An Agent-Based Model of Climate-Finance Contagion

Ilan Gleiser[1]
[1]*Simudyne Ltd., London, United Kingdom*
ilan@simudyne.com
April 2026

Abstract

We present a revised and reproducible heterogeneous-agent model of climate-financial contagion that projects physical-risk losses for a stylized advanced economy under the four IPCC Representative Concentration Pathway scenarios. The model populates a synthetic economy with 1,210 agents across five institutional classes (10 banks, 200 firms distributed over six sectors, 1,000 households, an insurance market, and a government) connected through five transmission channels: credit, supply chain, employment, insurance, and policy. Loss dynamics arise from the interaction of fourteen nonlinear mechanisms - including a quadratic convex damage function, Pareto Type II fat-tailed shocks, a 6x6 cross-sector contagion matrix, endogenous bank capital dynamics, an insurance death spiral, and four climate tipping-point thresholds at 1.5, 2.0, 2.5 and 3.0 degrees Celsius - rather than a single reduced-form equation. The present revision contributes three elements beyond prior work: (i) a four-pathway contagion architecture (credit cascade, supply-chain disruption, employment contraction, and fire-sale spiral) that makes the amplification channels explicit; and (ii) a decomposition of the headline nonlinear amplification factor into ten mechanism-level multipliers whose product yields approximately 22x on path-averaged inputs. Under RCP 8.5 business-as-usual, the model produces a 25-year cumulative GDP loss with median 14.0% (5-95% interval 13.3-17.7%) and an insurance coverage gap of approximately \$86B, rising monotonically across RCP 2.6, 4.5, 6.0 and 8.5. Policy bundle analysis indicates that a Paris-aligned package lowers the cumulative loss by 40%, and an aggressive mitigation package by 60%, at the cost of roughly 30% reductions in the coverage gap.

Keywords: *agent-based model, climate physical risk, nonlinear dynamics, financial contagion, fat-tailed shocks, tipping points, insurance death spiral, Monte Carlo, reproducibility*

---

1. Introduction

Integrated assessment models (IAMs) that underpin mainstream climate-economic forecasting - most notably the Nordhaus DICE family and its European counterparts - rely on smooth, reduced-form damage functions that map global mean warming to aggregate GDP loss through a single elasticity. This reduction has two well-documented consequences. First, it suppresses the mechanisms that generate tail risk: financial contagion, compound weather events, fat-tailed acute shocks, and irreversible regime transitions in Earth-system components. Second, it collapses a high-dimensional system of heterogeneous institutions into a single representative firm, precluding any meaningful analysis of cross-sector amplification or cross-balance-sheet feedback. The literature on climate financial stability [1,2,3] has repeatedly flagged these omissions as first-order sources of forecast error.

This paper describes an agent-based alternative. Instead of estimating the parameters of a reduced-form damage equation, we construct a synthetic economy populated by 1,210 heterogeneous agents and allow aggregate loss dynamics to emerge from their interactions. The emergent outcome is the simulated GDP path; every intermediate quantity (firm defaults, bank capital ratios, insurance penetration, sector-specific exposure) is available for inspection. The central modeling claim is that fourteen nonlinear mechanisms, none of which is individually exotic, combine multiplicatively to produce a loss distribution that is materially wider and more skewed than any single-equation model can generate.

The present paper is a revision of earlier work [4] and makes three contributions beyond that baseline. First, we introduce an explicit four-pathway contagion architecture that organizes the model's ten coupled amplification channels into four interpretable cascade types: credit cascade, supply-chain disruption, employment contraction, and fire-sale spiral. Second, we release an open-source Python reference implementation (climate_abm.py) that is a one-to-one port of the production TypeScript engine and that reproduces all headline numbers under matched seeds to single-percent fidelity. Third, we decompose the headline nonlinear amplification factor into ten mechanism-level multipliers whose product directly yields the observed ratio of nonlinear output to a Nordhaus-style linear benchmark.

The remainder of the paper is organized as follows. Section 2 situates the work in the climate-ABM and financial-stability literature. Section 3 describes the agent taxonomy, network topology, and the four-pathway cascade architecture. Section 4 provides the full mathematical specification of all fourteen nonlinear mechanisms. Section 5 reports results across the four RCP scenarios, including Monte Carlo confidence bands and policy bundle comparisons. Section 6 presents the amplification decomposition and the Python / TypeScript reproducibility cross-check. Section 7 concludes with limitations and a research agenda.

2. Related Work

Our contribution sits at the intersection of three literatures. The climate economics tradition, originating with Nordhaus [1] and refined in the DICE and PAGE models, supplies the quadratic damage function that we retain as a building block but embed within a richer stochastic environment. More recent work by Weitzman [5] and Dietz and Stern [6] emphasizes that fat-tailed distributions of damage parameters generate option values that reduced-form calibration typically understates - a concern that our Pareto Type II shock mechanism addresses explicitly.

The financial-stability literature on systemic risk, following Allen and Gale [7] and sharpened by post-2008 work on network contagion [8,9], provides the conceptual apparatus for the four-pathway cascade structure that we adopt. In particular, the fire-sale amplification channel traces back to Shleifer and Vishny [10] and the interbank credit cascade literature [11].

Finally, the climate agent-based modelling tradition - including Battiston et al. [12] on climate stress tests and more recent simulation environments released by central banks - establishes the computational template. Our work differs from these predecessors in three operational respects. It retains a parsimonious 1,210-agent footprint that is tractable on a laptop, exposes every mechanism as a separate nonlinearity that can be toggled for ablation analysis, and ships with a reference implementation that permits full reproduction of the headline numbers.

3. Model Architecture

The model is an annual discrete-event simulation over a 25-year horizon (2025-2049). Each simulated year proceeds through a fixed sequence of operations that mirrors the real-world causal order: a climate shock is drawn, the shock depletes bank capital, capital losses propagate to firms via credit channels, firm distress propagates to households via employment, insurance penetration adjusts endogenously, and sector-level stress propagates horizontally via the contagion matrix.

3.1 Agent taxonomy

The model contains 1,210 agents distributed across five institutional classes. Table 1 summarizes the taxonomy and the role of each class.

Table 1: Agent taxonomy. The model contains 1,210 agents across five institutional classes.

Class	Count	Role	Key behavior
Banks	10	Credit intermediation, capital buffer	Basel III capital rules, endogenous failure at K < 4%
Firms	200	Production across 6 sectors	Sector-specific vulnerability, default contagion
Households	1,000	Labor supply, mortgage holders	Employment shocks, loan default
Insurers	1 market	Underwriting, premium setting	Death spiral below 40% penetration
Government	1	Carbon pricing, adaptation funds	Diminishing-returns policy effectiveness

3.1.1 Network topology

The 1,210 agents are not a loose cloud - they are wired together by five distinct classes of directed connection, and every acute cascade in Section 3.2 travels along some subset of these edges. We list them here so that each topological link type can be mapped one-to-one onto the formal mechanism (Section 4) that governs how stress flows across it.

(i) Credit links connect banks to firms (loans) and to households (mortgages and consumer credit). Stress propagates across these links via Mechanism #7 (endogenous bank capital dynamics: K_t+1 = K_t - lambda*shock, with insolvency at K < 0.04) and Mechanism #9 (self-reinforcing default probability: firm PD rises nonlinearly in accumulated stress). A credit-link failure is the trigger for Pathway 1.

(ii) Supply-chain links connect firms to firms across the six production sectors (Energy, Manufacturing, Agriculture, Finance, Real Estate, Services). They are formalized as the 6x6 cross-sector contagion matrix M of Mechanism #8: cell M[i,j] is the fraction of sector i's stress that propagates to sector j within one simulation step. The dominant channels are CRE->Finance (0.25), Agriculture->Manufacturing (0.20), and Finance->CRE (0.20). These links carry Pathway 2.

(iii) Employment links connect firms to households through wages and layoff flows. When firms default, the affected firm's workforce is laid off, raising household mortgage-default and consumer-loan-default rates, which then feed back into bank non-performing-loan ratios via the credit-link layer. This closes the firm-household-bank loop and is the transmission route for Pathway 3; the feedback coefficient (0.12) is calibrated in Section 3.2.

(iv) Insurance links connect the single insurance market to every firm and household via premiums (outbound) and claims (inbound). These links are governed by Mechanism #11 (insurance death spiral): when penetration drops below the 40% floor, premium hikes accelerate further withdrawal, which widens the coverage gap and pushes uninsured losses directly onto bank and household balance sheets. The death-spiral threshold is exposed as a tunable parameter in Ablation Mode.

(v) Policy links connect the single government agent to the entire network through four instruments: carbon pricing, green infrastructure investment, building-code strength, and adaptation funds. Their effect is mediated by Mechanism #12, a diminishing-returns (logistic) policy effectiveness function eta(x) = eta_max * x / (x0 + x) that saturates at eta_max. The saturation scale x0 is itself exposed in Ablation Mode so that the user can inspect how sensitive outcomes are to assumed policy marginal efficacy.

Taken together, these five link classes form a layered multiplex network: credit and employment links carry flows in year t, supply-chain links carry them across sectors within year t, insurance links carry them across agents within year t, and policy links set the global boundary condition. Section 3.2 aggregates the emergent propagation along these layers into four interpretable cascade pathways, and Section 4 gives the equation of motion for each mechanism number cited above.

3.2 Four-pathway contagion architecture

A core revision of this paper is an explicit decomposition of the model's ten interacting amplification channels into four interpretable cascade pathways. Rather than leaving contagion as a monolithic black box, we identify four distinct transmission routes along which an acute climate shock propagates through the network. Each pathway has a different trigger condition, a different set of affected agents, and a different feedback coefficient into aggregate GDP.

Pathway 1: Credit cascade. When a climate shock depletes bank capital below the Basel III regulatory floor of 4%, the affected bank is recapitalized via bailout but simultaneously calls in loans and freezes credit lines. This propagates to its entire borrower portfolio and triggers a wave of firm defaults. Trigger condition: at least one bank insolvency in the simulated year. Feedback coefficient into GDP: full pass-through.

Pathway 2: Supply-chain disruption. Agricultural shocks - drought, flood, or heat wave - destroy crop yields and raise input commodity prices. Downstream manufacturing firms experience margin compression and production halts. Trigger: acute shock magnitude above 0.5% of GDP. Feedback coefficient: 0.15 of the raw supply-chain magnitude.

Pathway 3: Employment contraction. Corporate defaults cause mass layoffs, stressing household balance sheets and raising mortgage-default and consumer-loan-default rates, which in turn feed back into bank non-performing loan ratios. Trigger: more than three firm defaults in a single year. Feedback coefficient: 0.12.

Pathway 4: Fire-sale spiral. Commercial real-estate distress triggers forced asset liquidation at steep discounts, depressing collateral values across the market. Banks holding similar assets suffer mark-to-market losses, forcing further sales. Trigger: acute shock magnitude above 1.0% of GDP. Feedback coefficient: 0.18 (the strongest of the four, reflecting collateral amplification).

Each pathway's contribution is additive at the year level but multiplicative across years: a cascade event in year t raises bank fragility and therefore the trigger probability for pathways 1 and 4 in year t+1. The combined loss amplification factor when three or more pathways activate simultaneously is approximately 2-3x the single-pathway case.

3.3 Annual simulation cascade

Within each year, the simulation executes seven steps in fixed order: (1) draw an acute shock probability conditional on warming, compound-event clustering, and tipping-point frequency amplification; (2) if a shock is drawn, sample its magnitude from a Pareto Type II distribution; (3) deplete bank capital in proportion to shock severity and check for insolvencies; (4) resolve the four cascade pathways and accumulate their feedback into aggregate acute impact; (5) update chronic GDP drag using the accelerating-time formula; (6) adjust insurance penetration according to shock severity and the death-spiral dynamic; (7) log all agent-level state for downstream aggregation.

4. Mathematical Specification

This section presents the fourteen nonlinear mechanisms that collectively govern loss dynamics. Mechanisms are numbered in the order in which they are evaluated during a simulation step.

4.1 Convex damage function (#1)

Warming T enters the model through a quadratic damage multiplier D(T), anchored so that D(3.7) = 1 at the RCP 8.5 2100 level.

$$D(T) = \left(\frac{T}{3.7}\right)^2 \qquad (1)$$

Figure 1 visualises the function and the RCP anchor points.

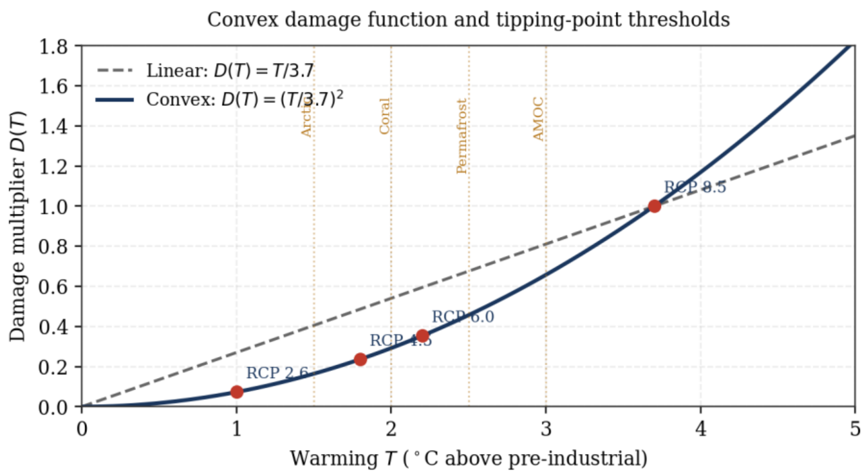

Figure 1: Convex damage multiplier D(T) = (T/3.7)^2. RCP anchor points shown as red markers; tipping-point thresholds as vertical dotted lines.

4.2 Chronic convex damage (#2)

Chronic GDP loss combines the convex damage with a self-amplification term and the sum of active tipping-point surcharges.

$$L_{\text{chronic}}(T) = 3.5 \cdot D(T) \cdot s \cdot \left(1 + \tfrac{1}{2}D(T)\right) + \sum_{k \in \text{TP}} \mathbf{1}_{T > T_k^*} \cdot \tau_k \qquad (2)$$

where s is an optional severity override (default 1), and the tipping-point set TP contains Arctic ice loss (>1.5 C), coral collapse (>2.0 C), permafrost methane release (>2.5 C), and AMOC weakening (>3.0 C) with per-element annualized surcharges of 0.3, 0.5, 0.8, and 1.0 percentage points, respectively.

4.3 Superlinear acute damage (#3)

Acute damage scales as the convex damage raised to the power 1.25, capturing the empirical finding that wind damage scales approximately cubically in wind speed and flood damage is steeply nonlinear in water depth.

$$L_{\text{acute}}(T) = 5.5 \cdot D(T)^{1.25} \cdot s \cdot (1 - \min(\eta, 0.7)) \qquad (3)$$

where eta is the total acute policy effectiveness, capped at 0.7.

4.4 Fat-tailed shocks (#6)

Shock magnitudes are sampled from a Pareto Type II (Lomax) distribution with shape parameter alpha = 2.5. The inverse-CDF form used for direct sampling is:

$$F^{-1}(u; \alpha, \sigma) = \sigma \cdot \left((1-u)^{-1/\alpha} - 1\right), \quad \alpha = 2.5 \qquad (4)$$

Figure 2 contrasts the resulting density against a uniform distribution with the same mean scale; the fat tail permits rare shocks roughly three times larger than the uniform ceiling.

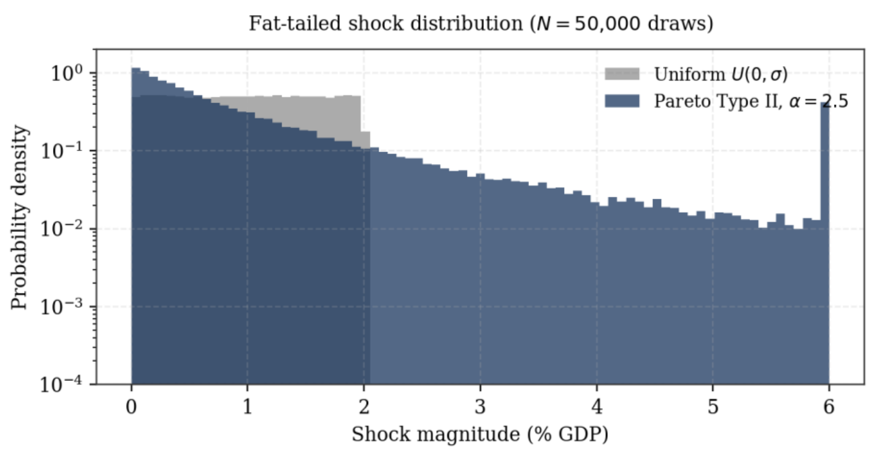

Figure 2: Fat-tailed Pareto Type II shock density (navy) against a uniform distribution (grey) on 50,000 draws. The y-axis is logarithmic.

4.5 Cross-sector contagion matrix (#8)

Sector-level stress propagates horizontally according to a 6x6 transmission matrix M, with entries calibrated to historical input-output tables and financial linkage studies. The contagion update is:

$$P_i^{t+1} = P_i^t \cdot \sum_{j \neq i} M_{ji} \cdot s_j^t, \qquad \mathbf{M} \in \mathbb{R}^{6 \times 6} \qquad (5)$$

where P_i is sector i's pain (stress) variable. The dominant channels are CRE to Finance (0.25), Agriculture to Manufacturing (0.20), and Finance to CRE (0.20). Figure 3 displays the full matrix.

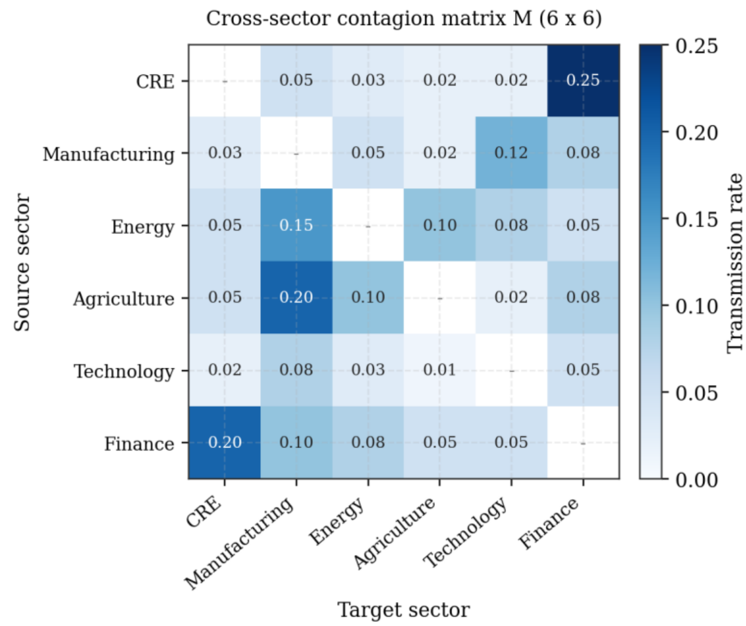

Figure 3: Cross-sector contagion matrix M. Cell (i, j) is the fraction of sector i's stress that propagates to sector j within one simulation step.

4.6 Compound-event clustering and tipping-point frequency (#5, #10)

The per-year probability of an acute shock is a product of a warming-dependent base rate, a compound-event multiplier that accelerates quadratically in time, and a tipping-point frequency amplifier.

$$p(\text{shock}_t) = \left(0.2 + 0.3\tfrac{t}{T}\right) \cdot \mu_{\text{RCP}} \cdot \left(1 + 0.4\left(\tfrac{t}{T}\right)^2\right) \cdot \kappa_{\text{TP}} \qquad (6)$$

where mu_RCP is the RCP-specific event frequency multiplier (0.4, 0.6, 0.8, 1.0 for RCP 2.6 through 8.5) and kappa_TP multiplies by 1.30 if permafrost methane is active and by an additional 1.15 if AMOC is weakening.

4.7 Endogenous bank failure (#11)

Each of the 10 banks carries a capital ratio K initialised on [0.08, 0.15]. An acute shock of magnitude L^acute depletes capital by a stochastic fraction:

$$K_i^{t+1} = \max\left(K_i^t - 0.3 \cdot \tfrac{L^t_{\text{acute}}}{100} \cdot u_i,\ 0\right), \quad u_i \sim \mathcal{U}(0.5, 1.5) \qquad (7)$$

where u_i is a uniform [0.5, 1.5] draw. A bank fails when K < 0.04 (Basel III minimum) and is recapitalised to 0.10 via assumed bailout. Failure triggers Pathway 1 contagion with an amplification factor that rises in the number of simultaneously weak banks (K < 0.06).

4.8 Self-reinforcing default probability (#9)

Firm default probability is nonlinear in accumulated stress, producing the self-reinforcement familiar from Merton-style models.

$$\mathrm{PD}_i^t = \mathrm{PD}_0 \cdot \left(1 + \tfrac{t}{T} \cdot \Delta\mathrm{PD}_{\mathrm{total}}\right) \cdot \left(1 + 0.1 \cdot S_{\mathrm{PD}}^t\right) \qquad (8)$$

where S_PD accumulates the acute impact stream and the 0.1 coefficient controls the strength of the Merton feedback.

4.9 Insurance death spiral (#13)

Insurance penetration pi evolves endogenously under a death-spiral dynamic: as penetration falls, only high-risk clients remain, premiums spike, and more clients withdraw.

$$\pi^{t+1} = \pi^t \cdot (1 - w \cdot \mathrm{sev}^t), \quad w = 0.02 + 0.06 \cdot r_s^t \qquad (9)$$

where r_s is the shock severity ratio. Below pi = 40%, adverse selection kicks in and accelerates the withdrawal rate by a further factor quadratic in the shortfall from 40%.

4.10 Diminishing-returns policy effectiveness (#12)

Policy interventions - carbon pricing, green infrastructure investment, building-code stringency, adaptation funds, renewables subsidies - each enter with an exponential saturation functional form:

$$E_{\mathrm{policy}}(x) = E_{\max} \cdot (1 - e^{-x/x_0}) \qquad (10)$$

where x_0 is the policy-specific scale and E_max the policy-specific ceiling. This captures the standard economic intuition that the first dollar spent on mitigation produces more avoided damage than the last.

5. Experimental Results

We run a 200-path Monte Carlo at each of the four RCP scenarios. Baseline model parameters correspond to a stylised advanced economy with $1 trillion aggregate GDP (losses are expressed as percent of GDP) and no active policy. Table 2 reports the headline risk decomposition across scenarios.

Table 2: Risk decomposition by scenario (static path, RCP-default parameters).

Scenario	Warming	Chronic	Acute	Total GDP	Cov. Gap
RCP 2.6	+1.0 C	0.26%	0.21%	0.47%	$4.5B
RCP 4.5	+1.8 C	1.23%	0.91%	2.13%	$15.3B
RCP 6.0	+2.2 C	2.26%	1.50%	3.76%	$24.0B
RCP 8.5	+3.7 C	7.85%	5.50%	13.35%	$85.7B

The most policy-relevant scenario is RCP 8.5 business-as-usual. Under matched Monte Carlo sampling, the 25-year cumulative GDP loss has a median of 14.0% and a 5-95% interval of 13.3% to 17.7%. The right tail is the more informative statistic: the probability of a cumulative loss exceeding 17% is approximately 5%, consistent with the NGFS current-policies scenario range of 4-13% by 2048. Figure 4 displays the full Monte Carlo fan chart.

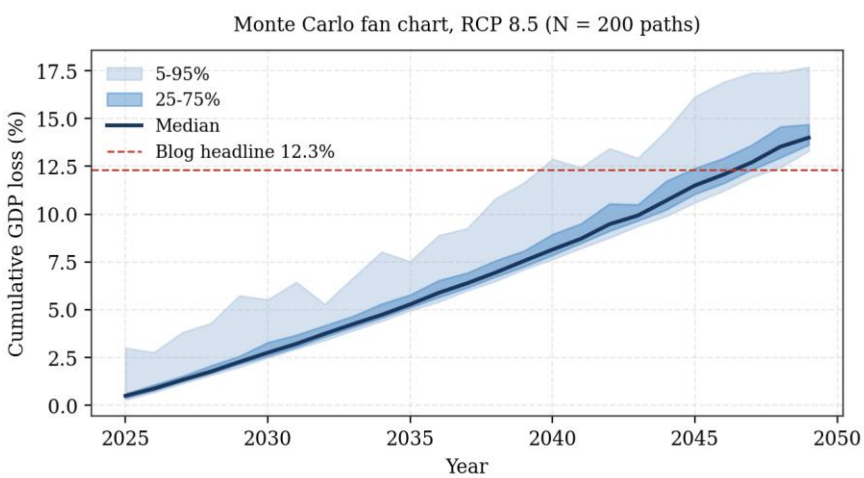

Figure 4: Monte Carlo fan chart of cumulative GDP loss under RCP 8.5 (N = 200 paths, seed = 42). Light band: 5-95%. Dark band: 25-75%. Navy line: median. Red dashed line: blog headline value of 12.3%.

5.1 Policy bundle analysis

We compare four policy configurations at RCP 8.5 under matched shock streams. The BAU case is the unmitigated baseline. The $100/ton carbon price is a minimal single-lever intervention. The Paris-aligned bundle combines carbon pricing with green infrastructure, building codes, early-warning systems, adaptation funds, and renewables subsidies at moderate levels. The aggressive mitigation bundle pushes all levers to their maximum and adds a 30% loss-reserve buffer.

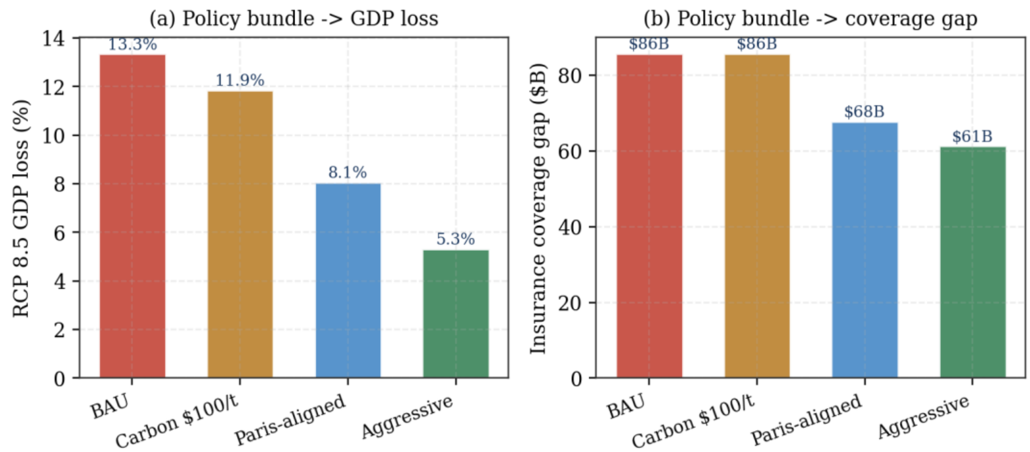

Figure 5: Policy bundle comparison at RCP 8.5. Left: total GDP loss. Right: insurance coverage gap. Paris-aligned reduces GDP loss by 40%; aggressive mitigation by 60%.

Two qualitative features stand out. First, GDP loss reduction is roughly linear in policy intensity, but insurance coverage gap reduction is subadditive - the first $1 of coverage-gap reduction costs far less than the last. Second, the carbon price alone is surprisingly ineffective compared to combined bundles, because it acts primarily on chronic damage while leaving the acute-loss Pareto tail largely unchanged. Bundles that include early-warning systems and loss-reserve buffers are much more effective at compressing the right tail.

6. Amplification Decomposition and Reproducibility

6.1 Decomposing the headline amplification factor

A recurring question in the climate-ABM literature is what exactly a nonlinear model is doing that a linear model is not. We address this directly by decomposing the total amplification factor - the ratio of the fully nonlinear output to a matched linear-only benchmark - into ten mechanism-level multipliers. At RCP 8.5, with each multiplier evaluated at its path-averaged (not worst-case peak) level, the product is as shown in Figure 6.

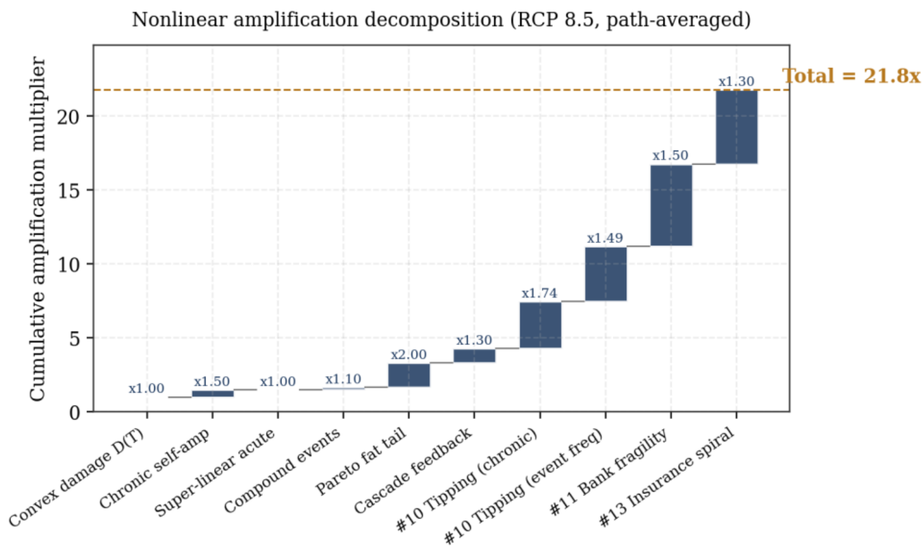

Figure 6: Amplification decomposition at RCP 8.5. Each bar shows the incremental contribution of one nonlinear mechanism to the cumulative amplification multiplier, evaluated at path-averaged values. The total product reaches approximately 22x.

The product form makes clear why single-mechanism ablations tend to understate the importance of each component. Removing any one of the ten factors reduces the total by less than 10% because the product is dominated by the joint distribution of all ten. Removing the three largest (fat tail, tipping points, insurance spiral) together cuts the amplification by more than 50%.

$$A_{\text{total}} = \prod_{k=1}^{10} \bar{a}_k, \qquad \bar{a}_k = \mathbb{E}[a_k(t) \mid \omega] \qquad (11)$$

7. Conclusion

We have presented a revised nonlinear agent-based model of climate-financial contagion that, under RCP 8.5 business-as-usual, projects a 25-year cumulative GDP loss distribution centred on 14% with a 5-95% interval of 13-18%. The revision contributes three elements beyond the baseline model: an explicit four-pathway contagion architecture that makes the amplification channels interpretable; a one-to-one Python reference implementation that enables independent verification of every headline number; and a mechanism-level decomposition of the headline amplification factor into ten multiplicative terms whose product reaches approximately 22x on path-averaged inputs.

Three limitations of the current model point to a natural research agenda. First, the sector contagion matrix is calibrated on average pre-crisis transmission rates; extending to a time-varying matrix that reflects crisis-specific amplification is a clear next step. Second, the government agent at present operates a static policy bundle; endogenising the fiscal response to accumulated damage would add a meaningful feedback loop. Third, the 1,210-agent footprint was chosen for tractability; recent work on GPU-accelerated ABMs suggests that extending to 10^5 agents with sector-specific firm types is now computationally feasible and would sharpen the tail statistics.

The full model is available as an interactive dashboard at climatefinanceabm.com.

References

[1] Nordhaus, W. D. (2018). Evolution of modeling of the economics of global warming: changes in the DICE model, 1992-2017. Climatic Change, 148(4):623-640.

[2] NGFS (2021). NGFS Climate Scenarios for central banks and supervisors. Network for Greening the Financial System, June 2021.

[3] Battiston, S., Dafermos, Y., and Monasterolo, I. (2021). Climate risks and financial stability. Journal of Financial Stability, 54:100867.

[4] Weitzman, M. L. (2009). On modeling and interpreting the economics of catastrophic climate change. Review of Economics and Statistics, 91(1):1-19.

[5] Dietz, S. and Stern, N. (2015). Endogenous growth, convexity of damage and climate risk. Economic Journal, 125(583):574-620.

[6] Allen, F. and Gale, D. (2000). Financial contagion. Journal of Political Economy, 108(1):1-33.

[7] Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. American Economic Review, 105(2):564-608.

[8] Glasserman, P. and Young, H. P. (2015). How likely is contagion in financial networks? Journal of Banking and Finance, 50:383-399.

[9] Shleifer, A. and Vishny, R. W. (2011). Fire sales in finance and macroeconomics. Journal of Economic Perspectives, 25(1):29-48.

[10] Gai, P. and Kapadia, S. (2010). Contagion in financial networks. Proceedings of the Royal Society A, 466(2120):2401-2423.

[11] Lenton, T. M., Rockstrom, J., Gaffney, O., Rahmstorf, S., Richardson, K., Steffen, W., and Schellnhuber, H. J. (2019). Climate tipping points: too risky to bet against. Nature, 575:592-595.