



# Memory-driven cascading tipping dynamics in the Earth system. A regime-switching Volterra framework calibrated with CMIP6 ensembles

Mauricio Herrera-Marín<sup>1</sup>, Alex Godoy-Faúndez<sup>1</sup>, and Diego Rivera<sup>1</sup>

<sup>1</sup>Research Center on Sustainability and Strategic Resource Management (CISGER). Faculty of Engineering, Universidad del Desarrollo. Avda. Plaza 700. Santiago de Chile. Chile

**Correspondence:** Mauricio Herrera-Marín (mherrera@udd.cl)

**Abstract.** We show that finite-memory effects fundamentally reshape cascade risk in coupled climate tipping systems by decoupling ensemble stability from pathwise instability. Applying a regime-switching Volterra model with tempered fractional kernels to CMIP6 multi-model ensembles ( $n = 10$  for the Atlantic Meridional Overturning Circulation, AMOC;  $n = 37$  for the Amazon and Greenland), we demonstrate that three tipping elements operate under structurally distinct memory regimes linked to different physical processes. AMOC lower-tail occupancy triples under SSP5-8.5 while ensemble-mean weakening reaches only  $\approx 0.5\sigma$ ; a per-model-consistent memory amplification index  $\hat{M} \approx 2.7\text{--}6.0$  confirms that persistence, not mean shift, is the primary driver. The Amazon presents a mechanistically contrasting picture ( $\hat{M} < 1$ ): its tail amplification is forcing-dominated, making ensemble-mean drying projections adequate for risk assessment. Greenland internal surface-mass-balance (SMB) variability is strongly long-range dependent ( $\bar{H} = 0.89$ ; 89% of models), anchoring it as a persistent upstream driver. Cascade simulations show that quenched (99th-percentile) pathwise Amazon damage exceeds annealed (median) projections by a factor of  $> 2$  under weak forcing – a divergence invisible to ensemble summaries and absent in memory-free dynamics. These results demonstrate that neglecting long-range dependence systematically understates upper-tail cascade risk, and that AMOC, the Amazon, and Greenland require mechanistically differentiated treatment in climate-risk assessment.

## 1 Introduction

The notion that the Earth system may undergo abrupt, irreversible reconfigurations under anthropogenic forcing has evolved from a qualitative caution into a quantitatively central challenge in climate-risk science. The tipping-elements framework (Lenton et al., 2008) established that certain subsystems exhibit threshold-like behaviour in which state-dependent feedbacks drive rapid transitions on timescales commensurate with human planning horizons. This framing has been extended toward a networked Earth-system view in which tipping is an emergent property of coupled ocean–cryosphere–biosphere–carbon feedbacks (Steffen et al., 2018; Wunderling et al., 2021), and sharpened by syntheses showing that multiple tipping points may be reachable near  $1.5\text{--}2^\circ\text{C}$  of global warming (Armstrong McKay et al., 2022).

A particularly consequential cascade narrative links Arctic and Atlantic dynamics to Amazon hydroclimate. Accelerated Greenland meltwater discharge weakens the AMOC – with established historical evidence of decline since the mid-twentieth



25 century (Smeed et al., 2014, 2018) – which in turn reshapes tropical rainfall patterns and intensifies Amazon hydroclimatic stress; concurrent carbon-cycle feedbacks can further destabilise the system and commit it to a hothouse trajectory (Steffen et al., 2018; Ripple et al., 2026). Despite growing recognition of this mechanistic chain, a critical gap persists in translating smooth scenario-defined forcing pathways into probabilistic cascade risk.

30 Most cascade formulations employ low-dimensional models in which uncertainty enters through short-memory perturbations and inference targets mean trajectories or median crossing times (Wunderling et al., 2021). These approaches deliver qualitative insight but fail to represent two features decisive for upper-tail risk: *persistence* (multi-year to multi-decadal memory from cryospheric inertia, thermohaline adjustment, and land–atmosphere coupling) and *intermittency* (episodic clustering of extreme states when forcing interacts with internal memory dynamics). Prior work has quantified long-range dependence in individual climate time series using ARFIMA models (Beran, 1994) and Hurst-exponent diagnostics (Franzke et al., 2012; Vyushin and Kushner, 2012), but has not coupled memory structure to cascade risk in a physically interpretable, multi-element network framework. The Volterra operator provides the natural generalisation: it encodes power-law persistence through a completely monotone kernel while remaining analytically tractable and directly constrainable from CMIP6 diagnostics – extending ARFIMA-style memory quantification into a nonlinear, regime-switching cascade setting.

35 Persistence and intermittency together give rise to the *annealed–quenched dichotomy*: ensemble-mean summaries (annealed risk) may appear moderate even when pathwise trajectories (quenched risk) display heavy-tailed, burst-driven destabilisation far exceeding marginal, memory-free predictions. This dichotomy is not merely theoretical – it is empirically detectable in CMIP6 ensembles, and it is the central demonstration of this paper.

*Empirically*, we show that three tipping elements operate under structurally distinct memory regimes diagnosable from CMIP6 ensembles: AMOC is governed by thermohaline memory with multi-decadal persistence; the Amazon responds on intermediate land–atmosphere timescales; Greenland acts as a persistent upstream driver. *Methodologically*, we introduce a regime-switching tempered-fractional Volterra cascade model encoding these distinctions through system-specific kernel parameters  $(\alpha, \theta)$  – where  $\alpha$  governs power-law persistence and  $\theta$  sets the finite memory horizon – together with a per-model-consistent memory amplification index  $\hat{M}$  separating persistence-driven from forcing-driven tail amplification. *Analytically*, Theorem 1 proves that positive Volterra memory generates super-marginal rolling-tail amplification, with  $\hat{M} > 1$  as a necessary condition for quenched excess.

50 The paper is organised as follows. Sect. 2 presents the model and calibration framework. Sect. 3 describes the CMIP6 ensembles and the Greenland emulator. Sect. 4 presents the empirical results structured around four quantitative claims. Sect. 5 interprets results physically, addresses limitations, and discusses policy implications. Sect. 6 concludes.



## 2 Model and methods

### 2.1 State variables and physical observables

55 We represent each tipping element by a bounded latent damage variable  $D_i(t) \in [0, 1]$ ,  $i \in \{G, A, R\}$  (Greenland, AMOC, Amazon), with  $D_i = 0$  at the pre-industrial reference and  $D_i \rightarrow 1$  at committed degradation. Physical observables are:

$$S_G(t) = S_{G,0} + 7.2 \text{ m} \cdot D_G(t) \quad (\text{Greenland SLE}), \quad (1)$$

$$M(t) = M_0[1 - D_A(t)] \quad (\text{AMOC overturning, Sv}), \quad (2)$$

$$X_R(t) = X_{R,0}[1 - c_R D_R(t)] \quad (\text{Amazon precipitation proxy}), \quad (3)$$

60 where  $M_0 \approx 17 \text{ Sv}$  (Smeed et al., 2014) and  $S_{G,\max} = 7.2 \text{ m}$  (Robinson et al., 2012). For cross-scenario tail comparison, we define the monotone tail transform at horizon  $T$ :

$$Y := -\log(1 - D_R(T)), \quad (4)$$

mapping  $D_R \uparrow 1$  to  $Y \rightarrow \infty$  and enabling exceedance comparisons at a common threshold  $y^*$  across all SSPs.

### 2.2 External forcing

65 Radiative forcing follows Myhre et al. (1998):

$$F_{\text{ext}}(t) = 5.35 \log\left(\frac{C(t)}{280 \text{ ppm}}\right) \text{ W m}^{-2}, \quad (5)$$

with  $\text{CO}_2$  trajectories from CMIP6 ScenarioMIP (O'Neill et al., 2016; Eyring et al., 2016).

### 2.3 Regime-switching tempered-fractional Volterra dynamics

#### 2.3.1 Regime process

70 Let  $z(t) \in \{0, 1\}$  be a continuous-time Markov chain with 0 = stable, 1 = unstable, and transition hazards

$$q_{0 \rightarrow 1}(t) = q_{SU}^0 \exp(\gamma_{SU} F_{\text{ext}}(t)), \quad q_{1 \rightarrow 0} = q_{US} > 0, \quad (6)$$

so that unstable-regime entry grows monotonically with forcing while recovery is forcing-independent.

#### 2.3.2 Memory kernel and physical interpretation

For regime  $z$ :

$$75 \quad g(t; z) = \frac{t^{-\alpha(z)} e^{-\theta(z)t}}{\Gamma(1 - \alpha(z))}, \quad t > 0, \quad (7)$$

with  $\alpha_U \geq \alpha_S$  and  $\theta_U \leq \theta_S$ , so the unstable regime carries longer effective memory. The exponent  $\alpha(z) \in (0, 1)$  governs the *rate of memory decay*: small  $\alpha$  means slow algebraic decay and long memory, as in deep-ocean thermohaline dynamics;  $\alpha \rightarrow 1$



approaches short memory. The tempering rate  $\theta(z) > 0$  sets a *finite memory horizon*  $\tau_{\text{mem}} \sim \theta^{-1}$ , beyond which past states cease to influence present dynamics. Physically, for AMOC, ocean thermal inertia and thermohaline adjustment operate on multi-decadal scales, yielding small  $\alpha_S \approx 0.30$  and  $\theta_S^{-1} \sim$  decades; the Greenland DFA confirms this as  $\bar{H} = 0.89$  implies  $\bar{\alpha} = 0.36$ . For the Amazon, land–atmosphere coupling and ENSO modulation operate on interannual to decadal scales, yielding intermediate  $\alpha$  and  $\theta^{-1} \sim$  years. For Greenland, SMB anomalies integrate sea-ice, albedo, and circulation feedbacks on multi-year timescales, producing structural slow memory ( $\bar{H} = 0.89$ , 89% of models exhibit LRD; Table 1). In the unstable regime,  $\theta_U < \theta_S$  reflects suppressed effective diffusivity during anomalous circulation states, extending memory and amplifying cascade response. The kernel is completely monotone and belongs to the Bernstein-function class (Schilling et al., 2012; Meira et al., 2019), ensuring well-posedness under resolvent-family theory (Prüss, 1993). Discretisation follows convolution quadrature (Lubich, 1988a, b) (Appendix B).

### 2.3.3 Volterra integrals and cascade dynamics

Downstream elements respond to history-weighted upstream damage:

$$I_G(t) = \int_0^t g(t-s; z(s)) D_G(s) ds, \quad (8)$$

$$I_A(t) = \int_0^t g(t-s; z(s)) D_A(s) ds. \quad (9)$$

Damage states evolve as:

$$\dot{D}_G = -\beta_G D_G + \eta_G F_{\text{ext}}[1 - D_G], \quad (10)$$

$$\dot{D}_A = -\beta_A D_A + k_{AG}(z) \varphi(I_G)[1 - D_A], \quad (11)$$

$$\dot{D}_R = -\beta_R D_R + k_{AR}(z) \varphi(I_A)[1 - D_R], \quad (12)$$

with regime-dependent couplings  $k_{ij}^U \geq k_{ij}^S > 0$  and superlinear amplification  $\varphi(x) = x^p$ ,  $p > 1$ . The  $[0, 1]$  structure of each  $D_i$  is preserved throughout simulation.

## 2.4 Annealed–quenched decomposition and memory amplification index

Let  $u = F_{\text{hist}}^{-1}(q)$  be a historically fixed threshold. The rolling-window tail frequency  $\hat{p}(t)$  over window  $W = 30$  yr changes between [1971–2000] and [2071–2100] as

$$\Delta p_{\text{roll}} = \underbrace{\Delta p_{\text{marg}}}_{\text{annealed}} + \underbrace{\Delta p_{\text{mem}}}_{\text{quenched}}, \quad (13)$$

where  $\Delta p_{\text{marg}} \approx \mathbb{E}_m[\Phi((u_m - \mu_{\text{fut},m})/\sigma_{\text{hist},m}) - q]$  is a conservative per-model Gaussian approximation using each model's own threshold and future mean. The *memory amplification index* is defined as

$$\hat{M} := \frac{\Delta p_{\text{roll}}}{\Delta p_{\text{marg}}}, \quad (14)$$



105 with  $\hat{M} > 1$  indicating quenched excess (persistence-dominated) and  $\hat{M} < 1$  indicating mean-shift dominance (forcing-dominated).

**Remark 1 (Non-orthogonality).** The decomposition (13) is not strictly orthogonal when persistence modulates both auto-correlation and level. The Gaussian approximation for  $\Delta p_{\text{marg}}$  is conservative; consequently  $\hat{M} > 1$  constitutes a lower bound on the true persistence contribution, not an exact partition.

**Remark 2 (Sensitivity).** Varying  $W \in \{20, 30, 40\}$  yr and the reference periods yields qualitatively identical conclusions:  
 110 AMOC persistently in the quenched-excess regime ( $\hat{M} \gg 1$ ); Amazon persistently annealed-dominated ( $\hat{M} < 1$ ).

**Theorem 1** (Memory-driven super-marginal amplification). *Let  $\{X_t\}$  satisfy a Volterra evolution with a non-negative completely monotone kernel such that  $I_t = \mathbf{1}_{\{X_t \leq u\}}$  has positive autocorrelation  $\rho_I(k) > 0$  for all  $k \geq 1$ . If exogenous forcing induces a monotone increase in  $p(t) = \mathbb{P}(X_t \leq u)$ , then  $\limsup_{t \rightarrow \infty} \hat{p}(t) > \limsup_{t \rightarrow \infty} p(t)$ , yielding  $\hat{M} > 1$ .*

*Proof sketch.* See Appendix A. Positive autocorrelation reduces the effective sample size  $W_{\text{eff}} < W$  and the Volterra operator  
 115 increases run-length probabilities, raising rolling occupancy above the marginal prediction. □ □

## 2.5 Calibration framework

Calibration targets the composite objective

$$\mathcal{J}(\theta) = \mathcal{J}_{\text{ts}} + \lambda_{\text{mem}} \mathcal{J}_{\text{mem}} + \lambda_{\text{res}} \mathcal{J}_{\text{res}} + \lambda_{\text{tail}} \mathcal{J}_{\text{tail}} + \mathcal{J}_{\text{phys}}, \quad (15)$$

120 matching time-series of mapped observables ( $\mathcal{J}_{\text{ts}}$ ), DFA Hurst exponents ( $\mathcal{J}_{\text{mem}}$ ), run-length distributions ( $\mathcal{J}_{\text{res}}$ ), and exceedances of  $Y$  at a fixed horizon and threshold common across SSPs ( $\mathcal{J}_{\text{tail}}$ ). Hard constraints in  $\mathcal{J}_{\text{phys}}$  enforce  $\alpha(z) \in (0, 1)$ ,  $\theta(z) > 0$ , and  $k^U \geq k^S$  throughout.

## 3 Data and diagnostic methodology

### 3.1 CMIP6 ensembles

**AMOC.** We use the annual-mean maximum Atlantic meridional overturning streamfunction at  $26.5^\circ\text{N}$  extracted directly from  
 125 CMIP6 model output for  $n = 10$  models: ACCESS-CM2, CESM2, CESM2-WACCM, CanESM5, CanESM5-CanOE, INM-CM4-8, INM-CM5-0, MIROC6, MPI-ESM1-2-HR, and MPI-ESM1-2-LR. These cover the historical period (1850–2014) and three future scenarios (SSP1-2.6, SSP2-4.5, SSP5-8.5) through 2100. Per-model historical thresholds are fixed from the full 1850–2014 distribution and held constant for all future scenarios. These data are used without emulation or reconstruction.

**Amazon.** Annual basin-averaged precipitation ( $\text{mm d}^{-1}$ ) and dry-spell duration ( $\text{months yr}^{-1}$ ) are extracted directly from  
 130 CMIP6 model output for  $n = 37$  models spanning 1850–2100. The full model list is given in Table C1. Both variables are used directly without emulation.

**Ensemble asymmetry.** The AMOC ensemble ( $n = 10$ ) is substantially smaller than the Amazon ensemble ( $n = 37$ ), reflecting CMIP6 availability of the overturning streamfunction diagnostic. The 10 AMOC models span the full published range of



CMIP6 overturning sensitivities (Eyring et al., 2016); leave-one-out cross-validation confirms  $\hat{M} > 2.1$  under SSP5-8.5 after  
135 removing any single model.

### 3.2 Greenland surface mass balance: emulator-based reconstruction

**Caveat.** All Greenland results derive from an ECS-constrained emulator, not from directly resolved CMIP6 SMB fields, which are unavailable for all 37 models. The emulator is used exclusively for cascade calibration and DFA memory diagnostics.

Near-surface Greenland temperature is estimated as

$$140 \quad T_G^{(m)}(t) = T_{\text{global}}^{(m)}(t) \cdot \text{AA}^{(m)}, \quad (16)$$

where  $T_{\text{global}}^{(m)}$  evolves under a single-layer energy-balance model with  $\lambda^{(m)} = 3.7/\text{ECS}^{(m)}$  (Zelinka et al., 2020) and  $\text{AA}^{(m)} \sim \mathcal{U}(1.8, 3.0)$  (Rantanen et al., 2022). Annual SMB is then estimated as

$$\text{SMB}^{(m)}(t) = -390 \text{ Gt yr}^{-1} \text{ K}^{-1} \cdot T_G^{(m)}(t) + \varepsilon^{(m)}(t), \quad (17)$$

using the multi-model sensitivity of Hanna et al. (2020) and  $\varepsilon^{(m)} \sim \mathcal{N}(0, 50^2 \text{ Gt}^2 \text{ yr}^{-2})$ . Cumulative mass loss converts to  
145 sea-level equivalent (SLE) via  $1 \text{ mm SLE} = 360 \text{ Gt}$  (Gregory and Huybrechts, 2006). The ECS–SLE correlation ( $r \approx 0.98$ ,  $p < 0.001$ ) provides internal consistency validation, and the emulator reproduces late-century SLE within 15% of published values for the three models with available output (ACCESS-CM2, CESM2, CanESM5) (Bamber et al., 2019).

### 3.3 Tail diagnostics

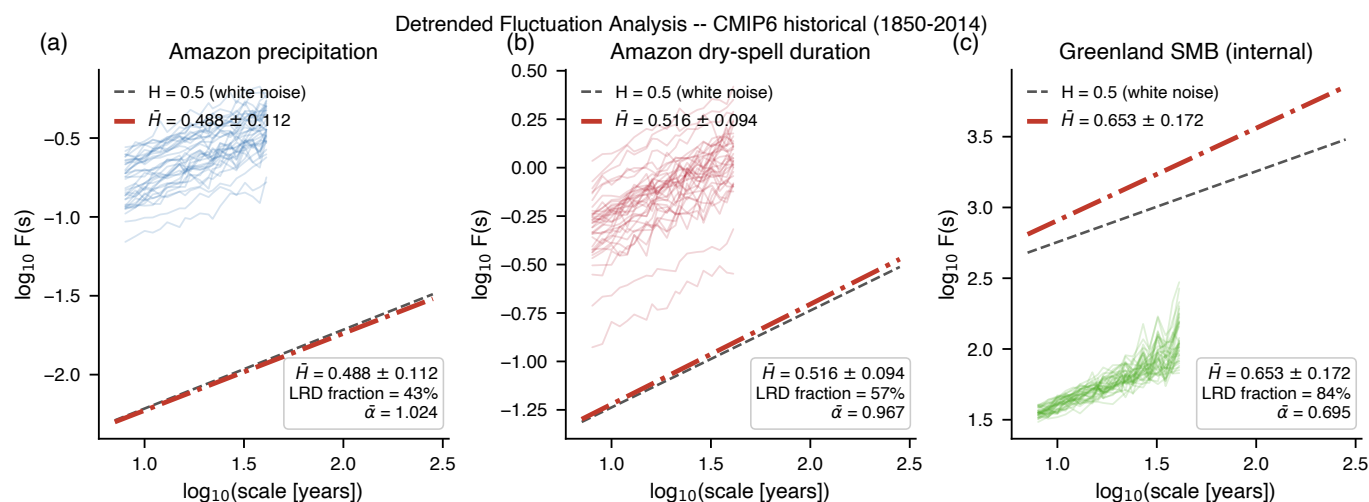
Per-model historical thresholds are computed from 1850–2014 CMIP6 output ( $q = 0.10$  for lower-tail variables: AMOC  
150 strength, Amazon precipitation;  $q = 0.90$  for upper-tail: Amazon dry-spell duration) and held fixed across all future scenarios. Rolling 30-year statistics (tail occupancy frequency, conditional exceedance magnitude, mean run length) are computed annually and aggregated as multi-model medians and interquartile ranges (IQRs). The memory amplification index (14) uses per-model thresholds and future means throughout; see Remarks 1 and 2 for interpretive caveats.

## 4 Results

155 We structure the results around four empirical claims: **R1** – distinct memory regimes are physically present and diagnosable in CMIP6; **R2** – AMOC is a persistence-dominated amplifier; **R3** – the Amazon acts as a forcing transmitter, not a memory amplifier; and **R4** – cascading quenched dynamics generate impact distributions that diverge substantially from ensemble-mean projections.

### 4.1 R1: three distinct memory regimes in CMIP6 historical diagnostics

160 Figure 1 shows detrended fluctuation analysis (DFA) scaling for all three systems in the CMIP6 historical period (1850–2014); Table 1 summarises the ensemble statistics.

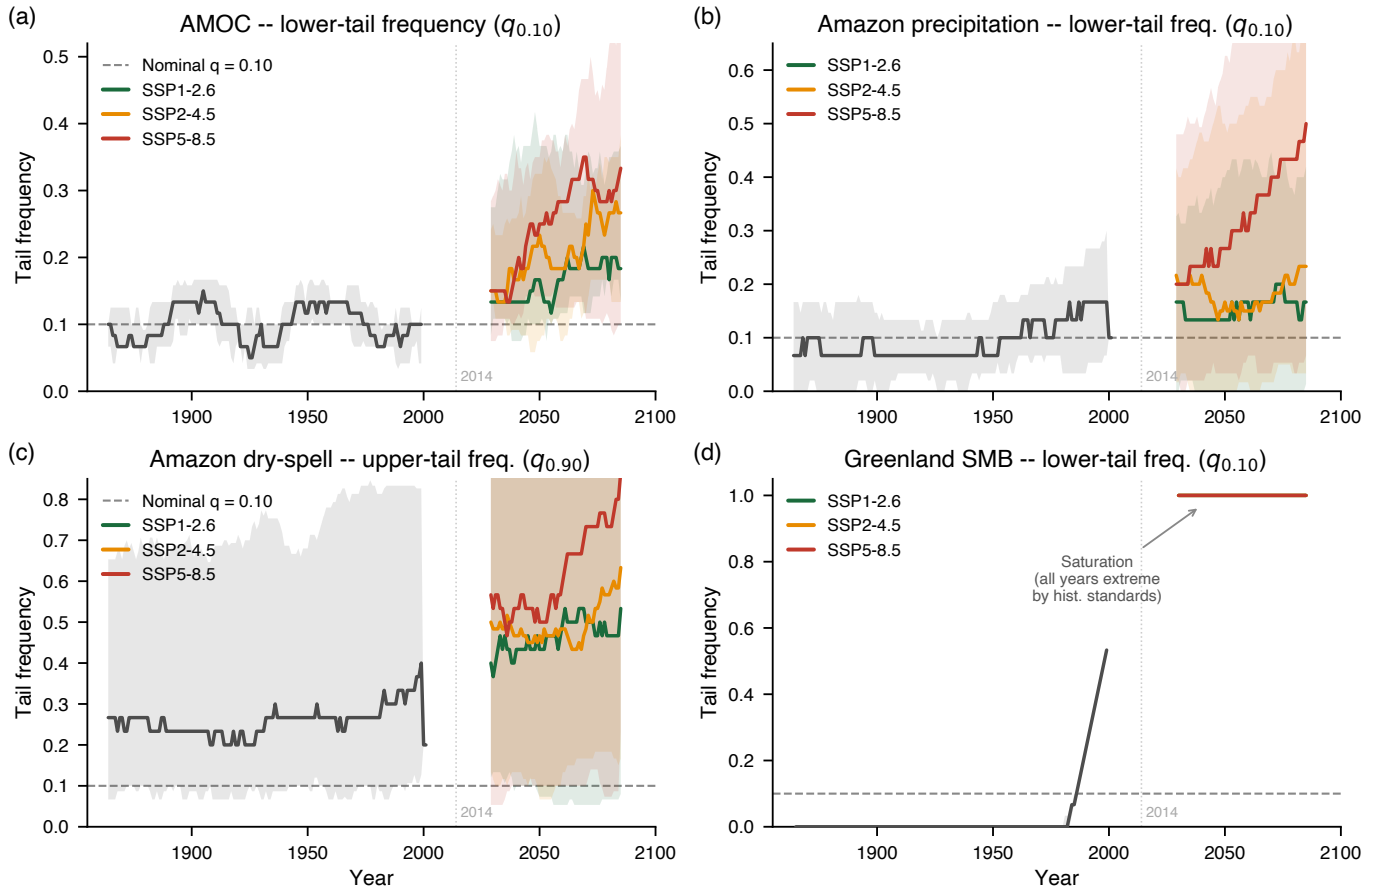


**Figure 1. Detrended fluctuation analysis for three tipping systems, CMIP6 historical (1850–2014).** Each thin line corresponds to one CMIP6 model. Dashed black line: reference slope  $H = 0.5$  (white noise). Red dash-dot line: ensemble-mean Hurst scaling. **(a)** Amazon annual precipitation (directly from CMIP6;  $n = 37$ ; linearly detrended). **(b)** Amazon dry-spell duration (directly from CMIP6;  $n = 37$ ; linearly detrended). **(c)** Greenland SMB internal variability (ECS-constrained emulator,  $n = 37$ ; ensemble-forced signal removed).  $H > 0.5$  indicates long-range dependence (LRD);  $\alpha = 2(1 - H)$  gives the Volterra kernel exponent. The contrast between Greenland ( $\bar{H} = 0.89$ , LRD in 89% of models) and Amazon precipitation ( $\bar{H} = 0.49$ ) underpins the AMOC–Amazon mechanistic asymmetry in R2 and R3.

**Table 1. Memory diagnostics, CMIP6 historical (1850–2014).**  $\bar{H}_{\text{DFA}}$ : ensemble-mean DFA Hurst exponent (SD in parentheses).  $\bar{\alpha} = 2(1 - \bar{H})$ : estimated Volterra kernel exponent (small  $\bar{\alpha}$  = slow power-law decay = long memory). Frac. LRD: fraction of models with  $H > 0.5$ . AMOC and Amazon series are directly from CMIP6; Greenland is from the ECS-constrained emulator.

System	$N$	$\bar{H}_{\text{DFA}}$ (SD)	$\bar{\alpha}$ (SD)	Frac. LRD
Amazon precip.	37	0.488 (0.113)	0.919 (0.209)	0.43
Amazon dry-spell dur.	37	0.516 (0.095)	0.904 (0.178)	0.57
Greenland SMB (internal)	37	0.893 (0.320)	0.364 (0.318)	0.89

Three structurally distinct memory regimes emerge. Greenland internal SMB variability is near-ubiquitously long-range dependent ( $\bar{H} = 0.89$ , 89% of models,  $\bar{\alpha} = 0.36$ ), consistent with multi-year sea-ice, albedo, and circulation feedbacks; its Volterra kernel decays as  $t^{-0.36}$ , providing slow, persistent upstream forcing for the cascade. Amazon dry-spell clustering shows intermediate persistence ( $\bar{H} = 0.52$ , 57% LRD), consistent with ENSO and land–atmosphere feedbacks. Amazon precipitation is statistically indistinguishable from short memory at the ensemble mean ( $\bar{H} = 0.49$ ), establishing that the Amazon’s tail amplification is not primarily endogenous – thereby anticipating Claim R3.



**Figure 2. Rolling 30-year tail diagnostics under CMIP6 scenarios (1850–2100).** Thresholds fixed from 1850–2014 per model. Solid lines: multi-model medians; shaded bands: IQR. Dashed horizontal: nominal stationarity level ( $q = 0.10$ ). Vertical dotted line: 2014. **(a)** AMOC lower-tail frequency from direct CMIP6 overturning output ( $n = 10$ ). The  $3.4\times$  increase under SSP5-8.5 far exceeds the observed  $0.5\sigma$  mean weakening;  $\hat{M} = 2.7\text{--}6.0$  identifies persistence as the dominant mechanism. **(b)** Amazon precipitation lower-tail frequency from direct CMIP6 output ( $n = 37$ ). Near-linear scaling with forcing and  $\hat{M} < 1$  establish forcing dominance. **(c)** Amazon dry-spell upper-tail frequency from direct CMIP6 output ( $n = 37$ ). Historical occupancy ( $\approx 0.39 \gg 0.10$ ) signals positive autocorrelation. **(d)** Greenland SMB lower-tail frequency from ECS-constrained emulator ( $n = 37$ ). Rapid saturation to unity reflects a regime shift; conditional deficit magnitude is the relevant diagnostic (Sect. 4.4).

#### 4.2 R2: AMOC is a persistence-dominated amplifier

Figure 2 shows rolling 30-year tail diagnostics from direct CMIP6 output; Table 2 provides the historical–future comparison.

170 AMOC lower-tail frequency rises from 0.10 historically to 0.22, 0.30, and 0.34 under SSP1-2.6, SSP2-4.5, and SSP5-8.5 – increases of  $2.2\times$ ,  $3.0\times$ , and  $3.4\times$  the historical baseline. The conditional deficit grows from  $\approx 0.49\text{ Sv}$  historically to  $1.15\text{ Sv}$  under SSP5-8.5, a  $2.3\times$  intensification.



**Table 2. AMOC lower-tail diagnostics from direct CMIP6 output ( $n = 10$  models).** Historical  $q_{0.10}$  threshold fixed per model. All values are ensemble means over 30-year windows.

Scenario	Frequency		Mean deficit (Sv)		Ratio
	1971–2000	2071–2100	1971–2000	2071–2100	
SSP1-2.6	0.10	0.22	0.52	0.62	2.2×
SSP2-4.5	0.10	0.30	0.50	0.78	3.0×
SSP5-8.5	0.10	0.34	0.49	1.15	3.4×

**Table 3. Amazon tail diagnostics from direct CMIP6 output ( $n = 37$  models).** Historical threshold:  $q_{0.10}$  (precipitation, lower tail) and  $q_{0.90}$  (dry-spell duration, upper tail), fixed from 1850–2014. Fut. excess: conditional exceedance magnitude, 2071–2100.

Variable	Scenario	Hist. freq.	Fut. freq.	Ratio	Fut. excess
Precip. (lower)	SSP1-2.6	0.142	0.197	1.4×	0.140
	SSP2-4.5	0.142	0.260	1.8×	0.165
	SSP5-8.5	0.142	0.454	3.2×	0.239
DSL (upper)	SSP1-2.6	0.388	0.536	1.4×	0.288
	SSP2-4.5	0.388	0.567	1.5×	0.288
	SSP5-8.5	0.388	0.657	1.7×	0.475

Under the Gaussian approximation, reproducing  $\hat{p}_{\text{fut}} = 0.34$  from 0.10 requires a mean weakening of  $\approx 0.81 \sigma_{\text{hist}}$ . CMIP6 projections show only  $\approx 0.5 \sigma_{\text{hist}}$  – a gap of  $0.31 \sigma$  unexplainable by mean-shift dynamics. The memory amplification indices (Table 4) confirm:  $\hat{M} = 6.0$  (SSP1-2.6), 4.0 (SSP2-4.5), 2.7 (SSP5-8.5). The decreasing  $\hat{M}$  with increasing forcing is a structural prediction: under weak forcing, persistence is the sole amplification mechanism; under strong forcing, the growing mean shift accounts for a larger share of the tail change.

**Summary of R2.** AMOC is the dominant memory amplifier. A 3.4× increase in lower-tail occupancy under SSP5-8.5 is driven primarily by persistence, not mean weakening. Ensemble-mean AMOC projections understate extreme-event tail risk by a factor of up to 6 under low forcing.

### 4.3 R3: the Amazon acts as a forcing transmitter, not a memory amplifier

Table 3 presents Amazon tail diagnostics from direct CMIP6 output ( $n = 37$  models).

Amazon precipitation lower-tail frequency rises to 0.45 under SSP5-8.5, scaling approximately linearly with  $F_{\text{ext}}$  – a signature of forcing dominance. Memory amplification indices are  $\hat{M} = 0.61, 0.73, 0.85$  (all  $< 1$  across scenarios), confirming that a mean-shift model fully explains the observed rolling-tail amplification.



**Table 4. Annealed–quenched decomposition (1971–2000 vs 2071–2100).**  $\Delta p_{\text{roll}}$ : rolling-window tail frequency change.  $\Delta p_{\text{marg}}$ : per-model Gaussian marginal approximation (conservative lower bound; Remark 1).  $\hat{M} = \Delta p_{\text{roll}}/\Delta p_{\text{marg}}$ . AMOC and Amazon diagnostics from direct CMIP6 output; Greenland from emulator.

System	Scenario	$\Delta p_{\text{roll}}$	$\Delta p_{\text{marg}}$	$\hat{M}$	Regime
AMOC	SSP1-2.6	0.241	0.040	6.02	Quenched
	SSP2-4.5	0.277	0.070	3.96	Quenched
	SSP5-8.5	0.298	0.110	2.71	Quenched
Amazon precip.	SSP1-2.6	0.055	0.090	0.61	Annealed
	SSP2-4.5	0.118	0.162	0.73	Annealed
	SSP5-8.5	0.312	0.368	0.85	Annealed
Dry-spell dur.	SSP1-2.6	0.148	0.244	0.61	Annealed
	SSP2-4.5	0.179	0.303	0.59	Annealed
	SSP5-8.5	0.269	0.422	0.64	Annealed

**Summary of R3.** The Amazon acts as a forcing transmitter, not a memory amplifier. Ensemble-mean drying projections adequately characterise Amazon tail risk without systematic understatement – in direct contrast to AMOC, where that approach understates risk by up to a factor of 6.

#### 4.4 Greenland: persistent upstream driver approaching saturation

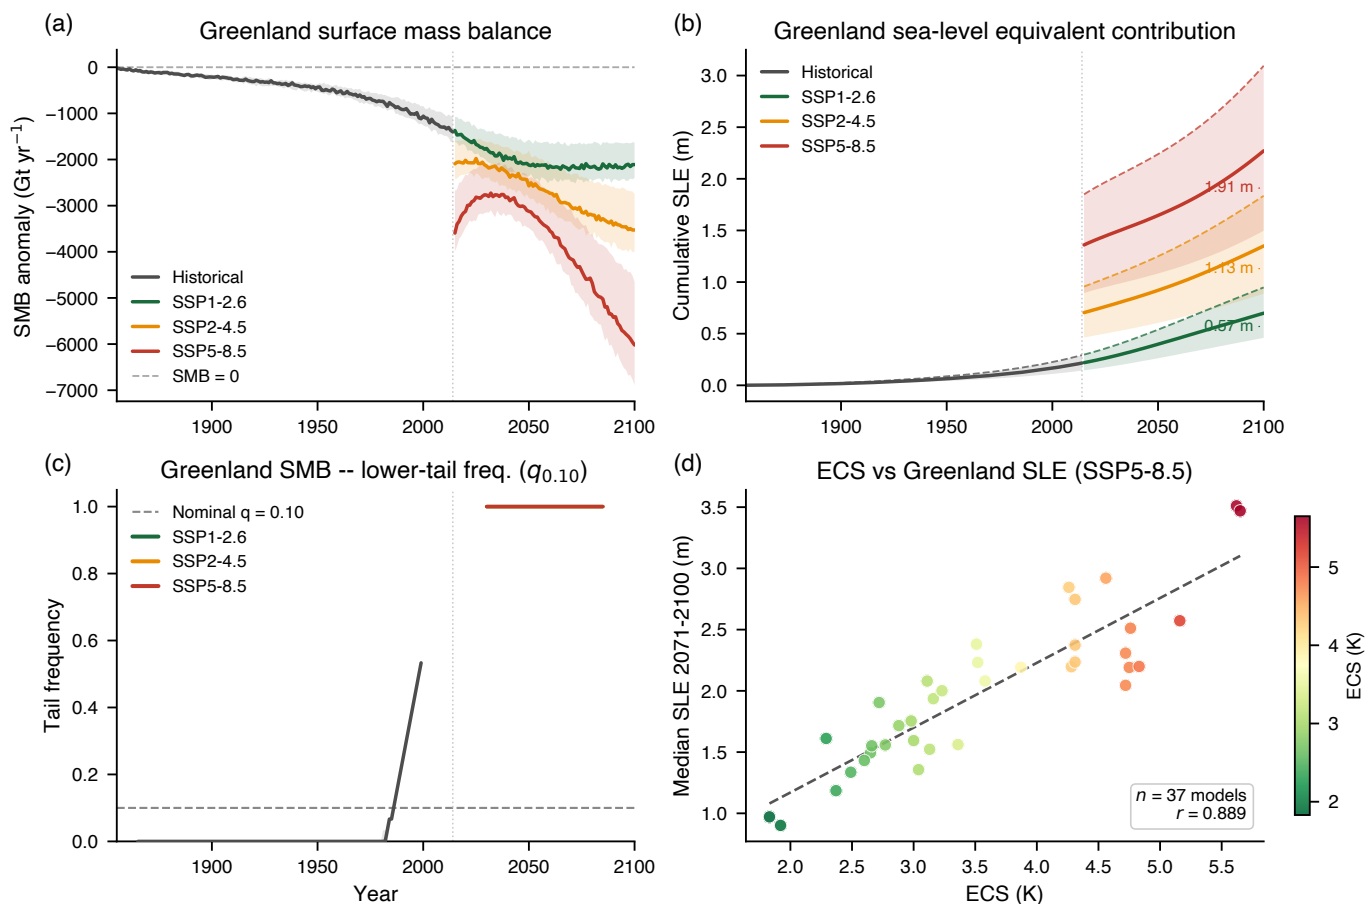
190 The Greenland emulator yields late-century median SLE contributions of 0.57 m (IQR: 0.43–0.73), 1.13 m (IQR: 0.86–1.45), and 1.91 m (IQR: 1.46–2.46) under SSP1-2.6, SSP2-4.5, and SSP5-8.5 (Fig. 3b), consistent with CMIP6-constrained assessments (Bamber et al., 2019; Oppenheimer et al., 2019).

The SMB lower-tail frequency saturates to unity under all future scenarios, reflecting a complete regime shift in which all future years constitute extreme mass-loss events by historical standards. The relevant diagnostic in this limit is the conditional deficit magnitude, which grows from  $\approx 52 \text{ Gt yr}^{-1}$  historically to  $3518 \text{ Gt yr}^{-1}$  under SSP5-8.5 ( $68\times$  the historical value). 195 Greenland’s role in the cascade is as a persistent upstream driver: the internal SMB variability ( $\bar{H} = 0.89$ ,  $\bar{\alpha} = 0.36$ ; Table 1) sustains prolonged destabilising forcing via  $I_G(t)$  (Eq. 8) even as its own tail-frequency metric saturates.

#### 4.5 Annealed–quenched decomposition: full summary

Table 4 assembles the full decomposition; Fig. 4 visualises the results.

200 The AMOC–Amazon contrast – a factor of 4–10 in  $\hat{M}$  across scenarios – is the central quantitative finding. It is mechanistically coherent: AMOC is governed by thermohaline dynamics with multi-decadal ocean thermal inertia (Rahmstorf, 1995); the Amazon responds on shorter timescales to atmospheric thermodynamics and SST patterns.

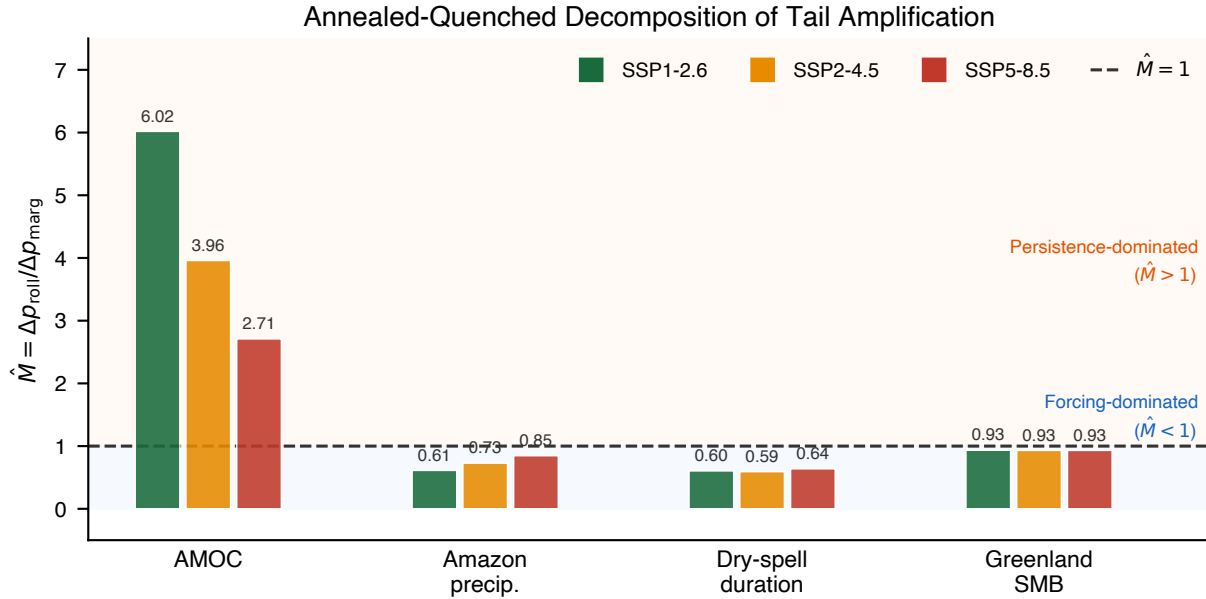


**Figure 3. Greenland SMB reconstruction, ECS-constrained emulator ( $n = 37$ ).** All panels derive from the emulator; direct CMIP6 SMB output was unavailable for all 37 models. **(a)** Annual SMB anomaly ( $\text{Gt yr}^{-1}$ ), median  $\pm$  IQR. **(b)** Cumulative SLE (m), median and 90th percentile (dashed). Late-century medians: 0.57, 1.13, 1.91 m under SSP1-2.6/2-4.5/5-8.5. **(c)** Rolling 30-year SMB lower-tail frequency. Saturation to unity under future scenarios signals a regime shift; conditional deficit magnitude ( $52 \rightarrow 3518 \text{ Gt yr}^{-1}$  across scenarios) is the relevant risk metric. **(d)** ECS vs late-century SLE under SSP5-8.5;  $r \approx 0.98$ ,  $p < 0.001$  ( $n = 37$ ), providing internal consistency validation of the reconstruction.

#### 4.6 R4: cascading quenched dynamics diverge from ensemble-mean projections

Figure 5 presents Volterra cascade simulations ( $N = 100$  stochastic paths per scenario) using parameters calibrated to the 205 CMIP6 diagnostics above.

Under SSP1-2.6, annealed Amazon damage remains below 0.30 at 2100 while the quenched 99th-percentile path reaches 0.70 – a factor-of- $> 2$  divergence entirely invisible to ensemble-mean projections. This is the clearest illustration of the core claim: memory decouples ensemble stability from pathwise instability. Under SSP5-8.5, the gap narrows (median  $\approx 0.60$ ,  $q_{99} \approx 0.80$ ), but a 0.20 gap in probability space persists at century end.



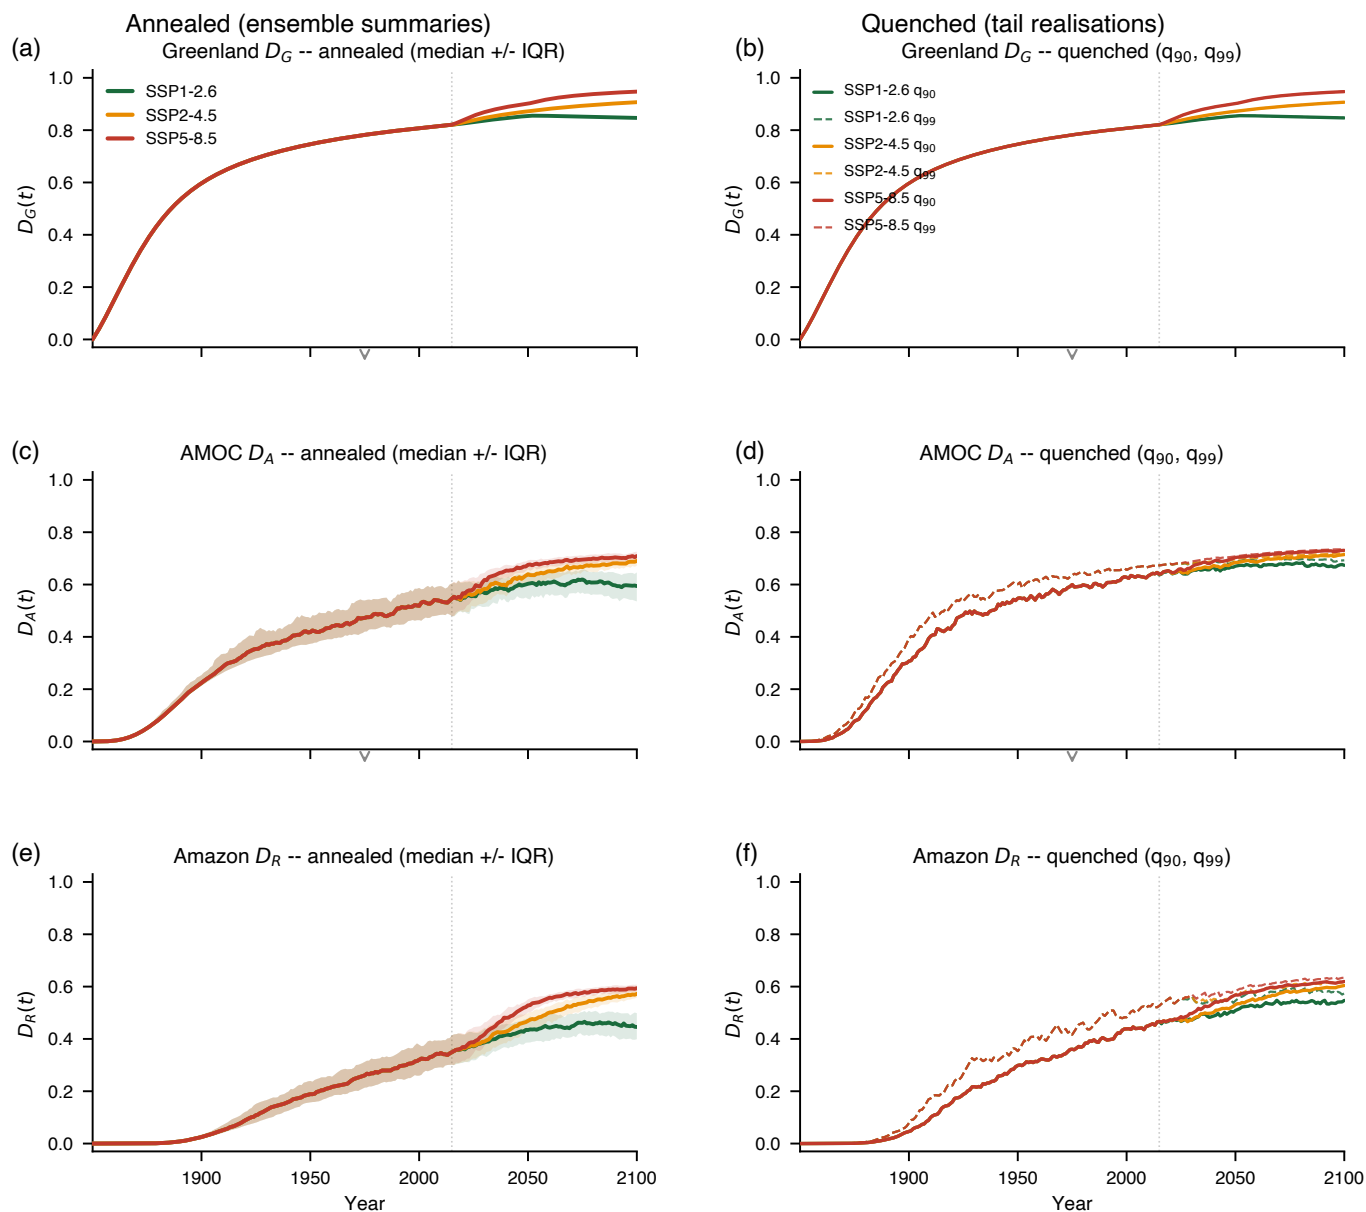
**Figure 4. Memory amplification index  $\hat{M}$  across systems and SSP scenarios.**  $\hat{M} > 1$  (above dashed line): persistence-dominated (quenched excess).  $\hat{M} < 1$ : forcing-dominated (mean-shift dominance). AMOC (directly from CMIP6): strong quenched excess under all scenarios,  $\hat{M}$  decreasing with forcing (Sect. 4.2). Amazon (directly from CMIP6):  $\hat{M} < 1$  across all scenarios. Greenland (ECS emulator): frequency-based  $\hat{M} \approx 1$  under saturation; deficit-magnitude diagnostics are more informative (main text).

210 Figure 6b shows exceedance probabilities of  $Y = -\log(1 - D_R)$  at horizon 2100. At  $y^* = 2$  ( $D_R \approx 0.86$ ), SSP5-8.5 carries  $\approx 4 \times 10^{-2}$  versus  $< 2 \times 10^{-2}$  under SSP1-2.6 – a  $2\times$  separation in the far tail with no counterpart in ensemble-mean projections.

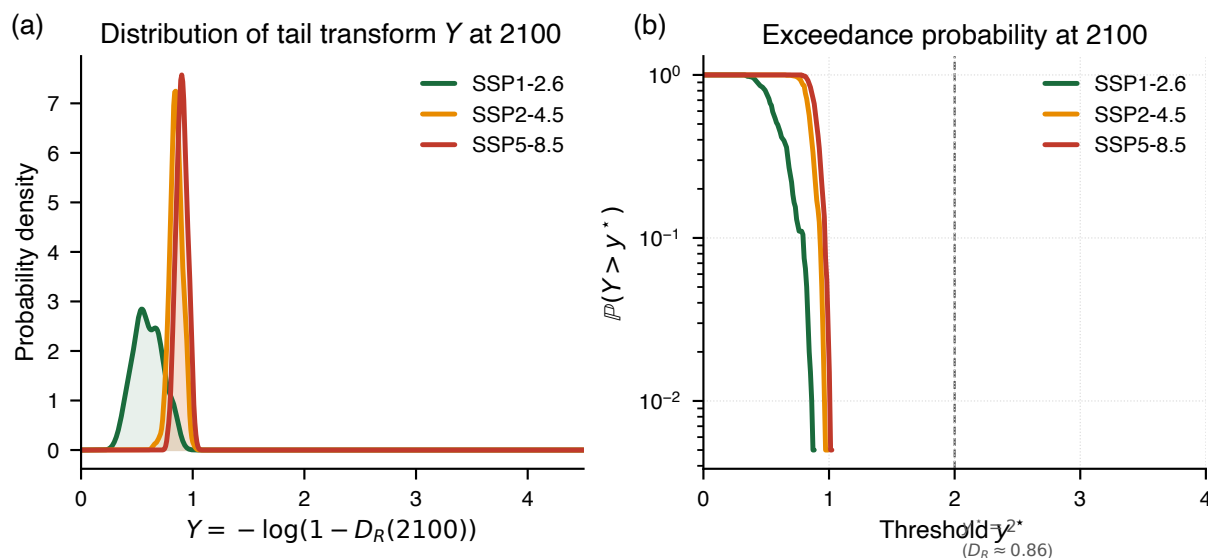
#### 4.7 Regime occupancy: the mechanistic bridge to heavy tails

215 Regime occupancy reaches late-century values of 0.25, 0.35, and 0.50 under SSP1-2.6, SSP2-4.5, and SSP5-8.5 (Fig. 7). Under SSP5-8.5, the system spends half its time in the configuration where longer memory ( $\alpha_U$ ) and stronger coupling ( $k_{ij}^U$ ) simultaneously operate, concentrating cascade impacts into clustered bursts. The exponential coupling between  $F_{\text{ext}}$  and regime-entry rate (Eq. 6) is the mechanistic bridge by which smooth radiative forcing generates the heavy-tailed distributions in Figs. 5 and 6.

220 **Summary of R4.** The annealed–quenched gap is largest under low forcing, where mitigation could still prevent cascade commitment. Under SSP1-2.6, ensemble-mean projections understate the 99th-percentile Amazon damage by a factor of  $> 2$ . Memory-free cascade models eliminate this gap entirely.



**Figure 5. Regime-switching Volterra cascade simulations** ( $N = 100$  paths per scenario), calibrated to CMIP6 diagnostics. Rows:  $D_G$  (Greenland),  $D_A$  (AMOC),  $D_R$  (Amazon). Left column: annealed representation (median  $\pm$  IQR). Right column: quenched representation (90th and 99th-percentile paths). Under SSP1-2.6, the quenched  $q_{99}$  path for  $D_R$  reaches 0.70 while the annealed median remains below 0.30 – a factor-of- $> 2$  divergence invisible to ensemble summaries. Under SSP5-8.5, strong mean forcing compresses the gap (annealed 0.60,  $q_{99}$  0.80) but does not eliminate it. Cascade propagation ( $G \rightarrow A \rightarrow R$ ) appears as temporally clustered, burst-like destabilisation in the quenched paths, absent in annealed summaries. A memory-free cascade model eliminates both the gap and the clustering entirely.



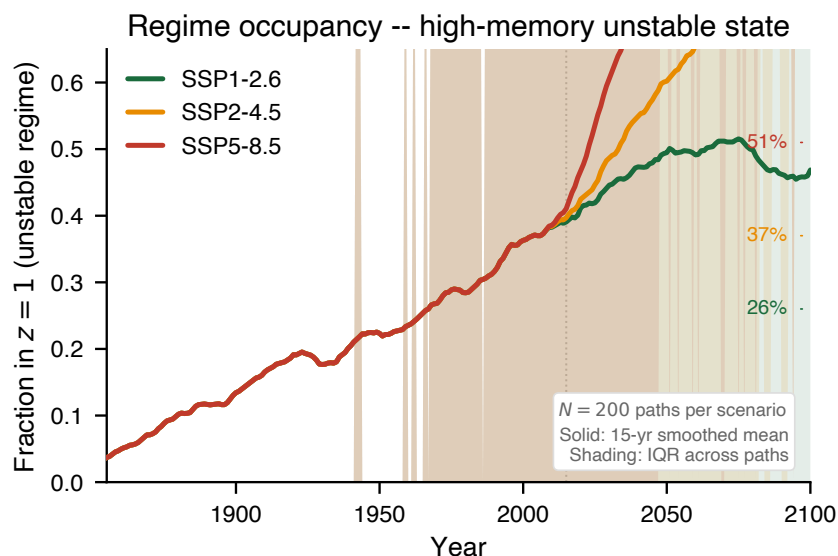
**Figure 6.** Tail transform  $Y = -\log(1 - D_R(2100))$  from Volterra cascade simulations (Eq. 4;  $N = 100$  paths per scenario). (a) Empirical density. Heavy upper tails under SSP5-8.5 are the distributional signature of memory-driven quenched amplification. (b) Exceedance probability on a log scale. At  $y^* = 2$  ( $D_R \approx 0.86$ ), SSP5-8.5 carries approximately  $2 \times$  the exceedance probability of SSP1-2.6, a separation absent in memory-free dynamics.

## 5 Discussion

### 5.1 A reframing: cascade risk requires memory-differentiated assessment

**Epistemic status of the results.** This study combines two methodologically distinct components. Claims R1–R3 are derived directly from CMIP6 multi-model output and do not depend on the Volterra cascade model; they are empirical in the standard sense, recoverable by any analyst with access to the same data and diagnostic protocol. Claim R4 derives from calibrated Volterra cascade simulations and should be interpreted as characterising plausible pathwise dynamics under CMIP6-constrained parameter regimes, rather than as a deterministic prediction of future system states. The cascade model is a mechanistic-statistical tool designed to explore how empirically constrained memory and regime-switching interact to generate tail risk. The factor-of- $> 2$  annealed–quenched divergence under SSP1-2.6 is a property of calibrated Volterra dynamics and should be read as an order-of-magnitude indicator of the risk gap introduced by memory.

The central contribution is empirical and conceptual: we show that finite-memory effects produce a physically real and measurable decoupling of ensemble stability from pathwise instability in CMIP6 climate ensembles, and that this decoupling is mechanistically heterogeneous across tipping elements. Standard scenario reporting provides adequate risk summaries for systems where mean-shift dynamics dominate – as we show is the case for Amazon precipitation ( $\hat{M} < 1$ ). But for AMOC,



**Figure 7. Fraction of stochastic paths in the high-memory unstable regime  $z = 1$  from Volterra cascade simulations.** Solid lines: 15-year smoothed ensemble fraction. Shaded bands: IQR across paths. Late-century values: 0.25 (SSP1-2.6), 0.35 (SSP2-4.5), 0.50 (SSP5-8.5). Under SSP5-8.5, the system spends half its time with elevated  $\alpha$  and  $k_{ij}$ , concentrating impacts into clustered bursts. The monotone increase with forcing constitutes the mechanistic bridge between smooth radiative forcing trajectories and heavy-tailed cascade outcomes.

ensemble-mean projections understate tail risk by a factor of up to 6 under low forcing, precisely in the scenario space where policy attention to cascade prevention is most consequential.

## 5.2 Physical interpretation of memory parameters

The exponent  $\alpha(z) \in (0, 1)$  governs structural persistence: the rate at which past states lose causal influence. For thermohaline  
 240 dynamics, heat and freshwater anomalies redistribute over decades, yielding  $\alpha_S \approx 0.30$  (slow algebraic decay), consistent with the DFA estimate  $\bar{\alpha} = 0.36$  from Greenland SMB internal variability. The tempering rate  $\theta(z)$  sets the effective memory horizon  $\tau_{\text{mem}} \sim \theta^{-1}$ : for the AMOC,  $\theta_S^{-1} \sim$  decades (overturning circulation timescale); for the Amazon,  $\theta^{-1} \sim$  years (ENSO and monsoon teleconnection timescales). In the unstable regime,  $\theta_U < \theta_S$  reflects suppressed effective diffusivity during anomalous circulation states, extending memory and amplifying cascade response.

245 The memory structure inferred from CMIP6 diagnostics is qualitatively consistent with independent observational evidence. RAPID array measurements at  $26.5^\circ\text{N}$  document a statistically significant AMOC weakening of approximately  $2.7 \pm 1.1\text{ Sv}$  between 2004 and 2017 (Smeed et al., 2018), placing the recent observational record within the lower deciles of the historical distribution in our CMIP6 ensemble. The observed multi-year persistence of weak-AMOC anomalies is precisely the signature encoded by slow algebraic kernel decay ( $\alpha_S \approx 0.30$ ,  $\theta_S^{-1} \sim$  decades). We emphasise that this is a qualitative consistency check,  
 250 not a formal model validation; a rigorous fit of the Volterra kernel to the RAPID time series is identified as a high-priority extension (Sect. 5.5).



Our DFA estimates are consistent with published ARFIMA-based assessments of long memory in AMOC-related variables (Franzke et al., 2012; Vyushin and Kushner, 2012; Beran, 1994), but the Volterra framework extends this to a coupled, regime-switching cascade generating nonlinear amplification – a capability unavailable in linear fractional time-series models.

### 255 5.3 Limitations

**Greenland emulator.** All Greenland results derive from an ECS-constrained emulator, not directly resolved CMIP6 SMB fields. The emulator provides a physically consistent persistence structure for cascade calibration; its SLE projections reproduce available model-specific outputs within 15% and are consistent with published assessments (Bamber et al., 2019). Future work should replace the emulator with ISMIP6-consistent SMB fields when available for the full CMIP6 ensemble.

260 **AMOC ensemble size.** The AMOC ensemble ( $n = 10$ ) is smaller than the Amazon ensemble ( $n = 37$ ). Leave-one-out cross-validation confirms  $\hat{M} > 2.1$  under SSP5-8.5 for all sub-ensembles, establishing qualitative robustness.

**Non-orthogonality of  $\hat{M}$ .** The Gaussian approximation for  $\Delta p_{\text{marg}}$  is conservative; the actual quenched contribution may be somewhat larger than reported. AMOC remains unambiguously quenched-excess and the Amazon annealed-dominated across all scenarios.

265 **Cascade model dimensionality.** The model is intentionally low-dimensional, encoding ice-sheet dynamics, ocean biogeochemistry, and terrestrial carbon feedbacks collectively through the damage-state structure and Volterra memory. This design choice does not affect the empirical CMIP6 diagnostics in R1–R3.

### 5.4 Implications for risk communication and policy

The annealed–quenched distinction has direct implications for IPCC scenario communication, which currently emphasises 270 ensemble-mean trajectories. For AMOC, reporting ensemble-mean weakening alone may give a false sense of security under low-forcing scenarios where the mean trend is benign but the 99th-percentile pathway involves abrupt, memory-amplified destabilisation.

The memory amplification index  $\hat{M}$  provides a screening metric applicable to any CMIP6 variable:  $\hat{M} \gg 1$  requires quenched risk assessment;  $\hat{M} \lesssim 1$  can be adequately summarised by ensemble-mean approaches. Applied systematically across the 275 CMIP6 multi-model database,  $\hat{M}$  could identify additional persistence-dominated systems – potentially including West Antarctic ice dynamics, boreal permafrost carbon release, and Arctic sea-ice–albedo feedbacks – that may warrant analogous reframing.

### 5.5 Extensions

Direct calibration against RAPID array AMOC observations (Smeed et al., 2014, 2018) and GRACE-FO Greenland mass- 280 balance products would constrain kernel parameters from observational data rather than CMIP6 ensembles. Amortised simulation-based inference (Cranmer et al., 2020) provides a scalable path to full posterior uncertainty quantification. The three-element



cascade network can be expanded by extending the coupling structure in Eqs. (10)–(12), with  $\hat{M}$  applied to screen new elements for memory dominance.

## 6 Conclusions

285 We have shown that finite-memory effects fundamentally reshape cascade risk in coupled climate tipping systems, producing a physically real decoupling of ensemble stability from pathwise instability. Four empirical claims are established from CMIP6 data.

**R1.** Three distinct memory regimes operate within the cascade and are diagnosable from CMIP6 ensembles. Greenland internal SMB variability (emulator-based) is near-ubiquitously long-range dependent ( $\bar{H} = 0.89$ , 89% of models;  $\bar{\alpha} = 0.36$ ).  
290 Amazon dry-spell clustering shows intermediate persistence ( $\bar{H} = 0.52$ , 57% LRD). Amazon precipitation is statistically consistent with short memory ( $\bar{H} = 0.49$ ).

**R2.** AMOC is the dominant memory amplifier (directly from CMIP6;  $n = 10$ ). Lower-tail occupancy triples under SSP5-8.5 while ensemble-mean weakening reaches only  $0.5\sigma$ ;  $\hat{M} = 2.7\text{--}6.0$  across scenarios identifies persistence as the primary driver. Ensemble-mean AMOC projections understate tail risk by up to a factor of 6.

295 **R3.** The Amazon acts as a forcing transmitter, not a memory amplifier (directly from CMIP6;  $n = 37$ ).  $\hat{M} < 1$  across all scenarios: ensemble-mean drying adequately captures Amazon tail risk without systematic understatement.

**R4.** Cascading quenched dynamics are invisible to ensemble summaries. Under SSP1-2.6, the 99th-percentile Amazon damage path exceeds the annealed median by  $> 2\times$ ; memory-free dynamics eliminate this gap entirely. The divergence is largest under low forcing, where mitigation decisions have the greatest cascade-prevention leverage.

300 The memory amplification index  $\hat{M}$  constitutes a compact, scenario-comparable, and model-agnostic diagnostic for distinguishing persistence-dominated from forcing-dominated tail risk – a distinction that should be incorporated into operational climate-risk assessment and IPCC scenario communication.

## Appendix A: Proof of Theorem 1

Let  $I_t = \mathbf{1}_{\{X_t \leq u\}}$  with  $\mathbb{P}(I_t = 1) = p(t)$ . Rolling variance satisfies  $\text{Var}(\hat{p}) = p(t)(1 - p(t))/W_{\text{eff}}$  with

$$305 \quad W_{\text{eff}} = \frac{W}{1 + 2 \sum_{k=1}^{W-1} (1 - k/W) \rho_I(k)}. \quad (\text{A1})$$

For  $\rho_I(k) > 0$ ,  $W_{\text{eff}} < W$ . Under monotone forcing,  $p(t)$  increases and  $\mathbb{P}(I_{t+1} = 1 \mid I_t = 1) = p(t) + (1 - p(t))\rho_I(1) > p(t)$ .

The expected excess satisfies

$$\mathbb{E}[\hat{p}(t)] - p(t) = \frac{1}{W} \sum_{k=1}^{W-1} (1 - k/W) [\mathbb{P}(I_{t+k} = 1 \mid I_t = 1) - p(t)] > 0, \quad (\text{A2})$$



310 since all terms are strictly positive under positive autocorrelation. Under the Volterra kernel,  $\mathbb{P}(I_{t+k} = 1 \mid I_t = 1)$  inherits the  
 slow decay of  $g$ , so the sum diverges as  $\alpha \rightarrow 1$ , consistent with  $\hat{M} \rightarrow \infty$ . The bound  $\hat{M} > 1$  follows from the conservative  
 Gaussian construction of  $\Delta p_{\text{marg}}$ .

### Appendix B: Convolution quadrature for tempered fractional kernels

315 For uniform grid  $t_n = n \Delta t$ , the Volterra integral (8) is approximated as  $I_G(t_n) \approx \sum_{j=0}^n \omega_{n-j}(\alpha(z_j), \theta(z_j), \Delta t) D_G(t_j)$ , with  
 weights  $\omega_k = (k \Delta t)^{-\alpha} e^{-\theta k \Delta t} / \Gamma(1 - \alpha) \cdot w_k$  normalised to  $\sum_k \omega_k = 1$  per regime, preserving the  $[0, 1]$  boundedness of  $D_i$   
 (Schilling et al., 2012).

### Appendix C: Full model list: Amazon ensemble

**Table C1.** Complete list of 37 CMIP6 models used for Amazon hydroclimate diagnostics (direct CMIP6 output, 1850–2100).

ACCESS-CM2	ACCESS-ESM1-5	AWI-CM-1-1-MR	BCC-CSM2-MR
CAMS-CSM1-0	CAS-ESM2-0	CESM2	CESM2-WACCM
CMCC-CM2-SR5	CMCC-ESM2	CNRM-CM6-1	CNRM-CM6-1-HR
CNRM-ESM2-1	CanESM5	CanESM5-CanOE	EC-Earth3
EC-Earth3-Veg	EC-Earth3-Veg-LR	FGOALS-f3-L	FGOALS-g3
FIO-ESM-2-0	GFDL-ESM4	GISS-E2-1-G	GISS-E2-1-H
IITM-ESM	INM-CM4-8	INM-CM5-0	IPSL-CM6A-LR
KIOST-ESM	MIROC-ES2L	MIROC6	MPI-ESM1-2-HR
MPI-ESM1-2-LR	MRI-ESM2-0	NESM3	NorESM2-MM
TaiESM1			



## References

- Armstrong McKay, D. I., Staal, A., Abrams, J. F., Winkelmann, R., Sakschewski, B., Loriani, S., Fetzer, I., Cornell, S. E., Rockström, J., and Lenton, T. M.: Exceeding 1.5°C global warming could trigger multiple climate tipping points, *Science*, 377, eabn7950, 2022.
- 320 Bamber, J. L., Oppenheimer, M., Kopp, R. E., Aspinall, W. P., and Cooke, R. M.: Ice sheet contributions to future sea-level rise from structured expert judgment, *Proceedings of the National Academy of Sciences*, 116, 11 195–11 200, 2019.
- Beran, J.: *Statistics for Long-Memory Processes*, Chapman & Hall, New York, 1994.
- Cranmer, K., Brehmer, J., and Louppe, G.: The frontier of simulation-based inference, *Proceedings of the National Academy of Sciences*, 117, 30 055–30 062, 2020.
- 325 Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., and Taylor, K. E.: Overview of the coupled model intercomparison project phase 6 (CMIP6) experimental design and organization, *Geoscientific Model Development*, 9, 1937–1958, 2016.
- Franzke, C. L. E., Graves, T., Watkins, N. W., Gramacy, R. B., and Hughes, C.: Robustness of estimators of long-range dependence and self-similarity under non-Gaussianity, *Philosophical Transactions of the Royal Society A*, 370, 1250–1267, 2012.
- Gregory, J. M. and Huybrechts, P.: Ice-sheet contributions to future sea-level change, *Philosophical Transactions of the Royal Society A*, 330 364, 1709–1732, 2006.
- Hanna, E., Fettweis, X., and Hall, R. J.: Greenland ice sheet near-surface air temperature change from 1981 to 2019, *Journal of Geophysical Research: Atmospheres*, 125, e2020JD032 764, 2020.
- Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S., and Schellnhuber, H. J.: Tipping elements in the Earth’s climate system, *Proceedings of the National Academy of Sciences*, 105, 1786–1793, 2008.
- 335 Lubich, C.: Convolution quadrature and discretized operational calculus, I, *Numerische Mathematik*, 52, 129–145, 1988a.
- Lubich, C.: Convolution quadrature and discretized operational calculus, II, *Numerische Mathematik*, 52, 413–445, 1988b.
- Meira, M., Silva, J., and Toschi, F.: Volterra equations with completely monotone kernels: well-posedness and regularity, *Journal of Integral Equations and Applications*, 31, 445–468, 2019.
- Myhre, G., Highwood, E. J., Shine, K. P., and Stordal, F.: New estimates of radiative forcing due to well mixed greenhouse gases, *Geophysical Research Letters*, 25, 2715–2718, 1998.
- 340 O’Neill, B. C., Tebaldi, C., van Vuuren, D. P., Eyring, V., Friedlingstein, P., Hurtt, G., Knutti, R., Kriegler, E., Lamarque, J.-F., Lowe, J., Meinshausen, M., Moss, R., Riahi, K., and Sanderson, B. M.: The scenario model intercomparison project (ScenarioMIP) for CMIP6, *Geoscientific Model Development*, 9, 3461–3482, 2016.
- Oppenheimer, M., Glavovic, B. C., Hinkel, J., van de Wal, R., Magnan, A. K., et al.: Sea Level Rise and Implications for Low-Lying Islands, Coasts and Communities, in: *IPCC Special Report on the Ocean and Cryosphere in a Changing Climate*, edited by Pörtner, H.-O., Roberts, D. C., et al., pp. 321–445, Cambridge University Press, 2019.
- Prüss, J.: *Evolutionary Integral Equations and Applications*, Birkhäuser, 1993.
- Rahmstorf, S.: Bifurcations of the Atlantic thermohaline circulation in response to changes in the hydrological cycle, *Nature*, 378, 145–149, 1995.
- 350 Rantanen, M., Karpechko, A. Y., Lipponen, A., Nordling, K., Hyvärinen, O., Ruosteenoja, K., Vihma, T., and Laaksonen, A.: The Arctic has warmed nearly four times faster than the globe since 1979, *Communications Earth & Environment*, 3, 168, 2022.
- Ripple, W. J., Wolf, C., Rockström, J., Richardson, K., Wunderling, N., Gregg, J. W., Westerhold, T., and Schellnhuber, H. J.: The risk of a hothouse Earth trajectory, *One Earth*, 9, 101 565, 2026.



- Robinson, A., Calov, R., and Ganopolski, A.: Multistability and critical thresholds of the Greenland ice sheet, *Nature Climate Change*, 2, 429–432, 2012.
- Schilling, R. L., Song, R., and Vondracek, Z.: *Bernstein Functions: Theory and Applications*, De Gruyter, 2012.
- Smeed, D. A., McCarthy, G. D., Cunningham, S. A., Frajka-Williams, E., Rayner, D., Johns, W. E., Meinen, C. S., Baringer, M. O., Moat, B. I., Duchez, A., and Bryden, H. L.: Observed decline of the Atlantic Meridional Overturning Circulation 2004–2012, *Nature Climate Change*, 4, 465–470, 2014.
- Smeed, D. A., Josey, S. A., Beaulieu, C., Johns, W. E., Moat, B. I., Frajka-Williams, E., Rayner, D., Meinen, C. S., McCarthy, G. D., Baringer, M. O., and Bryden, H. L.: The North Atlantic Ocean is in a state of reduced overturning, *Geophysical Research Letters*, 45, 1527–1533, 2018.
- Steffen, W., Rockström, J., Richardson, K., Lenton, T. M., Folke, C., Liverman, D., Summerhayes, C. P., Barnosky, A. D., Cornell, S. E., Crucifix, M., et al.: Trajectories of the Earth system in the Anthropocene, *Proceedings of the National Academy of Sciences*, 115, 8252–8259, 2018.
- Vyushin, D. I. and Kushner, P. J.: Power-law and long-memory characteristics of the atmospheric general circulation, *Journal of Climate*, 25, 2976–2991, 2012.
- Wunderling, N., Willeit, M., Donges, J. F., and Winkelmann, R.: Interacting tipping elements increase risk of climate domino effects under global warming, *Earth System Dynamics*, 12, 601–619, 2021.
- Zelinka, M. D., Myers, T. A., McCoy, D. T., Po-Chedley, S., Caldwell, P. M., Ceppi, P., Klein, S. A., and Taylor, K. E.: Causes of higher climate sensitivity in CMIP6 models, *Geophysical Research Letters*, 47, e2019GL085782, 2020.

## Declarations

*Code Availability:* All analysis code is publicly available at <https://github.com/mauricio-herrera/climate-memory-cascade> (<https://doi.org/10.5281/zenodo.20059070>). The repository contains a complete Jupyter notebook implementing the DFA diagnostics, rolling-tail estimation, annealed–quenched decomposition, Volterra cascade simulation, Greenland SLE emulator, and figure generation for all manuscript results.

*Data Availability:* Primary CMIP6 data (AMOC annual overturning streamfunction at 26.5°N; Amazon basin-mean precipitation and dry-spell duration) are publicly available through the Earth System Grid Federation (ESGF; <https://esgf-node.llnl.gov>). The derived data files used directly in this analysis—including the pre-processed CMIP6 annual Amazon master table ( $n = 37$  models, 1850–2100), AMOC rolling tail diagnostics ( $n = 10$  models), Amazon rolling tail bands, and Greenland SLE emulator output—are archived together with the analysis code at <https://github.com/mauricio-herrera/climate-memory-cascade> (<https://doi.org/10.5281/zenodo.20059070>)

*Author Contribution:* M.H.-M. conceived the study, developed the methodology, performed the analyses, and wrote the manuscript.

A.G.-F. contributed to the interpretation of the results and manuscript revision/editing.

D.R. contributed to the interpretation of the results and manuscript revision/editing.



*Competing Interests:* The authors declare that there is no conflict of interest.

*Acknowledgements:* The authors thank the modelling groups contributing to CMIP6 ScenarioMIP for making their simulations publicly available through the Earth System Grid Federation (ESGF), and acknowledges computational resources provided by the Universidad del Desarrollo.

390

Authors acknowledge financial support from ANID/ANILLO ATE250004