Map Impact

A Data Journey for Physical Climate Risk under SS5/25

Connecting Hazard Data Design, Regulatory Expectations, and Management Action

Map Impact Whitepaper, February 2026

Foreword: How to Use this Document & Its Purpose

The purpose of this 'Data Journey' is to support regulated firms in understanding, evidencing, and governing the use of physical climate hazard data within the expectations set out in PRA SS5/25. It is intended to provide a clear explanatory bridge between data design choices, regulatory workflows (including asset identification, materiality assessment, and scenario analysis), and the practical use of hazard insight to inform proportionate management action.

This document does not prescribe modelling approaches, replace firm-owned judgement, or introduce new risk methodologies. Rather, it is designed to improve explainability, transparency, and consistency in how hazard-led data is interpreted and used alongside existing risk management frameworks.

The 'Data Journey' is a supplementary narrative that connects (i) the technical construction of hazard data, (ii) the practical steps firms must undertake under PRA SS5/25, and (iii) how hazard insight can be translated into proportionate management action and increased portfolio resilience.

It is not designed to be read linearly by all readers. Instead, it functions as a reference 'journey map' that different stakeholders can dip into depending on their role. The Methodology Pack remains the authoritative source for technical method details, controls, and limitations; our prior whitepaper – 'Supporting Climate Risk Materiality Assessment under PRA SS5/25: A Hazard-Led Approach' remains the external positioning and context. Where referenced by firms in governance or supervisory dialogue, this document should be read in conjunction with the Map Impact Methodology Pack.

Recommended reading routes:

- **Board / ExCo (10–15 minutes)**: Read Sections 1, 5, 8, 9 and 10 for the strategic and governance implications.
- **First Line (risk owners / model owners)**: Read Sections 1–7 to understand how data choice affects attribution, materiality, and scenario behaviour; then Sections 8–10 for management action and guardrails.
- **Second Line (MRM / ERM / Compliance)**: Read Sections 1–7 with particular focus on Sections 4–7; validate that scope boundaries, interpretation guidance and guardrails are explicit.
- **Internal Audit / Assurance**: Focus on Sections 1, 4–7 and 10; confirm the audit trail of claims and limitations; validate consistency with the Methodology Pack.

- **PRA / supervisory specialist engagement**: Use Sections 1, 4–7 and 10 as the narrative bridge explaining why the data is fit for purpose in asset-referenced workflows, and how false precision is avoided.

How to use in governance:

- Use this document to support the 'why' behind data selection, especially when documenting materiality decisions and the rationale for property-level attribution approaches.
- Use the Map Impact Methodology Pack to support the 'how' (data lineage, transformations, QA/QC, versioning, limitations).
- Where a firm uses risk-embedded composite vendor outputs, use the comparison table to clarify classification, validation expectations, and whether embedded assumptions align to internal frameworks.

Key principle: the document is intended to improve explainability and decision quality, not to increase apparent precision. Where the document describes benefits of native or hazard-led approaches, those benefits relate to interpretability, governance, and actionability, rather than claims of predicting outcomes.

Lines of defence and PRA review map

Table 1 indicates the primary reviewers and typical review focus by section. (Sections as labelled in this Data Journey whitepaper: Section 1–10.)

Table 1. Primary reviewers and review focus.

Section	Primary reviewers	Typical questions	Evidence / checks	How to use
1. Regulatory problem	2LoD, PRA, Audit, 1LoD	Is the framing SS5/25-aligned and proportionate?	Consistency with SS5/25 themes; no over-claims	Use as the framing for governance packs and committee papers
2. Native vs downscaled	1LoD modellers, 2LoD MRM, PRA	What is the provenance of 'asset-level' outputs?	Method consistency; artefact risk recognised	Use when justifying dataset choice and limitations
3. Scale & tails	1LoD, 2LoD, PRA	Does the data preserve hotspots / extremes appropriately?	Distributional checks; tail behaviour explained	Use to justify materiality sensitivity to extremes
4. Asset attribution	1LoD, 2LoD, Audit, PRA	Can we defend how hazards are assigned to secured assets?	Lineage (asset→location→grid); spatial QA	Use in model inventory entries and attribution documents

Section	Primary reviewers	Typical questions	Evidence / checks	How to use
5. Materiality & concentration	2LoD, PRA, 1LoD	Are we likely to miss concentrations? Are tails compressed?	Materiality logic; sensitivity analysis approach	Use for documenting materiality decisions and thresholds
6. Scenario analysis	1LoD, 2LoD, PRA	Is this exploratory (not forecasting)? How do distributions shift?	Scenario provenance; interpretability; governance notes	Use in climate scenario narratives (ICAAP/ORSA)
7. Hazard to loss integration	1LoD, 2LoD MRM, Audit, PRA	Where do loss assumptions sit? Any embedded risk duplication?	Clear boundaries; firm-owned vulnerability/ transmission	Use to defend model architecture and split of responsibility
8. Management action	1LoD, 2LoD, ExCo/CRO	How does insight translate into action and resilience?	Proportionality; mitigation pathway; monitoring loop	Use in risk appetite and portfolio steering discussions
9. Strategic implications	Board/ExCo, CRO, 2LoD	Does this enable selectivity vs blunt overlays?	Consistency with risk appetite; governance clarity	Use as Board narrative for why data choice matters
10. Guardrails & conclusions	2LoD, Audit, PRA, Board	Are limitations explicit? Are we avoiding false precision?	Non-claims; permitted/prohibited uses; clarity	Use as the supervisory 'comfort' close and audit anchor

Legend: 1LoD = First Line (business/model owners). 2LoD = Second Line (MRM/ERM/Compliance). Audit = Third Line (Internal Audit/Assurance).

Document taxonomy and how the pack fits together

To reduce 'mode confusion' and streamline onboarding, the recommended document set is structured as a tiered suite. Each document has a distinct purpose, audience and governance role.

Table 2. Document taxonomy, audience and roles

Document	Primary purpose	Primary audience	Use / governance role
Technical Whitepaper – A hazard-led approach	Market + regulatory context; why hazard-led matters	Clients, partners, mixed technical	External-facing; should not introduce methodology details beyond what the Methodology Pack supports.
Map Impact Methodology Pack	Authoritative 'how'; lineage, transformations, QA/QC, limitations	1LoD modellers, 2LoD MRM, Audit, PRA	MRM-grade governance artefact; referenced in procurement and model inventory.
Technical Whitepaper – 'Data Journey'	Connect data choice → SS5/25 workflows → actionability	Risk, modelling, 2LoD, PRA	Reference narrative; supports explainability and fitness-for-purpose justification.
Data Journey (Executive)	Strategic implications, resilience and governance	Board, ExCo, CRO, NEDs	Board-safe narrative; complements the technical journey without technical depth.
MRM Onboarding Pack	Concise onboarding wrapper; scope, LoD, classification, sign-off	Procurement, 2LoD, Audit	Accelerates approvals; attaches to model inventory entries.
Comparison Tables	Quick discriminators and governance implications	2LoD, Audit, procurement	Used in vendor selection notes and audit trails
FAQs	Consistent answers for challenge and onboarding	All; especially procurement and 1LoD users	Reduces misinterpretation and ensures consistent messaging
Change Log + Versioning Policy	Governance over updates	2LoD, Audit	Supports ongoing monitoring and control requirements
Data Dictionary / Schema	Field-level definitions and technical integration	1LoD engineers/modellers	Implementation support; reduces integration ambiguity

Practical packaging recommendation:

- External share (technical): Whitepaper – Hazard-led approach + Whitepaper - Data Journey + Methodology Pack (core) + Comparison Tables/FAQ.
- Board pack: Whitepaper - Data Journey (Executive) + one-page summary of Methodology Pack scope/limitations.
- Onboarding with 2LoD: MRM Onboarding Pack + Methodology Pack + Whitepaper - Data Journey, with change log/versioning included.

Table of Contents

1. The Regulatory Problem SS5/25 Is Actually Trying to Solve 9

2. Two Technical Approaches to Climate Hazard Data: Native and Downscaled 10

2.1. Native hazard data 10

2.2. Downscaled hazard data 11

2.3. Why the distinction matters technically 12

2.4. A note on appropriateness rather than superiority 12

2.5. What is meant by "compressed tails" in downscaled hazard data? 13

2.6. The statistical intuition 14

2.7. A simplified example 14

2.8. Why this matters for physical climate hazards 14

2.9. Why native hazard data behaves differently 15

2.10. Governance and modelling implications 15

3. Why Spatial Scale and Distribution Shape Matter for Physical Climate Hazards 16

3.1. Spatial heterogeneity is not noise 16

3.2. Distribution shape: why tail behaviour drives materiality 17

3.3. Why tail behaviour matters for risk management 18

3.4. Implications for interpretation and governance 19

4. Secured Asset Identification and Hazard Attribution Under SS5/25 19

4.1. Asset location as the foundation 20

4.2. Hazard attribution using native data 21

4.3. Hazard attribution using downscaled data 21

4.4. Implications for materiality and control 22

4.5. Proportionality and fitness for purpose 22

4.6. Illustrative example: neighbouring secured properties 23

5. Materiality Assessment and Concentration Risk Under SS5/25 24

5.1. The effect of hazard construction on materiality outcomes 25

5.2. Implications for concentration risk 26

5.3. Interaction with the neighbouring-property example 26

5.4. Forward-looking materiality under scenario analysis 27

5.5. Governance and evidencing materiality decisions 27
5.6. Regulatory context: materiality under SS5/25 27
6. Scenario Analysis and Forward-Looking Assessment............................. 28
6.1. Scenario analysis under SS5/25: purpose and expectations 28
6.2. How hazard distributions change under scenarios......................... 29
6.3. Native hazard data and scenario interpretability 30
6.4. Governance considerations for scenario outputs............................. 30
7. From Hazard to Loss: Integration with Credit Risk and Capital Frameworks 31
7.1. Maintaining clear ownership and governance boundaries 31
7.2. How hazard data is typically used in credit risk frameworks 33
7.3. Why hazard construction matters for downstream modelling 34
7.4. Implications for ECL, stress testing, and capital assessment 34
7.5. Avoiding over-interpretation and false precision 35
7.6. Market context: composite hazard scores and embedded risk assumptions 35
7.7. Why this distinction matters for SS5/25 36
8. From Risk Identification to Management Action, Mitigation, and Resilience... 36
8.1. Why hazard-led insight is a prerequisite for action 37
8.2. The role of native hazard data in targeted mitigation...................... 38
8.3. Limitations of risk-embedded composite scores for management action... 38
8.4. Supporting resilience rather than defensive retrenchment 39
8.5. Governance benefits under SS5/25... 39
8.6. Illustrative example: management action under different data approaches 39
8.7. Using hazard-only, native data 40
8.8. Using a risk-embedded composite score 40
8.9. Governance and SS5/25 implications.. 40
8.10. Why this example matters.. 41
9. Competitive and Strategic Implications of Data Choice.......................... 41
9.1. Regulatory confidence as a strategic asset 41
9.2. Selectivity versus blunt risk responses 42

9.3. Enabling portfolio resilience over time 42
9.4. Strategic implications for capital and risk appetite 42
9.5. Data choice as a long-term capability decision 43
10. Limitations, Guardrails, and Concluding Principles 43
10.1. Recognising the limits of hazard data 43
10.2. Guardrails against false precision 44
10.3. Governance and accountability under SS5/25 44
10.4. Principles emerging from the journey 44
10.5. Concluding reflection 45

1. The Regulatory Problem SS5/25 Is Actually Trying to Solve

Supervisory Statement SS5/25 does not prescribe specific climate models, datasets, or methodologies. Instead, it sets expectations for how firms should think about, govern, and act on physical climate risk within their existing risk management frameworks.

At its core, SS5/25 is concerned with decision quality under uncertainty. Firms are expected to demonstrate that they can:

- Identify where physical climate risks arise within their portfolios,
- Assess whether those risks are material now and over relevant future horizons,
- Incorporate forward-looking climate considerations proportionately into risk management, and
- Explain and evidence how risk insight informs governance, strategy, and management action.

The emphasis is not on precision or prediction, but on appropriateness, explainability, and control.

This has important implications for the role of data. Under SS5/25, the question is not simply whether a firm has access to climate risk data, but whether the data it uses is fit for the decisions it supports, and whether its limitations are understood and governed.

In the context of secured lending, physical climate risk assessment introduces an additional layer of complexity. Unlike many traditional credit risk drivers, physical hazards such as heat stress, drought, and wildfire are:

- Spatially heterogeneous,
- Driven by environmental conditions that vary over short distances, and
- Expected to evolve over time in ways that are uncertain but directionally understood.

As a result, the choice of hazard data becomes a foundational decision. It determines:

- How credibly physical risk can be attributed to individual assets,
- Whether concentrations and hotspots can be identified with confidence,
- How materiality judgements are formed and defended, and
- How smoothly hazard insight can be integrated into downstream modelling and governance processes.

SS5/25 implicitly recognises this by placing strong emphasis on proportionality, transparency, and avoidance of false precision. Firms are expected to understand not only what their models and data show, but why they show it, and where those outputs should and should not be relied upon.

This creates a practical challenge. Many climate datasets present outputs at high apparent resolution, often labelled as 'property-level' or 'asset-level'. However, the technical approach used to generate those outputs varies significantly, and those differences materially affect how the data behaves when used for regulatory purposes. For firms selecting between data providers, understanding these differences is essential to ensuring that 'asset-level' outputs are genuinely fit for the regulatory purposes for which they are intended.

This document explores how two broad technical approaches to climate hazard data - native and downscaled - shape what firms can reasonably do under SS5/25, and how those choices influence not only compliance, but also the ability to manage and mitigate physical climate risk in practice.

2. Two Technical Approaches to Climate Hazard Data: Native and Downscaled

Climate hazard data used within financial risk management is often described as "high resolution" or "asset level". However, outputs that appear similar on the surface can be generated using fundamentally different technical approaches. These differences are not cosmetic; they directly influence how the data behaves when applied to secured portfolios, materiality assessment, and downstream modelling under SS5/25.

Broadly, climate hazard datasets fall into two categories; native and downscaled.

2.1. Native hazard data

Native hazard data is generated at, or close to, the spatial resolution at which the relevant physical processes are directly observed or modelled. In this approach, the hazard signal arises from locally specific environmental conditions, rather than being inferred from broader regional averages.

In practice, native hazard datasets are typically constructed using combinations of high-resolution inputs such as:

- Earth observation data (e.g. land surface temperature, vegetation condition, moisture proxies),
- Terrain and topographic information,
- Soil and geological characteristics,
- Land cover, habitat type, and land management indicators.

Each location is assessed independently based on its own physical attributes. Spatial variation in hazard therefore reflects measured or directly derived differences in environmental conditions, rather than mathematical interpolation across space.

This approach preserves:

- Sharp spatial gradients,
- Local extremes and hotspots,
- The influence of land cover, soils, and built environment, and
- Heterogeneity between neighbouring locations.

From a technical perspective, native hazard data produces outputs whose spatial structure is grounded in observed or physically modelled processes.

2.2. Downscaled hazard data

Downscaled hazard data takes a different approach. It begins with coarse-resolution climate model outputs, often produced at grid sizes of tens of kilometres, and applies statistical or mathematical techniques to generate values at finer spatial resolution.

Common downscaling techniques include:

- Linear or spline interpolation,
- Bias correction,
- Regression against proxy variables, and
- Statistical disaggregation of coarse fields.

The resulting fine-resolution outputs do not introduce new physical observations. Instead, they redistribute a regional-scale climate signal across smaller grid cells according to the assumptions embedded in the downscaling method.

Downscaled data can be valuable for understanding broad climatic trends, particularly over large areas or long time horizons. However, its fine-scale detail is inherently inferred rather than observed.

2.3. Why the distinction matters technically

The difference between native and downscaled approaches becomes most apparent when data is used at the level of individual assets or secured properties.

Native hazard data:

- Reflects local physical conditions that can change over short distances,
- Allows neighbouring assets to legitimately exhibit different hazard characteristics, and
- Can be interrogated in terms of underlying environmental drivers.

Downscaled hazard data:

- Inherits the smoothness of the original coarse climate model grid,
- Tends to dampen local extremes and compress tails, and
- Can introduce artefacts that appear precise but are not physically grounded at that scale.

From a validation and governance perspective, this has important consequences. Fine-scale outputs derived from downscaling cannot be independently validated at that same scale, because the underlying signal did not originate there. As a result, distinguishing between genuine local variation and mathematically induced patterns can be challenging.

2.4. A note on appropriateness rather than superiority

It is important to note that this distinction is not about one approach being universally 'better' than the other. Each has contexts in which it is appropriate.

Downscaled climate data is often well suited to:

- Macro-level scenario narratives,
- Regional or national trend analysis, and
- Strategic, high-level exploration of future climate pathways.

Native hazard data becomes increasingly important where:

- Decisions depend on asset-level differentiation,
- Concentrations and hotspots are relevant, and
- Outputs are expected to support materiality assessment, governance, and action under SS5/25.

The relevance of this distinction therefore depends not only on the dataset itself, but on how and where the data is intended to be used within the firm's risk framework.

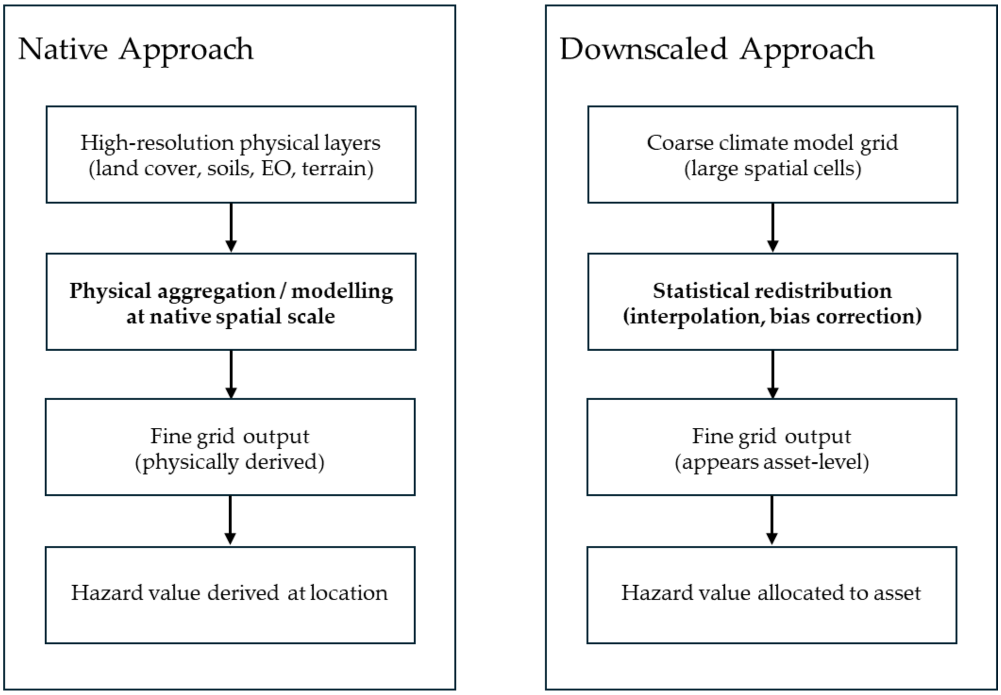

Figure 1. Conceptual comparison of downscaled and native approaches to producing asset-referenced hazard data. Downscaled outputs redistribute a coarse regional climate signal to finer grids, whereas native outputs are derived directly from high-resolution physical inputs.

As illustrated in Figure 1, datasets described as "asset-level" may be produced through fundamentally different technical pathways, with important implications for how hazard values should be interpreted and governed.

A further distinction that is often critical when selecting between climate data providers is whether a dataset presents hazard in isolation, or whether hazard has already been blended with assumptions about exposure, vulnerability, probability, or loss. Where these elements are combined into a single "asset score," the output no longer represents a physical hazard condition but a partially modelled risk construct. This distinction becomes important later when considering model governance, materiality assessment, and integration into firm-owned risk frameworks.

2.5. What is meant by "compressed tails" in downscaled hazard data?

In this context, "compressed tails" refers to the tendency of downscaled climate hazard data to under-represent extreme local conditions at the high (or low) end of the hazard distribution when values are inferred from coarse-resolution inputs.

This effect arises from the mechanics of downscaling rather than from any single modelling error.

2.6. The statistical intuition

Most downscaling methods begin with a coarse-resolution climate field, in which each grid cell represents an average over a large spatial area. When that field is redistributed to finer resolution, the process typically preserves the mean of the coarse cell, and sometimes its variance, but it does not preserve local extremes that were never observed at the coarse scale. As a result, the inferred fine-scale values tend to cluster closer to the regional mean. The upper and lower extremes of the distribution, the "tails", are therefore narrower than they would be if the hazard were measured or modelled natively at that scale.

This is what is meant by "tail compression".

2.7. A simplified example

Consider a coarse climate model grid cell covering a 25 × 25 km area. Within that area some locations may experience very high heat stress due to urban density and lack of vegetation, whilst others may remain relatively cooler due to shading, elevation, or land cover. If the coarse model outputs a single average value for that entire cell, that value already suppresses local extremes.

When that average is downscaled to a 50 m or 100 m grid, the resulting fine-scale values are variations around the same underlying mean. Even if statistical adjustments are applied, the most extreme local conditions cannot exceed what the coarse model "saw" in the first place.

The distribution of fine-scale hazard values therefore has less mass in the extreme tails than a distribution derived from physically observed or high-resolution inputs.

2.8. Why this matters for physical climate hazards

Physical climate hazards such as heat stress, drought susceptibility, and wildfire exposure are often tail-driven:

- A small proportion of locations may account for a disproportionate share of exposure or management concern.
- Concentration risk emerges from extremes, not averages.

If tail behaviour is compressed:

- Hotspots appear less severe,
- Concentrations appear weaker,
- Materiality thresholds may not be crossed when they should be.

This does not mean that downscaled data is "wrong", but it does mean that it can systematically understate the severity and concentration of local hazard extremes when used at asset-level resolution.

2.9. Why native hazard data behaves differently

Native hazard data does not start from a spatially averaged climate signal. Instead, it is constructed from locally varying physical inputs (e.g. land surface temperature, vegetation stress, soil characteristics).

As a result:

- Extremes emerge naturally where physical conditions support them,
- The distribution of hazard values exhibits longer and more differentiated tails, and
- High-hazard locations remain distinguishable from the rest of the portfolio.

This leads to a fuller representation of the hazard distribution, particularly at the upper end where risk management attention is typically focused.

2.10. Governance and modelling implications

From a modelling and regulatory perspective, compressed tails matter because they can:

- Reduce apparent concentration risk,
- Weaken the signal used for materiality assessment,
- Flatten relationships when hazard variables are used for segmentation,
- Make it harder to justify why certain assets warrant closer scrutiny.

Under SS5/25, where firms are expected to demonstrate proportionate, forward-looking assessment, understanding whether tails are preserved or compressed is a key part of assessing data fitness for purpose.

The implications of this behaviour for spatial heterogeneity, materiality assessment, and risk governance are examined in the next section.

Understanding these two technical approaches is a necessary foundation. The next step is to examine why spatial scale and physical heterogeneity matter specifically for physical climate hazards, and how this influences the credibility of hazard attribution in secured lending portfolios. Downscaled datasets may exhibit internal mathematical consistency at fine spatial resolution; however, the fine-scale variation they display remains conditioned on the scale and information content of the driving model.

3. Why Spatial Scale and Distribution Shape Matter for Physical Climate Hazards

Physical climate hazards differ from many traditional risk drivers in that they are inherently spatial and heterogeneous. The relevance of this behaviour lies not in the mechanics of downscaling itself, but in how distribution shape influences concentration, materiality assessment, and management judgement under SS5/25. Heat stress, drought susceptibility, and wildfire exposure do not vary smoothly across geography. Instead, they are shaped by combinations of local factors such as land cover, soil type, topography, vegetation condition, and the built environment. These factors can change materially over short distances.

As a result, the spatial scale at which hazard data is constructed has a direct influence on how accurately it represents real-world conditions, particularly at the level of individual assets or secured properties.

3.1. Spatial heterogeneity is not noise

In the context of physical climate hazards, local variation should not be treated as statistical noise to be smoothed away. It is often the primary driver of exposure and concentration.

For example:

- Heat stress can differ markedly between shaded, vegetated areas and densely built urban environments only a few hundred metres apart.
- Drought susceptibility is strongly influenced by soil type and local hydrology, which can vary within the same town or postcode.
- Wildfire exposure is shaped by fine-scale patterns of fuel load, land management, and the urban–rural interface.

Where hazard data preserves this heterogeneity, it enables a more faithful representation of where environmental stress is genuinely elevated. Where it does not, important local extremes can be diluted.

This heterogeneity is not incidental. For hazards such as heat stress, drought susceptibility, and wildfire exposure, land cover, habitat condition, soil characteristics, and surface properties are first-order drivers of how environmental stress manifests locally. Datasets that incorporate these factors directly into hazard construction are able to reflect meaningful short-distance variation grounded in physical conditions, rather than relying primarily on redistribution of broader climatic signals.

This linkage between land, habitat condition, and physical hazard is increasingly relevant in the context of emerging regulatory thinking around the climate–nature nexus. Industry and supervisory workstreams have begun to highlight that nature condition is not an external overlay to climate risk, but a factor that directly influences how hazards such as heat, drought, and wildfire manifest locally. Where hazard construction already incorporates these drivers, firms are less likely to face future challenges around overlapping metrics, double counting, or the need to reconcile separate "nature risk" overlays with existing climate hazard assessments.

3.2. Distribution shape: why tail behaviour drives materiality

As discussed in the previous section, downscaled hazard data tends to compress the upper and lower tails of the hazard distribution as a structural consequence of inferring fine-scale values from coarse climate model outputs.

The importance of this behaviour lies not in the mechanics of downscaling itself, but in how distribution shape influences risk identification, materiality assessment, and management attention under SS5/25.

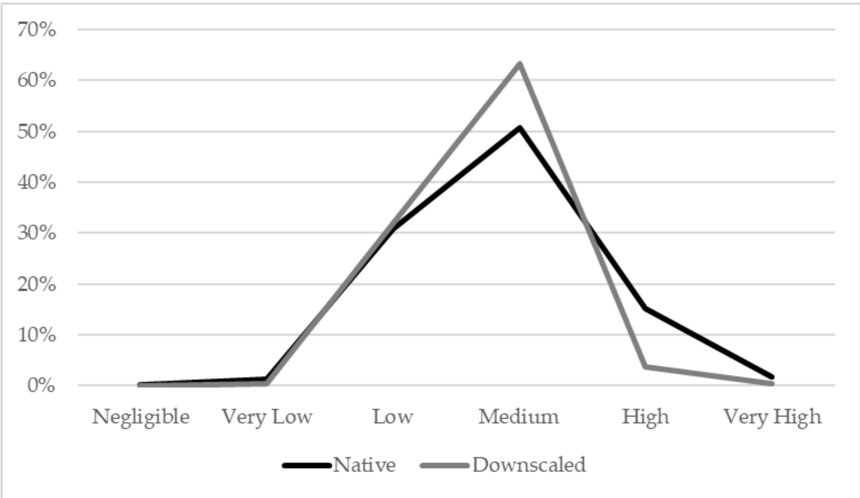

Figure 2. Illustrative comparison of hazard distributions. Native hazard data preserves local extremes, resulting in a longer upper tail. Downscaled data tends to compress extremes towards the regional mean. Percentages indicate the proportion of assets falling within each relative hazard category. Categories reflect relative hazard intensity only and do not represent probability of damage, loss, or impairment.

Figure 2 provides a conceptual illustration of this effect, showing how downscaling tends to compress extremes towards the regional mean, whereas native hazard data preserves a more differentiated upper tail.

In this context, compressed tails arise because downscaling starts from a spatially averaged climate signal produced at coarse resolution. Even when that signal is redistributed to finer grids using statistical techniques, the most extreme local conditions cannot exceed what was present in the original coarse cell. The redistribution process therefore preserves the regional mean, but pulls fine-scale values back towards that mean. For discussion of distributional effects and inherited constraints in statistical downscaling, see von Storch et al. (1999); Maraun et al. (2010); Maraun (2013).

The practical effect is that:

- Extreme local hazard values are dampened,
- The mass of the distribution shifts away from the tails and towards the centre, and
- The distinction between the highest-hazard locations and the rest of the portfolio is reduced.

This is not a flaw in implementation; it is a structural consequence of inferring fine-scale detail from coarse inputs.

3.3. Why tail behaviour matters for risk management

For physical climate hazards, tail behaviour matters disproportionately. Materiality, concentration risk, and management attention are typically driven by a relatively small subset of locations where hazard intensity is highest. If those extremes are compressed:

- Hotspots appear less severe,
- Concentrations appear weaker,
- Threshold-based assessments may fail to identify risks that are genuinely significant.

This can lead to systematic understatement of where physical climate risk is most acute, particularly when analysis is conducted at asset or collateral level.

Native hazard data behaves differently in this respect. Because it is constructed from locally varying physical inputs, extreme values emerge naturally where environmental conditions support them. The upper tail of the distribution is therefore more differentiated, allowing high-hazard locations to remain distinct from the wider population.

3.4. Implications for interpretation and governance

From a governance and validation perspective, the treatment of distribution tails has several implications:

- Materiality assessment: Compressed tails can delay or weaken the identification of risks as material, particularly where materiality depends on concentration rather than average exposure.
- Segmentation and sensitivity analysis: Flattened distributions reduce the explanatory power of hazard variables when used to segment portfolios or explore sensitivity.
- Explainability: It becomes harder to justify why certain assets should be prioritised for further analysis or management action if hazard values cluster tightly around a regional mean.

Under SS5/25, firms are expected to demonstrate that their assessment of physical climate risk is proportionate, forward-looking, and grounded in an understanding of the underlying drivers. Recognising how spatial scale and distribution shape influence hazard outputs is therefore a critical part of evaluating whether a dataset is fit for purpose.

Having established why spatial scale and distribution shape matter, the next section examines how these technical differences translate into practical challenges and opportunities when identifying and attributing physical climate hazard to secured assets within regulated lending portfolios.

4. Secured Asset Identification and Hazard Attribution Under SS5/25

A central expectation within SS5/25 is that firms can identify where physical climate risks arise within their portfolios, and do so in a way that is transparent, proportionate, and defensible. For secured lending portfolios, this requirement has a very practical interpretation; firms must be able to demonstrate that hazard attributes have been applied appropriately to the assets securing their exposures.

This step, moving from abstract climate risk to asset-referenced hazard, is the point at which differences in data construction become most consequential.

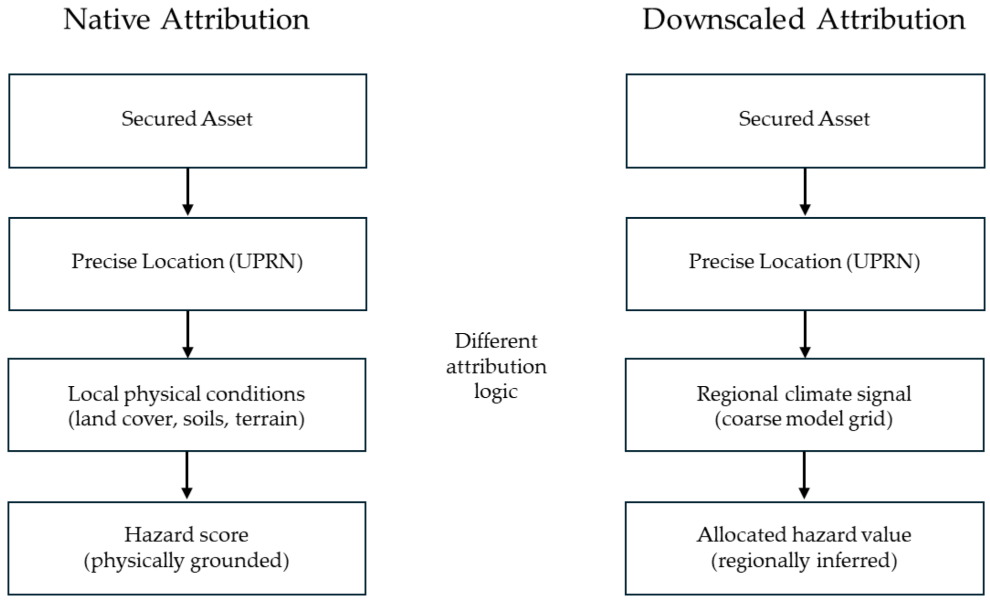

Figure 3. Asset-level hazard attribution pathways. Native approaches derive hazard from local physical conditions, while downscaled approaches allocate a regional signal to individual assets.

Figure 3 summarises the two attribution pathways, highlighting the distinction between hazard values derived from local physical conditions and those allocated from a regional climate signal. This distinction underpins how asset-level hazard assignments can be explained to second line functions and supervisors.

4.1. Asset location as the foundation

Secured lending risk assessment is fundamentally anchored in asset location. Properties are immovable, and their exposure to physical climate hazards is determined primarily by the environmental conditions of the location in which they sit.

As a result, the starting point for physical climate risk assessment under SS5/25 is typically:

- Accurate identification of the secured asset,
- Reliable geolocation (for example, via UPRN or equivalent reference), and
- Consistent linkage between that location and hazard data.

This creates an implicit expectation of traceability. Firms should be able to explain, at least in principle, how a hazard attribute assigned to an exposure was derived from its physical location.

This traceability is strengthened where hazard datasets are designed specifically around UK property referencing, land systems, soils, and environmental data structures. Datasets developed to operate consistently across many countries must necessarily generalise methods across differing climates and data regimes, whereas UK-native constructions can align more directly to the physical, geographic, and property characteristics of secured lending portfolios governed under PRA expectations.

For firms undertaking vendor due diligence or model onboarding, traceability becomes a key test of whether 'asset-level' claims are technically defensible.

4.2. Hazard attribution using native data

Where native hazard data is used, the attribution process aligns naturally with this expectation.

In this approach:

- Each secured asset is linked to a specific location,
- That location is associated with locally derived hazard indicators,
- The resulting hazard score reflects physical conditions observed or modelled at, or close to, that scale.

The firm can therefore articulate a clear attribution chain:

Secured asset → location → local environmental conditions → hazard score

This supports governance in several ways. First, it allows differences in hazard exposure between neighbouring assets to be explained in terms of observable physical drivers, rather than statistical allocation. Second, it enables internal challenge and review functions to interrogate the plausibility of hazard assignments by reference to known features of the environment.

From a supervisory perspective, this strengthens confidence that hazard attribution is physically grounded, even where uncertainty remains about future outcomes.

4.3. Hazard attribution using downscaled data

When downscaled hazard data is used, the attribution process takes on a different character.

Although fine-resolution outputs may be available, the underlying hazard signal originates from a coarse spatial representation of climate conditions.

The fine-scale values assigned to individual assets are therefore the result of statistical redistribution, not locally observed or modelled conditions.

In this case, hazard attribution is more accurately described as:

Secured asset → location → allocation of a regional hazard signal

This does not make the data unusable, but it does change how it should be interpreted and governed. Differences in hazard scores between nearby assets may reflect artefacts of the downscaling process rather than meaningful physical variation. Conversely, genuinely higher-risk locations may not stand out clearly from their surroundings.

Under SS5/25, firms using downscaled data at asset level therefore face an additional governance burden; they must be clear that the hazard attribute represents regional exposure allocated to assets, rather than asset-specific hazard conditions.

4.4. Implications for materiality and control

The distinction between these two attribution pathways has knock-on effects throughout the risk management process. Where hazard attribution is physically grounded:

- Concentrations of exposure can be identified more confidently,
- Materiality assessments are less reliant on judgement alone, and
- Internal challenge focuses on interpretation rather than data credibility.

Where hazard attribution is based on downscaled allocation:

- Materiality judgements may need to be framed more cautiously,
- Sensitivity analysis becomes more important,
- Firms may choose to aggregate exposures before drawing conclusions.

SS5/25 does not prohibit either approach, but it does require firms to understand and articulate the implications of their choice. Using asset-level outputs derived from downscaled data without acknowledging these limitations risks overstating precision and weakening governance.

4.5. Proportionality and fitness for purpose

A recurring theme in SS5/25 is proportionality. Hazard attribution does not need to be perfect, but it does need to be fit for purpose.

Native hazard data tends to be most appropriate where:

- Asset-level differentiation matters,
- Concentration risk is a concern,
- Hazard insight is expected to inform management action.

Downscaled data may be sufficient where:

- The objective is high-level screening,
- Analysis is conducted at regional or portfolio level,
- Asset-level precision is not relied upon.

Making this distinction explicit allows firms to align their data choices with their stated use cases, reducing the risk of challenge from second line, internal audit, or supervisors.

4.6. Illustrative example: neighbouring secured properties

Consider two residential properties located on the same street, approximately 200 metres apart, and securing similar mortgage exposures.

- Property A is adjacent to unmanaged vegetation, sits on south-facing ground, and is surrounded by low tree canopy and heat-retaining surfaces.
- Property B is bordered by managed green space, benefits from shading, and sits on soils with higher moisture retention.

Using native hazard data, these differences are reflected directly in the hazard attribution. Each property is assigned a hazard score derived from locally observed or modelled physical conditions. The firm can therefore explain why Property A exhibits higher heat stress or wildfire susceptibility than Property B, despite their proximity.

Using downscaled hazard data, both properties may be assigned similar or identical hazard values. The underlying regional climate signal is the same for both locations, and fine-scale variation reflects statistical redistribution rather than measured environmental differences. Any differentiation between the two properties is therefore limited, even where local conditions diverge materially.

From a governance perspective, this distinction matters. In the first case, the firm can justify asset-level differentiation and targeted management action. In the second, the firm must recognise that the hazard attribution represents a regional exposure applied to individual assets and should avoid over-interpreting small differences between neighbouring properties.

Having established how hazard attribution differs at the asset level, the next section examines how these differences play out when firms undertake materiality assessment and concentration analysis, which sit at the heart of SS5/25's expectations for physical climate risk management. As with all location-referenced analysis, residual uncertainty may arise from address matching, boundary definitions, and spatial reference systems, and should be reflected in firm governance and controls.

5. Materiality Assessment and Concentration Risk Under SS5/25

Materiality assessment sits at the centre of SS5/25. Firms are expected to determine which physical climate risks are, or are likely to become, material to their business, and to justify those conclusions in a way that is forward-looking, proportionate, and evidence-based. The technical approach used to construct hazard data has a direct influence on how this assessment behaves in practice. Materiality is driven by distributions, not averages. Materiality remains a management judgement informed by evidence, rather than a mechanical output of any single dataset.

In secured lending portfolios, physical climate risk is rarely evenly distributed. Materiality typically emerges from:

- Concentrations of exposure in higher-hazard locations, and
- Tail behaviour, where a relatively small subset of assets accounts for a disproportionate share of potential concern.

As a result, materiality assessment is sensitive not only to average hazard levels, but to the shape of the hazard distribution across the portfolio. Where hazard data preserves local extremes and heterogeneity, the upper tail of the distribution remains differentiated. This allows firms to identify:

- Clusters of higher exposure,
- Geographic hotspots,
- Segments of the portfolio that warrant deeper analysis.

Where hazard data smooths or compresses extremes, those signals are weakened.

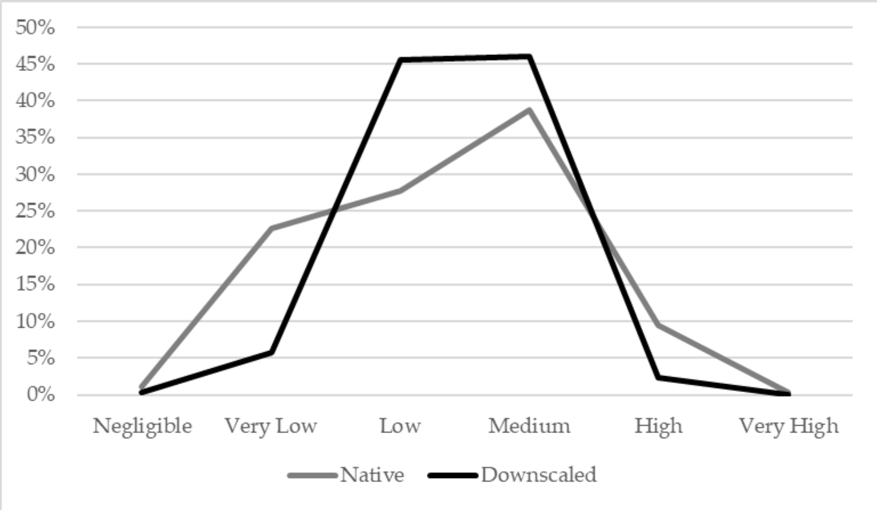

Figure 4. Materiality in physical climate risk is typically driven by concentrations in the upper tail of the hazard distribution, rather than portfolio averages. Percentages indicate the proportion of assets falling within each relative hazard category. Categories reflect relative hazard intensity only and do not represent probability of damage, loss, or impairment.

As illustrated in Figure 4, materiality in physical climate risk typically emerges from concentrations in the upper tail of the hazard distribution, rather than from portfolio-wide averages.

5.1. The effect of hazard construction on materiality outcomes

Using native hazard data, materiality assessment benefits from:

- Clearer separation between typical and high-hazard assets,
- More stable identification of hotspots over time,
- Reduced reliance on judgement to infer where risk may be concentrated.

In this context, a firm can demonstrate that its assessment of materiality is grounded in physically observed or modelled variation, rather than inferred averages.

Using downscaled hazard data, materiality assessment often behaves differently:

- Hazard values cluster more tightly around regional means,
- The upper tail of the distribution is compressed,
- Local concentrations may be diluted or missed entirely.

This does not invalidate the analysis, but it changes the nature of the conclusions that can be drawn. Firms may find that hazards appear less material at asset level, even where underlying environmental conditions suggest otherwise.

5.2. Implications for concentration risk

SS5/25 places emphasis on understanding not just whether risks exist, but also where they are concentrated. Native hazard data supports this by allowing concentration to be identified through:

- Genuinely elevated hazard scores in specific locations,
- Persistence of those elevations across time horizons,
- Coherence with known environmental drivers.

This enables firms to articulate concentration risk in a way that is spatially and physically credible.

With downscaled data, concentration risk often appears more diffuse. Regional patterns may be visible, but fine-scale clustering is harder to distinguish from noise. As a result, firms may be forced to rely on broader geographic aggregation, or conservative overlays applied uniformly across large areas.

Whilst both approaches are defensible, they are less targeted.

5.3. Interaction with the neighbouring-property example

Returning to the example of two neighbouring properties on the same street, the implications for materiality become clear.

Where native hazard data reveals that Property A consistently falls into a higher hazard category than Property B, that differentiation contributes to the formation of a longer upper tail in the portfolio distribution. As similar cases accumulate, a cluster of higher-hazard assets emerges, supporting a conclusion that the hazard may be material for that segment of the book.

Where downscaled data assigns similar values to both properties, those distinctions are absorbed into the average. Even if many such pairs exist across the portfolio, the distribution may fail to exhibit a sufficiently pronounced tail to trigger materiality thresholds. The difference lies not in the number of assets, but in the ability of the data to surface extremes.

Where hazard, vulnerability, and loss assumptions are combined into single asset metrics, tail behaviour can be further obscured, because the resulting distribution reflects not only environmental variation but also embedded modelling assumptions. In such cases, it can become difficult for firms to determine whether apparent concentrations arise from genuine physical hazard, or from the way vulnerability has been parameterised within the tool. This adds an additional layer of interpretation when undertaking materiality assessment under SS5/25, particularly where materiality is driven by localised extremes rather than portfolio averages.

5.4. Forward-looking materiality under scenario analysis

SS5/25 expects firms to assess materiality not only on a current basis, but under plausible future climate pathways. Here again, the behaviour of the hazard distribution matters. Native hazard data allows firms to observe:

- How the upper tail of the distribution expands over time,
- Whether new hotspots emerge or existing ones intensify, and
- How concentration risk migrates geographically.

This supports a more nuanced discussion of when and how a risk may become material, rather than a binary present/not-present judgement.

Downscaled data can still support forward-looking assessment, but scenario-driven changes may appear more uniform across regions, making it harder to identify which parts of the portfolio drive future materiality.

5.5. Governance and evidencing materiality decisions

From a governance perspective, SS5/25 requires firms to be able to evidence and defend their materiality assessments. Where hazard data preserves heterogeneity and tail behaviour:

- Materiality decisions can be linked to observable portfolio characteristics,
- Internal challenge focuses on thresholds and interpretation,
- Audit trails are clearer.

Where hazard data compresses extremes:

- Materiality decisions rely more heavily on judgement,
- Sensitivity analysis becomes critical,
- Firms should be explicit about the limitations of the underlying data.

Neither approach is prohibited, but failing to recognise these differences increases the risk of supervisory challenge.

5.6. Regulatory context: materiality under SS5/25

SS5/25 makes clear that firms are expected to undertake proportionate, risk-based materiality assessments that reflect the nature, scale, and complexity of their exposures, and that are capable of identifying concentrations and emerging risks over time, rather than relying solely on average or aggregate measures.

In particular, the supervisory expectation is that materiality judgements should be:

- Grounded in an understanding of where risk is concentrated,
- Sensitive to tail exposures and non-linear effects,
- Capable of supporting forward-looking assessment under plausible climate pathways.

In this context, the ability of hazard data to preserve or suppress extremes is not a technical detail but a substantive consideration. Where the construction of hazard data compresses the upper tail of the distribution, firms may find it more difficult to evidence why certain physical climate risks should be considered material, particularly when materiality is driven by localised concentrations rather than portfolio-wide averages.

Conversely, hazard data that preserves heterogeneity and tail behaviour provides a clearer evidential basis for materiality decisions that are consistent with the intent of SS5/25, even where uncertainty remains.

Having established how data construction influences materiality and concentration risk, the next section examines how these differences affect scenario analysis and forward-looking assessment, which SS5/25 treats as a core component of effective physical climate risk management.

6. Scenario Analysis and Forward-Looking Assessment

Scenario analysis plays a central role in SS5/25, not as a forecasting exercise, but as a structured way for firms to explore how physical climate risks may evolve over time and how those changes could affect the distribution of exposures within their portfolios. The technical characteristics of hazard data materially influence how informative and defensible this exercise becomes.

6.1. Scenario analysis under SS5/25: purpose and expectations

SS5/25 is explicit that scenario analysis should be exploratory and forward-looking, recognising uncertainty rather than attempting to predict specific outcomes. Firms are expected to use scenarios to test resilience, identify vulnerabilities, and inform strategic and risk management decisions.

In this context, the quality of scenario analysis is judged less on numerical precision and more on whether it:

- Captures plausible pathways of change,
- Highlights where risk may intensify or emerge,

- Supports informed discussion and action.

Hazard data therefore needs to behave in a way that allows meaningful interpretation of change, rather than simply producing different numbers under different scenarios.

In datasets derived primarily from climate model outputs, even present-day hazard is inherently dependent on model baselines and scenario framing. By contrast, where hazard is constructed from current environmental condition, present hazard can be explained independently of scenario selection, with future scenarios acting to stress an already physically grounded baseline. This distinction improves interpretability when explaining both current exposure and forward-looking change.

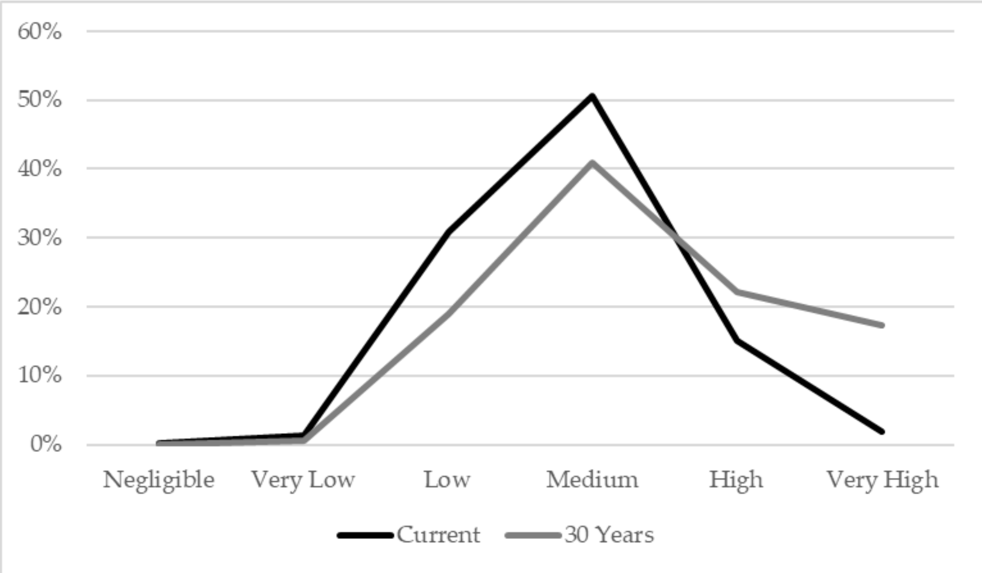

Figure 5. Scenario analysis explores how the distribution of hazard exposure across a portfolio may change over time. It is directional and exploratory, not predictive. Percentages indicate the proportion of assets falling within each relative hazard category. Categories reflect relative hazard intensity only and do not represent probability of damage, loss, or impairment.

Figure 5 illustrates how scenario analysis under SS5/25 is intended to explore changes in the distribution of exposure over time, rather than to forecast specific outcomes.

6.2. How hazard distributions change under scenarios

A key output of scenario analysis is not the absolute level of hazard at a given future date, but how the distribution of hazard across a portfolio evolves. When hazard data preserves spatial heterogeneity and tail behaviour, scenario analysis can reveal:

- Expansion of the upper tail of the distribution,
- Migration of assets into higher hazard categories,
- Emergence of new hotspots, and
- Increasing concentration of exposure in specific locations.

These distributional shifts are often more informative for risk management than changes in portfolio averages.

By contrast, where hazard data is derived from downscaled regional signals, scenario-driven changes tend to appear more uniform. Hazard levels may increase across broad areas, but differentiation between locations remains limited. As a result, it becomes harder to identify which parts of the portfolio drive future risk.

6.3. Native hazard data and scenario interpretability

Native hazard data supports scenario analysis by enabling physically coherent changes in hazard intensity across space.

As climate drivers evolve under different pathways:

- Locations with characteristics that amplify hazard respond more strongly,
- Local extremes intensify rather than simply shifting the mean,
- Spatial patterns of risk become clearer over time.

This allows firms to interpret scenario outputs in terms of:

- Where risk is growing fastest,
- Which assets transition into higher hazard bands,
- How concentration risk evolves.

Importantly, this does not imply greater certainty. Rather, it provides a more informative structure for uncertainty, which is aligned with the intent of SS5/25.

6.4. Governance considerations for scenario outputs

From a governance perspective, SS5/25 requires firms to be explicit about how scenario outputs should and should not be used. Regardless of the underlying data, scenario-adjusted hazard outputs:

- Do not represent probabilities,
- Do not model events or losses,
- Should not be interpreted as forecasts.

However, the interpretability of those outputs differs depending on data construction.

Where native hazard data is used, firms are better placed to:

- Explain why certain assets exhibit larger changes under scenarios,
- Link changes to known physical drivers,
- Justify why particular segments warrant closer scrutiny or management attention.

Where downscaled data is used, firms may need to rely more heavily on aggregation and narrative explanation, recognising that fine-scale variation reflects redistribution rather than observed physical change. Making these distinctions explicit reduces the risk of misinterpretation and strengthens internal and supervisory confidence.

Scenario analysis under SS5/25 is not an end in itself. Its value lies in how insights are carried forward into risk modelling, loss estimation, and capital frameworks, without overstating the role of hazard data or blurring governance boundaries.

The next section examines how hazard-led insights can be integrated into credit loss, stress testing, and capital assessment processes in a disciplined and regulator-safe manner.

7. From Hazard to Loss: Integration with Credit Risk and Capital Frameworks

Under SS5/25, firms are not expected to predict losses arising from physical climate hazards. They are, however, expected to demonstrate that hazard insights are coherently integrated into their existing risk management frameworks, including credit risk, stress testing, and capital assessment, in a way that is proportionate, explainable, and governed. This distinction is critical - physical climate hazard data is an input, not an outcome.

7.1. Maintaining clear ownership and governance boundaries

This separation becomes particularly important when comparing hazard-led datasets with platforms that provide composite asset scores combining hazard, vulnerability, and financial loss logic. In those cases, firms are not consuming a hazard input but a pre-modelled risk output. This can create additional complexity for model risk management, as firms must either accept embedded assumptions, attempt to reverse-engineer them, or risk double counting when integrating with internal credit and capital models.

A foundational principle under SS5/25 is that firms retain ownership of their risk models, assumptions, and decisions. Third-party hazard data should inform, but not replace, internal modelling and judgement.

A regulator-safe integration pathway therefore separates responsibilities clearly:

- Hazard data identifies relative exposure to physical climate drivers.
- Vulnerability assumptions translate hazard into potential damage mechanisms (firm-owned).
- Transmission channels link damage to credit outcomes such as arrears, default, recovery, or collateral value (firm-owned).
- Loss and capital models quantify impacts within governed frameworks such as IFRS 9, stress testing, ICAAP, or Pillar 2 (firm-owned).

This separation is not a limitation; it is a governance strength. This distinction is reinforced where hazard datasets are constructed as deterministic surfaces representing relative environmental intensity, rather than being derived from probabilistic event-based modelling traditions. Some climate risk platforms originate from catastrophe modelling approaches in which hazard is closely linked to assumptions about event frequency, return periods, and damage mechanisms, even where these are not presented explicitly. Deterministic hazard surfaces, by contrast, avoid importing implicit probability or loss logic, allowing firms to retain full control over how hazard translates into vulnerability, credit impact, and capital treatment.

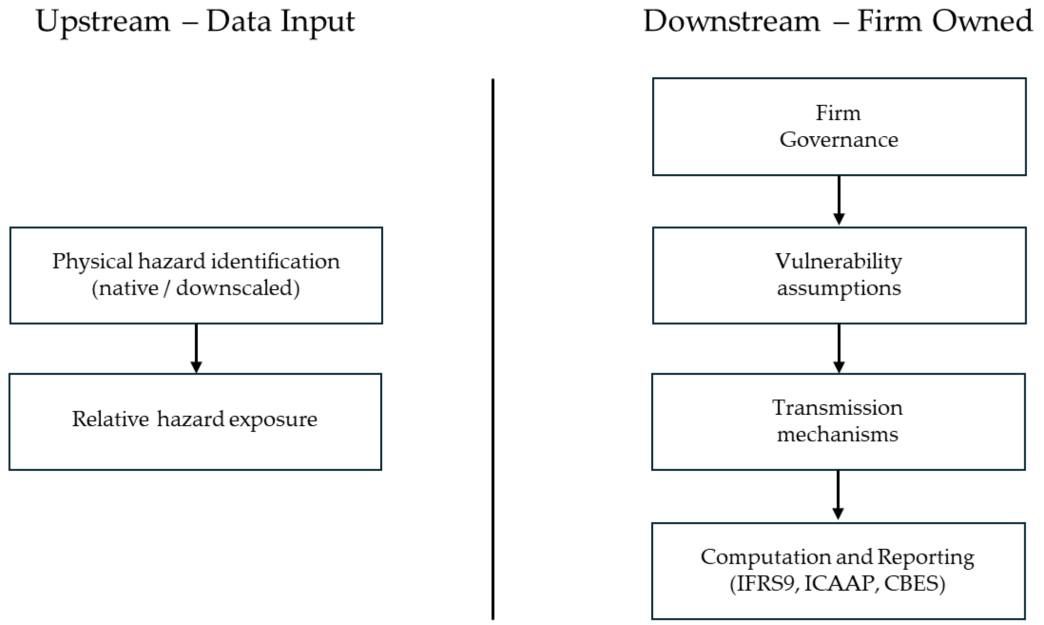

Figure 6. Governance separation between hazard identification and firm-owned modelling of vulnerability, loss, and capital impact.

Figure 6 highlights the governance separation between upstream hazard identification and downstream, firm-owned modelling of vulnerability, loss, and capital impact. Maintaining this separation supports model risk management by avoiding the inadvertent import of third-party loss assumptions.

7.2. How hazard data is typically used in credit risk frameworks

In practice, physical hazard information is most commonly used to condition, rather than determine, credit risk outcomes. Typical applications include:

- Segmentation of exposures into hazard bands,
- Sensitivity analysis to explore how credit metrics respond under different hazard assumptions, and
- Scenario overlays to test resilience under forward-looking pathways.

In all cases, hazard data acts as a contextual driver, not a direct predictor of loss.

7.3. Why hazard construction matters for downstream modelling

The technical properties of hazard data influence how effectively it can be integrated into these frameworks. Where hazard data preserves spatial heterogeneity and tail behaviour:

- Segmentation produces clearer differentiation between exposures,
- Relationships between hazard bands and credit outcomes are more interpretable,
- Consistent directional movement and stability are easier to assess, and
- Validation focuses on assumptions rather than data credibility.

This does not imply stronger statistical relationships, but rather cleaner modelling behaviour that is easier to govern and explain.

Where hazard data is derived from downscaled regional signals:

- Hazard bands may be less differentiated at asset level,
- Relationships can appear flatter or noisier,
- Firms may need to aggregate exposures to achieve stability.

Again, this is not incorrect, but it constrains how the data can be used without overstating precision.

7.4. Implications for ECL, stress testing, and capital assessment

Within expected credit loss (ECL), stress testing, and capital assessment processes, hazard data is most defensible when used to:

- Identify segments of the portfolio warranting differentiated assumptions,
- Inform scenario narratives and overlays,
- Support qualitative and quantitative sensitivity analysis.

Native hazard data can support these steps by providing more informative segmentation and by enabling clearer explanation of why certain exposures are treated differently under stress.

Downscaled data may remain appropriate where:

- Analysis is conducted at regional or portfolio level, and
- The firm explicitly avoids asset-level inference.

SS5/25 does not mandate one approach, but it does expect firms to align the resolution and construction of data with the resolution at which decisions are made.

7.5. Avoiding over-interpretation and false precision

A recurring supervisory concern is the risk of false precision - the appearance of analytical sophistication without a commensurate increase in decision quality. This risk is heightened where fine-resolution outputs are treated as asset-specific hazard measures without regard to how they were constructed. Firms using hazard data, whether native or downscaled, should therefore be explicit that:

- Hazard metrics do not represent probabilities or expected losses,
- Differences in hazard scores do not translate mechanically into credit outcomes, and
- All downstream impacts are contingent on firm-owned assumptions.

Being explicit about these limitations strengthens credibility rather than weakening it.

7.6. Market context: composite hazard scores and embedded risk assumptions

It is recognised that some climate risk data products available in the market present composite scores that incorporate elements beyond physical hazard, such as estimated probability of occurrence, event frequency, historical claims experience, or modelled loss metrics. These approaches are often designed to support specific use cases, particularly within insurance or portfolio-level risk pricing, where integration of hazard, vulnerability, and loss is appropriate and explicitly intended. However, such composite scores differ materially in character from hazard-only datasets. Where probability, frequency, or loss information is embedded within a composite metric:

- The output reflects a modelled risk construct, not a physical hazard condition,
- Assumptions about exposure, vulnerability, or damage mechanisms are implicitly embedded,
- Separation between upstream data and downstream modelling decisions is reduced.

This has governance implications under SS5/25. In particular, firms using composite metrics with embedded risk assumptions may need to:

- Treat the data as a form of risk model input, rather than a hazard indicator,
- Subject it to enhanced model risk management and validation,
- Carefully assess whether embedded assumptions align with internal frameworks and risk appetite.

By contrast, hazard-only composite scores, such as those described in this document, are constructed solely to reflect relative physical hazard intensity, without reference to probability, event frequency, historical loss, or financial impact. This preserves a clear boundary between:

- Identification of environmental stressors, and
- Firm-owned modelling of vulnerability, loss, and capital impact.

Under SS5/25, this separation supports proportional governance by allowing firms to integrate hazard insight into their existing risk frameworks without importing third-party assumptions about how hazards translate into losses.

7.7. Why this distinction matters for SS5/25

SS5/25 does not prohibit the use of composite or risk-based metrics. However, it does expect firms to:

- Understand what assumptions are embedded in the data they use,
- Ensure those assumptions are appropriate for their business model,
- Govern tools proportionately to their influence on decision-making.

Making the distinction between hazard-led composites and risk-embedded composites explicit helps firms demonstrate that:

- They are not double-counting vulnerability or loss assumptions,
- They retain control over critical modelling judgements,
- They can explain data use clearly to internal audit and supervisors.

Integrating hazard insight into modelling frameworks is only one part of the SS5/25 journey. The supervisory expectation extends beyond measurement and modelling to how firms use risk insight to inform management action. The next section explores how hazard-led approaches enable more targeted mitigation, adaptation, and resilience-building within secured portfolios.

8. From Risk Identification to Management Action, Mitigation, and Resilience

A consistent theme in SS5/25 is that firms should do more than identify and measure physical climate risks. Supervisors expect firms to demonstrate how risk insight informs management action, supports mitigation where feasible, and contributes to the long-term resilience of the business. The technical characteristics of hazard data - and whether it is hazard-only or risk-embedded -materially affect a firm's ability to move from analysis to action.

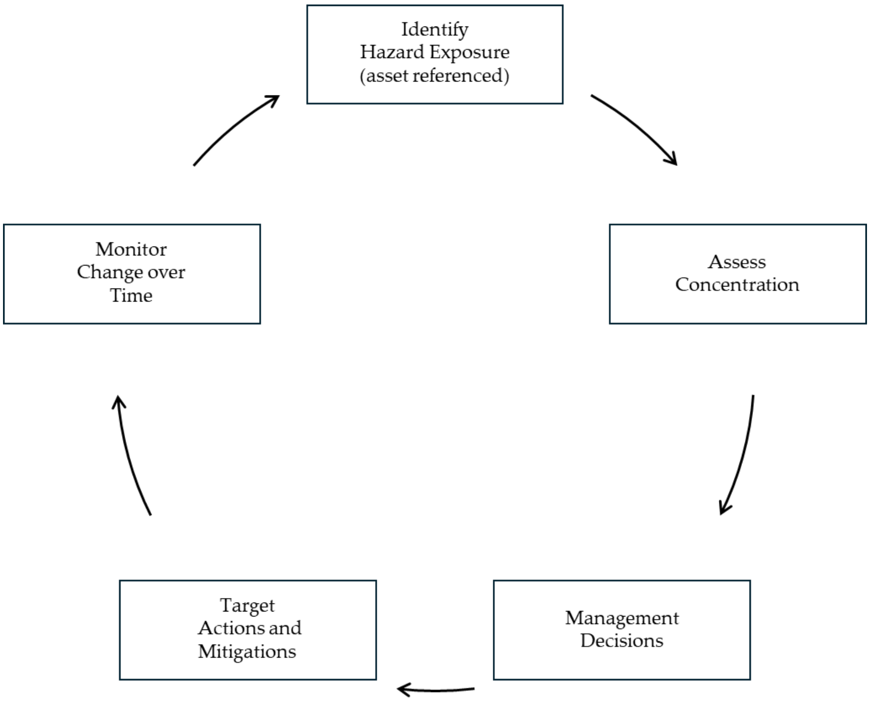

Figure 7. Hazard-led insight supports an iterative cycle of identification, action, and monitoring, consistent with SS5/25 expectations for ongoing risk management.

As shown in Figure 7, hazard-led insight supports an iterative cycle of identification, action, and monitoring, consistent with SS5/25's emphasis on ongoing risk management rather than one-off assessment.

8.1. Why hazard-led insight is a prerequisite for action

Management action requires clarity about where risk arises and what drives it. Hazard-only data supports this by:

- Isolating environmental stressors from financial outcomes,
- Allowing firms to understand which locations are subject to elevated physical conditions,
- Enabling targeted responses that are proportional to the nature of the hazard.

Because hazard-only datasets do not embed assumptions about probability, damage, or loss, they provide a neutral diagnostic layer. This allows firms to decide, within their own governance frameworks, how hazard exposure should translate into underwriting standards, borrower engagement, or portfolio strategy.

This separation is particularly important where management actions involve judgement, trade-offs, or customer interaction.

8.2. The role of native hazard data in targeted mitigation

Where hazard data is both native and hazard-only, firms are better positioned to design targeted mitigation and adaptation strategies. Examples include:

- Prioritising borrower engagement in locations subject to elevated heat stress or wildfire exposure,
- Encouraging or incentivising specific adaptation measures aligned to the hazard driver (e.g. shading, cooling, vegetation management),
- Refining underwriting or monitoring practices for clearly identified hotspots,
- Tracking whether interventions reduce relative hazard exposure over time.

Because the hazard signal is physically grounded and spatially differentiated, these actions can be justified as risk-sensitive rather than blunt or exclusionary.

8.3. Limitations of risk-embedded composite scores for management action

Some climate risk products available in the market provide composite scores that embed probability, event frequency, or modelled loss, often drawing on historical claims data or insurance-style loss models. While such metrics may be appropriate for certain use cases, they can constrain management action in a secured lending context.

Where risk-embedded composites are used:

- Assumptions about vulnerability and loss are already baked into the score,
- It can be difficult to disentangle why a location is flagged as higher risk,
- Management options may appear binary (accept / decline / price) rather than graduated.

This can limit a firm's ability to:

- Design targeted mitigation measures,
- Engage constructively with borrowers, or

- Demonstrate proportionality under SS5/25.

From a governance perspective, such composite metrics often need to be treated as risk model inputs, with corresponding validation and oversight. This may be appropriate, but it changes the nature of how the data can be used.

8.4. Supporting resilience rather than defensive retrenchment

One of the implicit objectives of SS5/25 is to encourage firms to manage physical climate risk in ways that support resilience, rather than defaulting to broad risk withdrawal. Hazard-led approaches, particularly when based on native data, support this objective by allowing firms to:

- Distinguish between genuinely high-risk locations and those that appear high-risk due to averaging or embedded assumptions,
- Avoid unnecessary tightening across entire regions,
- Align risk appetite with physical reality rather than proxy measures.

In this way, data choice becomes a strategic enabler. Firms with clearer, more granular hazard insight are better placed to manage risk selectively, maintain credit availability where appropriate, and adapt portfolios over time.

8.5. Governance benefits under SS5/25

From a supervisory perspective, this approach aligns closely with SS5/25 expectations that firms should:

- Understand the drivers of physical climate risk,
- Retain control over critical modelling and decision assumptions,
- Demonstrate how risk insight feeds into management decisions.

By separating hazard identification from vulnerability and loss modelling, firms can show that:

- Management actions are informed by transparent inputs,
- Decisions remain subject to internal governance and challenge,
- Uncertainty is acknowledged rather than hidden within composite metrics.

This strengthens both regulatory confidence and internal accountability.

8.6. Illustrative example: management action under different data approaches

Consider a lender with a residential mortgage portfolio concentrated in the outskirts of a large town, where summer heat stress and seasonal wildfire risk are expected to increase over time.

8.7. Using hazard-only, native data

Using hazard-only, native hazard data, the firm identifies a subset of properties that consistently fall into higher hazard categories due to a combination of local drivers, including limited tree cover, south-facing orientation, and proximity to unmanaged vegetation. Importantly, the data indicates where the environmental stress arises, but does not imply probability of damage or financial loss. The firm is therefore able to:

- Flag these assets for enhanced monitoring,
- Engage borrowers with targeted guidance on adaptation measures (e.g. shading, defensible space, vegetation management),
- Consider modest adjustments to underwriting or review frequency for new lending in those locations,
- Track over time whether mitigation actions reduce relative hazard exposure.

The hazard data supports graduated, proportionate intervention, with credit decisions remaining subject to internal vulnerability and affordability assessments.

8.8. Using a risk-embedded composite score

Using a composite score that embeds probability or loss assumptions, the same subset of properties is flagged as 'high risk' based on an aggregate metric. However:

- The drivers of the score are not easily separable,
- It is unclear whether the elevation reflects hazard, assumed vulnerability, or historical loss patterns, and
- The score implies an outcome rather than a condition.

As a result, management options narrow. The firm may feel compelled to:

- Apply broad pricing uplifts,
- Restrict new lending in the area, or
- Escalate exposures for exclusionary treatment,

even where mitigation or adaptation could plausibly reduce underlying exposure.

From a governance perspective, the composite score must be treated as a risk model input, increasing validation and oversight requirements and reducing flexibility in how the insight can be used.

8.9. Governance and SS5/25 implications

In both cases, the firm is responding to physical climate risk. The difference lies in how much control and discretion the firm retains.

Hazard-only, native data:

- Supports explainable, asset-referenced decision-making,
- Enables mitigation and resilience-building actions,
- Aligns with SS5/25's emphasis on proportionality and management action.

Risk-embedded composite scores:

- May be appropriate for certain use cases,
- But tend to drive binary outcomes,
- Require stronger justification that embedded assumptions align with the firm's internal frameworks.

8.10. Why this example matters

This example illustrates that the distinction between data approaches is not about sophistication, but about how insight translates into action.

Under SS5/25, firms are expected not only to identify risk, but to demonstrate that they can manage it credibly and proportionately. Data that preserves clarity around physical drivers enables a wider and more defensible range of responses.

Having explored how hazard-led insight supports management action and resilience, the next section draws these threads together to examine the competitive and strategic implications of data choice under SS5/25.

9. Competitive and Strategic Implications of Data Choice

While SS5/25 is framed as a supervisory statement, its implications extend beyond regulatory compliance. The way firms assess and manage physical climate risk increasingly shapes strategic flexibility, portfolio resilience, and competitive positioning. In this context, the choice of hazard data is not merely a technical or compliance decision; it influences how confidently a firm can act in the face of uncertainty.

9.1. Regulatory confidence as a strategic asset

Firms that can clearly explain what their climate risk data measures, how it is used within decision-making, and where its limitations lie, are better positioned in supervisory dialogue. This confidence reduces the likelihood of repeated data challenges from second line or audit, supervisory requests for rework or additional overlays, and overly conservative responses driven by uncertainty rather than evidence.

Over time, this translates into lower friction in regulatory engagement and greater internal confidence in climate risk processes.

9.2. Selectivity versus blunt risk responses

One of the strategic risks associated with physical climate risk is the temptation to respond defensively, for example, by applying broad geographic exclusions or uniform tightening of credit standards. Such responses may reduce exposure in the short term, but they can:

- Unnecessarily constrain lending in areas where risk is heterogeneous,
- Disadvantage customers who could reasonably adapt or mitigate,
- weaken portfolio diversification.

Hazard-led approaches, particularly those based on native data, support selective rather than blunt responses. By distinguishing genuinely elevated hazard from regional averages, firms can:

- Refine risk appetite with greater precision,
- Differentiate between exposures within the same area,
- Maintain access to lending where it remains appropriate.

This selectivity is increasingly important as climate considerations intersect with broader objectives around growth, inclusion, and sustainability.

9.3. Enabling portfolio resilience over time

Resilience is not achieved through a single modelling exercise. It is built through iterative assessment, action, and review. Where hazard data provides clear, physically grounded insight:

- Firms can monitor how exposure evolves under different climate pathways,
- Assess whether mitigation actions are having the intended effect,
- Adapt portfolio strategy as conditions change.

This dynamic capability supports a longer-term view of climate risk that aligns with SS5/25's emphasis on forward-looking management rather than static compliance.

9.4. Strategic implications for capital and risk appetite

As climate risk considerations become more embedded in capital planning and risk appetite discussions, firms that rely on coarse or opaque data may find themselves constrained by uncertainty. By contrast, firms using hazard data that is transparent in construction, aligned to asset-level decision points, and clearly separated from loss assumptions, are better able to articulate how physical climate risk is reflected in risk appetite, justify differentiation within portfolios, and engage with supervisors on proportional and optimised capital responses. This does not remove the need for judgement, but it provides a stronger foundation for exercising it.

9.5. Data choice as a long-term capability decision

Ultimately, the implications of data choice extend beyond any single regulatory cycle. Selecting hazard data that supports explainability, governance, and management action enables firms to:

- Respond more flexibly as supervisory expectations evolve,
- Integrate climate considerations into core risk disciplines, and
- Avoid repeated retooling of frameworks as new requirements emerge.

In this sense, hazard-led, physically grounded data should be viewed as part of a long-term risk management capability, not a short-term compliance solution.

The final section of this document draws together the technical, regulatory, and strategic strands of the journey, and sets out the limitations, guardrails, and concluding principles that underpin responsible use of hazard-led climate data under SS5/25.

10. Limitations, Guardrails, and Concluding Principles

Effective management of physical climate risk under SS5/25 does not depend on eliminating uncertainty. It depends on understanding where uncertainty lies, governing it appropriately, and avoiding false precision. This section sets out the key limitations and guardrails that underpin responsible use of hazard-led climate data, and summarises the principles that emerge from the preceding journey.

10.1. Recognising the limits of hazard data

Physical climate hazard data, whether native or downscaled, provides insight into environmental conditions, not outcomes. In particular, hazard data:

- Does not represent probability of occurrence,
- Does not estimate damage or loss,
- Does not capture asset-specific vulnerability,
- Does not predict credit impairment or capital impact.

Even where hazard data is constructed at high spatial resolution and grounded in physical inputs, it remains an upstream diagnostic tool.

Outcomes depend on how hazards interact with asset characteristics, borrower behaviour, management action, and broader economic conditions. SS5/25 implicitly recognises this by emphasising proportionality and judgement rather than deterministic modelling.

10.2. Guardrails against false precision

A recurring supervisory concern is the risk that climate risk analysis appears more precise than it truly is. This risk is heightened when fine-grained outputs are interpreted without sufficient regard to how they were constructed. Responsible use of hazard data therefore requires clear guardrails, including:

- Explicit statements of what hazard scores do and do not represent,
- Alignment between data resolution and decision-making resolution,
- Avoidance of mechanical translation from hazard metrics to credit outcomes,
- Clear separation between third-party data inputs and firm-owned assumptions.

Being explicit about these guardrails strengthens, rather than weakens, the credibility of climate risk frameworks.

10.3. Governance and accountability under SS5/25

SS5/25 places responsibility for climate risk management firmly with the firm. Data providers, methodologies, and tools support that responsibility, but do not displace it. Firms should therefore ensure that:

- The role of hazard data within their risk framework is clearly defined,
- Model ownership and accountability remain internal,
- Challenge and review functions understand the limitations of upstream inputs,
- Management decisions informed by climate risk are properly governed and documented.

Hazard-led approaches that preserve transparency and separation of roles make this governance easier to sustain.

10.4. Principles emerging from the journey

Several principles emerge from the analysis set out in this document:

1. Data choice shapes decision quality: The technical construction of hazard data directly affects how confidently firms can identify, assess, and manage physical climate risk.

2. Resolution should match use case: Asset-level decisions require data that is credible at asset level. Regional trends require less granularity, but greater caution in interpretation.
3. Separation supports governance: Keeping hazard identification distinct from vulnerability and loss modelling preserves internal control and reduces model risk.
4. Uncertainty should be structured, not hidden: Native, hazard-led data does not remove uncertainty, but it allows uncertainty to be explored in a more informative and transparent way.
5. Insight should enable action: Climate risk analysis is most valuable when it supports proportionate management action and resilience-building, not just disclosure.

10.5. Concluding reflection

SS5/25 does not require firms to predict the future. It requires them to demonstrate that they are thinking clearly, acting proportionately, and governing uncertainty responsibly in the face of evolving physical climate risks.

The journey set out in this document shows how technical choices made at the data level propagate through regulatory compliance, modelling practice, and management action. By understanding those connections, firms can make more informed decisions about how to assess physical climate risk, how to integrate it into existing frameworks, and how to build resilience over time.

Ultimately, better data does not eliminate risk. It enables better control over how risk is identified, understood, and managed. For firms operating under SS5/25, this means that the technical design of hazard data is not a secondary consideration, but a primary determinant of how effectively physical climate risk can be governed, evidenced, and acted upon in practice. Use of this document does not remove the need for firm-specific validation, challenge, and approval processes.