity would not imply equal monetary contribution per unit of catalogue weight, because process transmission, site values and organisational multipliers also enter. Nor is the model globally linear in λ : the process weights ϕ_{jp} themselves contain normalization denominators.

For scale only, an explicitly hypothetical equal-contribution completion can be calculated by imposing $E_U/0.37 = E_O/0.63$. Under this additional assumption, full-catalogue exposure would be $331.6/0.63 = \text{EUR } 526.3$ million, approximately 46.6 per cent of production value. This is an illustrative accounting scenario, not an estimate or an identified upper bound; the available evidence does not establish the assumption or the actual size of the omitted contribution. Completing the catalogue requires the missing physical inputs and the disclosed process-transmission weights.

Let O denote the six populated hazards. The pilot retains the process weights ϕ_{jp} implied by the full 17-hazard catalogue and computes pressure from the observed subset, $H_{sp\tau}^{(q,O)} = \left(\sum_{j \in O} \phi_{jp} h_{j\sigma\tau}^q \right)^{1/q}$. The weights are not re-normalized over the six observed hazards. Hence the missing 37 per cent is not redistributed to the populated hazards: its contribution is zero because the physical inputs are unavailable, not because those hazards are physically absent. This convention makes the current result a partial-coverage output and preserves the meaning of the original catalogue weights.

The available material reports the combined declared weight of the six populated hazards, 0.63, but does not disclose the 17 individual λ_j values, their elicitation basis and date, or their uncertainty. The case results are therefore conditional on the declared expert-weight vector. The supplied information does not establish whether the same local physical information is reused in both λ_j and $h_{j\sigma\tau}$. This is an information-provenance limitation, not a requirement that economic importance and physical intensity be statistically uncorrelated: correlation can arise without duplicated inputs, and an absence of correlation would not exclude duplication. Likewise, fixing weights before extracting the current hazard series would document the order of operations but would not, by itself, establish information separability.

The existing perturbation experiment nevertheless provides a conditional sensitivity diagnostic. At the highest joint severity in Table 9, which combines band changes and lognormal weight noise with $\sigma = 0.30$, the reported intensity-rank correlation is 0.748, the absolute-rank correlation is 0.980, and the dominant process area is preserved in 77 per cent of runs. These results describe the specified joint perturbations; they neither isolate the effect of weight error nor bound systematic information overlap. To complete the assessment,

the replication package must disclose all 17 weights, their normalization, evidence sources, elicitation date and uncertainty, together with an input-by-input account distinguishing conditional economic severity, general likelihood, historical losses, engineering evidence and expert judgement. Any local physical information also used in h_{jst} must be identified. Alternative documented weight vectors should then be propagated through levels, rankings and dominant-process assignments. These disclosures and computations remain necessary for a reproducible assessment of the weight specification.

Two fifteen-year windows are used: a current-climate window and a window around 2050. The available synthesis reports one scenario pathway. The paper uses the future values to demonstrate relative and absolute aggregation, not to claim a probabilistic forecast of a particular site's 2050 loss.

The forward exercise holds the 2026 portfolio fixed and varies only the climate. This static-portfolio assumption is standard in climate stress testing, but it is declared here rather than left implicit because one site in the sample is a financial-services facility whose commercial relationship ended in 2014 and whose last audit dates from 2015. Projecting that site to 2050 does not assert that it will exist; it isolates the climate channel from portfolio turnover. The consequences for the aggregate result are reported later in this section and in Section 5.

For each site, physical variables are normalized to six levels, mapped to processes, combined with $q = 1$, multiplied by the audit-based M_{sp} and sector correction, and scaled by production value. The case compares full assessment with screening, changes the power-mean exponent, varies the reference population and economic weight basis, and performs 2,000 Monte Carlo perturbations.

The robustness experiment perturbs one band by one level for a fraction of observations, adds lognormal noise to economic weights, and adds noise to the organisational multiplier. Three severities are reported: 10, 20, and 30 per cent band perturbation, paired with weight-noise standard deviations of 0.10, 0.20, and 0.30. Outcomes are rank correlation for absolute exposure, rank correlation for intensity, and preservation of the dominant component.

These perturbations are researcher-specified stress scenarios for model inputs. Their distributions are not estimated from geocoding error, audit reliability, climate-model dispersion, or sampling variation, so the resulting ranges are sensitivity measures rather than confidence intervals or posterior uncertainty. The operational material reports that changing the reference population can move hazard levels by as much as 22 per cent, but it does not provide full portfolio results under the three benchmarks suggested in the comment – the 17-site sample, a sector–country peer group, and a fixed external population. A complete robustness appendix should report all three because they answer different questions and may alter both levels and rankings.

The case also tests whether public management-system certifications can proxy the audit. Ten standards are mapped to the eleven process areas, and the proxy's scale is varied. This is a useful negative control because certifications are public and scalable, but may

contain little information about site-specific implementation quality.

The case is theoretical and diagnostic. It does not estimate the causal effect of climate on firm performance, because there is no untreated comparison group or exogenous shock design. It does not estimate predictive accuracy, because observed losses are absent. It does not estimate a statistical tail quantile, because hazard frequencies and damage functions are not jointly calibrated. The empirical contribution is to show how the framework behaves, where measurement error enters, and which conclusions are stable under specified perturbations.

Case-study results

Under current climate windows and six populated hazards, the portfolio model-based monetary exposure (the current GCPML output) is EUR 331.6 million, equal to 29.4 per cent of aggregate production value. Monetary results are reported to the nearest tenth of a million euro; reporting the portfolio total to the unit would suggest a precision that the underlying bands, weights, and expert scores do not support. Site intensities range from 23.6 to 47.3 per cent, with a median of 31.1 per cent. These magnitudes should not be read as predicted realised loss. They are the monetary output of the model at its current normalization and coverage.

A model-implied covariance decomposition of log normalized GCPML over the $17 \times 11 = 187$ site–process cells yields shares of 47.5 per cent for the physical factor, 52.7 per cent for the organisational factor, and -0.1 per cent for the sector-process correction. These are algebraic covariance shares of a multiplicative model and sum to 100 per cent by construction. They are not explanatory power, causal attribution, or a decomposition of observed losses. The earlier total of 100.6 per cent arose because the sector term was omitted and different degrees-of-freedom conventions were used; both issues are corrected here.

The 187 cells are not 187 independent empirical observations. The eleven process cells within a facility share the same site-level hazard vector, location, and production-value scale, and may also share common measurement error. No standard errors or population inference are attached to the decomposition. The correlation between $\log H$ and $\log M$ is 0.038 in the constructed cells; this is a descriptive property of the pilot inputs, not evidence that physical exposure and organisational quality are independent in a broader population. The decomposition is useful only for showing how the model’s own output is assembled in this sample.


Only five sites are correct at civic-address level. Six are displaced by 157–824 metres. Another six are displaced by 1.168–6.631 kilometres. The large errors are systematic rather than random data noise: a same-name street in the wrong place, a different firm’s plant, a registered office instead of the facility, a district centroid for an address without a street, a municipal centroid when the street is missing from the database, and the wrong segment of a ring road.

Table 6: Model-implied logarithmic covariance decomposition of normalized GCPML

Component	Value	Share of $\text{Var}(\log r)$
$\text{Var}(\log H)$	0.0337	46.1%
$\text{Var}(\log M)$	0.0370	50.6%
$\text{Var}(\log(1 + k))$	0.0001	0.2%
Covariance terms	0.0022	3.0%
$\text{Var}(\log r)$	0.0731	100.0%
Covariance share, physical hazard		47.5%
Covariance share, organisational mediation		52.7%
Covariance share, sector correction		−0.1%

Notes: The 187 entries are 17×11 constructed site–process cells, not independent observations. With $r = \text{GCPML}_{SP}/(V_S/P)$, covariance shares are $\text{Cov}(\log x, \log r)/\text{Var}(\log r)$ and sum to unity by construction. All moments use the same degrees-of-freedom convention; the table has no inferential interpretation.

Four of the six errors above one kilometre are reported as high precision by the geocoding service. In this sample, service confidence is therefore not a reliable substitute for building verification. The result is consistent with the broader geocoding literature, which finds that geocoding method and reference data affect positional accuracy and exposure assignment (Cayo and Talbot, 2003; Zandbergen, 2008; Kinnee et al., 2020).



Notes: Cases are anonymised and reconstructed from the attached case synthesis. Four of the six large errors were labelled high precision by the geocoding service.

Figure 2: Six address-derived points more than one kilometre from the verified facility. Four were labelled high precision by the service. Cases are anonymised and reconstructed from the attached synthesis.

The strongest conclusion is methodological but bounded. A verified facility point is essential when a decision depends on a specific site's hazard classification. The current evidence does not prove that all portfolio rankings are biased by address error. To quantify

that effect, the physical layer must be re-run at both the address and verified coordinates and the difference in $h_{js\tau}$ and GCPML reported hazard by hazard. That recomputation, set out in (27), is the first item of the research agenda in Section VIII: it is the step that converts the present observation about distance into a measurement of exposure error, and it uses data the study already holds.

Table 7: Address-geocoding outcome for the 17 sites

Outcome	Number of sites
Correct at civic-address level	5
Approximate, error 157–824 metres	6
Incorrect, error 1.168–6.631 kilometres	6
of which service reported high precision	4

The intensity ranking is the more informative measure of how much organisational information coordinate-only screening preserves: its Spearman correlation with the full assessment is 0.917. The absolute-exposure correlation is higher, at 0.983, but both modes multiply by the same site-level production value V_s , whose range spans three orders of magnitude. Part of the stability of the absolute ranking is therefore mechanical and should not be interpreted as evidence that the organisationally adjusted and unadjusted intensities are nearly identical. The dominant physical hazard agrees at all 17 sites by construction because the physical layer is shared; this is a consistency check, not a screening-performance result. The informative diagnostic is the dominant process area, which agrees at only five sites.

At portfolio level, screening produces EUR 335.6 million compared with EUR 331.6 million under full assessment, a difference of -1.2 per cent when measured as full relative to screening. At individual sites, the difference ranges from -21.3 to $+32.7$ per cent. The combination of a 0.917 intensity-rank correlation and agreement on the dominant process at only five sites suggests that screening is useful for coarse triage but not for intervention design. The audit contributes process diagnosis and site-specific adjustment; the pilot does not establish that either mode predicts realised loss.

Holding site-level model outputs and the original 17-site normalization calibration fixed, the exact leave-one-out aggregate for a site k with current model-exposure share θ_k and current-to-2050 change g_k is $\bar{g}_{(-k)} = (\bar{g} - \theta_k g_k) / (1 - \theta_k)$, and its masking gap is $\Gamma_{(-k)} = \text{median}_{s \neq k}(g_s) - \bar{g}_{(-k)}$. These weights are exposure shares, not production-value shares. Re-estimating the sample-based normalization after each exclusion would also change the remaining sites' scores and constitutes a separate benchmark-sensitivity exercise. The available summary reports that excluding the dominant site raises the aggregate change from 2.1 to 9.2 per cent, while the reported median remains 11.2 per cent and the gap falls from 9.1 to 2.0 percentage points. These are the supplied deletion summaries; they are not a reconstruction of the full 17-site exercise.

Concentration can be assessed further from the reported aggregates. With 17 sites, total

Table 8: Coordinate-only screening versus full assessment

Comparison outcome	Result
Spearman rank correlation, intensity	0.917
Spearman rank correlation, absolute exposure	0.983; partly stabilised by common V_s
Dominant hazard	17 of 17 by construction; not a performance statistic
Dominant process area agrees	5 of 17 sites
Portfolio total	EUR 335.6m screening; EUR 331.6m full
Portfolio difference	−1.2%
Site-level difference	−21.3% to +32.7%

Notes: The comparison is internal to the model and does not establish which mode predicts realised loss more accurately. Intensity is the more informative rank diagnostic because absolute exposure contains the same production-value scale in both modes.

production value EUR 1,128.45 million, and observed minimum and maximum values EUR 0.45 million and EUR 450 million, the Herfindahl index of production-value shares lies approximately between 0.183 and 0.357. The lower endpoint assigns the other fifteen sites EUR 45.2 million each; the upper endpoint assigns two sites EUR 450 million, one EUR 222.15 million and fourteen EUR 0.45 million. The corresponding inverse-Herfindahl equivalent is between 2.8 and 5.5 equally weighted sites. This describes value concentration; it is not an estimate of exposure concentration. Inverting the reported dominant-site deletion gives $\theta_d = (9.2 - 2.1)/(9.2 + 9.3) \simeq 0.384$, compared with a production-value share of $450/1128.45 \simeq 0.399$. It implies a current model-exposure intensity of approximately 28.3 per cent at that site, consistent with the reported intensity range. These calculations use rounded published summaries.

The direction of each deletion follows from $\bar{g}_{(-k)} - \bar{g} = \theta_k(\bar{g} - g_k)/(1 - \theta_k)$. Thus removing a site above the aggregate lowers the aggregate, while removing one below it raises the aggregate. A median site change of 11.2 per cent against an aggregate of 2.1 per cent implies that at least nine of the seventeen deletions lower the aggregate. A conservative analytical envelope can also be established. Treating the rounded reported extrema $a = -0.093$, $b = 0.351$, mean $\mu = 0.021$ and dominant-site deletion 0.092 as inputs, every deletion lies approximately between -9.3 and $+9.2$ per cent. The lower endpoint is a convexity bound and need not be attained; the upper endpoint is the reported deletion of the dominant site.³

The complete vectors of current site exposures and site-level climate changes are not disclosed in the available case material. The full leave-one-site-out table, including the

³Let $t = \theta_d = (0.092 - \mu)/(0.092 - a)$ and $g_d = a$. For $k \neq d$, writing $\theta = \theta_k$, the remaining weighted mean is bounded above by both $(\mu - a\theta)/(1 - \theta)$ and $b - (b - a)t/(1 - \theta)$. The first expression increases in θ and the second decreases. Their intersection, $\theta^* = (b - \mu)/(b - a) - t$, gives an upper bound of approximately 0.085 for every deletion other than d , below the reported 0.092 for d . Every remaining weighted mean is at least a . The numerical envelope is conditional on the rounded summaries and is not the exact observed range.

monetary portfolio total, aggregate change, recomputed median and masking gap for every exclusion, therefore remains outstanding. The analytical envelope and the dominant-site diagnostic supplement that table; they do not replace it or establish that every deletion preserves the sign of deterioration.

This result suggests a tiered operating model. A lender or public agency can screen a large portfolio with coordinates and physical data, then commission verification and process audits for sites near decision thresholds, sites with low location confidence, or sites for which process intervention is contemplated. An industrial group that already owns the facilities gains more from the full mode because it can act on the process map.

Under the specified perturbation design, outputs change gradually rather than discontinuously. At the intermediate severity, the intensity-rank correlation is 0.828, the absolute-rank correlation is 0.988, and the dominant process area is preserved in 79 per cent of the 2,000 runs, compared with 73 per cent for the dominant hazard. At the highest severity, the intensity correlation falls to 0.748 while the absolute correlation remains 0.980. The contrast again reflects the stabilising role of the common production-value scale in absolute GCPML. These figures describe model sensitivity to the chosen disturbances; they are not estimated probabilities or statistical coverage rates.

Table 9: Sensitivity under 2,000 researcher-specified perturbations

Perturbation severity	Absolute rank	Intensity rank	Dominant process preserved
Bands ± 1 at 10%; weight $\sigma = 0.10$	0.993	0.907	81%
Bands ± 1 at 20%; weight $\sigma = 0.20$	0.988	0.828	79%
Bands ± 1 at 30%; weight $\sigma = 0.30$	0.980	0.748	77%

Notes: The final column reports preservation of the dominant *process area*. At the intermediate severity the dominant *hazard* is preserved in 73.0 per cent of runs against 78.9 per cent for the process area. The disturbances are imposed by the researchers and are not calibrated to empirical measurement-error distributions; the frequencies are sensitivity summaries, not confidence levels.

The pilot uses the 17-site sample as its reference population Ω , so its normalization depends on portfolio composition. The proposed alternatives answer different questions: a pre-defined sector–country peer group supplies a peer benchmark, while a fixed external distribution supports longitudinal and cross-portfolio comparison. The available material does not contain the hazard-specific thresholds and complete case outputs for these alternatives. The reported movement of up to 22 per cent in some hazard levels therefore cannot establish their effects on site rankings, portfolio exposure or the masking gap.

The current sensitivity experiment remains informative about its declared random perturbations, but changing the reference population may induce systematic, correlated band reassignments. Even if no more than 30 per cent of cells changed by one band, the experiment would not guarantee rank correlations of at least 0.748 or 0.980, or dominant-

process preservation of at least 77 per cent. The pattern and location of the changes matter, as well as their number. One algebraic invariance can be stated separately: with all other inputs fixed, an admissible common multiplicative rescaling of every hazard score by the same positive factor in both climate windows rescales model exposures proportionally and leaves rankings, dominant components, site growth rates, exposure shares and the masking gap unchanged. A change of benchmark need not have this form, especially under discrete bands; this invariance is not evidence of robustness to the proposed benchmarks.

The minimum completion exercise should document each benchmark's source, composition, calibration window and hazard-specific thresholds, using a consistent reference calibration for the current and future windows. For each benchmark it should report the fraction and pattern of band transitions, site intensities and monetary exposures, rank changes, dominant process areas, portfolio exposure, aggregate 2050 change and the recomputed median and masking gap. Until those outputs are available, all levels and rankings remain conditional on the sample-based reference population. This normalization limitation is not resolved by the existing Monte Carlo results.

The interpretation is asymmetric. Absolute exposure is heavily influenced by production value, so its stability is partly a scale result. Intensity removes that common factor and is therefore the more demanding diagnostic of the physical and organisational layers. Near-neighbour intensity rankings are less stable, and dominant components are preserved roughly four times in five under the imposed perturbations. The model can support triage and the identification of large exposures, but the pilot does not justify fine discrimination among sites separated by only a few intensity points.

The portfolio monetary exposure generated by the current GCPML specification rises by only 2.1 per cent between the current and 2050 windows. Read alone, this appears reassuring. The median site rises by 11.2 per cent, however, and the largest increase is 35.1 per cent. The masking gap is therefore $\Gamma = 9.1$ percentage points. The low aggregate is driven by a single site whose GCPML falls by 9.3 per cent and which represents 39 per cent of production value; because site intensities lie in the narrow band 23.6 to 47.3 per cent, its exposure share θ_d in (22) is close to the same figure.

The sample satisfies the masking condition (23) but not the reversal condition (24), and the distance between the two is worth reporting. With $\theta_d \simeq 0.39$ the dominant site contributes -3.6 percentage points to the aggregate while the remaining sites contribute $+5.7$; reversal would require the dominant site to improve by 14.7 per cent rather than 9.3, or to hold an exposure share of 61.6 per cent rather than 39. The case therefore illustrates how an aggregate conceals the typical site well before it changes its sign, which is the more common and more easily overlooked configuration.




Figure 3: Changes in site GCPML between current and 2050 windows. The site distribution deteriorates more than the portfolio aggregate suggests because one high-value site improves. The masking gap Γ between the median site and the aggregate is 9.1 percentage points; the aggregate does not reverse sign. Values are reconstructed from the attached case chart; sites are anonymised.

Because the result rests on a single component, both readings should be published. Excluding the dominant site – the facility described in Section V, whose relationship ended in 2014 – the aggregate rises by 9.2 per cent against an unchanged median of 11.2 per cent, and the masking gap falls from 9.1 to 2.0 percentage points. The direction of the finding survives the exclusion; its magnitude does not. Reported together, the two readings show both that the phenomenon is real and that its size in this sample is governed by one facility.

Table 10: The 2050 aggregate under both readings

Reading, absolute change to 2050	Aggregate	Median site	Gap Γ
Full sample, 17 sites	+2.1%	+11.2%	9.1 pp
Excluding the dominant site, 16 sites	+9.2%	+11.2%	2.0 pp

Notes: Both readings hold the 2026 portfolio fixed and vary only the climate. The excluded site is a financial-services facility whose commercial relationship ended in 2014.

The result does not imply that value-weighting is wrong. Capital allocation requires monetary weighting. It implies that the aggregate answers a different question from the typical-site change. A board, lender, or supervisor should see both.

The two reported readings also serve different empirical purposes. The 17-site reading is a fixed-portfolio climate stress counterfactual that preserves the original exposure set, including a relationship that ended in 2014. The 16-site reading is closer to a current-portfolio diagnostic. Neither is a forecast of the portfolio's actual composition in 2050; presenting both prevents a static-sample convention from being mistaken for a business prediction.

Mapping public management-system certifications to process areas does not improve the dominant-process match: it remains five of 17 sites, the same as simple screening. Even with an implausibly large scale parameter, the match reaches at most seven sites. The reason is signal amplitude. The audit-based multiplier ranges from 0.259 to 1.536, a span of 1.277. The certification proxy ranges from 0.980 to 1.000, a span of 0.020 – roughly 64 times narrower.

The negative result is informative. Certifications may indicate that a management system exists, but not how well it operates at a particular facility or how its protections compare with process-specific aggravating conditions. They can be a prior or screening flag, not a substitute for the audit in the current specification.




Figure 4: Range of the audit-based organisational multiplier and the public-certification proxy. The proxy contains too little cross-process variation to recover the audit’s dominant process area.

5. Implications for industrial risk management and finance

For an industrial group, the framework separates two decisions. The first is portfolio triage: which sites account for the largest monetary exposure and which hazards dominate them? Screening performs this task with relatively little loss of ranking information. The second is intervention design: which process inside the site amplifies the hazard, and how much exposure could be reduced by improving it? That decision requires the organisational layer.

Because $\gamma = 1$, the local sensitivity of process GCPML to the organisational multiplier is

$$\frac{\partial \text{GCPML}_{sp\tau}}{\partial M_{sp}} = \frac{\text{GCPML}_{sp\tau}}{M_{sp}}. \quad (26)$$

A feasible reduction in M_{sp} can therefore be translated directly into avoided model exposure, holding the other layers fixed. This supports a mitigation ranking by avoided GCPML per unit cost. The calculation is a comparative-static management tool, not a causal estimate of realised savings.

The process map also disciplines resilience language. “Improve climate resilience” becomes a set of specific actions in logistics, supply-chain redundancy, business continuity, information technology, health and safety, or environmental management. The same site can be physically exposed but organisationally robust, or physically moderate but operationally fragile.

Financial transmission from facilities to valuation. The finance literature separates news about future cash flows from changes in discount rates. In the Campbell–Shiller representation, valuation ratios encode expectations of both future payouts and required returns (Campbell and Shiller, 1988); time variation in discount rates is itself a central source of asset-price movements (Cochrane, 2011). Physical hazards can affect the cash-flow side through downtime, repair costs, adaptation expenditure, lost sales, and asset impairment. Climate and transition uncertainty can affect the discount-rate side through risk premia, collateral haircuts, and financing conditions. The persistent response of price–dividend ratios in Guidolin and Ionta (2026) illustrates this financial amplification at aggregate and sectoral levels.

GCPML does not combine these channels into a valuation formula. Its defensible role is upstream: it supplies a documented facility–process exposure variable that can be linked empirically to realised cash flows and market prices. A finance extension should estimate whether GCPML predicts loan margins, covenant tightness, bond or CDS spreads, insurance premia, equity valuations, and implied volatility after controlling for emissions intensity, sectoral transition risk, and textual climate exposure. This preserves the distinction between a physical-risk measurement and the market’s price of that risk.

Banking guidance increasingly treats physical climate risk as a driver transmitted through credit, market, operational, and liquidity channels (Basel Committee on Banking Supervision, 2021b,a, 2022). The NGFS similarly emphasizes hazard, exposure, vulnerability, and the value of geospatial data (Network for Greening the Financial System, 2022, 2025). GCPML offers a practical hierarchy of data acquisition.

At the broadest level, coordinates and physical fields can rank counterparties or facilities. The next layer is location confidence: addresses that resolve to centroids, legal offices, or disagree across geocoders should be flagged. The most expensive layer is process audit and facility verification. The case suggests reserving it for material exposures and borderline decisions rather than requiring it for every borrower.

The model also clarifies why a single country coefficient is insufficient for collateral or industrial credit. Two facilities in the same municipality can occupy different floodplains, elevations, coastal zones, or urban heat environments. Conversely, site resolution alone is insufficient because the same physical hazard can produce different economic effects in organisations with different redundancies and controls.

Insurers may find the site and process decomposition useful for risk engineering and underwriting dialogue. However, the current GCPML should not be inserted into pricing as a return-period loss. It lacks hazard-frequency calibration, vulnerability curves, policy terms, deductibles, limits, and claims validation. Among the contractual parameters, the indemnity period for business interruption deserves particular attention, because it is the term most directly connected to the process logic of the model: the exposure attributed to a process is economically meaningful only over the horizon for which that process would remain impaired, and an indemnity period shorter than the plausible restoration time truncates precisely the loss the model attempts to represent. The proper path is to treat GCPML as a transparent exposure index and then estimate a mapping from the index to observed loss distributions.

A future actuarial extension could replace normalized bands with hazard exceedance curves and engineering damage functions. Organisational multipliers could then shift vulnerability or business-interruption duration rather than multiply a score directly. In the current management specification, “probable” recognises conditional physical risk and imperfect information about protective effectiveness, while “maximum” denotes the no-additional-action upper reference defined in (17). Neither term establishes a calibrated probability or an absolute loss bound. This explicit convention is needed when risk-engineering terminology is used for a diagnostic exposure measure; the engineering-

insurance guidance of Mbatha and Nkuna (2021) likewise emphasises documented scenarios and protection assumptions.

A monetary quantity equal to a substantial share of annual production value will be read by boards and lenders as an expected loss unless its accounting perimeter is stated explicitly. Three distinctions are necessary.

First, GCPML does not support the recognition of a provision. A provision requires a present obligation arising from a past event, whose settlement is probable and whose amount can be reliably estimated (IFRS Foundation, 1998b; Organismo Italiano di Contabilità, 2025). A forward-looking exposure conditioned on a climate scenario satisfies none of these conditions: there is no past obligating event and no counterparty. Treating the measure as a basis for accrual would misstate the financial position.

Second, the measure can legitimately inform impairment testing. A high and rising site intensity is an indicator that the recoverable amount of the assets deployed at that site may have fallen below their carrying amount, and it is relevant both to the cash-flow projections and to the discount rate used in a value-in-use computation (IFRS Foundation, 1998a; Organismo Italiano di Contabilità, 2016). This is a diagnostic use rather than a measurement input: the impairment test remains governed by its own standard, and GCPML enters as one of the external indicators that trigger or inform it.

Third, the measure is well suited to the narrative disclosure of principal risks and uncertainties in the management report (Repubblica Italiana, 1942). Its site and process decomposition maps naturally onto the disaggregation such disclosure requires, and its explicit limitations are themselves reportable.

A fourth consideration concerns financial uses. As reported in Section IV, the portfolio total rises by 12.9 per cent when the power-mean exponent moves from one to two – a modelling choice, not an observation. By analogy with the fair-value hierarchy, an output whose level depends materially on unobservable inputs and preparer judgement belongs to the third level of that hierarchy (IFRS Foundation, 2011). The analogy is not a claim that GCPML is a fair-value measurement; it is a statement about the evidential status of the number. The operational consequence, already implicit elsewhere in this paper, is that GCPML may be used for ranking, triage, and diagnosis, but should not enter covenants, pricing, or collateral valuation as a point estimate until it has been calibrated against observed losses.

European sustainability reporting requires precisely the object this framework produces. The climate standard obliges an undertaking to assess physical climate risk at the level of its own assets and their locations, over defined time horizons and under climate scenarios, and to describe the resilience of its strategy and business model (European Commission, 2023). A site-resolved, scenario-conditioned, process-decomposed exposure measure is a direct candidate for operationalising that requirement, and its process layer supports the connectivity expected between the sustainability statement and the financial statements.

The regulatory perimeter has narrowed. The reporting directive (European Parliament and Council of the European Union, 2022), transposed in Italy in 2024 (Repubblica Italiana,

2024), was amended in 2026 so that mandatory reporting applies only to undertakings above substantially higher size thresholds (European Parliament and Council of the European Union, 2026), with national transposition of the amending directive under way at the time of writing. Two consequences follow, and they point in opposite directions. Fewer undertakings are obliged to report; but those released from the obligation remain exposed to the same physical hazards and continue to face information requests from lenders and from customers that are themselves in scope. For that population a transparent, auditable exposure measure is more useful as a voluntary and value-chain instrument than as a compliance artefact.

The audience widens accordingly. Beyond the banks, insurers, and industrial groups discussed above, sustainability assurance practitioners require evidence that is documented, reproducible, and traceable to a method (International Auditing and Assurance Standards Board, 2024). The audit-based multiplier is a structured expert estimate in exactly that sense, and the governance requirements set out in Section VIII are what such an engagement would examine.

The common eleven-process architecture allows legal, geopolitical, and geoclimatic pressure to be displayed on one map. It does not justify summing their three monetary exposure quantities. The assessments may reuse the same economic value, cover overlapping damage channels, and rely on different scenarios; their dependence and joint occurrence are not modelled. A combined loss measure would therefore require reconciled economic perimeters and a joint event or loss specification. The current output is a process-level dashboard showing which risk family dominates each area, accompanied by separate totals and scenario narratives.

Validation, limitations, and research agenda

The next scientific step is a pre-registered blind test. A new portfolio should be selected without using known losses. The protocol should specify the reference population, hazard coverage, coordinate-verification rule, audit instrument, parameter values, and failure thresholds before outcomes are inspected. Candidate external outcomes include insured property loss, business-interruption days, emergency capital expenditure, production downtime, covenant deterioration, or facility closure.

With a larger panel, an empirical model could test whether GCPML predicts losses or performance beyond conventional country, sector, and balance-sheet controls. Because losses are likely zero-inflated and heavy tailed, hurdle, quantile, or extreme-value methods may be more appropriate than ordinary least squares. Facility and time fixed effects could absorb persistent unobserved heterogeneity, while event studies around floods, heatwaves, or storms could isolate acute shocks. The present paper does none of this and should not be read as if it did.

A finance-oriented validation should distinguish cash-flow outcomes from discount-rate and funding outcomes. The first group includes revenues, operating margins, emergency

capital expenditure, impairment charges, insured losses, and interruption days. The second includes bank loan spreads and maturities, covenant changes, bond and CDS spreads, equity valuation ratios, option-implied volatility, and insurance pricing. Incremental tests should compare GCPML with sectoral emissions, climate-policy uncertainty, and firm-level textual measures (Sautner et al., 2023; Li et al., 2024; Bolton and Kacperczyk, 2021; Guidolin and Ionta, 2026). Such a design would show whether verified site and process information adds explanatory content beyond the green–brown and firm-level classifications already used in climate finance, and whether physical and transition risks reinforce one another.

The six large coordinate errors motivate a direct re-computation exercise. For each site-hazard pair, calculate

$$\Delta h_{j\tau} = h_j(\tilde{\ell}_s, \tau) - h_j(\ell_s, \tau), \quad \Delta \text{GCPML}_{s\tau} = \text{GCPML}_s(\tilde{\ell}_s, \tau) - \text{GCPML}_s(\ell_s, \tau). \quad (27)$$

This would distinguish large positional error in a smooth field from smaller error at a sharp boundary. It would also allow the paper to estimate when address verification changes portfolio ranking, not merely when the address is geographically wrong.

A scalable confidence model should store both the address-derived and verified point, the address type, the confidence score assigned before verification, disagreement across independent geocoders, and the realised distance error. Every future on-site audit would then improve the calibration sample.

The current model populates six of 17 hazards and covers 63 per cent of declared hazard weight. Extending coverage is not simply a data-collection exercise. Each new hazard requires a physical variable, a presence rule, a normalization population, a damage-mechanism map, and evidence for its economic weight. Missing hazards should remain explicitly missing rather than coded as zero.

Only one future climate pathway is currently reported. A scientific version should show multiple scenarios, climate-model ensembles, and internal variability. Results should distinguish uncertainty from the climate model, hazard transformation, geocoding, normalization, transmission weights, organisational scoring, and economic value. The Monte Carlo exercise is a useful start but does not yet span all these layers.

Reference-population dependence is substantial. A stable external benchmark is preferable for time comparison, while a portfolio-specific benchmark may be useful for cross-sectional triage. Both can be reported, but they answer different questions. The paper should avoid presenting relative bands as absolute physical probabilities.

The weight-conservation identity in (10) shows that hazard importance is an input. The current economic weights require a documented elicitation or empirical basis. A defensible protocol could combine engineering severity, historical loss shares, expert elicitation with uncertainty intervals, and stakeholder-specific objectives. Alternative weight sets should be propagated through rank and dominance sensitivity.

The transmission matrices require governance as well. Each expert-assigned entry should have a rationale, reviewer, version, and evidence class. Entries derived mechanically

from row-stochastic decomposition should be distinguishable from expert inputs; in the present implementation every entry of F is derived and every entry of D is assigned, so the distinction is structural. The apparent 170-versus-187 cell discrepancy is reconciled in Section III: the two counts refer to D and to F respectively.

The audit is the model’s most valuable and least scalable layer. Future work should report inter-rater reliability, test-retest stability, item-level missingness, and the mapping from audit questions to a_{sp} and b_{sp} . The multiplier’s functional form should be validated against outcomes and compared with monotone alternatives.

Two further properties of the archive bear on the interpretation of M_{sp} . The first is instrument stability: a twelve-year archive supports intertemporal comparison only if the underlying items, anchors, and scoring conventions have remained constant, or if changes have been documented and back-mapped. The second is the age of the individual observation. The multiplier entering a current-window computation is the one recorded at the last audit, which for some sites predates the assessment by several years; the financial-services facility discussed in Section V was last audited in 2015, so its multiplier is eleven years old at the reference date. Stale organisational information is distinct from the fixed-portfolio convention, because it affects the mediating layer rather than the composition of the sample. A production implementation should record the audit date alongside M_{sp} and either re-audit or widen the reported uncertainty beyond a declared validity horizon. The requirements that auditing practice imposes on structured expert estimates – documented anchors, competence of the preparer, evidence of consistency, and independent review – apply directly here (International Auditing and Assurance Standards Board, 2018).

A legacy dashboard formula used in earlier GPML implementation is not monotone in adverse findings and can rank observations in the opposite direction from (14). This paper excludes that legacy series. Historical continuity should not take precedence over sign consistency; a harmonised back-cast is required before old and new outputs are compared.

6. Conclusion

This paper develops a site-resolution framework for translating macro climate conditions into corporate risk-management quantities. The central idea is simple: a climate field reaches a facility, the facility’s processes transmit or absorb the pressure, and economic value turns the resulting intensity into a monetary exposure. The contribution lies in making each step explicit and testable.

Three methodological conclusions stand out. First, location is part of the risk model rather than a data-cleaning footnote. Address-derived coordinates can be wrong by kilometres and service confidence may not reveal the error. Second, a row-stochastic transmission architecture distributes hazard pressure but cannot identify hazard importance; the weights remain external modelling choices. Third, compensatory aggregation can reverse the message delivered by the typical site when economic value is concentrated. Portfolio totals should therefore be accompanied by the dominant component and the site

Table 11: Current limitations and required next steps

Limitation	Consequence	Required next step
Six of 17 hazards populated	Current output omits 37% of declared hazard weight	Extend the physical layer and keep missing hazards explicit
17 sites; nine in one group	No population inference or sector generalisation	Blind, multi-firm validation sample
No observed losses	No predictive or actuarial validation	Link to claims, downtime, and financial outcomes
One future scenario	Scenario uncertainty understated	Multi-scenario, multi-model ensemble
Static portfolio; 2050 result driven by one site	Aggregate change depends on a single component	Publish both readings and the masking gap; refresh portfolio composition
Sample-based normalization	Bands move with the comparison set	Pre-specify stable external and portfolio benchmarks
Aggregation exponent set to $q = 1$	Monetary level rises 12.9% at $q = 2$; ordering unaffected	Pre-specify q per application and publish the level under at least two exponents
Address error measured, hazard error not fully re-run	Distance does not equal climate-score error	Recompute all hazards at proxy and verified coordinates
Expert weights and audit scores	Model risk and subjectivity	Elicitation protocol, uncertainty intervals, reliability tests
Legacy multiplier inconsistency	Historical series not directly comparable	Harmonise formula and back-cast
Address error measured only as distance	Priority ordering of the agenda	Execute (27) first; it requires no new fieldwork
Confidential raw data unavailable in this draft	Limited reproducibility	Release anonymised data dictionary, code, and permitted aggregates

distribution.

The 17-site case supports a tiered, but carefully qualified, use of the model. The intensity-rank correlation of 0.917 is the more informative measure of organisational information retained by coordinate-only screening. The higher absolute-exposure correlation of 0.983 is partly mechanical because both modes share the same production values, and agreement on the dominant physical hazard is exact by construction rather than an empirical success. Screening can therefore support coarse portfolio triage, while the process audit is needed when management must identify where to intervene within a site. Public certifications do not contain enough process-level variation to replace that audit in this pilot.

The terminology is deliberately conditional. “Probable” recognises uncertain physical conditions and imperfect knowledge after auditing; “maximum” denotes the no-additional-action upper reference relative to the stated risk-reducing interventions, retaining existing protections. Hazard averaging and value-weighted portfolio intensity are transparent summaries, not evidence of a calibrated loss tail. The evidence is preliminary. GCPML is not yet an actuarially calibrated PML, a credit-risk model, or an asset-pricing model. Its scientific value at this stage is to organize the macro-to-micro transmission problem, expose measurement and aggregation choices, and define a falsifiable validation programme. For finance, it should be interpreted as an exposure state variable upstream of cash-flow and discount-rate outcomes. Its added value will be demonstrated if verified site and process information predicts realised losses, financing conditions, or valuations beyond emissions and sector classifications. A blind portfolio test with verified coordinates, multiple climate scenarios, documented weights, audit reliability, observed losses, and financial outcomes is the necessary next step.

References

- Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.
- Addoum, J. M., Ng, D. T., and Ortiz-Bobea, A. (2020). Temperature shocks and establishment sales. *Review of Financial Studies*, 33(3):1331–1366.
- Alfieri, L., Burek, P., Dutra, E., Krzeminski, B., Muraro, D., Thielen, J., and Pappenberger, F. (2013). GloFAS – global ensemble streamflow forecasting and flood early warning. *Hydrology and Earth System Sciences*, 17:1161–1175.
- Ardia, D., Bluteau, K., Boudt, K., and Inghelbrecht, K. (2023). Climate change concerns and the performance of green versus brown stocks. *Management Science*, 69(12):7607–7632.
- Auffhammer, M. (2018). Quantifying economic damages from climate change. *Journal of Economic Perspectives*, 32(4):33–52.
- Aven, T. (2016). Risk assessment and risk management: Review of recent advances on their foundation. *European Journal of Operational Research*, 253(1):1–13.

- Barnett, M., Brock, W., and Hansen, L. P. (2020). Pricing uncertainty induced by climate change. *Review of Financial Studies*, 33(3):1024–1066.
- Barrot, J.-N. and Sauvagnat, J. (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *Quarterly Journal of Economics*, 131(3):1543–1592.
- Basel Committee on Banking Supervision (2021a). Climate-related financial risks – measurement methodologies. Technical report, Bank for International Settlements, Basel.
- Basel Committee on Banking Supervision (2021b). Climate-related risk drivers and their transmission channels. Technical report, Bank for International Settlements, Basel.
- Basel Committee on Banking Supervision (2022). Principles for the effective management and supervision of climate-related financial risks. Technical report, Bank for International Settlements, Basel.
- Battiston, S., Mandel, A., Monasterolo, I., Schuetze, F., and Visentin, G. (2017). A climate stress-test of the financial system. *Nature Climate Change*, 7:283–288.
- Bernstein, A., Gustafson, M. T., and Lewis, R. (2019). Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics*, 134(2):253–272.
- Bilal, A. and Känzig, D. R. (2026). The macroeconomic impact of climate change: Global versus local temperature. *Quarterly Journal of Economics*, 141(2):889–944.
- Boehm, C. E., Flaaen, A., and Pandalai-Nayar, N. (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tohoku earthquake. *Review of Economics and Statistics*, 101(1):60–75.
- Bolton, P. and Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics*, 142(2):517–549.
- Bolton, P. and Kacperczyk, M. (2023). Global pricing of carbon-transition risk. *Journal of Finance*, 78(6):3677–3754.
- Bressan, G., Đuranović, A., Monasterolo, I., and Battiston, S. (2024). Asset-level assessment of climate physical risk matters for adaptation finance. *Nature Communications*, 15:5371.
- Burke, M., Hsiang, S. M., and Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, 527:235–239.
- Campbell, J. Y. and Shiller, R. J. (1988). The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, 1(3):195–228.
- Carleton, T. A. and Hsiang, S. M. (2016). Social and economic impacts of climate. *Science*, 353(6304):aad9837.

- Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the great east japan earthquake. *Quarterly Journal of Economics*, 136(2):1255–1321.
- Cayo, M. R. and Talbot, T. O. (2003). Positional error in automated geocoding of residential addresses. *International Journal of Health Geographics*, 2:10.
- Cochrane, J. H. (2011). Presidential address: Discount rates. *Journal of Finance*, 66(4):1047–1108.
- Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, 4(3):66–95.
- Dietz, S., Bowen, A., Dixon, C., and Gradwell, P. (2016). Climate value at risk of global financial assets. *Nature Climate Change*, 6(7):676–679.
- Domenichini, P. and Ravazzolo, F. (2026). *Geopolitical Probable Maximum Loss: The STAIRS Corporate Geopolitical Risk Model*. ISR Editions, Oslo and Verona.
- Engle, R. F., Giglio, S., Kelly, B., Lee, H., and Stroebe, J. (2020). Hedging climate change news. *Review of Financial Studies*, 33(3):1184–1216.
- European Commission (2023). Commission delegated regulation (EU) 2023/2772, annex i: ESRS E1 Climate Change. Official Journal of the European Union, L 2023/2772, 22 December 2023.
- European Parliament and Council of the European Union (2022). Directive (EU) 2022/2464 as regards corporate sustainability reporting. Official Journal of the European Union, L 322, 16 December 2022.
- European Parliament and Council of the European Union (2026). Directive (EU) 2026/470 amending certain corporate sustainability reporting and corporate sustainability due diligence requirements. Official Journal of the European Union, L 2026/470, 26 February 2026.
- Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., and Taylor, K. E. (2016). Overview of the coupled model intercomparison project phase 6 (CMIP6) experimental design and organization. *Geoscientific Model Development*, 9:1937–1958.
- Gavrilidis, K. (2021). Measuring climate policy uncertainty. SSRN working paper 3847388.
- Gavrilidis, K., Känzig, D. R., Raghavan, R., and Stock, J. H. (2026). The macroeconomic effects of climate policy uncertainty. NBER Working Paper 34762, National Bureau of Economic Research. Revised May 2026.

- Giglio, S., Kelly, B., and Stroebe, J. (2021a). Climate finance. *Annual Review of Financial Economics*, 13:15–36.
- Giglio, S., Maggiori, M., Rao, K., Stroebe, J., and Weber, A. (2021b). Climate change and long-run discount rates: Evidence from real estate. *Review of Financial Studies*, 34(8):3527–3571.
- Greco, S., Ishizaka, A., Tasiou, M., and Torrisi, G. (2019). On the methodological framework of composite indices: A review of the issues of weighting, aggregation, and robustness. *Social Indicators Research*, 141:61–94.
- Grossi, P. and Kunreuther, H., editors (2005). *Catastrophe Modeling: A New Approach to Managing Risk*. Springer, New York.
- Guidolin, M. and Ionta, S. (2026). Uncertain climate policy as a source of macro-financial shocks: Evidence from carbon futures volatility. Working paper, Baffi Centre, Bocconi University, 18 August 2026.
- Hain, L. I., Kölbel, J. F., and Leippold, M. (2022). Let's get physical: Comparing metrics of physical climate risk. *Finance Research Letters*, 46:102406.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horanyi, 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., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragani, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R. J., Holm, E., Janiskova, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., and Thépaut, J.-N. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049.
- IFRS Foundation (1989). IAS 29: Financial reporting in hyperinflationary economies. International Accounting Standard.
- IFRS Foundation (1998a). IAS 36: Impairment of assets. International Accounting Standard.
- IFRS Foundation (1998b). IAS 37: Provisions, contingent liabilities and contingent assets. International Accounting Standard.
- IFRS Foundation (2006). IFRS 8: Operating segments. International Financial Reporting Standard.
- IFRS Foundation (2011). IFRS 13: Fair value measurement. International Financial Reporting Standard.
- IFRS Foundation (2014). IFRS 15: Revenue from contracts with customers. International Financial Reporting Standard.

- International Auditing and Assurance Standards Board (2018). ISA 540 (revised): Auditing accounting estimates and related disclosures. Technical report, International Federation of Accountants.
- International Auditing and Assurance Standards Board (2024). ISSA 5000: General requirements for sustainability assurance engagements. Technical report, International Federation of Accountants.
- Kahn, M. E., Mohaddes, K., Ng, R. N. C., Pesaran, M. H., Raissi, M., and Yang, J.-C. (2021). Long-term macroeconomic effects of climate change: A cross-country analysis. *Energy Economics*, 104:105624.
- Kinnee, E. J., Tripathy, S., Schinasi, L., Shmool, J. L. C., Sheffield, P. E., Holguin, F., and Clougherty, J. E. (2020). Geocoding error, spatial uncertainty, and implications for exposure assessment and environmental epidemiology. *International Journal of Environmental Research and Public Health*, 17(16):5845.
- Li, Q., Shan, H., Tang, Y., and Yao, V. (2024). Corporate climate risk: Measurements and responses. *Review of Financial Studies*, 37(6):1778–1830.
- Li, X. (2025). Physical climate change exposure and firms’ adaptation strategy. *Strategic Management Journal*, 46(3):750–789.
- Mandel, A., Battiston, S., and Monasterolo, I. (2025). Mapping global financial risks under climate change. *Nature Climate Change*, 15:329–334.
- Mbatha, P. and Nkuna, N. (2021). PML assessment for engineering insurance. Munich Re. Technical presentation, 18 May 2021.
- Munda, G. and Nardo, M. (2009). Noncompensatory/nonlinear composite indicators for ranking countries: A defensible setting. *Applied Economics*, 41(12):1513–1523.
- Network for Greening the Financial System (2022). Physical climate risk assessment: Practical lessons for the development of climate scenarios with extreme weather events from emerging markets and developing economies. Technical report, NGFS, Paris.
- Network for Greening the Financial System (2025). Leveraging physical climate risk data. Technical report, NGFS, Paris.
- Noailly, J., Nowzohour, L., and van den Heuvel, M. (2022). Does environmental policy uncertainty hinder investments towards a low-carbon economy? NBER Working Paper 30361, National Bureau of Economic Research.
- OECD and European Commission Joint Research Centre (2008). *Handbook on Constructing Composite Indicators: Methodology and User Guide*. OECD Publishing, Paris.

- Organismo Italiano di Contabilità (2016). OIC 9: Svalutazioni per perdite durevoli di valore delle immobilizzazioni materiali e immateriali. Technical report, Fondazione OIC.
- Organismo Italiano di Contabilità (2024). OIC 12: Composizione e schemi del bilancio d'esercizio. Technical report, Fondazione OIC. Principle issued in 2016 and updated through 18 March 2024.
- Organismo Italiano di Contabilità (2025). OIC 31: Fondi per rischi e oneri e trattamento di fine rapporto. Technical report, Fondazione OIC. Principle issued in 2016 and updated through 8 December 2025.
- Palikhe, H., Schaur, G., and Sims, C. (2024). Environmental policy uncertainty. *Journal of the Association of Environmental and Resource Economists*, 11(5):1135–1163.
- Pankratz, N., Bauer, R., and Derwall, J. (2023). Climate change, firm performance, and investor surprises. *Management Science*, 69(12):7352–7398.
- Pástor, L., Stambaugh, R. F., and Taylor, L. A. (2022). Dissecting green returns. *Journal of Financial Economics*, 146(2):403–424.
- Portner, H.-O., Roberts, D. C., Tignor, M., Poloczanska, E. S., Mintenbeck, K., Alegría, A., Craig, M., Langsdorf, S., Loschke, S., Möller, V., Okem, A., and Rama, B., editors (2022). *Climate Change 2022: Impacts, Adaptation and Vulnerability*. Cambridge University Press, Cambridge.
- Repubblica Italiana (1942). Codice civile, articolo 2425: Contenuto del conto economico. As amended.
- Repubblica Italiana (2024). Decreto legislativo 6 settembre 2024, n. 125: Attuazione della direttiva (UE) 2022/2464 sulla rendicontazione societaria di sostenibilità. Gazzetta Ufficiale della Repubblica Italiana, n. 212, 10 September 2024.
- Sautner, Z., van Lent, L., Vilkov, G., and Zhang, R. (2023). Firm-level climate change exposure. *Journal of Finance*, 78(3):1449–1498.
- STAIRS Srl (2026). Il rischio geoclimatico d'impresa misurato sul fabbricato: Il modello, le verifiche, i risultati, gli artefatti e ciò che resta da fare. Documento 13, revisione 1.0, internal research synthesis.
- Zandbergen, P. A. (2008). A comparison of address point, parcel and street geocoding techniques. *Computers, Environment and Urban Systems*, 32(3):214–232.

A. Proofs

Proof of the weight-conservation identity. Using (9) and row stochasticity,

$$\begin{aligned}\sum_p F_{jp} &= \sum_p \lambda_j \sum_c D_{jc} I_{cp} \\ &= \lambda_j \sum_c D_{jc} \sum_p I_{cp} \\ &= \lambda_j \sum_c D_{jc} \\ &= \lambda_j.\end{aligned}$$

If $\sum_r \lambda_r = 1$, then

$$\omega_j = \frac{\sum_p F_{jp}}{\sum_r \sum_p F_{rp}} = \frac{\lambda_j}{\sum_r \lambda_r} = \lambda_j.$$

□

Proof of finite-power-mean compensability. For finite q ,

$$A_q(z) = \left(\sum_i w_i z_i^q \right)^{1/q}.$$

Differentiation yields

$$\frac{\partial A_q}{\partial z_i} = \frac{1}{q} \left(\sum_r w_r z_r^q \right)^{1/q-1} q w_i z_i^{q-1} = w_i z_i^{q-1} A_q^{1-q} > 0.$$

For $i \neq k$, the marginal rate of substitution is

$$\frac{\partial A_q / \partial z_i}{\partial A_q / \partial z_k} = \frac{w_i}{w_k} \left(\frac{z_i}{z_k} \right)^{q-1},$$

which is finite and positive for positive components and weights. Hence changes can compensate. As $q \rightarrow \infty$, $A_q \rightarrow \max_i z_i$. □

Proof of the exposure-share representation. From (21), $\Delta \bar{r} = \sum_s v_s \Delta r_s = \sum_s v_s r_s g_s$. Dividing by $\bar{r} = \sum_i v_i r_i$ and using $v_s r_s = \text{GCPML}_s / \sum_i V_i$,

$$\bar{g} = \frac{\Delta \bar{r}}{\bar{r}} = \frac{\sum_s v_s r_s g_s}{\sum_i v_i r_i} = \sum_s \frac{\text{GCPML}_s}{\sum_i \text{GCPML}_i} g_s = \sum_s \theta_s g_s,$$

which is (22). The weights are exposure shares; they reduce to value shares only if r_s is constant across sites. □

Proof of the masking identity. Since $\sum_s \theta_s = 1$, we may write $\bar{g} - g^{\text{med}} = \sum_s \theta_s (g_s - g^{\text{med}})$. Hence $\bar{g} < g^{\text{med}}$ if and only if $\sum_s \theta_s (g_s - g^{\text{med}}) < 0$, and in that case $\Gamma = g^{\text{med}} - \bar{g} > 0$. □

Proof of the reversal condition. From (22),

$$\bar{g} = \theta_d g_d + \sum_{i \neq d} \theta_i g_i.$$

Because $g_d < 0$, the condition $\bar{g} < 0$ is equivalent to

$$\theta_d |g_d| > \sum_{i \neq d} \theta_i g_i.$$

For the equivalence with (23), note that $\bar{g} < 0$ holds if and only if $\Gamma = g^{\text{med}} - \bar{g} > g^{\text{med}}$, so reversal is the case of masking whose gap exceeds the median change. $\square$

B. Implementation details

The eleven process areas.

Table 12: Common process architecture inherited from GPML

No.	Process area	No.	Process area
1	Management responsibility	7	Health and safety
2	Finance	8	Environment
3	Supply chain	9	Privacy / data protection
4	Sales and marketing	10	Logistics
5	Anti-bribery	11	Information technology
6	Human resources		

Recommended reproducibility package. A submission-ready replication package should contain: (i) an anonymised site identifier and sector; (ii) address-derived and verified coordinates at a disclosure-safe precision; (iii) raw hazard variables and climate windows; (iv) presence-screen decisions; (v) normalization thresholds and reference populations; (vi) D , I , F , λ , and ϕ matrices with evidence classes; (vii) audit items, scoring anchors, and M_{sp} ; (viii) sector corrections and production values; (ix) code for all tables, figures, and robustness exercises; and (x) a data-access statement for confidential items.