

NatPar: Natural Parametric Modeling*

Hirbod Assa[†]

January 20, 2026

Abstract

This paper proposes a minimum viable reporting template for natural parametric insurance. The framework is motivated by prevailing industry practice: under intensifying climate volatility, rapidly evolving exposures, and binding solvency constraints, natural catastrophe (NatCat) teams increasingly re-purpose the hazard–exposure–vulnerability–finance (HEVF) stack to engineer index-linked liabilities that are verifiable, bounded, and more capital-tractable. We refer to this supply-side design class as *Natural Parametric* (NatPar) modelling.

NatPar retains the NatCat architecture but alters what becomes contractual. Hazard modelling specifies observable indices and trigger logic; exposure and vulnerability are re-tasked into a basis-risk engine that quantifies mismatch between indemnity loss and parametric payout; and the finance block is adapted to short-tailed liabilities with limited development risk. Because NatPar is constructed from the NatCat stack, standard portfolio outputs carry over directly—including annual average loss (AAL), EP/AEP/OEP curves, return-period levels, and one-year tail metrics—enabling like-for-like comparison between indemnity and parametric programmes on a common hazard and exposure base.

We then introduce a minimal reporting and regulatory standard that complements these familiar tail objects with basis-risk governance diagnostics: average basis shortfall and overpayment, parametric–indemnity dependence (e.g., correlation), two-sided basis-exceedance probability curves (BEP $^\pm$), and portfolio-level basis AEP and OEP (BAEP/BOEP) that separate shortfall from overpayment. Collectively, these outputs support an “apples-to-apples” evaluation of NatCat and NatPar designs by making explicit how payout shape and basis risk propagate into risk and capital outcomes.

A frost laboratory illustrates the framework. Starting from a NatCat model of seasonal frost damage, we construct AAL-neutral parametric payoffs and show that payout shape is a first-order determinant of tail capital: bounded continuous schedules can reduce tail VaR by limiting exposure-driven amplification, whereas trigger-based event parametrisations can concentrate losses and increase VaR at regulatory confidence levels.

*©Hirbod Assa

[†]Model Library and UCD, email:assa.hirbod@gmail.com

Contents

1	Introduction	3
2	Literature review	5
3	NatCat machinery in practitioner terms	7
3.1	The four modeling blocks: hazard, exposure, vulnerability, financial terms	7
3.2	Portfolio loss	7
4	NatPar framework	8
4.1	NatPar payout classes and their implications for tails and basis	8
4.2	NatCat to NatPar through AAL-neutral calibration	9
4.3	Basis risk as an auditable interface	10
4.3.1	Individual risk basis.	10
4.3.2	Portfolio basis: aggregate and occurrence.	11
4.4	Minimal reporting template	12
5	Two design blocks and one reporting language for individual risk	13
5.1	Common notations	13
5.2	NatCat: Analytic EP and AAL representations	15
5.3	NatPar: Analytic EP and AAL representations	15
5.3.1	Continuous parametric payout.	15
5.3.2	Binary payout.	16
5.4	Basis risk diagnosis	17
5.4.1	Point diagnosis	17
5.4.2	Basis exceedance curve	18
5.5	Continuous NatPar basis analysis: analytic shortfall and overpayment .	19
5.6	Binary NatPar: analytic shortfall and overpayment	19
6	Reporting the frost risk example	20
6.1	Hazard-exposure-vulnerability and NatCat indemnity	20
6.2	Basis risk diagnostics	21
6.2.1	Individual risk	22
6.2.2	Portfolio risk	26
7	Results interpretation of the frost laboratory	30
8	Extensions and discussion	34
8.1	Multi-trigger and multi-peril NatPar programmes	34
8.2	NatPar and solvency/climate-risk regulation	35
8.3	Limitations and practical challenges	36
9	Summary and lessons from the frost laboratory	37

1 Introduction

Over the last decade, climate pressure alongside rapid economic growth has increased both exposure and hazard intensity across many regions. In several markets, climate-exposed lines have moved toward outcomes increasingly described—informally or explicitly—as *uninsurable*: coverage is withdrawn, terms are tightened, or prices rise to levels that are politically and socially untenable. At the same time, protection gaps for climate perils remain large and, in some settings, continue to widen.

From the perspective of insurers and reinsurers, the binding challenge is often not only higher mean losses, but reduced *tractability of the loss distribution*, especially in the tail. Hazard distributions drift, exposures evolve quickly, and extremes may cluster in space and time; these dynamics inflate uncertainty precisely where solvency frameworks and risk appetite concentrate. Long-tailed indemnity liabilities become harder to reserve, harder to explain, and harder to support under one-year capital constraints and target returns. In this sense, a line drifts toward uninsurability when markets cease to clear because one or more constraints becomes binding: *affordability* (premiums exceed willingness or ability to pay), *capital* (tail requirements render the line uneconomic), and/or *model uncertainty* (pricing and reserving assumptions are no longer defensible to management, regulators, or capacity providers). These constraints reinforce each other: *uncertainty in exposure and vulnerability* increases capital charges and risk loadings, pushing premiums higher and accelerating capacity withdrawal.

Against this background, parametric structures have expanded not merely as customer-facing innovations, but as a supply-side reconfiguration of NatCat practice. Under climate-driven uninsurability and capital pressure, NatCat teams increasingly redeploy the hazard–exposure–vulnerability–finance stack to produce index-linked liabilities that are auditable, bounded by contract, and more capital-tractable. We refer to this pipeline as *Natural Parametric (NatPar) modelling*. Conceptually, NatPar preserves the NatCat architecture but changes what becomes contractual: hazard modelling defines measurable indices and triggers; exposure and vulnerability are re-tasked into a basis-risk engine that quantifies the mismatch between realised indemnity loss and parametric payout; and the finance block is adapted to short-tailed liabilities with minimal development risk. NatPar therefore exchanges a portion of exposure-driven indemnity variability for a transparent, index-verified payout, reducing reserve uncertainty and improving capital tractability *at the cost of explicit basis risk*.

Design-first versus pipeline-first. It is useful to contrast our objective—to set standards for NatPar reporting and regulation— contrasting the dominant framing in much of the academic literature and industrial development.

- *Academic view: design-first.* A large literature treats parametric insurance as a contract-design problem: specify an index (or index vector) X and choose a payout function $I(X)$ to optimise a welfare or risk objective under a premium constraint.

This perspective is mathematically flexible and emphasises basis risk as the mismatch between realised loss and index-based payout, making it well suited for studying optimality, robustness, behavioural frictions, and uptake.

- *Industry view: pipeline-first.* In practice, however, many parametric products are developed inside the NatCat ecosystem. Underwriting, reinsurance/ILS structuring, accumulation control, and solvency reporting rely on a modular workflow (hazard, exposure, vulnerability, financial terms) and a standard set of portfolio outputs (AAL and EP/AEP/OEP curves, return-period levels, and one-year tail metrics). From this viewpoint, a parametric contract is not an arbitrary $I(X)$: it must be verifiable, auditable, and *comparable* within the same reporting language used for indemnity portfolios and capital.

This creates a practical gap between design-first theory and pipeline-first implementation—a gap that becomes acute when products must be justified to boards, supervisors, and capacity providers.

This paper introduces a framework based on three pillars.

1. **NatPar standard within the NatCat stack.** We formalise a NatPar standard in which parametric contracts are constructed and evaluated within the conventional catastrophe-modelling architecture—hazard, exposure, vulnerability, and financial terms—so that outputs remain directly comparable to NatCat portfolio outputs. We express NatPar design objectives using the portfolio language used in practice—AAL, EP/AEP/OEP curves, and one-year tail metrics—and introduce AAL-neutral calibration to isolate distribution-shape effects across contracts.
2. **Basis risk as a distributional object.** We treat basis risk not as a single summary statistic but as a distributional object, introducing two-sided basis-risk exceedance curves (“basis EP”) for *shortfall* and *overpayment* as governance diagnostics that sit naturally alongside standard EP reporting and provide a tail-sensitive view of mismatch. We create such basis tail analysis under a so-called AAL-neutrality calibration to make sure the shortfall/overpay is not due to changes in the average loss.
3. **NatPar vs NatCat capital.** We show that contracts with the same mean loss can have materially different tails depending on payout shape; hence AAL-neutrality does not imply capital-neutrality under VaR-style solvency measures.

We consolidate numerical illustrations into a single case-study section and apply a consistent reporting template throughout.

The remainder of the paper is organised as follows. Section 2 reviews the contract-design and multi-hazard index-insurance literatures and contrasts them with the pipeline and reporting language used in catastrophe practice. Section 3 summarises the NatCat hazard-exposure-vulnerability-finance stack and the standard portfolio outputs used for underwriting and capital (AAL, EP/AEP/OEP curves, return-period levels, and one-year tail

metrics). Section 4 then formalises the NatPar mapping—how triggers and payout schedules are constructed to remain NatCat-compatible—and introduces the minimal reporting template that pairs tail objects with auditable basis-risk diagnostics (point measures and two-sided basis exceedance curves, at individual and portfolio level). Sections 5, 6 and 7 implement the framework in a stylised frost laboratory: Section 5 develops the analytic representations used for comparability and basis diagnosis, Section 6 reports the empirical outputs for both individual risks and a two-region portfolio (including BAEP/BOEP), and Section 7 interprets the results with an emphasis on how payout shape propagates into tails and solvency-relevant capital. Section 8 discusses extensions (multi-trigger and multi-peril programmes), regulatory implications, and practical limitations, and Section 9 concludes with a summary of lessons for NatPar reporting and governance.

2 Literature review

Much of the academic literature treats parametric insurance as a generic contract-design problem, whereas the industry has largely evolved parametric programmes as pipeline-compatible overlays on NatCat machinery.

A standard academic entry point models insurance design as a choice of a payout schedule subject to pricing constraints, often derived from expected-utility considerations and classic optimal insurance results (e.g., generalized deductibles) [1, 2]. Index (parametric) insurance adopts a related contract-design perspective, but replaces loss adjustment by a payout determined by an observable index (or index vector) X , typically written as $I(X)$. Modern treatments formalize the choice of $I(\cdot)$ under premium constraints and study welfare, risk, and implementability of index contracts [3].

Across this literature, *basis risk*—the mismatch between realized indemnity-style loss and the index-based payout—is consistently identified as the central friction that can dominate performance and suppress demand. Two influential directions are (i) multi-scale indices to reduce mismatch across spatial/temporal aggregation and (ii) empirical measurement of index quality and basis-risk outcomes in the field [4, 5]. Complementary work emphasises implementation and uptake constraints, including regulatory challenges and behavioural frictions such as complexity aversion, which can reduce uptake even when expected values are favorable [6, 7].

A conceptually close strand to our motivation is the review by Benso et al. [8] on weather index insurance for multi-hazard resilience and food security. Beyond synthesizing the literature, they propose a practical three-module framework for index insurance design: *hazard identification*, *vulnerability assessment*, and *financial methods and risk pricing*. This decomposition is a useful bridge between a broad (and sometimes fragmented) academic literature and the practical steps required to build indices and payout rules.

A key message in the multi-hazard setting is that hazard representation and combination are modelling choices, not merely data choices. Benso et al. [8] emphasise that multi-

hazard products should not default to independence assumptions and highlight interaction taxonomies (independent, synergistic, cascading) that can materially affect loss modelling and premium adequacy [8, 9]. Related work reviews quantitative methodologies for multi-hazard interrelationships and provides a structured view of how hazard interactions can be represented and tested [10]. In climate-risk contexts, compound-event research provides additional conceptual and statistical tools for thinking about multivariate extremes and interacting drivers [11].

In industry practice, catastrophe risk is commonly represented via a modular supply chain: *hazard* (event sets/footprints), *exposure* (assets at risk and where), *vulnerability* (damage/loss functions), and *financial terms* (deductibles, limits, aggregates and layers). This architecture is central to how portfolios are priced (e.g., AAL) and how tails are communicated (e.g., EP/AEP/OEP curves and return-period losses), and it is described in practitioner-oriented treatments of catastrophe risk modelling [12].

Recent climate physical risk assessment (PRA) and climate change risk assessment (CCRA) work builds naturally on the same decomposition, but adds scenario conditioning, spatial-temporal resolution constraints, and climate-model uncertainty. In particular, portfolio-scale CCRA requires separating changes in *hazard* from changes in *exposure*, because their interaction can magnify or attenuate losses [13]. Integrating climate scenarios into NatCat-style workflows typically requires explicit choices around downscaling and bias correction, which can materially influence hazard projections and therefore pricing and capital metrics [14, 15, 16, 17]. At the same time, climate disclosure and supervisory stress testing have accelerated demand for portfolio-level methodologies that are transparent and auditable [18, 19, 20].

The contract-design literature explains parametric insurance in the abstract language of $I(X)$ and basis risk, and multi-hazard reviews provide useful design taxonomies [3, 8]. However, these literatures do not typically track how parametric products have evolved operationally inside the insurance ecosystem: namely, as contractual overlays on the catastrophe-modelling pipeline that already governs underwriting, portfolio aggregation, reinsurance/ILS structuring, and solvency reporting.


This paper therefore adopts a *pipeline-first* notion of parametric insurance. We define NatPar programmes as contracts whose triggers, payouts, and validation procedures are explicitly designed to be compatible with the NatCat supply chain and its portfolio outputs. In this view, basis risk is not merely a single summary statistic but the measurable interface created when $(\text{exposure} \times \text{vulnerability})$ is replaced by a low-dimensional, index-verifiable hazard proxy. This is also where model governance lives: NatPar design should be readable in the same reporting language as NatCat (AAL, EP/AEP/OEP, tail metrics), with explicit basis distributions and basis exceedance curves as diagnostics.

3 NatCat machinery in practitioner terms

This section summarises the catastrophe-modelling workflow that underpins NatCat underwriting and portfolio reporting. We focus on the modelling blocks and the portfolio outputs that are used in practice, because NatPar is defined in this paper as a pipeline-compatible overlay on the same machinery.

3.1 The four modeling blocks: hazard, exposure, vulnerability, financial terms

A standard NatCat model can be represented as a modular pipeline:



The purpose of this decomposition is operational: each block can be updated, validated, and governed with its own data sources and model risk controls, while still producing consistent portfolio outputs.

Hazard. The hazard block specifies an event set (or stochastic process) for the peril, including frequency, intensity, and spatial footprints where relevant. In climate settings, the hazard block may also be conditioned on scenarios or nonstationary assumptions.

Exposure. The exposure block specifies what is at risk and where: locations, sums insured/values, relevant asset attributes, and aggregation rules. Exposure is often the most rapidly changing component in practice, and it is also a major driver of accumulation risk. Our main message is that m

Vulnerability. The vulnerability block maps hazard intensity (and modifiers) to damage ratios or loss distributions. It captures construction/asset sensitivity, secondary uncertainty, and model error.

Financial terms. Financial terms transform ground-up loss into contractual loss: deductibles, limits, layers, aggregates, reinstatements, and other features that define the insurer or reinsurer liability.

3.2 Portfolio loss

Let Y_r denote the contractual loss after applying financial terms for the set of covered risk $r \in \mathcal{R}$. Here \mathcal{R} is a set of risks, that can be for example a deterministic set of regions (as

we use in this paper), or even a random set of events (such as $r \in \mathcal{R} \equiv 1 \leq r \leq N_t$ for a Poisson process N_t).

Portfolio-level reporting commonly distinguishes:

$$\text{Occurrence loss: } M := \max_r Y_r, \quad \text{Aggregate loss: } S := \sum_r Y_r.$$

These are the random quantities that generate standard NatCat outputs such as OEP and AEP curves.

NatCat portfolio reporting distinguishes between exceedance probabilities for annual *occurrence* loss (OEP) and annual *aggregate* loss (AEP). Using the annual variables the one-year curves are defined as

$$\text{OEP}(x) := \mathbb{P}(M > x), \quad \text{AEP}(x) := \mathbb{P}(S > x), \quad x \geq 0.$$

In general, these curves differ because a year can contain multiple loss-causing events: OEP captures the tail of the single largest event, while AEP captures the tail of the sum of events. They coincide only in special cases (e.g. when there is at most one loss-causing event per year in the model, or when financial terms collapse event losses into a single outcome).

Return periods. A common presentation uses return periods. For any liability X , define the T -year return level

$$x_T(X) := \inf \{x : \mathbb{P}(X > x) \leq 1/T\}.$$

Return-period curves $T \mapsto x_T(\cdot)$ provide an interpretable mapping from tail probability to severity. For a return period T (years), the corresponding OEP return-period loss level x_T^{OEP} satisfies OEP ($x_T^{\text{OEP}} = 1/T$), and similarly for AEP.

4 NatPar framework

NatPar programmes are defined here as parametric structures anchored in the catastrophe-modelling workflow and evaluated in the same portfolio language as NatCat (AAL, EP/AEP/OEP curves, return periods, tail metrics). The central additional object is *basis risk*, which we treat as an auditable interface rather than an afterthought.

4.1 NatPar payout classes and their implications for tails and basis

A practical advantage of a pipeline-first view is that NatPar designs can be organised into a small number of payout classes, each with predictable implications for tail metrics and basis risk. This section provides a compact taxonomy used later to interpret the case study.

(i) Continuous bounded payouts (graded contracts). Contracts of the form $P = I(X)$ with I continuous and bounded (e.g. $P = q(D(X) - d)^+$ with $D \in [0, 1]$) mechanically cap the insurer's liability and often reduce high-quantile tail metrics relative to exposure-scaled indemnity losses, even under AAL-neutral calibration. Their basis risk tends to be *diffuse*: frequent moderate mismatch can occur because payouts do not scale with realised exposure, but severe basis shortfalls can be controlled by the cap and by shaping $I(\cdot)$.

(ii) Piecewise-linear / tiered payouts. Piecewise-linear or tiered payouts provide additional flexibility while remaining transparent and auditable. They can be used to approximate layer-like behaviour (e.g. increasing payout rate as hazard severity increases) and to enforce constraints such as $\text{VaR}_{0.995}(P) \leq c$ while still targeting basis shortfall objectives. In a governance setting, the breakpoints of $I(\cdot)$ become explicit design parameters that can be stress-tested and versioned.

(iii) triggers-driven events. A trigger-driven event (or simply binary) $P = q\mathbf{1}_{\{X \in \text{event}\}}$, for a particular event set event and index X for a payout q , is maximally simple and highly verifiable, and they often align with operational objectives such as rapid cashflow upon a threshold event. However, their loss distribution is discrete, which can have unintuitive implications: for a trigger probability $1 - \alpha$ (i.e., $\mathbb{P}\{X \in \text{event}\} = 1 - \alpha$), tail metrics can "jump" at confidence levels above α (e.g. $\text{VaR}_{0.995}$ may equal q when $\alpha = 0.99$). As a result, AAL-neutrality does not guarantee capital advantages, and these type of binary designs can concentrate basis mismatch into rare but large shortfalls or overpayments depending on how q is calibrated.

(iv) Layer-like parametric designs. Between continuous and binary extremes lie parametric designs that emulate insurance layers (e.g. step functions with multiple thresholds, or capped linear segments). These are often the most directly comparable to indemnity layer structures and can be tuned to match both AAL and a tail constraint, while keeping the trigger logic auditable.

4.2 NatCat to NatPar through AAL-neutral calibration

Let Y denote the NatCat contractual loss (for a policy, region, or portfolio segment), and let T denote an index (typically a hazard proxy or near-hazard proxy) that is observable and verifiable. A NatPar design specifies a payout rule $P = I(T)$ together with contract terms such as caps, layers, and trigger thresholds.

A natural baseline calibration target is AAL-neutrality:

$$\mathbb{E}[P] = \mathbb{E}[Y].$$

This constraint ensures that differences in EP curves and tail metrics primarily reflect distributional shape rather than mean loss. In practice, this can be implemented at different granularities (policy-level, peril-region segment, or portfolio bucket), and it can be supplemented by loadings for expenses and profit. The key point is that AAL-neutrality is a *comparability device*, not a full design criterion.

4.3 Basis risk as an auditable interface

The fundamental trade-off is that P is transparent and bounded but cannot replicate Y pointwise; the resulting mismatch is basis risk. Therefore, define basis risk as

$$B := Y - P.$$

4.3.1 Individual risk basis.

Two useful metric to view a basis risk are

- **Point measures.** This means measures that can represent the basis risk through a single scalar such as $\mathbb{P}(B > 0)$, $\mathbb{E}[B^+]$, $\mathbb{E}[(-B)^+]$, $\text{corr}(Y, P)$, etc.
 - $\mathbb{P}(B > 0)$ distinguishes whether the design tends to miss loss states (high shortfall frequency) or to pay in non-loss states (high overpayment frequency).
 - The severities $\mathbb{E}[B^+]$, $\mathbb{E}[(-B)^+]$ are the relevant averages conditional on sign: they quantify the expected size of liquidity *missing* when loss occurs (shortfall) and the expected size of liquidity *excess* when it is not needed (overpayment).
 - The indemnity-parametric correlation $\text{corr}(Y, P)$ indicate how the statistical mismatch between the real pay-off and the parametric product can be measured.
- **Tail measures.** While point summaries are useful but can hide whether mismatch is driven by frequent small shortfalls or rare severe ones. To make basis risk governable, we treat it as a distributional object and report exceedance curves:

$$\text{BEP}^+(x) := \mathbb{P}(B > x), \quad \text{BEP}^-(x) := \mathbb{P}(-B > x).$$

These “*Basis EP*” curves sit naturally next to NatCat EP curves and provide a tail-sensitive diagnostic of *shortfall* risk and *overpayment* risk.

Basis EP curves provide a portfolio-style, tail-sensitive view of mismatch. The shortfall curve $\text{BEP}^+(x) = \mathbb{P}(B > x)$ answers: *how often does the indemnity benchmark exceed the parametric payout by more than x ?* The overpayment curve $\text{BEP}^-(x) = \mathbb{P}(-B > x)$ answers: *how often does the parametric payout exceed the indemnity benchmark by more than x ?*

These curves are operational for governance because they separate:

- **frequency** of mismatch (the level of BEP near $x = 0$),
- **severity** of mismatch (how slowly BEP decays for large x),
- **direction** of mismatch (shortfall vs overpayment, via BEP^+ and BEP^-), and they can be compared across candidate contract shapes even when AAL is matched.

Return-period interpretation (basis return periods). Just as EP curves are often communicated via return periods, basis EP curves can be summarised by basis return-period levels. For a return period T (years), define the shortfall-basis return-period level b_T^+ by

$$\text{BEP}^+(b_T^+) = \mathbb{P}(B > b_T^+) = \frac{1}{T},$$

and similarly define the overpayment-basis return-period level b_T^- by

$$\text{BEP}^-(b_T^-) = \mathbb{P}(-B > b_T^-) = \frac{1}{T}.$$

Reporting (b_T^+, b_T^-) for a small set of T values (e.g. $T \in \{10, 20, 50\}$) yields a compact and interpretable summary of basis-tail behaviour that can sit next to the standard loss return-period table.

4.3.2 Portfolio basis: aggregate and occurrence.

Let $B_r = Y_r - P_r$ denote the basis for the risk $r \in \mathcal{R}$. In a multi-region portfolio, basis reporting must distinguish between (i) the *aggregate* mismatch over all regions in the model year, and (ii) the *occurrence* (worst-region) mismatch that can dominate operational and regulatory concerns.

Given portfolio, define the *aggregate basis* as

$$S^B = \sum_{r \in \mathcal{R}} B_r = S^Y - S^P.$$

This quantity is natural when the objective is balance-sheet impact and annual aggregate liquidity, and it aligns with AEP-style reporting for liabilities. For basis regulation and suitability, we additionally introduce an *occurrence basis* that captures the worst regional mismatch within the year:

$$M^B = \max_{r \in \mathcal{R}} B_r.$$

The occurrence basis is the appropriate object when basis risk is interpreted as an operational shortfall/overpayment exposure that can be driven by a single region, even if the portfolio aggregate is moderate. This parallels the NatCat distinction between AEP and OEP, and it is particularly relevant for supervisory review because it measures basis accumulation: the most adverse mismatch that the index can generate in a given year.

Two useful metric to view a basis risk are

- **Point measures.** Similar to the above this means measures that can represent the portfolio basis risk through $\mathbb{P}(S^B > 0)$, $\mathbb{E}[(S^B)^+]$, $\mathbb{E}[(-S^B)^+]$, $\text{corr}(S^Y, S^P)$, $\mathbb{P}(M^B > 0)$, $\mathbb{E}[(M^B)^+]$, $\mathbb{E}[(-M^B)^+]$, $\text{corr}(M^Y, M^P)$.

- **Tail measures.** Similar to above we can also look at the tail:

$$\text{BAEP}^+(x) := \mathbb{P}(S^B > x), \quad \text{BAEP}^-(x) := \mathbb{P}(-S^B > x).$$

$$\text{BOEP}^+(x) := \mathbb{P}(M^B > x), \quad \text{BOEP}^-(x) := \mathbb{P}(-M^B > x).$$

Basis return levels $b_T^{\pm, \{\cdot\}}$ for the portfolio can be defined by

$$\text{BAEP}^{\pm}(b_T^{\pm}) = \mathbb{P}(S^B > b_T^{\pm}) = \frac{1}{T},$$

$$\text{BOEP}^{\pm}(b_T^{\pm}) = \mathbb{P}(M^B > b_T^{\pm}) = \frac{1}{T}.$$

4.4 Minimal reporting template

For each NatPar candidate design (within a region/segment/etc), report:

1. **Mean comparability:** AAL as the calibration target.
2. **Tail comparability:** EP/AEP/OEP, return-period levels.
3. **Basis interface:**
 - point diagnostics.
 - basis tails: BEP/BAEP/BOEP, return-period levels.
4. **Capital metrics:** VaR_{0.995} (and TVaR where feasible).

As we will see in Section 5 we implement this reporting template in a stylised frost laboratory. We compare a continuous bounded NatPar design to an indemnity benchmark and then compare an extreme indemnity layer to an binary design, illustrating how payout shape can dominate tail capital even under AAL-neutral calibration.

The frost laboratory in Section 5 uses a continuous and a binary as two archetypes that bracket the design space. The former illustrates how removing exposure-scaled variability and bounding payouts can shorten tails; the latter illustrates how digital triggers can shift probability mass into capital-relevant quantiles, making explicit why payout shape must be treated as a first-order design decision rather than a cosmetic contract feature.

5 Two design blocks and one reporting language for individual risk

This section operationalises the NatCat-to-NatPar mapping in a stylised frost laboratory (two regions, one seasonal outcome per year) and reports outputs in the same portfolio language used in catastrophe practice for an individual risk. The point is not to build the most realistic frost model, but to make the *workflow* explicit: starting from a NatCat loss specification, constructing a NatPar payout that is legible to underwriting and governance, calibrating it to a transparent target, then constructing the reporting framework and finally comparing the two using the same diagnostics. Note that in order to better focus on the reporting and regulatory framework, we assume there is only one season and one region. This implies that $AEP = OEP = EP$. In the numerical assessment we will use a portfolio in addition to the individual risk.

5.1 Common notations

We index regions by $r \in \{\text{FL, CA}\}$. In the frost laboratory each season produces at most one relevant frost outcome per region. Hence, within a region, annual aggregate and annual occurrence losses coincide, and the usual catastrophe reporting objects (AEP/OEP) reduce to a single exceedance curve. This “one-loss-per-year” structure is deliberate: it isolates the effect of contract design and basis risk without confounding from within-year event counts.

State variables and loss components. For each region r :

- T_r denotes the (seasonal) trigger e.g., minimum temperature as hazard proxy/index.
- A_r denotes seasonal exposure e.g., monetary value at risk.
- $D(\tau) \in [0, 1]$ denotes a non-increasing hazard-driven damage fraction (vulnerability curve). We assume $D(\tau) > 0 \Leftrightarrow \tau < \tau_t$, for a given τ_t and $D(\tau) < 1 \Leftrightarrow \tau > \tau_c$. The left inverse of D is denoted by D^- .
- L_r denotes ground-up loss. In this paper we consider the multiplicative form $L_r = A_r D_r$. We use the notation $D_r = D(T_r)$, throughout the paper.

Let us introduce the following exposure fucntions.

- The CDF being denoted by:

$$F_A(x) := \mathbb{P}(A \leq x),$$

- The call functional

$$C_A(x) := \mathbb{E}[(A - x)^+].$$

NatCat vs. NatPar payoffs. We write contractual (NatCat-type) indemnity loss as a transformation of ground-up loss,

$$Y_r := \psi(L_r),$$

where ψ captures financial terms (e.g. deductible/stop-loss layer). The parametric (NatPar-type) payout depends only on the hazard proxy,

$$P_r := I(T_r),$$

where $I(\cdot)$ is the payout function (continuous, piecewise, or binary).

Calibration convention (comparability device). Within each region and design block, we calibrate NatPar by an *AAL-neutral* principle:

$$\mathbb{E}[P_r] = \mathbb{E}[Y_r],$$

so differences in EP curves, tail metrics, and basis risk reflect *distributional shape* and *index mismatch* rather than differences in mean payout. This is a comparability device, not a claim of optimality.

Ground-up loss. The seasonal ground-up loss is multiplicative:

$$L_r := A_r D(T_r) = A_r D_r.$$

This is the canonical NatCat construction: hazard produces a state (T), vulnerability maps hazard to a fractional impact (D), and exposure scales impact into monetary loss ($A \times D$). In this form, uncertainty in A loads directly into the loss distribution whenever $D_r > 0$.

NatCat indemnity cover with deductible. We benchmark against a simple deductible (stop-loss) indemnity form with region-specific deductible $d_r \geq 0$:

$$Y_r := (L_r - l_r)^+ = \max(A_r D_r - l_r, 0).$$

Because the frost laboratory produces at most one relevant frost loss per season in each region, Y_r is simultaneously the event loss, the annual occurrence loss, and the annual aggregate loss for that region. This is precisely what makes the example a clean laboratory for comparing EP curves, return levels, and tail metrics under a single-period loss distribution.

5.2 NatCat: Analytic EP and AAL representations

The analytic structure used later for basis risk begins already at the NatCat benchmark. Fix a region r and a threshold $x \geq 0$. The exceedance probability of indemnity loss is

$$\text{EP}_r^{\text{Cat}}(x) = \text{EP}_r(x) := \mathbb{P}(Y_r > x) = \mathbb{E} \left[\bar{F}_{A_r} \left(\frac{l_r + x}{D_r} \right) \right] \quad (1)$$

where \bar{F}_{A_r} is the exposure survival in region r . This form makes explicit the supply-side sensitivity: tail exceedance is controlled by the probability that exposure exceeds a hazard-dependent threshold proportional to $\frac{d_r + x}{D_r}$.

Similarly, the average annual loss (AAL) of the deductible indemnity admits the conditional form

$$\text{AAL}_r^{\text{Cat}} = \mathbb{E} \left[D_r C_{A_r} \left(\frac{l_r}{D_r} \right) \right]. \quad (2)$$

Equations (1)–(2) show that, in the ground-up NatCat benchmark, exposure uncertainty enters the liability distribution through F_{A_r} (and through $\sigma_{A,r}$ equivalently). This is the structural mechanism that NatPar will later modify: once payout depends only on T_r (and is AAL-calibrated), exposure uncertainty drops out of the insurer’s liability distribution, but it reappears as a measurable residual mismatch in the basis variable $B_r = Y_r - P_r$.

5.3 NatPar: Analytic EP and AAL representations

We now construct a NatPar contract that is “supply-side natural” in the sense that it is designed directly from the hazard–vulnerability block and is priced and governed with the same loss-distribution diagnostics used in catastrophe practice. The key structural change relative to the NatCat benchmark is that the payout depends only on the observable hazard proxy T_r (through the damage function), not on exposure A_r . So in a similar manner the ground-up payout in parametric is defined as

$$\mathcal{L}_r^c = (D_r - d_r)^+, \quad \mathcal{L}_r^b = \mathbf{1}_{\{T_r < \tau_r\}} = \mathbf{1}_{\{D_r > d_r\}}$$

where $D(\cdot)$ is the same vulnerability curve used in the NatCat benchmark and $d_r = D(\tau_r)$ where τ_r is a region specific trigger. This payoff is (i) fully index-driven (depends only on T_r), (ii) bounded by construction ($0 \leq P_r^c \leq q_r$), and (iii) interpretable: q_r is the maximum payout and D_r controls how the payout increases as temperatures fall.

5.3.1 Continuous parametric payout.

Define a continuous payout NatPar as follows:

$$P_r^c := q_r^c \mathcal{L}_r^c = q_r^c (D_r - d_r)^+. \quad (3)$$

To compare distributional *shape* rather than mean level, we choose q_r^c such that NatPar is AAL-neutral to the NatCat benchmark indemnity Y_r :

$$\text{AAL}_r := \mathbb{E}[P_r^c] = \mathbb{E}[Y_r]. \quad (4)$$

Substituting (3) into (4) yields the explicit calibration

$$q_r^c = \frac{\mathbb{E}[(A_r D_r - l_r)^+]}{\mathbb{E}[(D_r - d_r)^+]} = \frac{\mathbb{E}\left[D_r C_A\left(\frac{l_r}{D_r}\right)\right]}{C_{D_r}(d_r)}. \quad (5)$$

Because P_r^c depends only on T_r , its EP curve has a simple hazard-only representation. For any threshold $x \geq 0$,

$$\text{EP}_r^c(x) := \mathbb{P}(P_r^c > x) = \bar{F}_{D_r}\left(\frac{x}{q_r^c} + d_r\right) \quad (6)$$

In the frost laboratory (one-loss-per-season), this EP curve is simultaneously the AEP and OEP curve.

5.3.2 Binary payout.

NatPar binary payout is a binary contract given as follows:

$$P_r^b = q_r^b \mathcal{L}_r^b = q_r^b \mathbf{1}_{\{T_r < \tau_r\}}, \quad (7)$$

with AAL-neutral calibration $\mathbb{E}[P_r^b] = \mathbb{E}[Y_r]$. By substituting into the equations (7) into the AAL neutrality we get:

$$q_r^b = \frac{\mathbb{E}[(A_r D_r - l_r)^+]}{\mathbb{P}\{D_r > d_r\}} = \frac{\mathbb{E}\left[D_r C_A\left(\frac{l_r}{D_r}\right)\right]}{\bar{F}_{D_r}(d_r)}. \quad (8)$$

Therefore, we easily can get the exceedance as follows:

$$\text{EP}_r^b(x) = \begin{cases} \bar{F}_{D_r}(d_r), & x \leq q_r^{\text{bin}} \\ 0 & x > q_r^{\text{bin}} \end{cases}.$$

As we mentioned earlier in the frost laboratory EP curve is simultaneously the AEP and OEP curve:

$$\text{AAL}_r := \mathbb{E}[P_r^b] = \mathbb{E}[Y_r].$$

5.4 Basis risk diagnosis

A central governance concern in parametric risk transfer is basis risk: the mismatch between the indemnity-style loss Y_r and parametric P_r . Let

$$B_r^{\{\cdot\}} := Y_r - P_r^{\{\cdot\}},$$

for $\cdot \in \{c, b\}$.

5.4.1 Point diagnosis

It is not difficult to show that

$$\mathbb{E} \left[B_r^{\{\cdot\},+} \right] = \mathbb{E} \left[D_r C_A \left(\frac{l_r + P_r^{\{\cdot\}}}{D_r} \right) \right]$$

On the other hand, given the AAL neutrality, we get $\mathbb{E} \left(B_r^{\{\cdot\}} \right) = 0$ which implies $\mathbb{E} \left(B_r^{\{\cdot\},+} \right) = \mathbb{E} \left(B_r^{\{\cdot\},-} \right)$. In addition, we can show that

$$\text{Cov} \left(Y_r, P_r^{\{\cdot\}} \right) = \mathbb{E} \left[P_r^{\{\cdot\}} D_r C_A \left(\frac{l_r}{D_r} \right) \right] - \left(\text{AAL}_r^{\{\cdot\}} \right)^2.$$

Continuous NatPar. With $P_{r,t}^c$, we get

1. $\mathbb{E} \left[B_r^{c,+} \right] = \mathbb{E} \left[D_r C_{A_r} \left(\frac{l_r}{D_r} \right) \mathbf{1}_{\{D_r \leq d_r\}} \right] + \mathbb{E} \left[D_r C_{A_r} \left(\frac{l_r}{D_r} + q_r^c \left(1 - \frac{d_r}{D_r} \right) \right) \mathbf{1}_{\{D_r > d_r\}} \right],$
2. $\text{Var} (P_r^c) = (q_r^c)^2 \left(\mathbb{E} \left[D_r^2 \left(\left(1 - \frac{d_r}{D_r} \right)^+ \right)^2 \right] - \mathbb{E} \left[D_r \left(1 - \frac{d_r}{D_r} \right)^+ \right]^2 \right),$
3. $\text{Cov} (Y_r, P_r^c) = q_r^c \mathbb{E} \left[(D_r)^2 \left(\left(1 - \frac{d_r}{D_r} \right)^+ \right)^2 C_A \left(\frac{l_r}{D_r} \right) \right] - (\text{AAL}_r)^2.$

Binary NatPar. With P_r^b , the same way we get,

1. $\mathbb{E} \left[B_r^{b,+} \right] = \mathbb{E} \left[D_r C_{A_r} \left(\frac{l_r}{D_r} \right) \mathbf{1}_{\{D_r \leq d_r\}} \right] + \mathbb{E} \left[D_r C_{A_r} \left(\frac{l_r + q_r^b}{D_r} \right) \mathbf{1}_{\{D_r > d_r\}} \right],$
2. $\text{Var} (P_r^b) = (q_r^b)^2 \mathbb{P}(D_r > d_r) (1 - \mathbb{P}(D_r > d_r)),$
3. $\text{Cov} (Y, P^b) = q_r^b \left(\mathbb{E} \left[D_r C_{A_r} \left(\frac{l_r}{D_r} \right) \right] - \text{AAL}_r \mathbb{P}(D_r > d_r) \right).$

5.4.2 Basis exceedance curve

In this part we treat basis risk as a *distributional object*, not just a point diagnostic. The appropriate analogue of an EP curve is the exceedance-basis curve

$$\text{BEP}_r^{\{\cdot\},+}(x) := \mathbb{P}\left(B_r^{\{\cdot\}} > x\right), \quad x \geq 0,$$

For symmetry we also report

$$\text{BEP}_r^{\{\cdot\},-}(x) := \mathbb{P}\left(-B_r^{\{\cdot\}} > x\right), \quad x \geq 0.$$

Here $B_r^{\{\cdot\}} > 0$ is *shortfall* (parametric pays less than indemnity by more than x), while $B_r < 0$ (or $-B_r > 0$) is *overpayment*. Under the AAL-neutral calibration $\mathbb{E}[P_r^{\{\cdot\}}] = \mathbb{E}[Y_r]$, we have $\mathbb{E}[B_r^{\{\cdot\}}] = 0$, but this does not control tails: the relevant governance question is how often and how severely the contract underpays (shortfall tails) and overpays (overpayment tails).

Fix $x \geq 0$ and condition on T_r . We distinguish two cases:

Deriving Shortfall exceedance. If $D_r = 0$, then $Y_r = P_r^c = 0$ and $B_r^c = 0$. If $D_r > 0$, then

$$B_r^{\{\cdot\}} > x \iff (A_r D_r - l_r)^+ > x + P_r^{\{\cdot\}} \iff A_r > \frac{l_r + x + P_r^{\{\cdot\}}}{D_r}.$$

Therefore, as $\bar{F}_{A_r}\left(\frac{l_r + x + P_r^{\{\cdot\}}}{D_r}\right) > 0$ implies $D_r > 0$ we get:

$$\text{BEP}_r^{\{\cdot\},+}(x) = \mathbb{E}\left[\bar{F}_{A_r}\left(\frac{l_r + x + P_r^{\{\cdot\}}}{D_r}\right)\right]. \quad (9)$$

Deriving Overpayment exceedance. If $D_r = 0$, then $Y_r = P_r^c = 0$ and $B_r^c = 0$. If $D_r > 0$, we can consider two cases.

- **Case 1:** $A_r D_r - l_r > 0$. This implies $-B_r^c = P_r^{\{\cdot\}} - Y_r > x \iff \frac{l_r}{D_r} < A_r < \frac{l_r - x + P_r^{\{\cdot\}}}{D_r}$. This also implies $P_r^{\{\cdot\}} > x$.
- **Case 2:** $A_r D_r - l_r \leq 0$. This implies $B_r^c < -x \iff P_r^{\{\cdot\}} > x$. This also implies $D_r > 0$.

So combining the two:

$$\text{BEP}_r^{\{\cdot\},-}(x) = \mathbb{E}\left[F_{A_r}\left(\frac{l_r - x + P_r^{\{\cdot\}}}{D_r}\right) \mathbf{1}_{\{P_r^{\{\cdot\}} > x\}}\right]. \quad (10)$$

5.5 Continuous NatPar basis analysis: analytic shortfall and overpayment

Under the baseline independence assumptions, basis exceedance admits a conditional reduction to the exposure CDF $F_{A,r}$, in the same spirit as the NatCat EP reduction in (1).

Shortfall exceedance. Therefore,

$$\text{BEP}_r^{c,-}(x) = \mathbb{E} \left[\bar{F}_{A_r} \left(\frac{l_r + x}{D_r} + q_r^{\text{cont}} \left(1 - \frac{d_r}{D_r} \right) \right) \mathbf{1}_{\{D_r > d_r\}} \right] + \mathbb{E} \left[\bar{F}_{A_r} \left(\frac{l_r + x}{D_r} \right) \mathbf{1}_{\{D_r \leq d_r\}} \right]. \quad (11)$$

This formula shows explicitly how exposure uncertainty enters shortfall tails: even though the NatPar payout is hazard-only, the probability of severe shortfall is controlled by the upper tail of exposure through $1 - F_{A_r}(\cdot)$.

Overpayment exceedance. Overpayment is piecewise because $Y_{r,t}$ is truncated by the deductible. We get the following result:

$$\text{BEP}_r^{c,-}(x) = \mathbb{E} \left[F_{A_r} \left(\frac{l_r - x}{D_r} + q_r^{\text{cont}} \left(1 - \frac{d_r}{D_r} \right) \right) \mathbf{1}_{\{D_r q_r^{\text{cont}} \left(1 - \frac{d_r}{D_r} \right) > x\}} \right]. \quad (12)$$

5.6 Binary NatPar: analytic shortfall and overpayment

Because P_r is *piecewise constant* in the hazard proxy, the basis distribution decomposes cleanly by whether the trigger fires:

Shortfall exceedance. For $x > 0$, we get

$$\text{BEP}_r^{b,+}(x) = \mathbb{E} \left(\bar{F}_{A_r} \left(\frac{l_r + x + q_r^{\text{bin}}}{D_r} \right) \mathbf{1}_{\{D_r > d_r\}} \right) + \mathbb{E} \left(\bar{F}_{A_r} \left(\frac{l_r + x}{D_r} \right) \mathbf{1}_{\{D_r \leq d_r\}} \right). \quad (13)$$

Each term admits the same exposure-CDF reduction as in (1), with thresholds $\frac{l_r + x}{D_r}$ and $\frac{l_r + x + q_r^{\text{bin}}}{D_r}$, respectively.

Overpayment exceedance. Similarly,

$$\text{BEP}_r^{b,-}(x) = \begin{cases} \mathbb{E} \left(F_{A_r} \left(\frac{l_r - x + q_r^{\text{bin}}}{D_r} \right) \mathbf{1}_{\{D_r > d_r\}} \right), & q_r^{\text{bin}} > x \\ 0, & q_r^{\text{bin}} \leq x \end{cases}. \quad (14)$$

6 Reporting the frost risk example

We begin with a stylised frost model in two regions $r \in \{\text{FL, CA}\}$ (representing Florida and California). A winter season t is characterised by the seasonal minimum temperature T_r (hazard proxy) and a seasonal exposure A_r (value at risk). The objective of this subsection is twofold: (i) to make the NatCat “hazard–exposure–vulnerability” loss construction explicit, and (ii) to express key distributional objects (EP curves and AAL) in a conditional form that will later carry over *verbatim* to basis-risk analysis.

6.1 Hazard–exposure–vulnerability and NatCat indemnity

Hazard (temperature). For each region r , seasonal minimum temperature is modelled as Gaussian and independent across seasons:

$$T_r \sim \text{Normal}(\mu_{T,r}, \sigma_{T,r}^2).$$

Region-specific parameters $(\mu_{T,r}, \sigma_{T,r})$ are reported in the input-parameter table.¹ This is the hazard component of the NatCat stack.

Exposure (value at risk). Exposure is modelled as lognormal:

$$A_r \sim \text{LogNormal}(\mu_{A,r}, \sigma_{A,r}^2), \quad A_{r,t} > 0.$$

Rather than parameterising by $(\mu_{A,r}, \sigma_{A,r})$ directly, we specify (by region) the mean $\bar{A}_r := \mathbb{E}[A_{r,t}]$ and the coefficient of variation

$$\text{CV}_{A,r} := \frac{\sqrt{\text{Var}(A_{r,t})}}{\bar{A}_r},$$

which is a natural underwriting summary of exposure uncertainty. These imply the log-normal parameters

$$\sigma_{A,r}^2 = \ln(1 + \text{CV}_{A,r}^2), \quad \mu_{A,r} = \ln(\bar{A}_r) - \frac{1}{2}\sigma_{A,r}^2.$$

Across seasons, $(A_{r,t})_t$ are independent. As a simplifying baseline we assume

$$T_r \perp A_r,$$

so exposure variability is not mechanically linked to hazard realizations. This independence is relaxed in extensions (e.g. copulas or $A_r \mid T_r$ regression) if dependence is empirically material.

¹The Gaussian assumption is used for transparency in the laboratory. Any alternative calibrated marginal for T can be substituted without changing the NatCat–NatPar logic.

Vulnerability (temperature–damage). Let $D : \mathbb{R} \rightarrow [0, 1]$ denote the frost damage fraction:

$$D(\tau) = \begin{cases} 0, & \tau > \tau_c, \\ \left(\frac{\tau - \tau_c}{\tau_t - \tau_c} \right)^\eta, & \tau_t \leq \tau \leq \tau_c, \\ 1, & \tau < \tau_t, \end{cases} \quad (\tau^1, \tau^0, \eta) = (28, 20, 1.5),$$

where τ_t is the critical temperature at which damage begins, $\tau_t < \tau_c$ is the temperature at which damage is total, and $\eta > 0$ controls curvature. In the following we illustrate the damage function in Figure 1.




Figure 1: Citrus fruit damage function from frost.

6.2 Basis risk diagnostics

This subsection compares the NatCat extreme-layer benchmark to two NatPar designs (continuous and binary) using exceedance and return-period diagnostics. Within each (r, β) block we enforce AAL-neutrality, so all differences reported below come from *tail shape* and *basis structure*, not from mean level.

For each region r we simulate N seasons and compute damage and loss as $D_r = D(T_r) \in [0, 1]$ and $L_r = A_r D_r$. For each $\beta \in \{0.90, 0.95, 0.99\}$ we define the NatCat extreme-layer benchmark via

$$l_{r,\beta} := \text{VaR}_\beta(L_r), \quad Y_{r,\beta} := (L_r - l_{r,\beta})^+,$$

and set the NatPar trigger on damage at the same percentile,

$$d_{r,\beta} := \text{VaR}_\beta(D_r).$$

We evaluate two NatPar archetypes:

$$P_{r,\beta}^c = q_{r,\beta}^c (D_r - d_{r,\beta})^+, \quad P_{r,\beta}^b = q_{r,\beta}^b \mathbf{1}_{\{D_r > d_{r,\beta}\}}.$$

Both designs are calibrated by AAL-neutrality relative to $Y_{r,\beta}$. Hence, within each (r, β) , the benchmark $Y_{r,\beta}$ and both NatPar payouts have equal mean, but potentially very different tails and mismatch (basis) behaviour.

6.2.1 Individual risk

In the following we will be looking into each individual risk associated with the two region.

Exceedance probability curves (liability tails) We compare the exceedance probability (EP) functions on a log scale. The binary design has a structural feature: $P_{r,\beta}^b$ places a point mass at $q_{r,\beta}^b$, so $EP^{b,+}(x)$ is essentially flat up to that level and drops to zero immediately after. This creates visible “digital” tail geometry relative to the smoother benchmark and the continuous design; see Figure 2.




Figure 2: Exceedance probability curves for $Y_{r,\beta}$, $P_{r,\beta}^{\{\cdot\}}$ (log scale), by region and β .

Point diagnosis of basis risk Before presenting tail curves, we report a set of *point diagnostics* that summarize the sign and average magnitude of basis mismatch at the

calibrated layer. For each region r and confidence level β , we construct the indemnity-layer benchmark $Y_{r,\beta} = (L - l_{r,\beta})^+$ and two NatPar designs, continuous $P_{r,\beta}^c$ and binary $P_{r,\beta}^b$, calibrated by mean matching $\mathbb{E}[P_{r,\beta}^{\{\cdot\}}] = \mathbb{E}[Y_\beta]$.

Table 1 reports the calibration levels and dependence diagnostics, including $(l_{r,\beta}, d_{r,\beta}, q_{r,\beta})$, $\mathbb{E}[Y_{r,\beta}]$, and $\text{Cov}(Y_{r,t,\beta}, P_{r,\beta}^{\{\cdot\}}) / \text{Corr}(Y_{r,\beta}, P_{r,\beta}^{\{\cdot\}})$. Table 2 reports the basis mismatch decomposition: $\mathbb{P}(B_{r,\beta}^{\{\cdot\}} > 0)$, $\mathbb{P}(B_{r,\beta}^{\{\cdot\}} < 0)$, $\mathbb{E}[(B_{r,\beta}^{\{\cdot\}})^+]$ (average shortfall), and $\mathbb{E}[(B_{r,\beta}^{\{\cdot\}})^-]$ (average overpayment). These point diagnostics capture the *frequency* and *average size* of mismatch, while the subsequent EP and return-period plots describe the *tail behavior*.

Region	β	Design	$l_{r,\beta}$	$d_{r,\beta}$	$q_{r,\beta}$	$\mathbb{E}[Y_\beta]$	$\text{Var}(B)$	$\text{Corr}(Y_\beta, P_\beta)$
FL	0.90	c	11.58	0.05	257.77	5.12	107.95	0.91
FL	0.90	b	11.58	0.05	51.16	5.12	368.93	0.62
FL	0.95	c	42.12	0.18	275.78	2.96	103.10	0.86
FL	0.95	b	42.12	0.18	59.20	2.96	224.00	0.64
FL	0.99	c	137.84	0.54	346.45	0.69	62.19	0.64
FL	0.99	b	137.84	0.54	68.82	0.69	66.34	0.57
CA	0.90	c	81.09	0.41	209.43	6.48	91.55	0.92
CA	0.90	b	81.09	0.41	64.76	6.48	255.81	0.77
CA	0.95	c	136.42	0.69	231.27	2.47	67.21	0.80
CA	0.95	b	136.42	0.69	49.36	2.47	84.22	0.74
CA	0.99	c	215.21	1.00	0.00	0.30	16.39	NaN
CA	0.99	b	215.21	1.00	0.00	0.30	16.39	NaN

Table 1: Point diagnosis: calibration levels and dependence.

Region	β	Design	$\mathbb{P}(P > 0)$	$\mathbb{P}(B > 0)$	$\mathbb{P}(B < 0)$	$\mathbb{E}[B^+]$	$\mathbb{E}[-B^+]$
FL	0.90	c	0.10	0.05	0.06	0.98	0.98
FL	0.90	b	0.10	0.04	0.07	2.13	2.13
FL	0.95	c	0.05	0.03	0.03	0.79	0.79
FL	0.95	b	0.05	0.02	0.03	1.25	1.25
FL	0.99	c	0.01	0.01	0.01	0.33	0.33
FL	0.99	b	0.01	0.01	0.01	0.36	0.36
CA	0.90	c	0.10	0.05	0.06	1.16	1.16
CA	0.90	b	0.10	0.05	0.06	2.12	2.12
CA	0.95	c	0.05	0.03	0.03	0.75	0.75
CA	0.95	b	0.05	0.03	0.03	0.88	0.88
CA	0.99	c	0.00	0.01	0.00	0.30	0.00
CA	0.99	b	0.00	0.01	0.00	0.30	0.00

Table 2: Point diagnosis: basis mismatch decomposition.

Basis exceedance curves (shortfall vs overpayment) We report the two-sided basis exceedance functions shortfall and overpayment. These separate consumer-protection tail risk (shortfall) from capital/pricing tail risk (overpayment). In the outputs, the binary design typically exhibits a sharper truncation in one basis tail (due to the fixed payout size), while the continuous design exhibits smoother decay see Figure 3.

Return-period levels and return-period curves A key feature in the results is the presence of many zeros at low return periods: for the layer $Y_{r,\beta} = (L - l_{r,\beta})^+$, we have




Figure 3: Basis exceedance curves for continuous and binary NatPar by region and β (log scale).

$\mathbb{P}(Y_{r,\beta} > 0) = 1 - \beta$, so when $1/T \geq 1 - \beta$ the return level is exactly zero. This is visible for $\beta = 0.90$ up to $T = 10$, for $\beta = 0.95$ up to $T = 20$, and for $\beta = 0.99$ up to much larger T in several region/design combinations.

The binary payout also produces plateau behaviour in $x_T(P_{r,\beta}^b)$ because of its point mass at $q_{r,\beta}^b$. Once T exceeds the reciprocal of the trigger probability, the return level cannot increase further and saturates at the fixed payout size; see Figure 4.

Basis return-period levels: T -year shortfall and overpayment We summarise basis tails via

$$b_T^+(B) := x_T(B), \quad b_T^-(B) := x_T(-B),$$

so b_T^+ is the T -year *shortfall* magnitude and b_T^- is the T -year *overpayment* magnitude. Results show that the binary design can generate substantially larger shortfall return levels at high T (e.g. for $\beta = 0.90$ and $\beta = 0.95$, $b_T^+(B^b)$ notably exceeds $b_T^+(B^c)$ at $T = 100, 200$), reflecting that fixed payouts can under-respond to the largest indemnity-layer realisations even under AAL matching; see Figure 5.

Capital metrics. To support solvency-style benchmarking, we additionally report one-year tail capital metrics for the NatPar liability P and, for reference, for the matched NatCat layer Y . Specifically, for each region r and layer confidence β , we compute $\text{VaR}_{0.995}$




Figure 4: Return-period curves $T \mapsto x_T(\cdot)$ for $Y_{r,\beta}$, $P_{r,\beta}^{\{\cdot\}}$, by region and β .




Figure 5: Return-period curves for basis tails: $T \mapsto b_T^+$ (shortfall) and $T \mapsto b_T^-$ (overpayment), for both continuous and binary NatPar, by region and β .

(and $\text{TVaR}_{0.995}$ where feasible) for Y_β and for each NatPar design P_β^c and P_β^b . These quantities summarize extreme one-year liability levels and enable like-for-like comparison between indemnity and parametric programmes on the same hazard/exposure base.

Table 3 reports $\text{VaR}_{0.995}$ and $\text{TVaR}_{0.995}$ for $Y_{r,\beta}$, $P_{r,\beta}^{\{\cdot\}}$ across $\beta \in \{0.90, 0.95, 0.99\}$. In the laboratory setting (at most one seasonal event per year), these correspond directly to EP/AEP return levels at exceedance probability 0.5%.

Region	β	$\text{VaR}_{0.995}(Y_\beta)$	$\text{TVaR}_{0.995}(Y_\beta)$	$\text{VaR}_{0.995}(P_\beta^c)$	$\text{TVaR}_{0.995}(P_\beta^c)$	$\text{VaR}_{0.995}(P_\beta^b)$	$\text{TVaR}_{0.995}(P_\beta^b)$
CA	0.90	156.77	184.43	123.27	123.27	64.76	64.76
CA	0.95	101.44	129.10	71.34	71.34	49.36	49.36
CA	0.99	22.65	50.31	0.00	0.00	0.00	0.00
FL	0.90	172.24	243.83	166.81	209.95	51.16	51.16
FL	0.95	141.69	213.29	143.11	189.27	59.20	59.20
FL	0.99	45.98	117.57	55.90	113.89	68.82	68.82

Table 3: Capital metrics for NatCat layer $Y_{r,\beta}$ and NatPar liabilities $P_{r,\beta}^{\{\cdot\}}$: $\text{VaR}_{0.995}$ and $\text{TVaR}_{0.995}$, by region and layer β .

6.2.2 Portfolio risk

In the following we will be looking into the protfolio risk associated with the two region.

Simulation assumptions for the portfolio. The numerical assessment is based on Monte Carlo samples of regional indemnity losses of the form ($L_r = A_r D(T_r)$). Within each region (r), the exposure factor (A_r) and the hazard driver (T_r) are taken to be independent, so that variation in severity arises from multiplicative exposure scaling and the damage mapping ($D(\cdot)$). At the portfolio level, regional sample pairs are combined by aligning draws across regions; unless explicitly stated otherwise, this corresponds to treating $(A_{\text{FL}}, T_{\text{FL}})$ and $(A_{\text{CA}}, T_{\text{CA}})$ as independent. If the regional samples are generated from common-year (shared-path) climate simulations, then any cross-region dependence in $(T_{\text{FL}}, T_{\text{CA}})$ is inherited from that data-generation step rather than imposed by the aggregation code.

Exceedance probability curves (liability tails) We compare the exceedance probability (EP) functions on a log scale for the liabilites in the portfolio. In more details we compare the OEP and AEP of the indemnity and the parametric products.

Point diagnosis of basis risk Let $B_{r,t,\beta}^{\{\cdot\}} = Y_{r,\beta} - P_{r,\beta}^{\{\cdot\}}$ denote the basis in region r at layer β (continuous or binary).

Given portfolio, define the *aggregate basis* as

$$S_\beta^{B^{\{\cdot\}}} = \sum_{r \in \mathcal{R}} B_{r,t,\beta}^{\{\cdot\}} = S_\beta^Y - S_\beta^{P^{\{\cdot\}}}.$$




Figure 6: AEP and OEP for the indemnity and parametric portfolio.

This quantity is natural when the objective is balance-sheet impact and annual aggregate liquidity, and it aligns with AEP-style reporting for liabilities.

For basis regulation and suitability, let:

$$M_\beta^{B^{\{\cdot\}}} = \max_{r \in \mathcal{R}} B_{r,\beta}^{\{\cdot\}}.$$

Point measures. Similar to the above this means measures that can represent the portfolio basis risk through $\mathbb{P}(S_\beta^{B^{\{\cdot\}}} > 0)$, $\mathbb{E}\left[\left(S_\beta^{B^{\{\cdot\}}}\right)^+\right]$, $\mathbb{E}\left[\left(-S_\beta^{B^{\{\cdot\}}}\right)^+\right]$, $\text{corr}(S_\beta^Y, S_\beta^P)$, $\mathbb{P}(M_\beta^{B^{\{\cdot\}}} > 0)$, $\mathbb{E}\left[\left(M_\beta^{B^{\{\cdot\}}}\right)^+\right]$, $\mathbb{E}\left[\left(-M_\beta^{B^{\{\cdot\}}}\right)^+\right]$, $\text{corr}(M_\beta^Y, M_\beta^P)$.

Scope	Design	β	$\mathbb{E}[Y]$	$\text{Corr}(Y, P)$	$\text{Var}(B)$
Agg	b	0.90	11.59	0.70	624.45
Agg	c	0.90	11.59	0.92	199.54
Agg	b	0.95	5.43	0.68	308.34
Agg	c	0.95	5.43	0.84	170.47
Agg	b	0.99	0.99	0.53	82.75
Agg	c	0.99	0.99	0.60	78.55
Occ	b	0.90	11.29	0.69	423.10
Occ	c	0.90	11.29	0.91	113.43
Occ	b	0.95	5.35	0.67	211.30
Occ	c	0.95	5.35	0.84	95.60
Occ	b	0.99	0.99	0.53	60.59
Occ	c	0.99	0.99	0.60	51.29

Table 4: Aggregate vs Occurrence: level and dependence.

Scope	Design	β	$\mathbb{P}(P > 0)$	$\mathbb{P}(B > 0)$	$\mathbb{P}(B < 0)$	$\mathbb{E}[B^+]$	$\mathbb{E}[(-B)^+]$
Agg	b	0.90	0.19	0.09	0.12	4.12	4.12
Agg	c	0.90	0.19	0.10	0.11	2.08	2.08
Agg	b	0.95	0.10	0.05	0.06	2.10	2.10
Agg	c	0.95	0.10	0.05	0.06	1.52	1.52
Agg	b	0.99	0.01	0.02	0.01	0.66	0.36
Agg	c	0.99	0.01	0.02	0.01	0.63	0.33
Occ	b	0.90	0.19	0.09	0.11	4.16	3.94
Occ	c	0.90	0.19	0.10	0.11	2.06	2.05
Occ	b	0.95	0.10	0.05	0.06	2.10	2.05
Occ	c	0.95	0.10	0.05	0.06	1.51	1.51
Occ	b	0.99	0.01	0.02	0.01	0.66	0.36
Occ	c	0.99	0.01	0.02	0.01	0.63	0.33

Table 5: Aggregate vs Occurrence: point diagnosis.

Tail measure. Then let $BAEP_{\beta}^{\pm}(x) = \mathbb{P}(S_{\beta}^{B^{\{\cdot\}}} > x)$ and $BOEP_{\beta}^{\pm}(x) = \mathbb{P}(M_{\beta}^{B^{\{\cdot\}}} > x)$, together with basis return levels $Ab_T^{\pm, \{\cdot\}}$ and $Ob_T^{\pm, \{\cdot\}}$ defined by $BAEP_{\beta}^{\pm}(x) = 1/T$ and $BOEP_{\beta}^{\pm}(x) = 1/T$.




Figure 7: BAEP and BOEP for the continuous and binary parametric portfolio.

Capital metrics. For each $\beta \in \{0.90, 0.95, 0.99\}$ we then compute $\text{VaR}_{0.995}$ and, where feasible, $\text{TVaR}_{0.995}$ for the benchmark NatCat layer $S_{Y,\beta}$ and $M_{Y,\beta}$ and for each NatPar design $S_{P,\beta}^c$, $M_{P,\beta}^c$, $S_{P,\beta}^b$, $M_{P,\beta}^b$. These capital metrics quantify one-year extreme liability levels under (i) *diversification across regions* in the aggregate view and (ii) *worst-region dominance* in the occurrence view, and they provide a direct, like-for-like benchmarking of parametric and indemnity-style tails on the same hazard/exposure base.




Figure 8: Return-period curves for liabilities (Aggregate vs Occurrence)




Figure 9: Return-period curves for basis (Aggregate vs Occurrence, cont & bin, +/-)

Scope	β	$\text{VaR}_{0.995}(Y_\beta)$	$\text{TVaR}_{0.995}(Y_\beta)$	$\text{VaR}_{0.995}(P_\beta^c)$	$\text{TVaR}_{0.995}(P_\beta^c)$	$\text{VaR}_{0.995}(P_\beta^b)$	$\text{TVaR}_{0.995}(P_\beta^b)$
Agg	0.90	123.89	173.81	119.25	149.64	55.24	55.24
Agg	0.95	100.32	150.47	100.78	133.39	41.44	42.21
Agg	0.99	33.71	82.65	39.13	79.72	48.17	48.17
Occ	0.90	120.57	170.68	116.77	146.97	35.81	35.81
Occ	0.95	99.19	149.30	100.18	132.49	41.44	41.44
Occ	0.99	33.33	82.51	39.13	79.72	48.17	48.17

Table 6: Capital metrics for portfolio of NatCat layer $Y_{r,\beta}$ and NatPar liabilities $P_{r,\beta}^{\{\cdot\}}$: $\text{VaR}_{0.995}$ and $\text{TVaR}_{0.995}$, by layer β .

7 Results interpretation of the frost laboratory

Figures 2–3 reveal a systematic pattern: in the far tail, the *shortfall* component of basis risk dominates the *overpayment* component. Interpreting this correctly requires separating an economic conclusion (who benefits) from a mathematical one (why the tail geometry looks the way it does). The plots primarily reflect a structural feature of the design: the indemnity-layer benchmark $Y_{r,\beta}$ is driven by *exposure–severity* through $L = AD$, while the parametric payouts are functions of the *damage index* D alone and are therefore bounded.

The core mechanism. In our setting $L = A_r D_r$, where the random exposure/severity multiplier $A_r > 0$ is independent of D_r and may exhibit substantial right-tail variability, while the damage ratio satisfies $D_r \in [0, 1]$. The NatPar designs are functions of D and inherit its bounded support:

$$0 \leq P_{r,\beta}^b \leq q_{r,\beta}^b, \quad 0 \leq P_{r,\beta}^c \leq q_{r,\beta}^c (1 - d_{r,\beta}),$$

where $d_{r,\beta} = \text{VaR}_\beta(D_r)$ and $q_{r,\beta}^{\{\cdot\}}$ are chosen to match the mean layer loss $\mathbb{E}[Y_{r,\beta}]$. In contrast, $Y_{r,\beta} = (A_r D_r - l_{r,\beta})^+$ can be large whenever A_r is large enough, even if D_r is only moderate. Consequently, in extreme states the basis behaves as

$$B_{r,\beta}^{\{\cdot\}} = Y_{r,\beta} - P_{r,\beta}^{\{\cdot\}} \approx Y_{r,\beta} \quad \text{whenever} \quad Y_{r,\beta} \gg \sup P_{r,\beta}^{\{\cdot\}},$$

so the far-right tail of $B_{r,\beta}^{\{\cdot\}}$ inherits the tail of Y_β while the overpayment side is capped. This is the mathematical reason the return-period shortfall curves $b_T^+(B_{r,\beta}^{\{\cdot\}})$ grow faster than the corresponding overpayment curves $b_T^-(B_{r,\beta}^{\{\cdot\}})$ in Figure 5, and why the basis exceedance curves BEP^+ remain comparatively heavy in Figure 3.

Unspanned severity. The dominance of shortfall is not merely “because P is bounded”; it arises because a D -only payout cannot span exposure-driven severity. Large values of $Y_{r,\beta}$ can occur via two mechanisms:

- (i) large exposure) A_r large with moderate D_r , (ii) large damage) D_r large with moderate A_r .

Binary and continuous NatPar designs respond only to mechanism (ii). Under independence of A and D , mechanism (i) occurs with non-negligible probability in the tail, producing states where $Y_{r,\beta} > 0$ but $P_{r,\beta}^{\{\cdot\}}$ is small or zero, hence $B_{r,\beta}^{\{\cdot\}}$ is large and positive (shortfall). In contrast, overpaymeY_{\beta}nt requires the reverse mismatch: D large (trigger) together with small A_r so that $Y_{r,\beta}$ is near zero while $P_{r,\beta}^{\{\cdot\}}$ is positive. Such states exist, but their severity is capped by $\sup P_{r,\beta}^{\{\cdot\}}$, and therefore overpayment does not dominate at long return periods.

Profitability. It is tempting to read “shortfall dominates overpayment” as an issuer advantage and therefore a reason to underwrite the product. That inference is incomplete. Overpayment corresponds to paying more than the benchmark layer and is indeed the issuer’s direct economic downside relative to $Y_{r,\beta}$. Shortfall, however, is not a cash loss to the issuer; it is a *coverage failure* borne by the policyholder and often a commercial or conduct risk for the issuer. A design that generates large tail shortfalls may be capital-tractable (bounded liability) but can be commercially fragile: it may underperform exactly when protection is most expected, raising reputational risk, suitability concerns, and disclosure requirements. Hence tail shortfall dominance is primarily evidence of *bounded-liability structure* rather than a direct profitability statement.

Design implication. We select $q_{r,\beta}^{\{\cdot\}}$ to match $\mathbb{E} \left[P_{r,\beta}^{\{\cdot\}} \right] = \mathbb{E} [Y_{r,\beta}]$ (AAL matching at the layer). Figures 2 and 4 show that this mean matching can coexist with large differences in tail behavior. The reason is structural: matching $\mathbb{E} [Y_{r,\beta}]$ does not control how Y_β decomposes into exposure-driven versus damage-driven extremes. When exposure variability dominates severity, a damage-only index necessarily under-replicates the far tail, producing heavy BEP⁺ and large b_T^+ .

In this numerical assessment, the observed dominance of tail shortfall reflects the combination of (i) exposure-driven severity in $L = A_r D_r$ with independent A and bounded D , and (ii) D -only payouts that cap the issuer’s liability. The results should not be read as “parametric is better,” but as a quantified statement about what this particular index can and cannot span. If the intended objective is tail risk transfer (rather than bounded liquidity), then either the index must incorporate exposure scaling (e.g. an exposure-adjusted index) or the contract must explicitly acknowledge its capped nature through limits, disclosures, and suitability framing.

Individual risk: what changes once FL and CA are pooled The individual capital table already shows that the two regional assets behave very differently in the far tail, and that the difference is design-dependent. For FL, the benchmark layer Y_β remains heavy-tailed: at $\beta = 0.90$ we have $\text{VaR}_{0.995}(Y_\beta) = 172.24$ and $\text{TVaR}_{0.995}(Y_\beta) = 243.83$, whereas the continuous design tracks the tail more closely with $\text{VaR}_{0.995}(P_\beta^c) = 166.81$

and $\text{TVaR}_{0.995}(P_\beta^c) = 209.95$. In contrast, the binary design is strongly capped, with $\text{VaR}_{0.995}(P_\beta^b) = \text{TVaR}_{0.995}(P_\beta^b) = 51.16$, i.e. the tail becomes essentially “flat” beyond the trigger. The CA asset is materially lighter in the far tail: at $\beta = 0.90$, $\text{VaR}_{0.995}(Y_\beta) = 156.77$ and $\text{TVaR}_{0.995}(Y_\beta) = 184.43$; and again the binary liability is capped (64.76) while the continuous liability is closer to the benchmark (123.27). At $\beta = 0.99$, CA becomes almost non-responsive in the tail under this index calibration: $\text{VaR}_{0.995}(P_\beta^c) = 0$ and $\text{VaR}_{0.995}(P_\beta^b) = 0$, while Y_β still has non-zero tail levels ($\text{VaR}_{0.995} = 22.65$, $\text{TVaR}_{0.995} = 50.31$). This is an explicit “tail non-spanning” signal for CA at very high β under a D -only trigger.

Portfolio versus individual risk Once we move to the two-asset portfolio, the point-diagnosis tables clarify *what pooling does and does not fix*. First, the mean layer size drops sharply as β increases, as expected for an exceedance layer: $\mathbb{E}[S_Y]$ is about 5.52 at $\beta = 0.90$, 2.81 at $\beta = 0.95$, and 0.57 at $\beta = 0.99$. This is not a “diversification” effect; it is the mechanical effect of pushing the attachment deeper into the tail. The portfolio story is instead in dependence and mismatch.

Second, the continuous design remains systematically more aligned with the benchmark than the binary design, and the gap is quantitatively large. In the aggregate (AEP-style) view, $\text{Corr}(S_Y, S_P)$ is 0.91 for design c versus 0.65 for design b at $\beta = 0.90$; 0.85 versus 0.65 at $\beta = 0.95$; and 0.63 versus 0.56 at $\beta = 0.99$. This is exactly what one expects from a “shape” argument: P_β^c is proportional to the excess damage layer $(D - d_\beta)^+$ and can respond continuously to severity once triggered, whereas P_β^b is essentially a two-point mass and cannot reproduce severity once it fires.

Third, and more importantly for governance, the portfolio basis dispersion is drastically higher for the binary design. In the aggregate view, $\text{Var}(S_B)$ is 203.74 (design b) versus 61.14 (design c) at $\beta = 0.90$, i.e. more than a threefold increase in mismatch variability; at $\beta = 0.95$, $\text{Var}(S_B)$ is 117.36 (b) versus 56.60 (c); even at $\beta = 0.99$, where both designs become sparse/degenerate, the binary still does not improve (33.99 vs 31.94). This is the numerical counterpart of what the basis-tail plots already suggest: the binary contract is not just “capped”; it is capped in a way that *concentrates* the unspanned component into the shortfall side B^+ whenever exposure-driven severity dominates.

The same conclusion holds under the occurrence (OEP-style) lens. Your portfolio tables show $\text{Var}(M_B)$ of 149.99 (b) versus 36.58 (c) at $\beta = 0.90$, and 84.31 (b) versus 32.46 (c) at $\beta = 0.95$. This matters because the occurrence view is the one that bites for suitability and operational stress: even if the aggregate mismatch nets out across assets in some years, a single region can still experience a large shortfall. That is why, in addition to reporting the aggregate basis $S_B = \sum_r B_r$, we recommend an *occurrence basis* as a regulatory object. In a two-asset portfolio this can be reported as

$$M_B := \max\{B_{\text{FL}}, B_{\text{CA}}\},$$

which directly measures the worst regional liquidity shortfall/overpay created by the index in any year. Importantly, this definition should be based on the per-asset basis terms (not portfolio-level quantiles of L and D), because attachment and triggers are contractual at the risk level; computing “portfolio” l_β^{port} and d_β^{port} mixes attachment/triggers across assets and can mask the very cross-asset tail asymmetries the portfolio analysis is meant to reveal.

Overall, the portfolio does not reverse the individual diagnosis; it *sharpens* it. The continuous design c preserves higher dependence with the benchmark and markedly lowers mismatch dispersion in both AEP- and OEP-style views. The binary design b systematically increases dispersion of basis, which means that pooling two assets does not create “free diversification” when the dominant unspanned driver is exposure severity in $L = AD$ and the payout depends only on D . In that sense, the portfolio results quantify a practical warning: under a damage-only index, adding a second asset may reduce aggregate variance in years where one region does not trigger, but it does not eliminate the structural tail shortfall mechanism, and the occurrence basis M_B remains the binding governance metric.

Regulatory and reporting implication. The diagnostics in Figures 3 and 5 provide information that cannot be inferred from AAL and standard EP curves alone. In particular, (i) basis exceedance curves BEP^\pm summarize the frequency of extreme shortfall/overpayment events as a function of severity, while (ii) basis return levels b_T^\pm translate those exceedance probabilities into T -year “basis severities.” These quantities make explicit whether a parametric design is best interpreted as (a) a loss proxy with acceptable tail mismatch, or (b) a bounded liquidity instrument whose tail protection is limited by construction. For NatPar reporting, we therefore view the pair $(\text{BEP}^\pm, b_T^\pm)$ as essential complements to the usual NatCat objects (AAL, EP curves, and return-period levels).

Market benefit: liquidity. A practical benefit of the bounded, index-linked structure is not that it perfectly replicates tail loss, but that it can function as a *contingent liquidity instrument* in markets where high exposure uncertainty creates insurability gaps. When exposure variability is large (or difficult to verify quickly), indemnity-style settlement becomes slow and costly, and insurers may ration capacity or impose high deductibles and strict wordings. A parametric contract tied to an observable hazard/damage index D can still be offered because the liability is auditable and capped by construction. In this interpretation, the product is best framed as *fast liquidity when conditions are adverse* rather than as a full loss proxy: it pays quickly upon an index trigger, stabilizing cash flows and enabling emergency response, operational continuity, and short-term financing, even though basis shortfall may remain material in the far tail. The relevant question is therefore not whether BEP^+ can be made negligible, but whether the trigger and scale deliver timely cash in the states where liquidity constraints bind, and whether residual loss risk is transparently disclosed and, where needed, complemented by additional layers (e.g. traditional cover, credit lines, or exposure-adjusted indices).

8 Extensions and discussion

The frost example shows, in a highly controlled setting, how NatPar emerges naturally from an existing NatCat model and how it changes the distribution of losses and capital-relevant metrics. In this section we sketch several extensions and discuss broader implications and limitations. The aim is not to exhaust the design space, but to show how the NatPar perspective can be carried into more realistic, multi-peril and multi-index portfolios and how it relates to current debates on solvency and climate-risk management.

8.1 Multi-trigger and multi-peril NatPar programmes

Real-world parametric programmes are rarely as simple as the single-index frost design considered above. They often involve multiple indices, multiple perils and layered structures across regions and client types. The NatPar framework accommodates these extensions, but also clarifies where complexity may undermine transparency and acceptance.

Multi-trigger within a peril. Within a single peril (e.g. tropical cyclone) it is natural to consider multi-trigger structures that use several aspects of the hazard, such as maximum wind speed *and* storm surge, or rainfall accumulation *and* wind. In NatPar notation this corresponds to payout functions

$$P_{e,m} = I_m \left(Z_e^{(1)}, \dots, Z_e^{(J)} \right),$$

with $J > 1$ indices per event. The hazard module already simulates the relevant fields, so constructing such indices is straightforward. The exposure and vulnerability blocks can then be used to assess whether adding indices meaningfully reduces basis risk or merely increases contractual complexity.

From a NatPar standpoint, multi-trigger designs are attractive when they capture qualitatively different damage channels that are difficult to encode in a single index (e.g. wind vs water vs rain-induced flooding), and when each index is itself contractible and robust. They become problematic when the resulting payout surface is opaque to clients and regulators or when small changes in parameterisation produce large, unintuitive changes in payouts.

Multi-peril portfolios. Insurers and reinsurers typically manage portfolios exposed to multiple perils across multiple regions. In a NatCat setting, this is handled by combining event sets and applying the same hazard–exposure–vulnerability–finance machinery across perils. NatPar generalises this by defining indices and parametric programmes for each peril and region, then aggregating the resulting parametric payouts.

If \mathcal{P} indexes perils (e.g. wind, flood, drought, wildfire) and \mathcal{R} indexes regions, a general NatPar portfolio can be written as

$$S^{\text{Par}} = \sum_{p \in \mathcal{P}} \sum_{r \in \mathcal{R}} S_{p,r}^{\text{Par}},$$

with each component $S_{p,r}^{\text{Par}}$ defined previously. The same event-level hazard simulations that drive the NatCat model are used to simulate indices and payouts across perils, allowing for consistent portfolio-level AAL and EP analysis.

Complexity vs transparency. The temptation in multi-trigger, multi-peril settings is to exploit the full flexibility of the $I(X)$ paradigm and to let parametric payouts be very complex functions of many indices. NatPar provides a counterweight: because the same NatCat models can be used to estimate the incremental reduction in basis risk from each added index or trigger, one can explicitly compare the benefit in terms of basis-risk metrics with the cost in terms of loss of transparency, operational complexity and regulatory acceptability. In many cases the NatPar lens suggests that a small number of well-chosen, simple triggers dominates more elaborate multi-trigger designs.

8.2 NatPar and solvency/climate-risk regulation

Solvency and climate-risk regulation increasingly require insurers and banks to demonstrate resilience under severe but plausible climate scenarios. NatCat models are already central to such exercises, providing loss distributions and stress-test results. NatPar programmes interact with this landscape in several ways.

Capital and risk measures. Regulatory capital regimes typically impose risk measures such as VaR or TVaR at high confidence levels over one-year or longer horizons. As shown in the frost example, NatPar portfolios often have shorter tails and lower TVaR for a given AAL compared with their indemnity counterparts, due to bounded, index-linked liabilities and the absence of a long development tail. Under capital regimes that are sensitive to tails rather than to AAL alone, NatPar programmes can therefore be more capital-efficient per unit of expected loss.

From a supervisory perspective, this raises both opportunities and questions: parametric structures may allow insurers to maintain or expand coverage in climate-exposed lines while respecting capital requirements, but they also introduce basis risk that must be understood and, in some cases, explicitly reported or capitalised.

Climate stress testing and scenario analysis. Climate-risk stress testing typically involves altering the hazard module (e.g. increasing the frequency or severity of certain perils, changing spatial patterns) and re-running the NatCat model. In a NatPar setting,

the same altered hazards will feed directly into index distributions and hence parametric payouts. Because NatPar portfolios are more tightly linked to hazard and less to exposure, they can provide sharper insight into how balance sheets respond to physical climate shifts under different adaptation and exposure scenarios.

Banks and public entities that use parametric structures (e.g. sovereign disaster risk-financing facilities, contingent credit lines with parametric triggers) can also incorporate NatPar portfolios into their own stress tests, using the NatCat/NatPar machinery to simulate joint behaviour of losses, payouts and macro-financial variables.

Regulatory views on parametric products. Supervisors and standard-setters have shown growing interest in parametric products, but also concern about their complexity and about consumer protection in the presence of basis risk. The NatPar framework offers a way to present parametric programmes in a language familiar to regulators: as variants of existing NatCat models with clear, quantifiable basis-risk profiles, rather than as opaque new instruments. It also suggests that regulatory guidance may wish to distinguish between simple, hazard-linked NatPar structures with well-understood basis risk and highly complex, multi-trigger designs that may be difficult to communicate or supervise.

8.3 Limitations and practical challenges

While the NatPar framework offers a coherent way of thinking about parametric insurance from a supply-side perspective, it has important limitations and faces practical challenges.

Data and model limitations. NatPar relies on the same underlying hazard, exposure and vulnerability models as NatCat. If these models are poor, biased or unstable, the resulting basis risk analysis and parametric designs will also be unreliable. In some markets hazard data are sparse, exposure data are incomplete, and loss histories are short; in others, non-stationarity is so pronounced that historical patterns may be a poor guide to future risks. NatPar does not remove these problems; it re-allocates them.

Governance of indices and triggers. Parametric payouts are only as good as the indices they rely on. This raises questions of governance: who controls the index, how revisions or errors are handled, how disputes are resolved, and what happens when measurement systems change (for example, replacement of weather stations or satellite products). Robust contractual and institutional arrangements are needed to ensure that indices remain trustworthy over the lifetime of a NatPar programme.

Market acceptance and client understanding. Even simple parametric structures can be unfamiliar to policyholders and intermediaries used to indemnity insurance. Basis risk, in particular, can be difficult to explain: clients may perceive a contract as unfair if it fails to pay in a year when they experience losses, even if the overall design is actuarially

sound and transparently documented. Education, clear documentation and careful choice of trigger structures are therefore essential for market acceptance.

Operational and legal considerations. Implementing NatPar programmes at scale requires robust operational processes: index calculation and verification, payout automation, integration with underwriting and claims systems, and coordination with reinsurance and capital market transactions. Legal frameworks for parametric products may be less well-developed than for traditional insurance in some jurisdictions, raising questions about enforceability, consumer protection and the boundary between insurance and derivatives.

9 Summary and lessons from the frost laboratory

The frost examples provided a concrete illustration of the NatPar thesis. Starting from a simple NatCat model of seasonal frost damage to citrus in two regions, we defined an indemnity cover and constructed two NatPar alternatives whose payouts depend only on temperature: a continuous design $P = qD(T)$ and a binary tail-trigger design $P = q\mathbf{1}_{\{T < \tau\}}$, calibrated by AAL-neutrality. Because the designs share the same hazard and vulnerability base, differences in EP curves, VaR, and loss-ratio behaviour are attributable to contract structure rather than to inconsistent modelling inputs.

Three lessons are generic and set the standard for NatPar evaluation.

Payout shape governs tail capital. “Parametric” is not automatically “capital-light”. Continuous bounded schedules tend to cap liabilities and dampen exposure-driven tails; digital schedules can concentrate losses into capital-relevant quantiles. Trigger probability and payout shape must be evaluated jointly against the chosen capital metric (e.g. $\text{VaR}_{0.995}$).

Basis risk is the operational interface between NatCat and NatPar. NatPar replaces the exposure–vulnerability component of indemnity loss by a low-dimensional hazard-only payout. The residual $B := Y - P$ is therefore not a nuisance term but a measurable interface. Exceedance-basis curves provide a tail view of shortfall and overpayment that is directly suitable for governance and model validation.

A disciplined reporting template supports both practice and supervision. By extending NatCat outputs (AAL, EP/AEP/OEP, return levels, tail metrics) to the parametric portfolio and pairing them with basis-risk exceedance diagnostics, NatPar programmes become comparable, auditable, and regulator-friendly. This discipline also discourages gratuitous complexity: additional triggers and indices should be justified by *measured* basis-risk reduction net of losses in transparency and operational robustness.

In summary, Natural Parametric Insurance is not a departure from NatCat practice but a reorganisation of it in response to climate-driven uninsurability, exposure uncertainty, and capital pressure. Making this reorganisation explicit provides a common language for theorists, practitioners, and supervisors to design, analyse, and govern parametric programmes in a way that is both mathematically coherent and operationally grounded in how catastrophe risk is actually managed.

Appendix

Proof of the equation (12).

Fix $x \geq 0$ and condition on D .

- *Case 1: no-loss-under-deductible ($AD \leq l$)*. Then $Y = 0$ and $P - Y = qD'$. Overpayment exceedance is

$$qD' > x,$$

which depends only on hazard.

- *Case 2: above deductible ($AD > l$)*. Then $Y = AD - l$ and

$$P - Y = qD' - (AD - l)^+ = l + qD' - AD.$$

The inequality $P - Y > x$ becomes

$$l + qD' - AD > x \iff A < \frac{l - x + qD'}{D},$$

together with the constraint $A > l/D$ (to ensure we are in Case 2).

Combining the two cases yields the conditional CDF form

$$\begin{aligned}
\text{BEP}^{c,-}(x) &= \mathbb{P}(P - Y > x) = \mathbb{E}[\mathbf{1}_{\{P-Y>x\}} \mathbf{1}_{\{D>0\}}] \\
&= \mathbb{E}[\mathbb{E}(\mathbf{1}_{\{P-Y>x\}}|D) \mathbf{1}_{\{D>0\}}] \\
&= \mathbb{E}[\mathbb{E}(\mathbf{1}_{\{P-Y>x\}} \mathbf{1}_{\{AD \leq l\}}|D) \mathbf{1}_{\{D>0\}}] \\
&\quad + \mathbb{E}[\mathbb{E}(\mathbf{1}_{\{P-Y>x\}} \mathbf{1}_{\{AD>l\}}|D) \mathbf{1}_{\{D>0\}}] \\
&= \mathbb{E}[\mathbb{E}(\mathbf{1}_{\{qD'>x\}} \mathbf{1}_{\{AD \leq l_r\}}|D) \mathbf{1}_{\{D>0\}}] \\
&\quad + \mathbb{E}\left[\mathbb{E}\left(\mathbf{1}_{\left\{A < \frac{l-x+qD'}{D}\right\}} \mathbf{1}_{\{AD>l\}}|D\right) \mathbf{1}_{\{D>0\}}\right] \\
&= \mathbb{E}[\mathbf{1}_{\{qD'>x\}} \mathbf{1}_{\{AD \leq l\}} \mathbf{1}_{\{D>0\}}] \\
&\quad + \mathbb{E}\left[\mathbf{1}_{\left\{A < \frac{l-x+qD'}{D}\right\}} \mathbf{1}_{\{AD>l\}} \mathbf{1}_{\{qD'>x\}} \mathbf{1}_{\{D>0\}}\right] \\
&= \mathbb{E}\left[\mathbf{1}_{\{q_r D'>x\}} \mathbf{1}_{\left\{A_r, t \leq \frac{l_r}{D}\right\}} \mathbf{1}_{\{D>0\}}\right] \\
&\quad + \mathbb{E}\left[\mathbf{1}_{\left\{\frac{l}{D} < A < \frac{l+qD'-x}{D}\right\}} \mathbf{1}_{\{qD'>x\}} \mathbf{1}_{\{D>0\}}\right] \\
&= \mathbb{E}\left[F_A\left(\frac{l}{D}\right) \mathbf{1}_{\{D>0, q_r D'>x\}}\right] \\
&\quad + \mathbb{E}\left[\left(F_A\left(\frac{l+qD'-x}{D}\right) - F_A\left(\frac{l}{D}\right)\right) \mathbf{1}_{\{D>0, qD'>x\}}\right] \\
&= \mathbb{E}\left[F_A\left(\frac{l+qD'-x}{D}\right) \mathbf{1}_{\{qD'>x\}}\right] \\
&= \mathbb{E}\left[F_A\left(\frac{l+q(D-d)^+ - x}{D}\right) \mathbf{1}_{\{(D-d)^+ > x/q\}}\right] \\
&= \mathbb{E}\left[F_A\left(\frac{l+qD - qd - x}{D}\right) \mathbf{1}_{\{D > x/q + d\}}\right] \\
&= \mathbb{E}\left[F_A\left(\frac{l - qd - x}{D} + q\right) \mathbf{1}_{\{D > x/q + d\}}\right].
\end{aligned}$$

Proof of equation (13).

$$\begin{aligned}
\text{BEP}^{\text{b},+}(x) &= \mathbb{P}(Y > x, D \leq d) + \mathbb{P}(Y > x + q, D > d) \\
&= \mathbb{P}(AD > x + l, D \leq d) + \mathbb{P}(AD > x + q + l, D > d) \\
&= \mathbb{E}\left(\bar{F}_A\left(\frac{x + l}{D}\right) \mathbf{1}_{\{D \leq d\}}\right) + \mathbb{E}\left(\bar{F}_A\left(\frac{x + q + l}{D}\right) \mathbf{1}_{\{D > d\}}\right)
\end{aligned}$$

Proof of equation (14).

$$\begin{aligned}
\text{BEP}^{\text{b},-}(x) &= \mathbb{P}(-B > x) = \mathbb{P}(q > x, D > d, Y = 0) \\
&\quad + \mathbb{P}(q - Y > x, D > d, Y > 0) \\
&= \mathbb{P}\left(q > x, D > d, A \leq \frac{l}{D}\right) \\
&\quad + \mathbb{P}\left(q > x, \frac{q + l - x}{D} > A > \frac{l}{D}, D > d\right) \\
&= \mathbb{P}\left(q > x, D > d, \frac{q_r + l_r - x}{D} > A\right) \\
&= \begin{cases} \mathbb{E}\left(F_A\left(\frac{q+l-x}{D}\right) \mathbf{1}_{\{D>d\}}\right), & x < q \\ 0, & x \geq q \end{cases}.
\end{aligned}$$

References

- [1] Arrow, K. J. (1974). Optimal insurance and generalized deductibles. *Scandinavian Actuarial Journal*, 1974(1), 1–42. 2
- [2] Raviv, A. (1979). The design of an optimal insurance policy. *American Economic Review*, 69(1), 84–96. 2
- [3] Zhang, J., Tan, K. S., & Weng, C. (2019). Index insurance design. *ASTIN Bulletin*, 49(2), 491–523. 2
- [4] Elabed, G., Bellemare, M. F., Carter, M. R., & Guirkinger, C. (2013). Managing basis risk with multiscale index insurance. *Agricultural Economics*, 44(4–5), 419–431. 2
- [5] Jensen, N. D., Barrett, C. B., & Mude, A. G. (2016). Index insurance quality and basis risk: Evidence from northern Kenya. *American Journal of Agricultural Economics*, 98(5), 1450–1469. 2

[6] Sonsino, D., & Mandelbaum, M. (2001). On preference for flexibility and complexity aversion: Experimental evidence. *Theory and Decision*, 51(2), 197–216. 2

[7] Sonsino, D., Benzion, U., & Mador, G. (2002). The complexity effects on choice with uncertainty—experimental evidence. *The Economic Journal*, 112(482), 936–965. 2

[8] Benso, M. R., Gesualdo, G. C., Silva, R. F., Silva, G. J., Castillo Rápalo, L. M., Navarro, F. A. R., Marques, P. A. A., Marengo, J. A., & Mendiondo, E. M. (2023). Review article: Design and evaluation of weather index insurance for multi-hazard resilience and food insecurity. *Natural Hazards and Earth System Sciences*, 23, 1335–1354. doi:10.5194/nhess-23-1335-2023. 2

[9] Gill, J. C., & Malamud, B. D. (2014). Reviewing and visualizing the interactions of natural hazards. *Reviews of Geophysics*, 52(4), 680–722. doi:10.1002/2013RG000445. 2

[10] Tilloy, A., Malamud, B. D., Winter, H., & Joly-Laugel, A. (2019). A review of quantification methodologies for multi-hazard interrelationships. *Earth-Science Reviews*, 196, 102881. doi:10.1016/j.earscirev.2019.102881. 2

[11] Zscheischler, J., van den Hurk, B., Ward, P. J., & Westra, S. (2020). Multivariate extremes and compound events. In J. Sillmann, S. Sippel, & S. Russo (Eds.), *Climate Extremes and Their Implications for Impact and Risk Assessment* (pp. 59–76). Elsevier. doi:10.1016/B978-0-12-814895-2.00004-5. 2

[12] Mitchell-Wallace, K., Jones, M., Hillier, J., & Foote, M. (2017). *Natural Catastrophe Risk Management and Modelling: A Practitioner's Guide*. John Wiley & Sons. 2

[13] Bourget, M., Boudreault, M., Carozza, D. A., Boudreault, J., & Raymond, S. (2024). A data science approach to climate change risk assessment applied to pluvial flood occurrences for the United States and Canada. *ASTIN Bulletin*, 54, 495–517. 2

[14] Chen, D., Rojas, M., Samset, B. H., Cobb, K., Diongue Niang, A., Edwards, P., Emori, S., et al. (2021). Framing, context, and methods. In *Climate Change 2021: The Physical Science Basis*. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press. doi:10.1017/9781009157896.003. 2

[15] Maraun, D., & Widmann, M. (2018). *Statistical Downscaling and Bias Correction for Climate Research*. Cambridge University Press. 2

[16] Martynov, A., Laprise, R., Sushama, L., Winger, K., Šeparović, L., & Dugas, B. (2013). Reanalysis-driven climate simulation over CORDEX North America domain using the Canadian Regional Climate Model, version 5: Model performance evaluation. *Climate Dynamics*, 41, 2973–3005. 2

- [17] Martel, J.-L., Mailhot, A., & Brissette, F. (2020). Global and regional projected changes in 100-yr subdaily, daily, and multiday precipitation extremes estimated from three large ensembles of climate simulations. *Journal of Climate*, 33(3), 1089–1103. 2
- [18] Financial Stability Board. (2017). *Final Report: Recommendations of the Task Force on Climate-related Financial Disclosures (TCFD)*. 2
- [19] Bank of England. (2019). *The 2021 biennial exploratory scenario on the financial risks from climate change*. 2
- [20] Office of the Superintendent of Financial Institutions (OSFI). (2023). *Guideline B-15*. Technical report, March. 2