SECTOR · REINSURANCE · CAT-MODEL SUBSTRATE

Your Cat Model is Right Inside Its Window. Bearing Reads the Three Layers Outside It

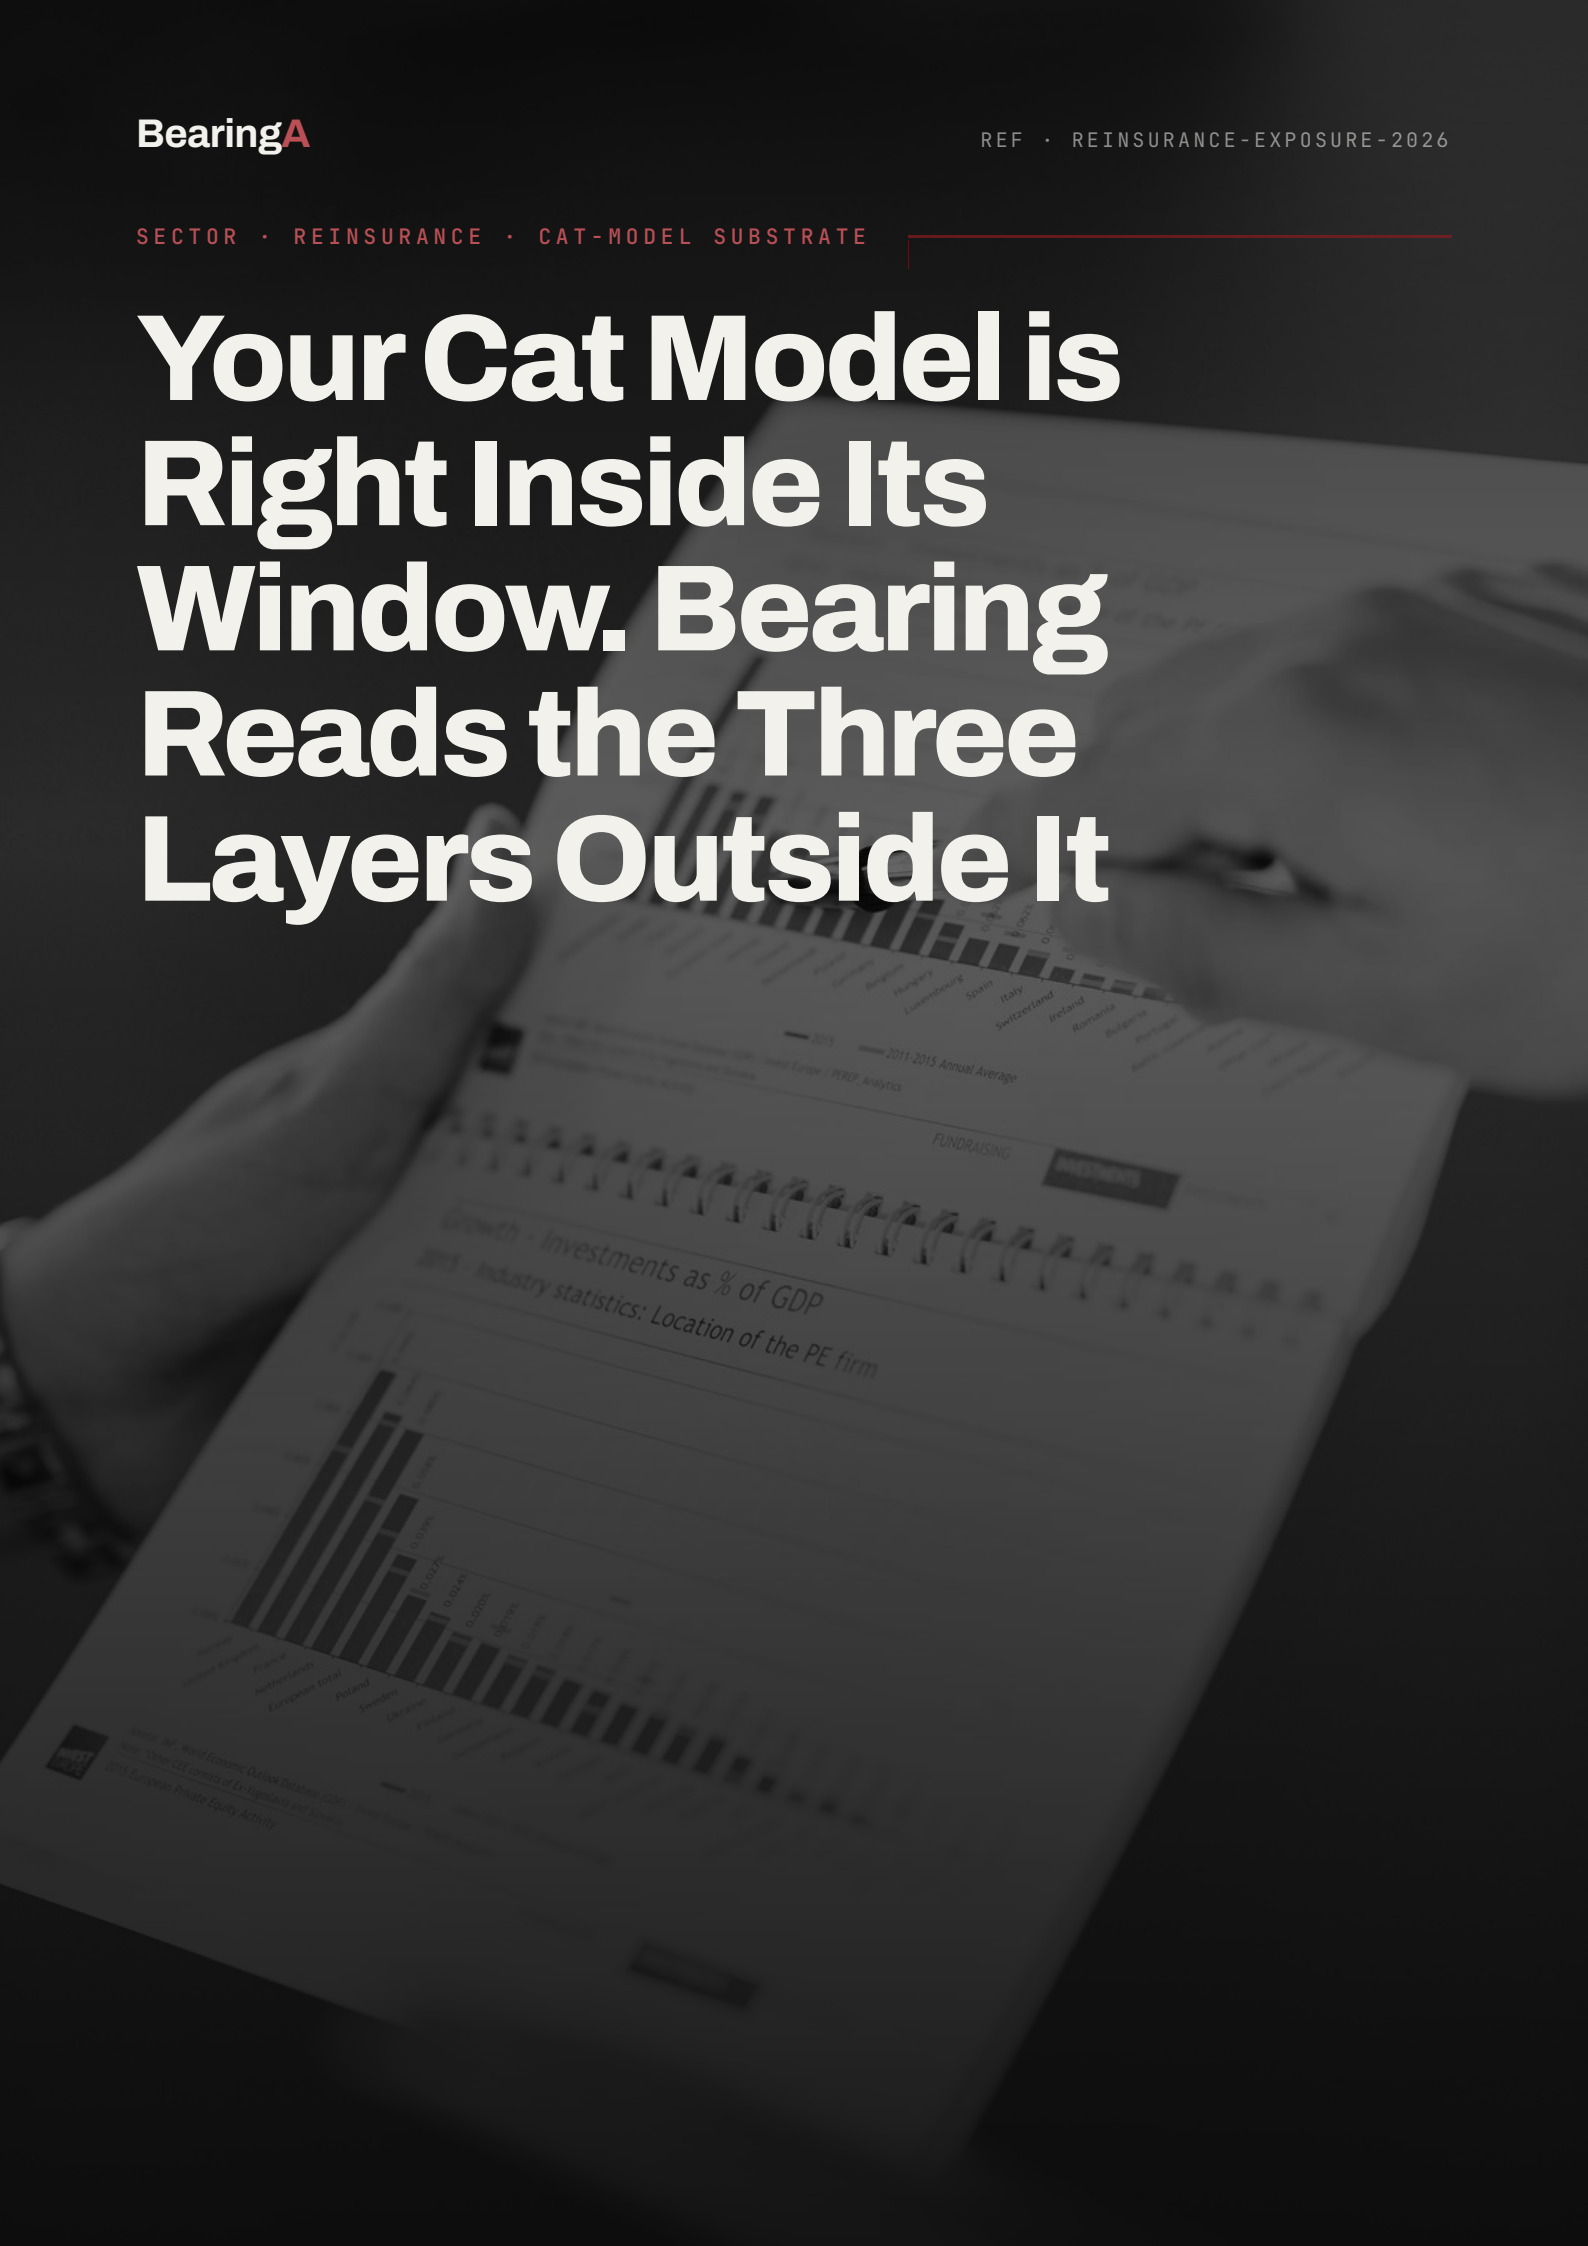

The exercise, the subject, the anchor, the method.

The 2026 SSM thematic prescribes both the shock and the depletion. Everything downstream is your composition.

i	CYCLE	€20–40B GWP
ii	RENEWAL	6 straight yrs $100bn+ insured loss
iii	WINDOW	Property cat −14.7% risk-adjusted
iv	TYPE	Catalog ~30–125 yrs vs record 150+

01

One soft market, three layers of concentration the cat model cannot include.

Your portfolio is a global multi-line reinsurer, property cat + casualty + specialty + L&H, ~€20–40B GWP, and the read addresses the Chief Underwriting Officer at portfolio-construction altitude . The methodology composes against the residual concentration the cat-model architecture structurally cannot include, not against the cat-model output itself. The cycle is the paradox: softening prices against escalating losses .

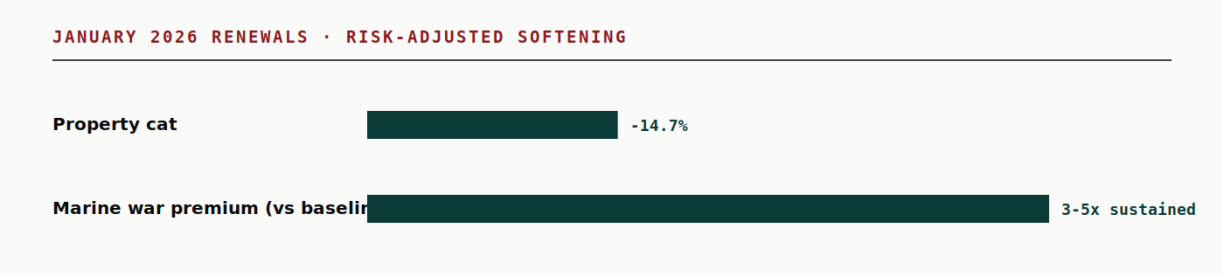

This read operates against the cat-model substrate your underwriting team already holds, and surfaces three layers of portfolio concentration that sit outside the model's calibration window. The methodology reads against the full instrumental record and resolved historical precedent the cat-model stochastic catalog structurally cannot include.

The cycle is the paradox: softening prices against escalating losses . Six reinsurers have filed four structurally distinct strategic responses against the same loss baseline and renewal-cycle softening.

The divergence is the methodology validation, published. The remainder of this read decomposes the three-layer concentration profile against the operator's portfolio reality, and surfaces where the cat-model frame structurally misses the concentration that is actually accumulating.

02

Your cat model is right inside its window. Three concentrations sit outside it.

The decision still open at your portfolio-construction window is not whether to follow the soft market down. It is which of the three layers your composition is reading, and whether the layer you are reading composes against the resolved precedents the long record carries.

Your January 2026 renewal book just priced at the largest year-on-year softening since 2014, property cat −14.7% risk-adjusted (Howden Re), −12% globally / −15% Europe (Guy Carpenter), retrocession −16.5%, against the sixth consecutive year of $100bn+ insured catastrophe losses.

The question is not whether your cat model is right within its window, it is which three layers of portfolio concentration sit outside the window your cat model reads.

Your cat-model stochastic catalog calibrates against ~30–125 years of instrumental record depending on peril (hurricane models to 1900; severe convective storm and wildfire catalogs ~30 years).

The methodology operates against the longer instrumental record, paleoclimate substrate, and historical sovereign-archive precedent, to surface analogue configurations whose structural shape matches current compound conditions and whose resolution is documented in the long record.

Three layers of concentration the cat model does not natively read. Your Layer 1 , direct cat losses, is the cat-model substrate; your team is already operational against it, and the methodology adds zero.

Your Layer 2 , cross-line correlation against the same climate driver, the cat model treats lines independently; the methodology reads the cross-line concentration (1906 SF earthquake → Knickerbocker Trust → 1907 Bankers' Panic is the canonical resolved precedent).

Six reinsurers, four strategic positions, one loss baseline. Munich Re cut April volume −18.5% on price discipline; Swiss Re held its combined ratio at 79.5%; SCOR grew P&C +5.4% while adding €300M to Best Estimate Liabilities; Hannover Re took growth below guidance; Berkshire Hathaway writes less as the cycle softens; Lloyd's Joint War Committee held specialty discipline independent of property softening.

Each position implicitly names which of the three layers that carrier reads as the dominant concentration.

03

The concentration map, and the honest zero at the bottom.

The map is the whole read in one frame. Each rung names where the concentration sits relative to the cat-model window, the resolved precedent it is anchored to, and the unmodelled magnitude band it carries, bounded against the long record, not asserted against the stochastic catalog.

The configuration regimes the catalog captures, hurricane to 1900, severe convective storm and wildfire ~30 years. The model simulates synthetic events within the window; your team operates against its EP curves. The methodology adds nothing here; discipline against the elevated baseline is the only move.

The same climate driver lands in property cat, D&O, marine, supply-chain BI, agriculture and political risk, each priced by its own unit, modelled by its own tools, rolled up at aggregate. Cross-line concentration against one configuration is structurally invisible to the per-line cat model.

Climate composing with geopolitical, sovereign-credit, supply-chain and energy configurations against the same underlying configuration. The model is single-peril, single-line, window-bounded by architecture; the compound sits outside the frame it is designed for, and is transmitting through your specialty book now.

Layer 1 is the honest zero: inside the cat-model window the methodology adds nothing. The concentration begins where the configuration composes across lines (Layer 2) and across climate-geopolitical compounds (Layer 3), the stochastic catalog reaches ~30–40 years for the non-peak perils; the instrumental record reaches 150, and the precedents that close these windows sit in the longer record.

04

What your cat model reads. The substrate your team already operates against.

Your cat-model substrate is the substrate every reinsurer holds. Munich Re NatCatSERVICE records $224bn overall / $108bn insured for 2025 , the sixth consecutive year above the $100bn insured threshold; Swiss Re's parallel estimate puts 2025 at $107bn insured. The 2025 cycle was front-loaded by the LA wildfires (Palisades + Eaton, Jan–Feb 2025), , the largest insured wildfire loss ever recorded. The cat model reads this layer.

The cat model reads this layer. Hurricane models calibrate against instrumental records back to 1900 (~125 years); severe convective storm and wildfire models calibrate against ~30 years.

The stochastic catalog simulates synthetic events within these windows; your underwriting team operates against the EP curves and expected-annual-loss outputs.

The methodology adds nothing at this layer, and acknowledges the addition as zero.

Where the methodology starts to operate. Munich Re's own 2025 commentary names what the cat model cannot fully read: "sheer luck spared the United States from hurricane landfalls in 2025" , yet non-peak perils (floods, severe convective storms, wildfires) accounted for $166bn overall / $98bn insured , far above the 30-year average ($90bn / $33bn).

(a) Hold cat-model-driven Layer 1 discipline against the structurally elevated baseline; the 14.7% global property cat softening at 1/1/2026 outpaces what the loss baseline justifies, and that gap is what the methodology composes against. (b) Mark the cat-model output as Layer 1 only, do not extend cat-model-derived probabilities to compound configurations the window does not include.

Confirming Layer 1 discipline is justified, 2026 H1 insured losses track at or above the $50–60bn pace; non-peak perils continue running ~50–80% above 30-year averages; at least one tier-1 reinsurer beyond Munich Re cuts capacity at mid-year.

Not landing, 2026 H1 insured losses come in below $40bn (a genuine soft loss year); cat-bond issuance continues at record levels without spreads widening; Munich Re's April cut is not followed at mid-year.

05

The cross-line correlation your cat model treats independently.

Your portfolio holds correlated exposure across lines against the same climate driver, and your cat-model architecture reads each line independently. The same compound climate driver lands in property cat (the model reads it), D&O (climate-disclosure litigation), marine (typhoon cargo, port disruption), supply-chain business interruption, agriculture (drought yield, crop reinsurance) and political risk (climate-induced sovereign credit pressure).

The 1906 San Francisco earthquake (18 April) caused approximately $400M in 1906-dollar losses , larger at modern exposure density.

Direct property losses were the visible Layer 1. The cascade ran through Layer 2: destruction of bank vaults, suspension of trade clearing through SF, transcontinental payment-system disruption, and the run on West Coast credit against an already-stressed US banking system.

The 2024–2025 cycle carries the same architecture in motion. The LA wildfires cascade through: direct property cat ($40bn insured); the (~$1bn, ~half landing on all California policyholders); seven of California's top twelve insurers limiting or withdrawing renewals since 2022, primary-insurer concentration cascading into reinsurance demand; the California Sustainable Insurance Strategy reform package restructuring the regulatory frame; potential D&O litigation on climate-disclosure positions; sovereign-credit pressure on California state finan

(a) Map your portfolio's cross-line concentration against named climate drivers, California wildfire: property cat + D&O + political risk + supply-chain BI rolling against one configuration; hurricane: property cat + marine + business interruption + agriculture in Gulf states.

Confirming the cross-line composition is operative, D&O claims activity from named primary/reinsurer cedents on climate-disclosure actions; state-level political-risk exposure (California, Florida, Louisiana) registering against the same drivers; aggregate combined ratio tracking worse than the per-line sum of expected.

Not landing, per-line books perform within expected ranges and the cascade does not materialize; FAIR Plan and residual-market mechanisms absorb the systemic pressure without rolling through your portfolio.

06

The compound configuration your cat model structurally cannot read.

Your portfolio sits inside compound configurations where climate composes with geopolitical, sovereign-credit, supply-chain and energy configurations against the same underlying configuration. The cat model is architected for single-peril, single-line, window-bounded simulation. Compound configurations, where the structural concentration emerges from the composition of climate + geopolitical + financial + supply-chain configurations operating against each other, sit outside the architectural frame the model is designed for.

The 1 November 1755 Lisbon earthquake (est. M 8.5–9.0), tsunami and fires destroyed Portugal's capital, killed 30,000–50,000, and caused damage estimated at 32–48% of Portuguese GDP .

The Layer 3 cascade ran: physical destruction → loss of the commercial fleet at port → disruption of Brazilian trade revenue → sovereign-credit pressure → forced restructuring under the Marquis of Pombal → political reorganization of the state.

The April 1815 eruption of Mount Tambora (VEI 7) cut global average temperature ~0.4–0.7°C through 1816. Europe experienced the "Year Without a Summer" , failed harvests across England, France, Germany and Switzerland; food riots; the 1816–19 typhus epidemic (~100,000 deaths); emigration cycles to North America.

The post-Napoleonic sovereign-credit configuration, already stressed from the 1815 reset, composed against the climate disruption to produce sustained credit pressure and multiple state-finance reorganizations through 1819.

The 1972 Sahel drought compounded with the 1972 Soviet harvest failure, the 1971 closure of the gold window and Bretton Woods collapse (1973), the October 1973 Yom Kippur War and the OPEC oil embargo.

The cascade ran through agricultural cost (grain roughly tripled 1972–74), energy cost (oil quadrupled), currency-regime restructuring, sovereign-credit stress across non-oil economies, and the "stagflation" structural shift through 1979.

The compound empirically active now is the Hormuz compound (PHM-CMP-0089). Following coordinated US and Israeli airstrikes on 28 February 2026, hull war premiums for Strait-of-Hormuz transits for 7 days of cover.

All 12 members of the International Group of P&I Clubs (~90% of oceangoing tonnage) cancelled certain war coverage with 72-hour notice; the expanded high-risk designation to the entire Persian Gulf.

Vessel transits cut ~95% from the 178/day pre-conflict baseline.

The compound is transmitting through your specialty book at quantifiable magnitude right now. Marine war pricing 4× baseline; cargo cover repriced or cancelled; P&I liability partially repriced; aviation war risk elevated for adjacent routes; energy reinsurance carrying Gulf-physical-asset exposure repriced; political-risk demand spiking against state-asset-confiscation and infrastructure-default scenarios.

Your cat model reads zero of this, none of these lines sit inside the property-cat calibration window.

The composition runs against substrates that do not revert. The post-COVID configuration carries a ~13% permanent upward price-level shift relative to pre-2019 trend, reserves calibrated against the pre-shock price level are short-covered against permanent adjustment, so the same Hormuz inflation cascade lands as structural rather than transitional.

(a) Map your portfolio's concentration against the active Hormuz compound: marine war, cargo, hull, P&I, aviation war, political risk, energy reinsurance, Gulf-state-asset exposure, the cross-specialty concentration is the load-bearing read; the per-line PML is the substrate.

07

Your peers are filing four different reads.

Reading Layer 1 as the dominant exposure. Cuts volume rather than chase price-softened business; protects per-policy economics; accepts absolute premium loss to defend portfolio quality. The strongest read on the elevated cat baseline; cycle discipline at tier-1 scale.

Reading Layer 1 as bounded within target. Accepts revenue compression to hold combined ratio; less explicit on cross-line concentration. Disciplined Layer 1 with implicit confidence the loss path stays within the model's calibration substrate.

Reading the soft market as an entry opportunity while building reserves against the elevated baseline. The €300m BEL buffer is the methodology-aligned read , taking premium but reserving against the baseline the cat model may under-read. Layer 1 discipline through reserves, not volume cut.

Reading the soft market as manageable through volume growth, missing the guidance trajectory. The gap between guidance and execution is the methodology-readable signal , the carrier read the softening as less severe than it landed; Layer 1 read more optimistic than the data substrate justified.

The cleanest cycle-discipline read in the set. No quarterly pressure, no growth target, no shareholder-rotation incentive. Reads the whole renewal cycle as a Layer 1 question and pulls back when pricing fails the discipline test. The read closest to the methodology's own posture on capacity-allocation timing.

The clearest Layer 3 read in the set. The Lloyd's structure separates specialty discipline from property-cat softening, and the Joint War Committee's expansion is the specialty market's read that the Hormuz compound carries unresolved concentration. The only major comparator explicitly reading the compound configuration as the dominant current exposure.

You are reading your peer reinsurers' Q1 2026 results and January renewal disclosures alongside your own portfolio. Six reinsurers have filed four structurally distinct strategic responses against the same loss baseline and renewal-cycle softening.

Each response implicitly names which of the three-layer concentration that carrier reads as the dominant exposure. The divergence is your methodology validation, published, and your calibration substrate.

(a) Track Munich Re's volume trajectory through mid-year renewals as your Layer 1 calibration, if the −18.5% April cut extends to further capacity reduction at mid-year, the elevated-baseline read is corroborated by the most disciplined tier-1 operator.

(b) Track SCOR's reserve trajectory, the €300m P&C BEL buffer is methodology-aligned; subsequent reserve movements name whether SCOR continues to read the Layer 1 elevation as warranting reserve discipline.

08

Your asymmetric exposure. The three-layer composite read.

Your three-layer concentration profile composes against a softening renewal market. The cat model reads Layer 1 within its window and your team operates against it, Layer 1 discipline defends +2–5pp combined ratio vs the soft-market follower. Layer 2 (cross-line correlation) sits outside the per-line architecture; the methodology reads it against the 1906 → 1907 precedent and the 2024–2025 California cascade.

(a) Hold cat-model-driven Layer 1 discipline against the elevated baseline; do not follow the renewal softening past the methodology-read floor, Munich Re's −18.5% April cut is the comparator anchor for disciplined Layer 1 at tier-1 scale. (b) Map cross-line concentration (Layer 2) against named climate drivers and stress-test the aggregate against the 1906 → 1907 cascade structure.

Confirming the composite move lands, your Layer 1 combined ratio outperforms the soft-market-follower peer group by 2–5pp through 2026; your Layer 2 reads correlate with named climate-driver events; your Layer 3 specialty book captures the compound pricing without absorbing the compound losses; Munich Re, Berkshire and Lloyd's specialty markets corroborate the direction.

09

What this read is not.

Not a cat-model replacement. Your cat-model output remains the canonical Layer 1 substrate. The methodology operates against the residual exposure the cat-model architecture structurally cannot include, not against the cat-model output itself.

Not a per-cedent underwriting recommendation. The three-layer composition is portfolio-construction substrate. Per-cedent pricing, capacity allocation and treaty-structure decisions compose against your underwriting team's substrate, not against this read.

Not a Solvency II SCR replacement. The composite bands are methodology-derived asymmetric-exposure surfacing, they do not substitute for the regulatory capital model. Your CRO's SCR and ORSA composition holds; the methodology composes an additional read against the residual.

Not a retrocession-structuring recommendation. Retro pricing and structure compose against your aggregate management and capital strategy. The Layer 3 compound read informs whether the retro composition reads against the compound shape, it does not specify the mechanism.

Not a substitute for the CUO's judgement. The methodology produces three-layer asymmetric-exposure surfacing substrate. Integration into portfolio construction, peril-mix, region-mix and cedent-concentration discipline composes against the CUO's judgement and the firm's capital-allocation strategy.

10

A Bearing Portfolio Exposure Read.

This artefact is a Bearing Reinsurance Portfolio Exposure Read , the deployed institutional product, composed at CUO portfolio-construction altitude for the underwriting leadership evaluating Bearing's analytical contribution alongside the cat-model substrate their team already holds. The operator's question is concentration-anchored: which three layers of portfolio concentration sit outside the window my cat model reads, and is the layer my composition is reading anchored to the precedents the long record carries.

This read composes against the corpus calibrated since March 2024, the catastrophe-cycle pack grounded in Munich Re NatCatSERVICE 2024–2025 factsheets, the January 2026 renewal disclosures across six named reinsurers, the resolved long-record precedents from the instrumental record and paleoclimate substrate (1755 Lisbon, 1815 Tambora, 1906–07, 1972–74), and the empirically active Hormuz compound.

Every load-bearing claim drills to source; every comparator action is publicly disclosed and citable.

The read does not replace the cat model. It surfaces the three layers of concentration the cat model's calibration architecture structurally cannot include, and grounds each against resolved precedents the long record carries.

Available for portfolio-construction discussion at the CUO's discretion · Bearing · BearingA · 23 May 2026

SOURCING

Every claim in this read traces to a primary source at the live web read. Sources are dated, versioned, and indexed continuously.

BOUNDARY

Anonymised archetype — no named institution's data enters this composition. The CET1 number is supervisor-prescribed, not a BearingA computation.

CONTACT

bearinga.com/reads/reinsurance-portfolio-exposure
bearinga.com
sheeringa@bearinga.com

Evidence · Charts and data exhibits

Primary sources cited at the live web read. Every claim in the body text traces to the exhibits below or to a source indexed in the BearingA corpus.

EXHIBIT 01

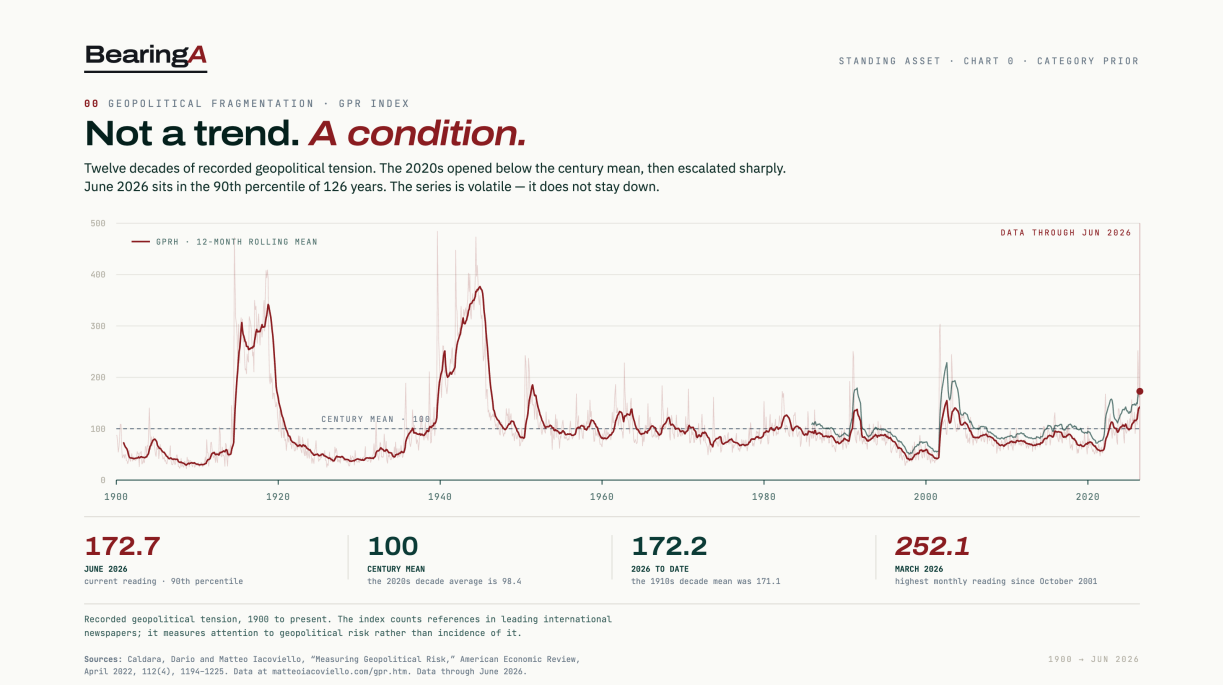

STANDING ASSET · CHART 0 — GEOPOLITICAL FRAGMENTATION, 1900–JUN 2026. 90TH PERCENTILE OF 126 YEARS.

EXHIBIT 02

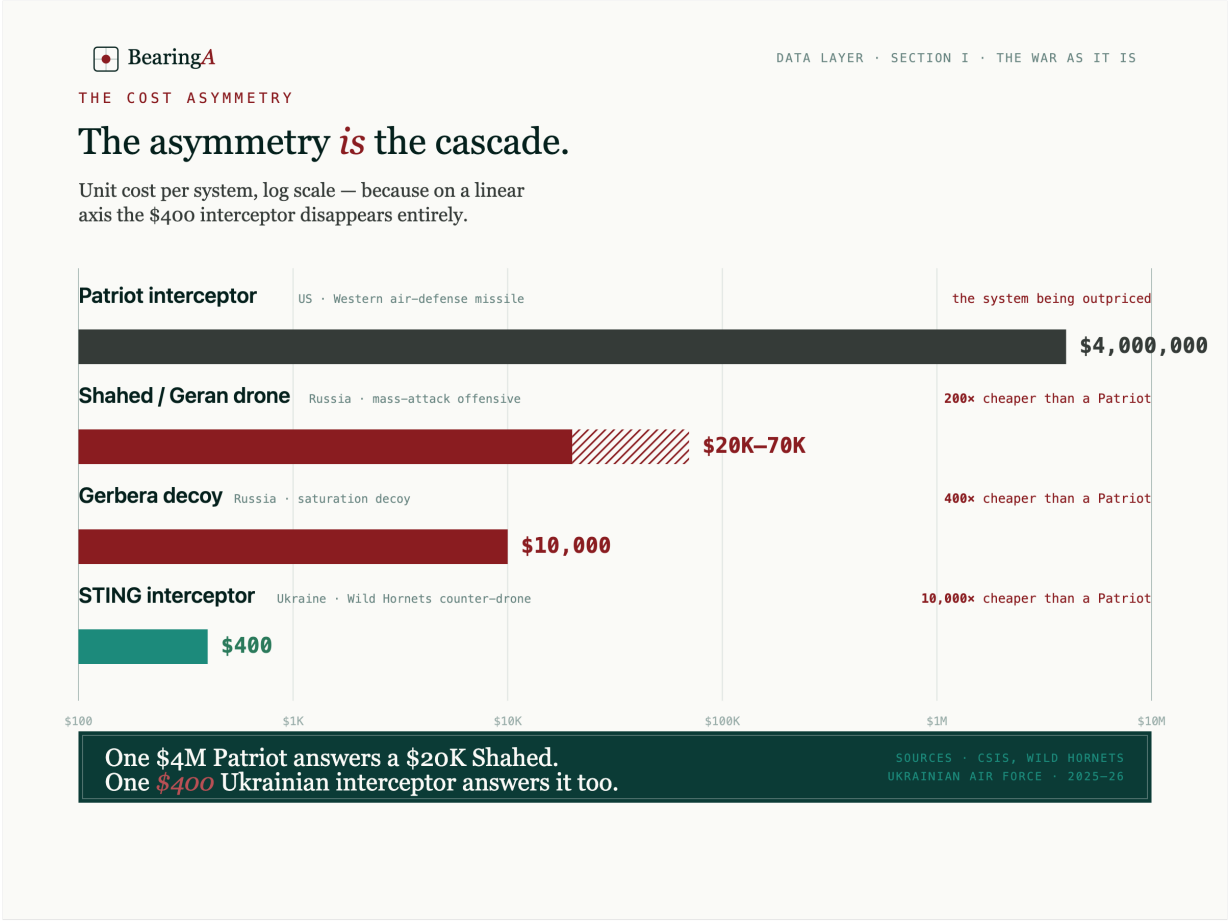

DATA LAYER · COST ASYMMETRY — PATRIOT $4M · SHAHED $20–70K · STING $400. THE ASYMMETRY IS THE CASCADE.

EXHIBIT 03

08 / CORRELATION MATRIX BREAKDOWN — COUNTDOWN TO REGIME FAILURE

Your VaR is about to lie to you. *In 19–23 days.*

Historical correlations are calibrated against periods between configurations — not against configuration activation. At every documented configuration boundary in the corpus, correlation regimes have shifted by 0.3–0.5 within 4–6 weeks. Your VaR's diversification benefit disappears precisely at the moment compound risk peaks. Bearing reads when each correlation pair will fail — named, dated, by mechanism.

CORRELATION PAIR	HISTORICAL (12M)	PROJECTED AT ACTIVATION	DAYS TO CROSSOVER	MECHANISM · PRECEDENT
Energy ↔ Credit (CVA)	0.21	0.61–0.74	19–23 days	*Refiners, airlines, utilities credit stress + energy price spike simultaneously.* **2008-Q3 documented: 0.18 → 0.71 over 5.5 weeks.** *2022 corpus: 0.22 → 0.64 over 4 weeks. Same configuration hits both books from the same direction.*
Gold ↔ DXY	-0.43	+0.05 to +0.25	Live now	*Standard inverse relationship overridden by safe-haven cascade. Gold rising despite DXY strength.* **2011 sovereign-stress corpus: gold-DXY moved positive for 9 weeks before reverting.** *Standard model is currently structurally wrong; configuration composition is the correct read.*
Energy ↔ Agri (protein)	0.12	0.42–0.55	28–42 days	*EU import substitution + input cost transmission.* **2022 cycle documented: 0.14 → 0.49 over 6 weeks.** *Energy desk and agri desk both running against the same underlying configuration. Cross-desk integrated read.*
DXY ↔ EM crosses (TWD/KRW/BRL/ZAR)	0.58	0.85–0.92	21–42 days	*Capital flight activation produces correlated EM cross moves.* **2014-15 and 2018 analogues: correlation peaked above 0.9 within 3 weeks of activation.** *EM diversification benefit disappears at exactly the activation moment.*
Copper ↔ LME base metals	0.67	0.45–0.55	42–70 days	**Correlation breakdown in the other direction.** *EU restructuring channel disproportionately benefits copper vs aluminium/zinc. Position-direction divergence opens within base metals complex.*
Cross-asset vol (energy ↔ gold)	0.34	Asymmetric divergence	14–28 days	*Energy vol expands first; gold vol follows but to a different ceiling.* **Cross-asset vol arbitrage opens.** *Long oil vol / short gold vol when Hormuz escalating but cross-compound sustaining is exactly the pair trade.*

CORRELATION MATRIX BREAKDOWN — YOUR VAR IS ABOUT TO LIE TO YOU. IN 19–23 DAYS.

ABOUT BEARINGA

BearingA

BearingA is a Ukrainian company. Headquartered in Kyiv. Operating across Kyiv and Amsterdam. Building and running the engine continuously since March 2024, through every geopolitical configuration the period has produced.

The compound geopolitical configurations the engine reads for institutional clients are, in Kyiv, operational conditions navigated daily while what reads them is being built. Not encountered first in the corpus. Navigated first. The reads compose with first-hand reference.

We are proud of Ukrainian resistance. Proud of the mental strength of the people building under these conditions. Proud of the depth of analytical talent operating here. This is where BearingA was built and this is where BearingA will remain headquartered.

SPEAK TO BEARINGA

BearingA does not run discovery calls, and does not gate the first conversation behind a form. Reach directly. You will hear back inside a working day.

Deploy against your position	sheeringa@bearinga.com
One specific read	sheeringa@bearinga.com
On behalf of a client or firm	enquiries@bearinga.com
Editorial · academic · media	enquiries@bearinga.com

THIS READ

WEB	bearinga.com/reads/reinsurance-portfolio-exposure
SOURCING	Every claim traces to a primary source at the live web read. Sources are dated, versioned, and indexed continuously from Kyiv since March 2024.
BOUNDARY	Anonymised archetype. No named institution's data enters this composition.