CROSS-DOMAIN RISK INTELLIGENCE

Why Cascading *Risk Models Fail*

Current cascading risk models assume linear propagation. This paper sets out why they systematically underestimate tail losses, and proposes the Abgalis Engine cross-domain transmission alternative.

SEVEN DOMAINS. ONE COUNSEL.

EXECUTIVE SUMMARY

Cascading risk has entered the supervisory vocabulary faster than the modelling discipline has caught up with it. The dominant analytical structures still in use — linear domino chains, simple Markov-style state transitions, correlation-matrix aggregation, and bolt-on contagion overlays — share three structural defects that bite hardest precisely where insurers most need the model to be reliable: in the tail, under stress, when channels become non-stationary. This paper sets out the three defects, explains why patches do not fix them, and describes the alternative cross-domain transmission approach that Abgalis has built into the Abgalis Engine.

§ 1

The cascading-risk vocabulary has run ahead of the maths

Read any 2025 supervisory speech on systemic or compound risk and you will find the word *cascade*. The Bank of England has used it of climate physical risk transmitting into financial stability. EIOPA has used it of climate stress transmitting into insurer solvency. The IAIS has used it of operational concentration transmitting into systemic risk in the insurance sector. The IFoA's Planetary Solvency work uses it as the load-bearing concept of the entire programme.

What is striking, when you read the underlying technical material that supports the speeches, is how thin the modelling infrastructure beneath the word actually is. In most cases the cascade is asserted narratively and aggregated quantitatively through a correlation matrix that was calibrated on data that did not contain a cascade. The result is a profession that talks about cascading risk in the executive summary and prices it as if cascades did not happen.

§ 2

Defect one — linear propagation

The simplest cascading-risk models propagate a shock as a sequence of weighted edges in a directed graph: shock arrives at node A, transmits to B with weight w_{AB}, then to C with weight w_{BC}, and so on. The total damage at C is approximately $w_{AB} \cdot w_{BC}$ times the original shock magnitude.

The structure is appealing, intuitive, and traceable. It is also wrong in three specific ways for cross-domain transmission in insurance:

- **Convexity in the receiving node.** The relationship between input shock magnitude and damage at the receiving node is rarely linear. A 2x shock typically generates more than 2x damage at the receiver, because the receiver's defences are themselves dependent on the absence of stress. Linear propagation systematically under-states tail loss for this reason alone.

- **Path dependence.** A shock arriving at C via two pathways — one direct from A, one indirect through B — does not produce additive damage. The two arrivals interact at C in ways that depend on timing, on whether the receiver has already been weakened, and on whether mitigation deployed against the first arrival is still available for the second.
- **Feedback.** Real cascades are rarely directed acyclic graphs. C's damage often increases B's susceptibility, which in turn increases the second wave from A. Linear propagation models cannot represent feedback at all.

> **THE DIAGNOSTIC**
>
> If a firm's cascading-risk model can be solved analytically with a single matrix multiplication, it is too linear. Real cascades require iterative, time-stepped, state-aware computation — and the firm's stress-testing infrastructure should be sized accordingly.

§ 3

Defect two — stationary calibration

The second structural defect is the assumption that the channel weights themselves are stable across regimes. They are not.

Calibrate any channel coefficient — whether by historical loss correlation, expert elicitation, or hybrid — and you have a number that describes how the channel behaved in the calibration window. Apply that number in stress, and you are implicitly asserting that the channel has the same transmission properties when the system is under load as when it is quiescent. Both empirical observation and elementary reasoning say otherwise.

Three sources of non-stationarity are common in practice:

1. **Mitigation depletion.** When the system is quiescent, mitigation is fresh — reinsurance is in place, capital buffers are full, third-party providers have spare capacity. Each of these dampens the channel coefficient. Under stress, mitigation depletes — first-event losses consume reinsurance, capital draws, third parties become bottlenecks. The same channel that transmits 30% of a shock in calibration may transmit 70% of the third shock in a sequence.
2. **Behavioural shift.** Counterparties, regulators, and competitors do not behave identically under stress. A regulator who defers to internal models in benign quarters may impose capital add-ons in stressed quarters. A reinsurer who pays claims promptly when its own losses are limited may dispute coverage when its own losses are large. The channel itself is reshaped.
3. **Technological obsolescence.** The historical data on which channels are calibrated reflect a different technological substrate. Cyber concentration, cloud-services consolidation, AI-driven third-party services, parametric trigger architectures — all of these are recent enough that the historical calibration window contains little or no signal about their channel behaviour under stress.

§ 4

Defect three — the missing seventh domain

The third defect is structural rather than mathematical. Most cascading-risk frameworks in production today were designed in the financial-stability tradition: nodes are firms or sovereigns, edges are claims or exposures, and the cascade transmits through the balance sheet. They handle insurance loss reasonably well and credit reasonably well. They handle operational risk poorly. They handle climate risk via bolt-on overlay. And they handle the regulatory channel — the channel by which a single shock generates supervisory action that constrains future capacity — barely at all.

The regulatory channel is the most under-modelled of the seven, and arguably the most consequential. The capital cost of a supervisory letter that constrains future underwriting capacity by 10% for two years is, in present value terms, often larger than the immediate insurance loss that triggered it. Most cascading-risk models do not have a node that represents this channel. Abgalis treats it as a first-class domain (Domain 7) for that reason.

> *A cascading-risk model that has insurance, credit, operational and climate & ESG but not strategic & emerging is missing the channel that determines whether the firm can write business next year.*

§ 5

What patching does not fix

The standard responses to these defects, as observed across the market, are patches: add a non-linear stress overlay, calibrate two correlation matrices (benign and stressed), append a regulatory-add-on scenario. Each patch addresses a symptom; none addresses the structure.

PATCH	DEFECT ADDRESSED	WHAT IT LEAVES UNSOLVED
Stressed-state correlation matrix	Stationary calibration (partial)	Does not capture path dependence, feedback, or mitigation depletion. Two static matrices is still two static states.
Non-linear loss function on receiver	Linear propagation (partial)	Convexity at the receiver does not fix convexity in the channel itself.
Bolt-on regulatory scenario	Missing regulatory domain	Treats regulatory action as exogenous rather than as a transmitted consequence of the originating shock.
Ad-hoc reverse stress	Tail under-statement	Generates plausible binding shocks but does not connect them to the channel structure that would deliver them.

The patches are not wrong; they are insufficient. They produce a model that is right at the average and lightly improved in the tail. The structural defects survive.

§ 6

The Abgalis Engine alternative

The Abgalis Engine treats transmission channels as first-class entities. Each channel has a direction, a latency, a state-dependent characteristic, and an explicit feedback profile. Channels are calibrated regime-aware, not regime-agnostic. The seven domains include the regulatory channel as a peer of insurance, credit, and operational rather than as a bolt-on. And the architecture is iterative rather than analytic — cascades are simulated in time, with mitigation depletion and feedback represented natively.

What this delivers, in operational terms, is not a different SCR number. It is a different account of *how* a given SCR number arose, and which mitigations actually reduce the path that produced it. That is what enables defensible reverse stress, defensible disclosure narrative, and defensible engagement with supervisors who have begun to ask "trace the shock through your business" as a standing question.

§ 7

Where to go from here

For firms reviewing their cascading-risk infrastructure, Abgalis suggests three diagnostic questions, in order of severity:

1. Does our cascading model represent the channel as a state-dependent object, or as a static coefficient calibrated on a benign window?
2. Does it include the regulatory channel as a domain, or treat regulatory action as exogenous?
3. Can it be re-run iteratively under mitigation depletion, or does it produce a one-shot answer?

A firm that answers *state-dependent / domain / iterative* is on solid ground. A firm that answers *static / exogenous / one-shot* is exposed precisely in the place where the supervisor will ask first. Abgalis engagements typically begin with this diagnostic and scale from there.

COLOPHON · SOURCES · TRADEMARKS · NOTICES

ABOUT ABGALIS

Abgalis is a climate and cross-domain risk intelligence platform for insurers and reinsurers. The name draws from the Sumerian *Abgal* — the seven divine sages sent by Enki to counsel the earliest kings of civilisation — re-instantiated as seven cooperating analytical domains. The civilisation that invented the insurance contract, the actuarial table, and the systematic record of risk.

TRADEMARKS & METHODOLOGY IP

References in this paper to Abgalis frameworks — the **Abgalis Engine™** (Abgalis's cross-domain transmission framework), **ICRIP™** (Integrated Climate Risk Intelligence Platform), and **The Cascade™** (Abgalis's monthly intelligence brief) — are trade marks of Abgalis Ltd, used and explained inline. The internal mathematical methodology is reserved to Abgalis Ltd and is not disclosed in this document. Reproduction outside Abgalis requires written permission.

VERIFICATION STANDARD

This paper is multi-source verified, hallucination-scanned, and editorial-panel reviewed at the four-gate NEXUS standard (citation, cross-check, sanity, currency). Where the text refers to specific regulatory instruments — Solvency II, SS3/19, SS1/23, SS1/21, SS2/21, CP10/25, DORA, ISSB / IFRS S2, IFoA Planetary Solvency — the regulator's or institute's primary text is authoritative; this paper does not paraphrase regulatory rules.

NO ADVICE; NO ALLEGATION

This paper is general thought leadership and does not constitute legal, regulatory, actuarial, investment, or compliance advice. Nothing in the paper alleges deficiency at any named firm. Where firms are referenced, the references are to publicly-available primary disclosures.

REGULATORY FRAMING & CONFIDENTIALITY

FPO Art. 19. Editorial commentary for insurance and reinsurance professionals and other investment professionals within the meaning of FPO Art. 19. Not a financial promotion; not investment advice; not a regulated actuarial opinion. Abgalis Ltd is not authorised by the FCA. Published commercially by Abgalis Ltd for distribution to identified industry recipients; onward redistribution permitted with attribution. No part may be re-purposed as the basis for a competing analytical framework without the express written permission of Abgalis Ltd.

Published by Abgalis Ltd, a company registered in England and Wales · Registered office: 128 City Road, London EC1V 2NX

"A cascading-risk model that has insurance, credit, operational and climate & ESG but not strategic & emerging is missing the channel that determines whether the firm can write business next year."

FROM §4 — THE MISSING SEVENTH DOMAIN

THE SERIES

This is one of a standing series of **Abgalis Position Briefs** — cross-domain risk-intelligence papers published by Abgalis Research. New papers are released as supervisory and market signals warrant.

SUBSCRIBE TO THE CASCADE

abgalis.com/cascade — Abgalis's monthly cross-domain climate-risk intelligence brief.

CONTACT & REPRINT

Abgalis Research · Abgalis Ltd · info@abgalis.com · linkedin.com/company/abgalis

Seven Domains. One Counsel.

ABGALIS

Published by Abgalis Ltd, a company registered in England and Wales
Registered office: 128 City Road, London EC1V 2NX