Supplementary Information

Institutions as Switching Circuits: Decision Synchrony under Rising Connectedness

Author: Ingo Piepers

Correspondence: ingopiepers@gmail.com

Purpose of this file. This SI accompanies *Institutions as switching circuits* and extends the main text by documenting the seven-step pipeline (S0–S7), robustness checks, and source data used to produce all figures and tables. It is designed as a continuation of the paper: each SI section mirrors the corresponding result in the main text and is cross-referenced there (e.g., "see SI S4").

Contents (overview).

Guide to appendices:

Appendix	Subject
A	Evidence spine (one-page map linking claims to figures/tables across S1–S7).
B	Design companion (operational templates: band-entry rule, ω–σ diagram, dashboards, treaty clauses, risk pricing).
C	North, revisited (extended synthesis: from transaction costs to switching risk).
D	Related work (compact literature map: North/Arthur/Williamson; cascades/EWS; networks/modularity).
E	Post-1945 prospective (DS2 note; no cycle claim; thresholds/leakage channels).
F	Leakage in RSP → next SW severity (empirical check; RSP2 anomaly).

Methods discipline and falsifiers (summary).

Discipline:

1. *Pre-whitening (L2)*: remove linear trend + AR(1); compute AC1/variance on residuals; z-score; composite = mean of z-scores on a 5-year end-year grid.
2. *Cadence/Severity*: estimate T_c from C1–C3 (C4 held out); fit $log S = c - p log(T_c - t)$ with T_c fixed.
3. *Ordering*: corr(κ_t, EWS_{t+L}), L = 0…40 y in 5-y steps; block-bootstrap 95% CIs (5-y blocks); conservative placebos (pseudo-schedules; per-series + window sensitivity).
4. *ABM (seed-locked)*: thresholds, topology, and update rule fixed; only post-onset ΔJ increments *J.* Onsets [54, 99, 134, 159]; intervals [45, 35, 25]; J at onsets [0.00, 0.11, 0.22, 0.33].
5. *Mapping*: $S = \rho \cdot Pop \cdot (J_0 + J)$ fitted on C1–C3; C4 held out; one historical uplift k at SW4 (WWII theatre fusion, 1941); no-tune EWS amplifier variant in SI.

Falsifiers:

(i) No shared T_c (cadence and severity do not converge to the same mid-1940s window).
(ii) κ does not lead the pre-whitened composite EWS in a broad 0–40 y scan, or placebo ramps routinely match/exceed observed rises.
(iii) Fixed-rule ABM fails to accelerate under plausible ω/θsettings or on simple alternative topologies (e.g., modest modularity).
(iv) The C1–C3 mapping collapses (wrong ranking or need for ad-hoc, cycle-specific tuning).

Data, notation, and reproducibility (at a glance).

- *Key datasets (SI cites files by name):* Levy SW BCD; OWID Europe population; κ (frozen, conservative L1 carrier on 5-y grid); pre-whitened L2 proxies (MILEX AC1; alliance-churn variance); COW alliance dyad-year (exposure matrices); ABM outputs (series/onsets).
- *Notation (see Glossary):* h(t) hazard/load; K(t) common knowledge (on-fraction); J actuation/rails; ΔJ ratchet; ω neighbour coupling; θ_i thresholds; κ L1 carrier; EWS composite (L2).
- *Provenance:* ABM seed = 20250912; onset rule ≥ ⌈0.83N⌉; refractory = 5; slope = 0.011; ΔJ = 0.11. Mapping parameters from C1–C3: ρ and J_0 as reported in SI S6; SW4 uplift marker k in figure/caption only (no re-fit).
- *Source data and scripts:* All CSVs and scripts referenced per section; a consolidated Data and artefact manifest closes the SI.

How to read this SI.

For a one-page map linking claims to figures/tables across S1–S7, see Appendix A (Evidence spine). For an extended synthesis linking North to the switching-risk mechanism, see Appendix C. Implementation templates and operating rules are collected in Appendix B (Design companion); a compact literature map is in Appendix D (Related work); a post-1945 note is in Appendix E (DS2); and the empirical leakage check is in Appendix F (RSP → next-SW).

Glossary (terms and symbols):

Terms:

Term / Symbol	Meaning	Where it appears / why it matters
L1 (carrier)	The slow, structural interaction layer of the system (e.g., connectivity, capacity, interdependence). In the paper we use a frozen index, κ, as a conservative proxy.	Drives long-run "pressure" toward criticality; we show κ leads behavioural signals by ~25 years (ordering).
L2 (dynamics)	Faster social–political processes that play out on L1 (alliances, mobilizations, expenditures, crises).	Source of early-warning texture (EWS) and visible churn.
κ (kappa)	A conservative, frozen measure of L1 connectedness/capacity on a 5-year grid.	Our structural driver; leads L2 EWS (ordering footprint).
EWS (early-warning signal)	Behavioural "texture" constructed from pre-whitened series: e.g., lag-1 autocorrelation (AC1) of aggregate MILEX, variance of alliance churn; standardized and averaged to a composite.	Rises before onsets; used to test ordering and band-entry.
Pre-whitening	Remove linear trend and AR(1) structure, then compute AC1/variance on residuals; z-score before compositing.	Standardizes L2 so EWS reflects approach to transition rather than drift.
Hazard band	A narrow regime where option sets are compressed, actuation is cheap, and correlated thresholds make synchronized "yes/no" decisions likely.	Where decentralized cabinet decisions effectively "line up."
(Unseen) Switch	The emergent tipping surface created when institutions compress options, build common knowledge, and lower actuation cost—so many thresholds cross together once the system enters the hazard band. Not a literal device.	Explains how decentralized decisions align without a metronome. Seen via phase-lock (band-entry), day-scale latencies, and fast cascades in 1914/1939. Central to our mechanism framing.
Alignment Engine (mechanism)	The four locks that produce synchrony: payoff (off-corridor options are costly), common-knowledge (public signals), actuation (protocols lower latency), threshold (correlated decision points).	Explains why decisions align without a metronome or central planner.
Actuation / rails (J)	Institutional "automaticity"—treaties, protocols, mobilization scripts—that reduce decision latency. Ratchet = J increases after each system reset.	Rising J shortens the next interval; core to the accelerating cadence.
Latency	The time lag between a trigger (e.g., ultimatum, invasion) and a cabinet's mobilization/war declaration. Low latency = cheap/fast actuation.	Measured in days for 1914/1939; shows that rails (treaties, timetables) enable rapid, synchronized action. Increases in J (ratchet) tend to shorten latency across punctuations.
Ratchet (ΔJ)	The post-event thickening of actuation rails (standing plans become faster/easier next cycle).	Small ΔJ changes have big effects on cadence and severity.
Drift (of h(t))	Slow increase in structural hazard (L1 pressure); modelled as a linear drift plus small noise.	The regular "push" that moves the system toward the band.
Safe leakage (leak)	Early, reversible venting of hazard (e.g., de-confliction tracks, arms-control queuing caps). In	Delays band entry, stretches cycles, and softens events.

Term / Symbol	Meaning	Where it appears / why it matters
	the model: reduce effective drift by a small constant.	
Neighbour coupling (ω)	Strength of cross-cabinet influence (doctrinal alignment, interoperability).	High ω → lock-step; lowering ω desynchronizes decisions.
Threshold heterogeneity (θ_i spread)	Diversity in national decision thresholds (doctrine, law, domestic checks).	Greater spread breaks cascades and reduces synchrony.
Backbone (activation)	Union of shortest activation paths (defence/guarantees) from a seed to each great power (1914, 1939).	Shows rails were in place before the trigger; enables fast cascades.
Tc (critical time)	Shared terminal window to which cycle lengths compress and severity diverges.	Empirical anchor for cadence and severity fits (≈ 1946).
RSP leakage	Share of casualties vented in non-systemic wars within a cycle; falls sharply in late cycles.	Explains why late cycles arrive "on time" and with higher severity.
Phase-lock / band-entry	Tests for clustering of onsets within a narrow phase range of L1/L2 band signals.	Used prospectively (DS2) to avoid war-prediction claims.
Block bootstrap	Resampling in blocks to keep time-dependence; used to add CIs on lag-curves.	Uncertainty wrapper for ordering footprint.
Population × rails mapping	$\hat{S}_c = \rho \cdot Pop_c \cdot (J_0 + J_c)$ links severity to capacity (population) and actuation (rails).	Explains amplitude with minimal parameters; WWII uplift via a single globalization factor.
Globalization factor (k)	One-off multiplier for SW4 reflecting 1941 theatre fusion (capacity expansion).	Aligns WWII amplitude without inventing a new mechanism.
J-trace	Time series of the actuation index J (with decay if reverse-ratchet is adopted).	Monitoring tool for automaticity/latency in institutions.
Bulkheads / lattice	Moving from a complete graph to a modular lattice (weak bridges), so contagion takes longer paths.	Reduces immediate system-wide cascades; a network design lever.
Path dependence and lock-in	Prior institutional choices create increasing returns that narrow future options; organizations adapt to (and reinforce) the rules, so the system tends to stay on its rails.	Interprets the ratchet (J↑ after resets) and the shortening of cycles; explains why late cycles show collapsed RSP leakage and faster actuation, consistent with North's institutional dynamics.
RSP (Relatively Stable Period)	The between-systemic wars segment of a cycle. Often contains non-systemic wars that maintain the status quo; can act as leakage that vents tension. RSP + SW = cycle.	Used in cadence fits (S1), leakage analysis (S2). Late cycles show low leakage (RSP share → ~0), so more tension arrives at the switch; helping explain acceleration and severity growth.
Systemic wars (SWs)	System-restructuring wars in which most/all great powers participate; they reset the order and thicken the actuation rails (J↑) for the next cycle. RSP + SW = cycle.	Define cycle endpoints (S1–S2). Their amplitude is mapped by Population × (J_0+J) (S6). SW4's uplift reflects globalization (1941 fusion), not a different mechanism.
Great Power (system)	Following Levy: states with major, security-relevant roles—distinguished by military power, interests, behaviour, how others perceive them, and some formal criteria. From 1495–1945 there were typically 4–8.	Determines who counts in systemic wars and the network/ABM composition (baseline N=6). "Systemic" implies GP participation, making rails & latency especially consequential.

Symbols:

This table covers all symbols, variables, and parameters used in Sections 1–7. Columns: Symbol/Name, Category, Definition / Meaning, Units / Scale, Typical / Baseline, Role / Where Used.

Symbol / Name	Category	Definition / Meaning	Units / Scale	Typical / Baseline	Role / Where Used
L_c	Geometry	Cycle c RSP length	years	C1–C4: 195, 167, 99, 27	Cadence fit $L = a(T_c - t)$ (S1)
T_c	Geometry	Critical time (terminal window)	year	≈1946	Shared horizon for cadence & severity (S1–S2)
t	Geometry	Systemic-war onset year	year	1618, 1792, 1914, 1939	Inputs for S1–S2
$d = T_c - t$	Geometry	Distance to terminal window	years		Regressor in S1–S2
S_c	Geometry/Severity	Systemic-war severity (Levy BCD)	persons		S2 power law; S6 mapping target
p	Severity exponent	Slope in log S = c - p log d	dimensionless	≈ 0.52 (95% CI ≈ [0.28, 0.76]), R^2 ≈ 0.98	S2 slope result
RSP/Total (leakage)	Geometry	RSP share of total cycle casualties	0–1	C1≈0.31–0.34; C2≈0.665; C3≈0.082; C4≈0.004	Interprets C2 anomaly; late collapse (S2)
$\Delta V_i(t)$	ABM	Cabinet i's decision margin	arbitrary		If >0 ⇒ x_i=1 (ON)
$\alpha, \beta, \gamma, \omega,$	ABM weights	Sensitivities to h, K, J, n_i	dimensionless	1.0, 0.35, 1.0, 1.0	ABM parameters (S5, S7)
h(t)	ABM (hazard)	Slow L1 structural pressure	steps	slope 0.011	Drives approach to band (S5)
K(t)	ABM (common knowledge)	Fraction ON (system visibility)	0–1		Synchrony channel
J(t)	ABM (actuation rails)	Institutional automaticity / latency reduction	index	rises post-onset	Ratchet variable (S5–S7)
ΔJ	ABM ratchet step	Post-onset thickening of rails	+index	+0.11 (slow-auto 0.06)	Shortens next cycle; lever in S7
n_i(t)	ABM neighbour term	Normalized mean of neighbours' ON states	0–1		Contagion channel
θ_i	ABM thresholds	Cabinet activation thresholds	$U[\mu \pm spread]$U	μ ≈ 0.65; spread ±0.07 (SI ±0.05–0.12)	Heterogeneity knob

Symbol / Name	Category	Definition / Meaning	Units / Scale	Typical / Baseline	Role / Where Used
$X_i(t)$	ABM state	1 if ON, else 0	0/1		System configuration
N	ABM size	Number of agents (GPs)	integer	6 (SI 4, 8)	Robustness (SI)
need	Onset quorum	Systemic onset threshold	integer	$\lceil 0.83N \rceil$ (≈ ≥ 5/6)	Onset rule (S5–S7)
REF	Refractory	Min gap between onsets	steps	5	Avoids double counting
σ	ABM common noise	Shared noise in h(t)	0–0.02	SI grid: 0, .005, .01, .02	Phase diagram (SI)
leak	ABM safe leakage	Early bleed: $h \leftarrow max(0, h + slope - leak)$	≈ 0.003–0.005	0.004 (S7)	Delays band entry; intervals ↑
J_{onset}	ABM readout	J at each event onset	index	[0.00, 0.11, 0.22, 0.33]	Maps to severity (S6)
κ**(t)**	Ordering (L1)	Structural carrier (frozen connectivity/capacity)	std units (z)		Leads EWS by ~25y (S4)
EWS_t	Ordering (L2)	Composite pre-whitened early-warning signal	std units (z)	mean of Z(MILEX AC1), Z(alliance var)	Behavioural texture (S4)
AC1	EWS component	Pre-whitened lag-1 autocorrelation of MILEX	unitless		EWS component (S4)
Var (alliance churn)	EWS component	Pre-whitened variance of alliance churn	unitless		EWS component (S4)
Lag peak	Ordering lead	Argmax corr (κ_t, EWS_{t+L})	years	~25 (composite corr ≈ 0.56)	Footprint of mechanism (S4)
Placebo p	Ordering placebo	p(null ≥ observed pre-onset slope)	p-value	composite ≈ 0.44 (per-series also in SI)	Conservative SI check (S4)
$S_c = \rho \cdot Pop_c \cdot (J_0 + J_c)$	Mapping	Severity predicted from rails × capacity	persons		S6 mapping formula
ρ	Mapping scale	Scale factor in mapping	persons/(pop·J)	≈ 0.01009	Fitted on C1–C3 (S6)
J_0	Baseline rails	Baseline actuation in mapping	index	≈ 1.465	Fitted on C1–C3 (S6)
J_c	Event rails	J at onset for cycle c	index	[0.00, .11, .22, .33]	From ABM (S6)
Pop_c	Capacity	Europe population at onset (OWID)	persons	interpolated	Capacity term (S6)
k	SW4 globalization	WWII 1941 theatre fusion factor	multiplier	≈ 1.366	One-off uplift for SW4 (S6)
g_{ews}	SI amplifier	EWS-derived global amplifier (no free k)	multiplier	data-derived	SI alternative to k (S6 SI)

Symbol / Name	Category	Definition / Meaning	Units / Scale	Typical / Baseline	Role / Where Used
share (BCD/Pop)	Severity share	BCD as share of population	%	~2–3% (C2 low)	Contextual check (S2/S6)
Backbone	Case studies	Union of shortest seed → GP activation paths	graph	1914 (seed A-H), 1939 (seed Germany)	Pre-trigger rails (S3 rev)
D / G / E	Edge types	Defence / Guarantee / Entente	categorical		Edge labels in S3 rev
Latency	Case metric	Trigger → action delay	days	median ≈ 2–3	Fast actuation (S3 rev)
Seed	Case root	Initial node for backbone	code	1914: 300; 1939: 255	Cascade root (S3 rev)
ABM seed	Const	Random seed (reproducibility)	value	20250912	All ABM runs
Slope	Const	Hazard drift per step	value	0.011	ABM
ΔJ (alt)	Const	Slower automaticity setting	value	0.06	S7 intervention
Onset rule	Const	Quorum & refractory	rule	$\geq \lceil 0.83N \rceil$, REF=5	ABM
Lag scan	Const	κ → EWS lags scanned	years	0–40 (5-year steps)	S4
Placebo window	Const	Pre-onset window length	years	25 (SI ~10, ~15)	S4, SI
Start date	Const	European system start used	year	1453 & 1495	Sensitivity (S1–S2)
Population source	Const	Europe population series	dataset	OWID	Mapping (S6)

Notes:

• *Pre-whitening:* fit $y_t = \mu + \tau t + \phi y_{t-1} + \varepsilon_t$; compute AC1/Var on residuals; z-score before compositing.
• *Phase diagram:* acceleration region documented over (ω, σ); modular lattices (bulkheads) shift the desync boundary to higher ω.
• *Leakage:* designed early and reversible; do not rely on late leak near the band apex.
• *SW4 uplift:* treated as one-off capacity fusion (k) in main; EWS-derived g_{ews} shown in SI without tuning.

S0. Methods discipline and falsifiers

What this section does. S0 states the rules of the road—how data are constructed, what is estimated where, the acceptance checks we require, and the falsifiers that would break the mechanism. Each later SI section (S1–S7) implements these rules with files and figures.

S0.1 Data and preprocessing (L1 κ; L2 EWS; grids):

- *Carrier (L1, κ; 5-year grid).* A frozen, conservative index (D6_kappa_verified.csv) built from: (i) ENS/state-lattice concentration (population-weighted, Europe-wide), (ii) rail km per capita (GP-weighted), (iii) urbanization share (GP-weighted). κ is z-scored and not re-tuned across analyses; it deliberately *excludes* information-intensity (telegraph/telephone), making it a lower bound.

- *Behavioural texture (L2, pre-whitened EWS; 5-year grid).* Two proxies: MILEX AC1 and alliance-churn variance. For each series we remove a linear trend and fit AR(1) (AIC-selected); compute the statistic (AC1/variance) on *residuals*; z-score; composite = mean of z-scores.
- *Cadence/severity geometry.* Systemic-war onsets t = {1618, 1792, 1914, 1939}; cycle lengths L; severities S (Levy BCD).
- *Cabinet ladders and exposure.* Matrices from COW alliances (defence "D") plus documented guarantees "G" (e.g., Belgian 1839; Anglo-Polish 1939); dated cabinet acts for 1914 (main) and 1939 (SI).
- *ABM artefacts.* abm_series.csv, abm_onsets.csv (seed-locked).

S0.2 Estimators and windows (what is tested where):

- S1 (Cadence → T_c). Fit $L = a(T_c - t)$ on C1–C3, C4 held-out. Grid $T_c \in [1938,1946]$ (0.01-y steps); OLS through origin on d = $T_c - t$. Bootstrap (2,000×) CIs.
- S2 (Severity with shared T_c). With T_c fixed from S1, $logS = c - plog(Tc_{-t})$ on C1–C4; report $\hat{p}$, small-sample t CI (df=2), and R^2.
- *S3 (Rails and latency).* Exposure matrices (D/G); ladders from trigger to dated cabinet acts (1914 main; 1939 compact SI).
- *S4 (Ordering).* Corr(κ_t, EWS_{t+L}) for L=0,5, ... ,40 years; block-bootstrap 95% CIs (5-year blocks). Placebos: 10k pseudo-schedules; series-specific slopes; ~10y/~15y windows; conservative composite.
- *S5 (ABM sufficiency).* Six cabinets on a complete graph; synchronous update $\Delta V_i = \alpha h + \beta K + \gamma J + \omega n_i - \theta_i, x_i(t+1) = 1\{\Delta V_i > 0\}$. Thresholds θ_i~ U[0.58,0.72] fixed; onset ≥ ⌈0.83N⌉; refractory = 5. Only post-onset ΔJ increments J. Baseline: slope = 0.011; ΔJ = 0.11; $\alpha = \gamma = 1.0, \beta$ = 0.35, $\omega = 1.0$; seed=20250912.
- *S6 (Mapping).* $S_c = \rho \cdot Pop_c \cdot (J_0 + J_c)$ fitted on C1–C3; C4 held-out. Report $\widehat{\rho}$, $\widehat{J_0}$. Display a single historical SW4 uplift marker k (WWII theatre fusion, 1941). SI also shows a no-tune EWS-derived amplifier variant.
- *S7 (Interventions).* Levers: safe leakage (reduces effective h pre-band), reverse-ratchet (lower ΔJ), decorrelation (lower ω/greater θ-spread). Readouts: Δ mean of first-three intervals; Δ total mapped severity (no re-fit of ρ, J_0).

S0.3 Acceptance checks (must pass):

- *Shared horizon.* S1 finds T_c in the mid-1940s with C4 held-out; S2 (with $\boldsymbol{T_c}$ fixed) yields $\hat{p}$ > 0, CI excluding 0, high R^2.
- *Year alignment.* κ and the composite EWS share the same 5-year end-year grid.
- *Ordering.* κ → EWS lag curve peaks near ~25 y; CI above zero around the peak; per-series placebos show elevated pre-onset ramps; composite placebo is conservative.
- *Rails and latency.* Exposure is spanning *pre-trigger*; ladders cluster at hours–days scale (1914 main).
- *ABM sufficiency.* Onsets [54, 99, 134, 159]; intervals [45, 35, 25]; J at onsets [0.00, 0.11, 0.22, 0.33] under fixed micro-rules; $N \in \{4, 8\}$ robustness preserves acceleration.
- *Mapping.* $\widehat{\rho}$, $\widehat{J_0}$ from C1–C3 reproduce rising shares; baseline under-predicts SW4; a *single* interpretable k aligns SW4; SI variant shows no-tune amplifier.

- *Interventions.* Safe leakage stretches cadence with flat ↓ severity; ΔJ ↓ stretches cadence and lowers mapped severity; decorrelation marks a regime boundary (becomes potent once below a threshold or with small idiosyncratic noise).

S0.4 Falsifiers (any one breaks the mechanism-to-design link):

1. *No shared* T_c: cadence and severity do not converge to the same mid-1940s window.
2. *No ordering*: κ fails to lead the pre-whitened composite in a broad 0–40 y scan, or placebo ramps routinely match/exceed observed rises.
3. *No sufficiency*: the fixed-rule ABM fails to accelerate for plausible ω, modest \theta-spread, or simple alternative topologies (i.e., acceleration is a tuning artefact).
4. *Mapping collapses*: C1–C3 fit fails (wrong ranking/magnitudes) or requires cycle-by-cycle tuning; uplift needs ad-hoc parameters beyond a single k.

S0.5 Reproducibility checklist (one pass):

1. Verify grids (κ, EWS 5-year end-years; onsets for geometry; dated cabinets for ladders).
2. Recreate S1–S2 fits (C1–C3 cadence with C4 held-out; severity with $\boldsymbol{T_c}$ fixed).
3. Run ordering scan + block-bootstrap CIs; export per-series and composite placebos.
4. Execute ABM (seed = 20250912) → confirm onsets, intervals, J.
5. Fit S6 mapping on C1–C3; export $\widehat{\rho}$, $\widehat{J_0}$, hold-out SW4; overlay k; render SI variants.
6. Run S7 levers; report Δ intervals and Δ total mapped severity.
7. Compile source CSVs + figures into the manifest.

S0.6 Notation (quick refresher; full Glossary at end)

h(t) hazard/load; K(t) common knowledge (on-fraction); J actuation/rails; ΔJ ratchet; ω neighbour coupling; θ_i thresholds; κ L1 carrier; EWS = pre-whitened composite (L2). T_c = critical time (mid-1940s).

All section-specific files and figures are listed at the end of each SI section and again in the final manifest.

S1. Cadence → T_c (fit on C1–C3; C4 held-out)

What we test and why:

We ask whether cycle lengths L shrink linearly with distance to a critical time T_c. If so, we can estimate T_c without using the last cycle and then check out-of-sample. This anchors the "terminal window" used across the paper (shared by severity in S2).

S1.1 Inputs:

- *Systemic-war onsets (C1–C4)*: t = {1618, 1792, 1914, 1939}.
- Cycle lengths: L= {195, 167, 99, 27} years.
- Source data: tc_cadence_stats.csv (contains observed L, best T_c, fitted $\hat{a}$, predicted $\hat{L}$, residuals).

S1.2 Estimator and procedure:

- Model. L = a($T_c - t$).
- Discipline. Fit on C1–C3 only; C4 held-out.
- Search. Grid T_c ∈ [1938, 1946] in 0.01-year steps; for each candidate T_c, estimate a by OLS through the origin on $d = T_c - t$. Choose T_c minimizing C1–C3 SSE.
- Uncertainty. Bootstrap (2,000×) over C1–C3 to form a 95% interval for T_c. Jackknife (leave-one-out over C1–C3) for stability.
- Diagnostics. Report C1–C3 fit and the C4 hold-out residual (observed minus predicted).

S1.3 Results:

- *Best T_c:* 1946.00 with 95% CI 1938.00–1946.00.
- *Slope $\hat{a}$:* 0.702 (years per year).
- *Hold-out (C4)*: residual +22.09 y (conservative miss expected in a small-n linear form).
- *Stability*. Each jackknife refit (drop C1/C2/C3) returns T_c =1946.0 at grid resolution. Using 1453 instead of 1495 as system start changes $\hat{a}$ modestly but leaves T_c in the same window.

(Exact values are in tc_cadence_stats.csv; jackknife and start-date sensitivity tables are provided below.)

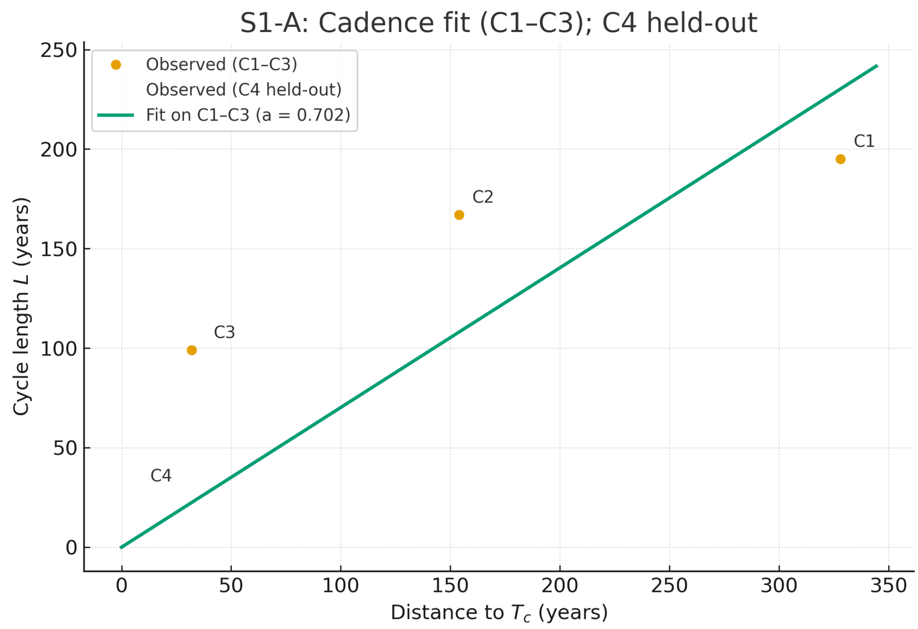

Figure S1-A: Cadence fit (C1–C3); C4 held-out. *Cycle length L versus distance to the terminal window d =* $T_c - t$ *with* T_c *estimated from cadence. Filled circles are C1–C3; the open circle is C4 (held-out). Line shows the C1–C3 fit through the origin L=* $\hat{a}$, *d (here* $\hat{a}$ *= 0.702); best* T_c *≈1946 (95% CI 1938–1946). Source data: tc_cadence_stats.csv.*

S1.4 Acceptance checks and falsifiers:

- *Checks.* (i) T_c lies in the mid-1940s with C4 held-out; (ii) jackknife stable; (iii) start-date sensitivity does not move T_c.
- *Falsifiers.* (i) No T_c near the WWII window; (ii) strong jackknife instability; (iii) C4 hold-out fails far outside the expected residual band.

S1.5 Sensitivities (reported):

- *Start date (1453 vs 1495)*. Re-run of S1 shows unchanged T_c (slope shifts, as expected).
  File: tc_cadence_1453_1495_sensitivity.csv.
- *Jackknife (C1/C2/C3 leave-one-out)*. All refits return T_c = 1946.0; per-fit $\hat{a}$, predictions, and residuals are tabulated.
  File: tc_cadence_jackknife.csv.

S1.6 Outputs (figures and tables):

- *Figure (main text reference)*: fig1_shared_tc_convergence.png (cadence and severity convergence; cadence part derived from S1).
- *Figure S1-A (SI)*: fig_tc_cadence.png — cadence fit on C1–C3 with C4 held-out (distance to T_c on x-axis).
- *Extended Data (ED1-A)*: ed1a_cadence_linear.png — linear view L vs d = $T_c - t$ (C1–C3 fit; C4 open marker).
- *Table (CSV)*: tc_cadence_stats.csv — observed L, best T_c, $\hat{a}, \hat{L}$, residuals.

S1.7 Reproducibility steps:

1. Load C1–C3 L and onsets; set the T_c grid [1938,1946].
2. For each T_c, compute d = $T_c - t$ and estimate a by OLS through the origin; retain C1–C3 SSE.
3. Choose best T_c; compute C1–C3 predictions and the C4 hold-out prediction/residual.
4. Bootstrap (2,000×) on C1–C3 to form the T_c 95% interval; run the three jackknife fits.
5. Export tc_cadence_stats.csv; render fig_tc_cadence.png and ED1-A.

Cross-references:

- Main text R1 (cadence to T_c); Figure 1 (shared-T_c convergence).
- S2 uses the same T_c to fit severity.
- S4 (ordering) and S5–S6 (ABM sufficiency and mapping) use the mid-1940s terminal window inferred here.

S2. Severity vs distance (shared T_c; diagnostics and linearized view)

What we test and why:

With T_c fixed from S1 (cadence), does systemic-war severity S follow a power-law in distance to the same horizon? A positive slope on log S vs log d (with $d = T_c - t$) confirms that compression (cadence) and divergence (severity) share one terminal window.

S2.1 Inputs:

- *Systemic-war onsets (C1–C4)*: t = {1618, 1792, 1914, 1939}.
- *Severities*: Levy systemic-war BCD S for C1–C4.
- *Critical time:* T_c fixed to S1's estimate (T_c = 1946.00; 95% CI 1938.00–1946.00).
- *Source data*: tc_logS_stats.csv (onset years, distances, *logS*, $\widehat{\log S}$, residuals).

S2.2 Estimator and procedure:

- *Model*. $\log S = c - p \log d$, $d = T_c - t$.
- *Discipline*. T_c is not re-estimated here; it is fixed from S1. Fit by OLS with intercept on C1–C4.
- *Diagnostics*. Report $\hat{p}$ with small-sample 95% CI (Student-t, df = 2), R^2, and per-cycle residuals.

S2.3 Results:

- *Slope* $\hat{p}$: 0.518 (95% CI 0.275–0.761).
- *Fit*: $R^2 \approx 0.977$.
- *Interpretation*. With T_c fixed to cadence, severity rises systematically as the terminal window is approached. C2's known idiosyncrasy affects amplitude but not the shared-T_c geometry (residuals tabulated in the CSV).

(Exact fields appear in tc_logS_stats.csv: log_distance, log_S, log_S_hat, and residual_log.)

A cross-cycle leakage check consistent with this geometry—high RSP leakage ↔ lower next-SW severity (notably RSP2→SW2)—is provided in Appendix F.

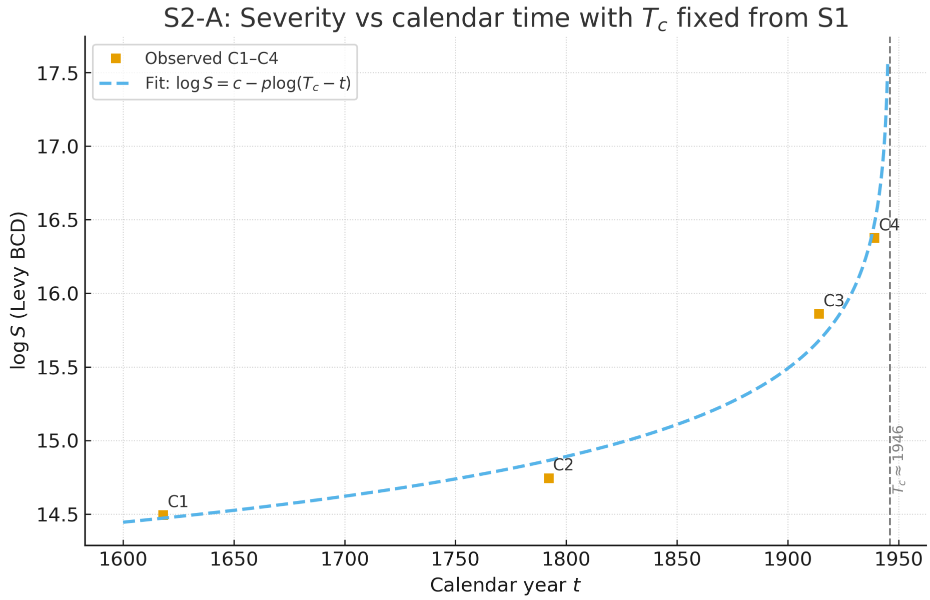

Figure S2-A: Severity vs calendar time with T_c fixed from S1. Observed log S (Levy BCD) at the four onsets plotted against calendar year; dashed curve is the fitted law log S = c $-\hat{p}\log(T_c - t)$ with T_c fixed to the cadence estimate (here $\hat{p}$ = 0.518, R^2 = 0.977). Vertical line marks T_c ≈ 1946. Source data: tc_logS_stats.csv.

S2.4 Acceptance checks and falsifiers:

- *Checks*. (i) $\hat{p} > 0$ with CI excluding 0; (ii) high R^2; (iii) residuals without pathological leverage.
- *Falsifier*. Slope indistinguishable from 0 when T_c is fixed to S1.

S2.5 Sensitivities (reported):

- Start-date sensitivity (1453 vs 1495). Changing the system start adjusts the cadence slope a (S1) but does not move T_c; the severity fit therefore remains anchored to the same T_c.
- Jackknife (by cadence). Because T_c is fixed from S1's jackknife-stable result, the severity estimate inherits that stability (see S1 and tc_cadence_jackknife.csv).

S2.6 Outputs (figures and tables):

- *Figure (main text reference)*: fig1_shared_tc_convergence.png — *Cadence (black) and severity (red) converge on the same T_c* (cadence part from S1; severity overlay from S2).

- *Figure S2-A (SI)*: fig_tc_logS.png — log S vs calendar year with the fitted law drawn using T_c from S1.
- *Extended Data (ED1-B)*: ed1b_severity_loglog.png — linearized diagnostic: log S vs log d with OLS line and C1–C4 labels.
- *Table (CSV)*: tc_logS_stats.csv — distances, log S, fitted values, and residuals.

S2.7 Reproducibility steps:

1. Load C1–C4 severities (Levy BCD) and S1's T_c.
2. Compute d = $T_c - t$; fit $log\ S = c - plog\ d$ (OLS with intercept).
3. Export $\hat{p}$, 95% small-sample CI, R^2, residuals (tc_logS_stats.csv).
4. Render fig_tc_logS.png and the ED diagnostic ed1b_severity_loglog.png.

Cross-references:

- S1 supplies the fixed T_c and its CI.
- S4 (ordering) motivates why severity's divergence should lag structural tightening.
- S6 uses the ABM read-out J with population to map to observed severity, showing SW4 uplift with a single historical marker k.

S3. Protocol exposure and decision latency (1914 main; 1939 SI)

What we test and why:

We test whether institutions furnished pre-existing rails - formal commitments and guarantees - that (i) create a spanning activation backbone before the trigger, and (ii) compress decision latency to hours–days once the trigger fires. This is the micro-behavioural footprint of the mechanism: rails manufacture common knowledge and cut actuation time.

S3.1 Inputs:

- *Alliance exposure (COW dyad-year)*. Formal alliances coded by type; we add documented legal guarantees not always tagged as "defense" in COW (e.g., Belgian neutrality, 1839; Anglo-Polish Agreement, 25 Aug 1939).
  Files: adjmatrix_1914.csv, adjmatrix_1939.csv (state × state, entries ∈ {D, G, 0});
  backbone_edges_1914.csv, backbone_edges_1939.csv.
- *Cabinet-dated acts (latency)*. Triggers and first formal acts by cabinets (mobilizations, declarations, initiating hostilities).
  1914 trigger: Austria-Hungary → Serbia (28 Jul 1914).
  1939 trigger: Germany → Poland (1 Sep 1939).
  Files: latency_1914.csv, latency_1939.csv.
- *Summary metrics*. GP reachability from the seed; max k-core; seed out-degree; median/percentile latencies.
  File: protocol_metrics_rev.csv.

S3.2 Construction and procedure:

- *Exposure matrices (D/G)*. Build state × state matrices for great powers and bridging nodes (Belgium, Serbia, Poland, etc.). Mark D = defence alliance; G = legal guarantee/obligation. Order rows/columns by seriation to reveal blocs and bridges.
- *Backbone*. Reduce to the minimum spanning activation routes (seed → other GPs) that are legally/organizationally feasible at the trigger.
- *Latency ladders*. Sort cabinet acts by days from trigger and plot a horizontal "ladder": one mark per cabinet event with label (e.g., "Russia GM (2d)").
- *Seeds*.
  1914: Austria-Hungary (→ Serbia).
  1939: Germany (→ Poland).

S3.3 Results:

- *Exposure is spanning pre-trigger*. In 1914, D/G ties produce a backbone from the trigger seed to the other great powers with a few legal bridges (e.g., Belgian and Serbian positions). In 1939, the network is denser and the relevant activation set smaller; the spanning route is present with fewer steps.
- *Latency compresses to hours–days*.
  1914 (main): decisions cluster at day-scale: Russia GM (2d), Germany declarations (4–6d), France GM (5d), UK declaration (7d).
  1939 (SI): UK and France at +2 days (tight cluster), consistent with thick rails and fewer necessary activations.

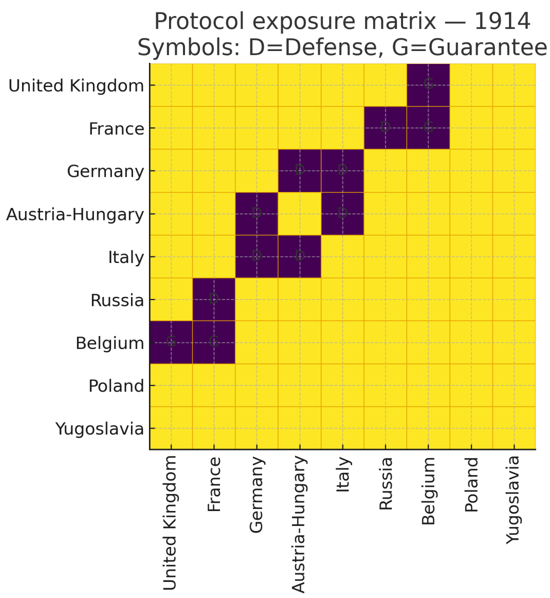

Figure S3-A: Protocol exposure matrix 1914. *Rows/columns list states; filled cells mark activation rails (D = defence alliance, G = legal guarantee). Ordering by seriation makes blocs and bridges (e.g., Belgium, Serbia) legible, showing a spanning activation route before the trigger. Source data: adjmatrix_1914.csv (guarantee overrides documented in SI text).*

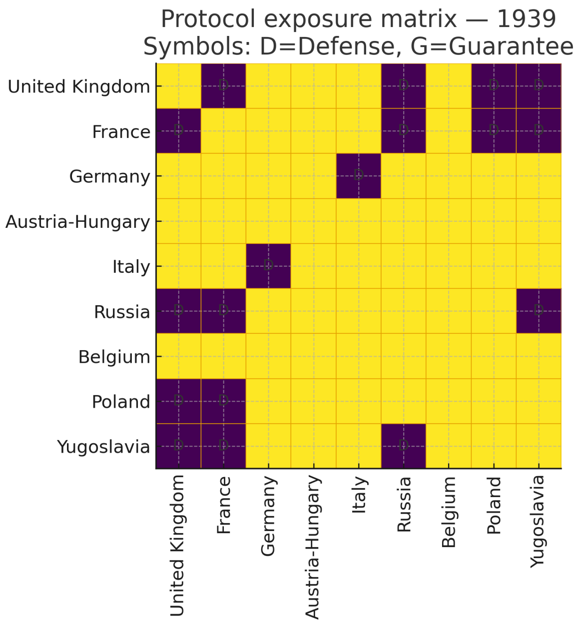

Figure S3-B: Protocol exposure matrix 1939. *As in S3-A, with bridges including Poland. The matrix shows a denser, more pre-wired order and a short spanning route from the trigger seed to other great powers (note the Anglo-Polish guarantee, 25 Aug 1939). Source data: adjmatrix_1939.csv (guarantee overrides documented in SI text).*

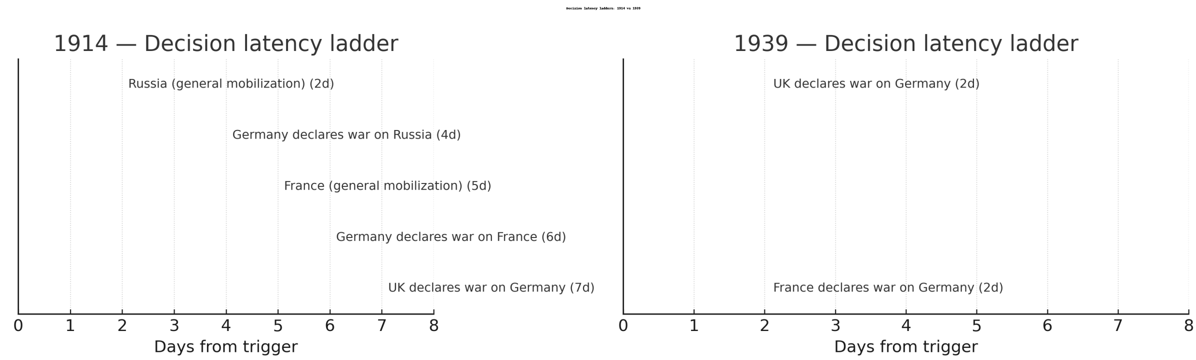

Figure S3-C: Decision latency ladders 1914 vs 1939. *Days from trigger to dated cabinet acts (1914: Austria–Hungary → Serbia, 28 Jul; 1939: Germany → Poland, 1 Sep). 1914 exhibits a stepped cascade (Russia GM 2d; Germany declarations 4–6d; France GM 5d; UK declaration 7d). 1939 shows a tight +2-day cluster (UK and France). The contrast illustrates the same rail-enabled logic—hours-to-days synchrony—under different activation sets. Source data: latency_1914.csv, latency_1939.csv.*

S3.4 Acceptance checks and falsifiers:

- *Checks*. (i) Exposure has a spanning route seed → other GPs *before* the trigger; (ii) cabinet acts cluster at hours–days (1914 main), with a compact 2-day cluster in 1939; (iii) summary metrics (reachability, k-core) consistent with a connected activation backbone.
- *Falsifiers*. (i) No spanning route pre-trigger; (ii) long, scattered latencies (multi-week) without documented exogenous obstacles; (iii) exposure depends on post-trigger additions rather than pre-existing rails.

S3.5 Figures and tables (this SI section):

- *Figure S3-A*. Protocol exposure matrix: *1914* (fig_1914_adjmatrix.png). Rows/columns are states; filled cells mark activation rails (D/G). Blocs and bridges are visible by seriation.
- *Figure S3-B*. Protocol exposure matrix*: 1939* (fig_1939_adjmatrix.png). As in S3-A, with bridges including Poland.
- *Figure S3-C*. Decision latency ladders: *1914 vs 1939 (two-panel)* (fig_latency_two_panel_1914_1939.png). Shared 0–8 day axis; 1914 shows a stepped cascade; 1939 shows a 2-day cluster, illustrating the same rail-enabled logic under fewer required activations.

S3.6 Notes on coding and overrides:

- *Guarantee overrides*. Some legally binding guarantees (G) are not captured by COW's "defence" tag but function as activation rails (e.g., *Treaty of London* 1839; *Anglo-Polish Agreement* 1939). We mark these as G and cite the archival/legal basis in the figure notes.
- *Precision asymmetry*. Early-modern cases (SW1/SW2) have coarser dating and slower coalition assembly; we therefore use 1914 as the main ladder and present 1939 as a compact consistency check. The SW trigger/entry table provides the broader historical context without forcing weak early-modern ladders.

S3.7 Reproducibility steps:

1. *Exposure*. Build adjmatrix_1914.csv / adjmatrix_1939.csv from COW + documented guarantees; render fig_1914_adjmatrix.png / fig_1939_adjmatrix.png.
2. *Ladders*. From latency_1914.csv / latency_1939.csv, compute days from trigger; render fig_1914_latency.png and the two-panel fig_latency_two_panel_1914_1939.png.
3. *Metrics*. Compute reachability, k-core, seed out-degree, and latency percentiles → protocol_metrics_rev.csv.
4. *SW trigger/entry table*. Compile SI_SW_triggers_GP_entries.csv and corresponding PNG pages (SI_table_SW1.png–SW4.png).
5. Verify that S3 checks pass (spanning pre-trigger; hours–days clustering); if any falsifier is tripped, flag in S0/S3 and halt downstream claims.

Cross-references:

- Main text R3 (rails and latency) cites Figure 1914 ladder and references this SI section.
- S4 (ordering) connects structural tightening to behavioural stickiness, explaining why ladders compress when κ is high.
- S6 maps actuation to observed harm; S7 dials (leakage, ΔJ, decorrelation) target the very rails S3 makes visible.

S4. Ordering: κ leads pre-whitened EWS (lag curve, CIs, and placebos)

What we test and why:

If institutions tighten the system's structural carrier first, then behavioural texture (stickiness, noise) should rise after structure. We therefore test whether κ (L1) leads a pre-whitened composite EWS (L2) by decades on the common 5-year grid, and we quantify uncertainty and null plausibility.

S4.1 Inputs:

- *Carrier (L1, 5-y grid)*: D6_kappa_verified.csv (frozen, conservative κ; see S0/SI Glossary).
- *Behavioural proxies (L2, 5-y grid; pre-whitened)*:
  - Military-expenditure AC1: D20_EWS_MILEX_AC1_prewhitened_5y.csv
  - Alliance-churn variance: D21_EWS_AllianceChurn_Var_prewhitened_5y.csv
    Each series is detrended + AR(1)-filtered; target statistic computed on residuals; z-scored; the composite is the mean of z-scores.
- *Outputs referenced here*:
  Lag curve with CI fig_ordering_lag_with_ci.png; placebo histograms fig_ordering_placebo_milex.png, fig_ordering_placebo_alliance.png, fig_ordering_placebo.png; tabulations ordering_lag_ci.csv, ordering_placebo_per_series.csv, ordering_placebo_sensitivity.csv, ordering_placebo_stats.csv.

S4.2 Estimator and procedure:

(A) Lag-correlation curve with block-bootstrap CIs:

- *Statistic*. For lags L ∈{0, 5, ..., 40\} years, compute $r(L) = corr(\kappa_t, EWS_{t+L})$, using the z-scored κ and the z-scored composite EWS; keep only years where both are observed.
- *Uncertainty*. Form 95% CIs at each L via block bootstrap (resampling paired (κ_t, EWS_{t+L}) in contiguous 5-year blocks to preserve local dependence).
- *Export*. ordering_lag_ci.csv (lag, point estimate, CI low/high); render fig_ordering_lag_with_ci.png.

(B) Placebos for pre-onset ramps (conservative composite):

- *Null construction*. Generate 10,000 pseudo-schedules of "onsets" that preserve serial structure; for each series compute the mean pre-onset slope in fixed windows (≈10y, ≈15y).
- *Per-series tests*. Report one-sided exceedance p for MILEX AC1 and Alliance variance separately.
- *Composite (conservative)*. Average z-scored series first, then repeat the slope test; this is intentionally conservative.
- *Export.* Per-series and composite histograms (fig_ordering_placebo), plus ordering_placebo_per_series.csv, ordering_placebo_sensitivity.csv, ordering_placebo_stats.csv.

S4.3 Results:

- *Lead (lag curve)*. The composite peaks near ~25 years (corr ≈ 0.56), with the 95% block-bootstrap band above zero around the peak. (fig_ordering_lag_with_ci.png, ordering_lag_ci.csv)
- *Placebos*. Series-specific pre-onset slopes lie on the upper tail of their nulls; the composite placebo is conservative (reported p does not over-state significance; see fig_ordering_placebo_milex.png, fig_ordering_placebo_alliance.png, fig_ordering_placebo.png). Window-length sensitivity (~10y/~15y) leaves conclusions unchanged.

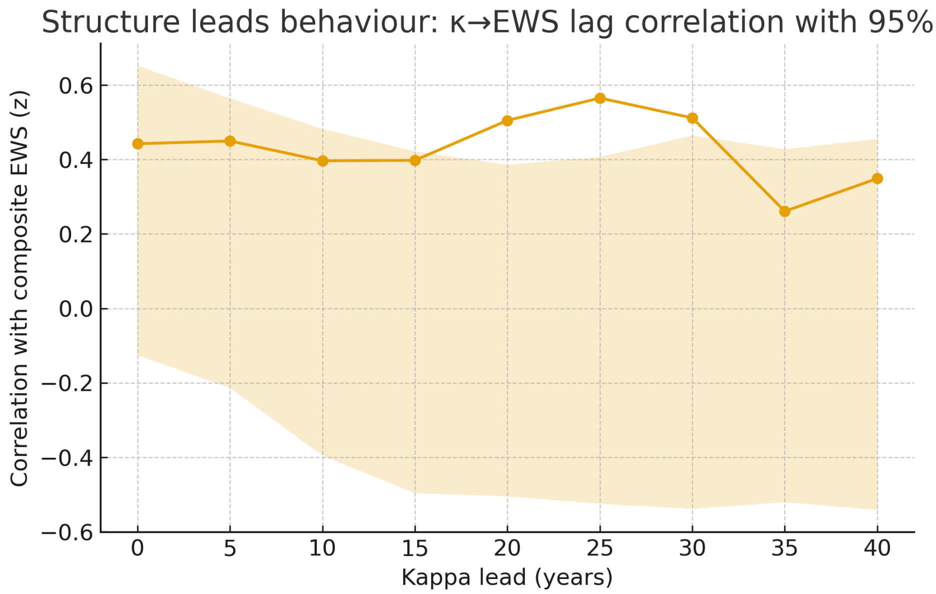

Figure S4-A: κ leads EWS (lag-correlation with 95% CI). *κ → EWS lag-correlation for κ leading by 0–40 years in 5-year steps. The curve peaks near ~25 years (corr ≈ 0.56); the shaded band is a 95% block-bootstrap CI (5-year blocks). EWS components are pre-whitened (detrend + AR(1) on residuals), z-scored, and averaged before the scan. Source data: ordering_lag_ci.csv.*

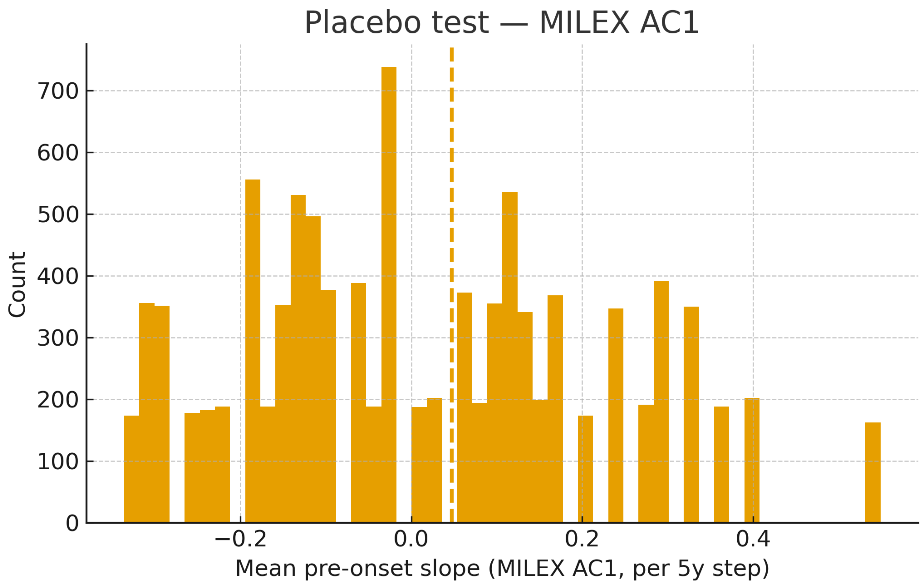

Figure S4-B: Placebo histogram (MILEX AC1). *Null distribution of pre-onset slope (≈10-year window) from 10,000 pseudo-schedules preserving serial structure; vertical line = observed slope. The observed MILEX AC1 ramp falls on the upper tail of its null, indicating a genuine rise ahead of onsets. Window sensitivity (≈10 / ≈15 y) in SI tables. Source data: ordering_placebo_per_series.csv, ordering_placebo_sensitivity.csv, ordering_placebo_stats.csv.*

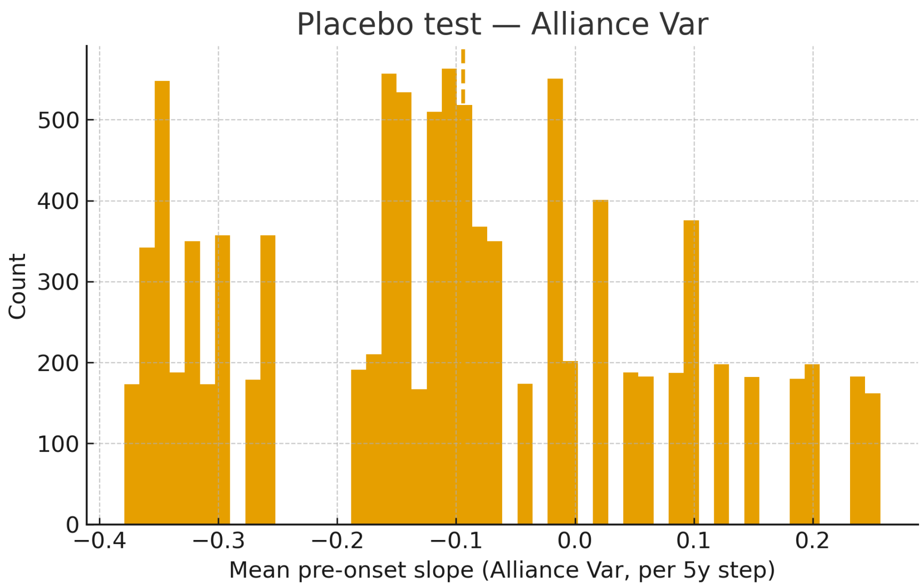

Figure S4-C: Placebo histogram (Alliance-churn variance). *As in S4-B, for the alliance-variance proxy. The observed pre-onset ramp lies on the upper tail of its null, consistent with a real rise before onsets. Window sensitivity confirmed in SI tables. Source data: ordering_placebo_per_series.csv, ordering_placebo_sensitivity.csv, ordering_placebo_stats.csv.*

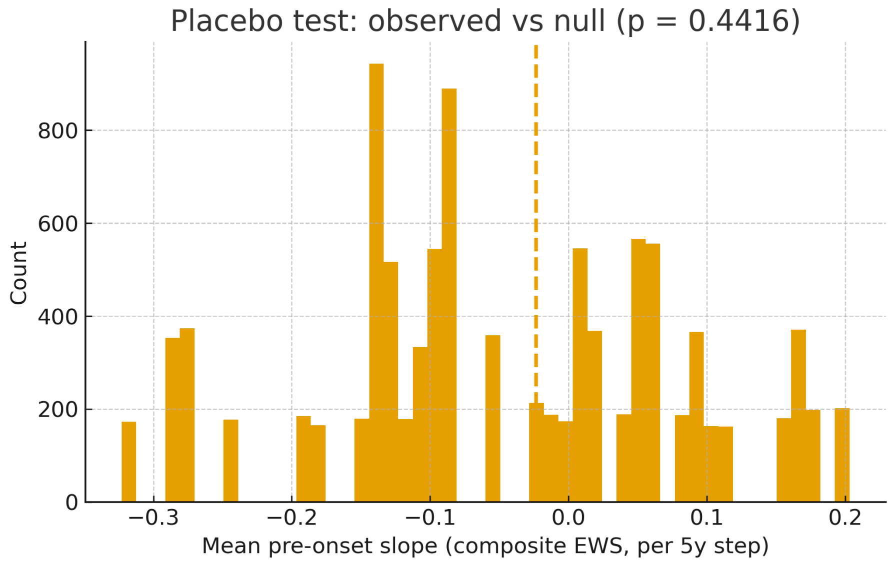

Figure S4-D: Placebo histogram (composite EWS — conservative). *Composite pre-onset slope (mean of pre-whitened, z-scored components) against its null from 10,000 pseudo-schedules; vertical line = observed composite slope. The composite test is conservative by design (averaging dampens series-specific ramps); exact one-sided p and window sensitivity are reported in SI tables. Source data: ordering_placebo_per_series.csv, ordering_placebo_sensitivity.csv, ordering_placebo_stats.csv.*

Interpretation. Structure tightens first; behavioural stickiness follows. This ordering is what the mechanism predicts if rails/connectedness compress latency and align thresholds before the system reaches the hazard band.

For a brief post-1945 (DS2) note on how the mechanism plausibly persists with altered thresholds and leakage channels, see Appendix E.

S4.4 Acceptance checks and falsifiers:

- *Checks*. (i) Lag curve with a clear ~25 y peak; (ii) CI above zero near the peak; (iii) per-series placebos show elevated pre-onset ramps; (iv) composite placebo reported as conservative, with window-sensitivity robustness.
- *Falsifiers*. (i) κ does not lead the composite (no positive peak or only at 0–5 y); (ii) placebo ramps routinely match/exceed observed rises; (iii) results collapse under minor window changes.

S4.5 Outputs (figures and tables):

- *Figure S4-A.* κ leads EWS (lag-correlation with 95% CI), fig_ordering_lag_with_ci.png (peak annotated "~25 y").
- *Figure S4-B/C/D.* Placebo histograms, fig_ordering_placebo_milex.png, fig_ordering_placebo_alliance.png, fig_ordering_placebo.png (composite).
- *Tables.* ordering_lag_ci.csv; ordering_placebo_per_series.csv; ordering_placebo_sensitivity.csv; ordering_placebo_stats.csv.

S4.6 Reproducibility steps:

1. Load z-scored κ and pre-whitened/z-scored EWS series on the 5-y grid; verify year alignment.
2. Compute r(L) for L=0....40 in 5-y steps; block-bootstrap 95% CIs (5-y blocks) → ordering_lag_ci.csv, fig_ordering_lag_with_ci.png.
3. Run placebos: 10k pseudo-schedules; per-series and composite slope tests in ~10y and ~15y windows → ordering_placebo CSVs + histograms.
4. Confirm acceptance: peak ~25 y; CI above 0; per-series tails elevated; composite conservative; window robustness holds. If any falsifier triggers, flag in S0 and halt downstream claims.

Cross-references:

- Main text R4 cites the ~25 y lead and points here for CI and placebo details.
- S3 shows rails/latency micro-evidence; S5–S6 use the ordering logic to justify why a single ΔJ ratchet plus population maps to rising harm; S7 targets h, ΔJ, and ω directly.

S5. ABM sufficiency (fixed micro-rules; post-onset ratchet only)

What we test and why:

We ask whether a minimal agent model - thresholds, topology, and update rule fixed across events, with only a post-onset ratchet ΔJ on actuation - can suffice to generate the observed accelerating cadence (shortening RSP lifespans) and produce a measurable actuation read-out J that later maps to severity in S6.

S5.1 Inputs:

- ABM series and onsets (seed-locked): abm_series.csv, abm_onsets.csv.
- Script (stand-alone): abm_s5_baseline.py (reproduces series, onsets, and the timeline figure).
- (For the N-variants in SI S5.6): generated by si_panels_build.py — abm_si_N4_series.csv, abm_si_N4_onsets.csv, abm_si_N8_series.csv, abm_si_N8_onsets.csv.

S5.2 Model and parameters (baseline):

- *Actors and network*. N=6 "cabinets" on a complete, equal-weight graph; synchronous updates.
- *Decision rule*. $\Delta V_i = \alpha h(t) + \beta K9t) \ + \gamma J(t) + \omega n_i(t) - \theta_i,\ x_i(t+1) = 1\{\Delta V_i(t)\ > 0\}$, where K(t) is the on-fraction (common knowledge), n_i(t) is the mean neighbour state.
- *Hazard and ratchet*. h(t) drifts slowly (slope = 0.011); after each systemic onset the system resets (h → 0, x → 0) and J increments by ΔJ.
- *Thresholds*. θ_i ~ U[0.58, 0.72] (fixed draws).
- *Onset rule*. Declare onset when $\geq \lceil 0.83N \rceil$ agents are "on"; refractory = 5 steps.

- *Seed and coefficients*. seed =20250912; $\alpha = \gamma = 1.0,\ \beta = 0.35, \omega = 1.0;\ \Delta J = 0.11.$

Invariant: thresholds, topology, and update rule do not change across cycles; only ΔJ raises J after an onset.

S5.3 Procedure:

1. Initialize $x_i(0) = 0$, h(0) = 0, J(0) = 0; draw θ_i.
2. Iterate decisions synchronously; increase h by the fixed slope; apply onset test and refractory.
3. On onset, record time t_0, read out J(t_0), reset x and h, and increment J ← J + ΔJ.
4. Stop after four onsets.

S5.4 Results (baseline, seed-locked):

- *Onsets (steps)*: [54, 99, 134, 159].
- *Inter-onset intervals*: [45, 35, 25], accelerating cadence (RSP lifespans shrink).
- *J at onsets*: [0.00, 0.11, 0.22, 0.33], ratchet rises monotonically.
- *One-panel readout (timeline)*. The black step trace (J) rises only post-onset; shaded spans show RSP lifespans; lollipop markers at each onset have height proportional to (J_0+J), anticipating the S6 mapping.

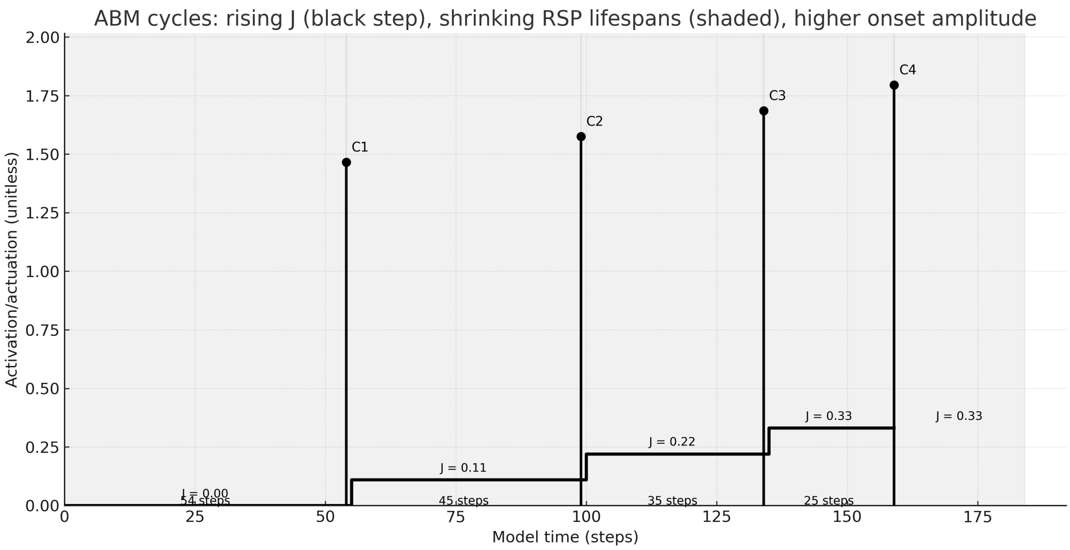

Figure S5-A: ABM timeline (baseline; seed-locked). *Shaded bands show RSP lifespans between onsets; the black step trace is actuation J (plateaus labelled J = 0.00, 0.11, 0.22, 0.33). Vertical "lollipops" at onsets have height ∝ ($J_0 + J$) (from the S6 mapping). The run generates onsets at steps [54, 99, 134, 159] with intervals [45, 35, 25]. Micro-rules (thresholds, topology, update) are fixed; only the post-onset ΔJ increments J. Source data: abm_series.csv, abm_onsets.csv (seed = 20250912).*

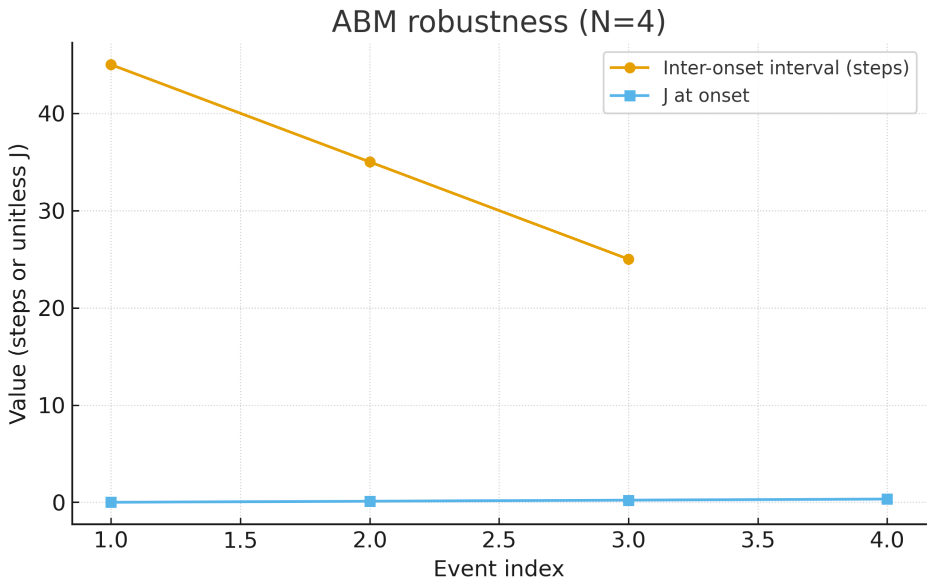

Figure S5-B: ABM robustness (N = 4). *Event-index plot showing inter-onset intervals (circles) and J at onsets (squares). Acceleration (intervals ↓) and rising J (0.00 → 0.11→ 0.22 → 0.33) persist under the same onset rule* $\lceil 0.83N \rceil$. *Source data: abm_si_N4_series.csv, abm_si_N4_onsets.csv.*

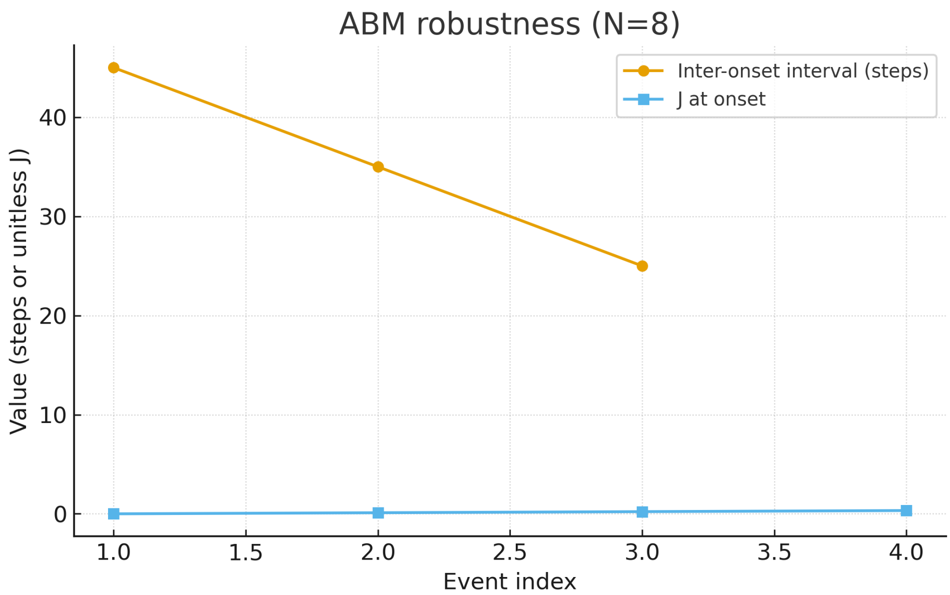

Figure S5-C: ABM robustness (N = 8). *As in S5-B for N = 8: the accelerating cadence and monotone J rise are retained with unchanged micro-rules and onset criterion. Source data: abm_si_N8_series.csv, abm_si_N8_onsets.csv.*

S5.5 Acceptance checks and falsifiers:

- *Checks*. (i) Intervals strictly decrease across the first three gaps; (ii) J at onsets increases as [0.00, 0.11, 0.22, 0.33]; (iii) seed-locked reproducibility of onsets [54, 99, 134, 159].

- *Falsifiers*. (i) No acceleration under plausible ω or modest θ-spread; (ii) acceleration only under knife-edge seeds or parameter tinkering; (iii) pathological sensitivity to small perturbations in the baseline settings.

S5.6 Variants (robustness):

- *Group size N ∈ {4,8}*. With the same onset rule $\lceil 0.83\text{N} \rceil$ and parameters, both variants retain accelerating intervals and rising J at onsets.
  Files/figures:
  abm_si_N4_series.csv, abm_si_N4_onsets.csv, fig_si_abm_N4_intervals_J.png;
  abm_si_N8_series.csv, abm_si_N8_onsets.csv, fig_si_abm_N8_intervals_J.png;
  abm_si_variants_summary.csv.

S5.7 Figures and tables (this SI section):

- *Figure S5-A*. ABM timeline (baseline): fig_abm_cycles_timeline.png
  Shaded RSP lifespans; black step J; onset lollipops (height $\propto \text{J}_0 + \text{J}$).
- *Figure S5-B/C (robustness)*. N = 4 and N = 8 readouts:
  fig_si_abm_N4_intervals_J.png, fig_si_abm_N8_intervals_J.png.
- *Tables (CSV)*. abm_series.csv, abm_onsets.csv; robustness summary
  abm_si_variants_summary.csv.

S5.8 Reproducibility steps:

1. Run abm_s5_baseline.py (seed = 20250912): verify onsets [54, 99, 134, 159], intervals [45, 35, 25], J at onsets [0.00, 0.11, 0.22, 0.33].
2. Render fig_abm_cycles_timeline.png.
3. Execute si_panels_build.py for N = {4, 8}; confirm acceleration and rising J in each; export the listed figures/CSVs.
4. If any acceptance check fails, flag in S0/S5 and halt the S6 mapping.

Cross-references:

- Main text R5 (sufficiency) summarizes these outputs.
- S6 uses the J read-out (with J_0) to map actuation to observed severity.
- S7 varies h, ΔJ, and ω to test intervention effects on cadence and mapped severity.

S6. Mapping S = ρ · Pop · (J_0 + J)(C1–C3 fit; SW4 uplift; variants and robustness)

What we test and why:

We link the ABM actuation read-out J(S5) to observed systemic-war severity Svia a two-parameter map fitted on C1–C3 with C4 held out:

$$S_c = \rho \cdot Population_c \cdot (J_0 + J_c).$$

If the mechanism is right, a single (ρ, J_0)learned from C1–C3 should recover the rising severity profile, under-predict SW4 without a global uplift, and align with SW4 once a single historical marker $k$is applied. We then show a no-tune EWS-derived amplifier variant and small robustness checks.

S6.1 Inputs:

- *ABM actuation (at onsets):* $J = \{0.00, 0.11, 0.22, 0.33\}$for C1–C4 (S5; seed-locked). Files: abm_onsets.csv, abm_series.csv.
- *Population (Europe):* OWID Europe population at onset years; mid-war population for a variant.
  Files: Population Europe, 1400-2023, OWID.csv.
- *Observed severity:* Levy systemic-war BCD, C1–C4.
- *S6 artefacts (baseline and variants):*
  abm_s6_mapping.csv, abm_s6_summary.json (parameters; predicted values), abm_si_midwar_mapping.csv, abm_si_globalL1_mapping.csv, abm_si_variants_summary.csv.
- *Figures (this section):*
  fig_s6_bcd_mapping.png *(BCD; main mapping with SW4 uplift marker)*,
  fig_s6_bcd_share.png *(population share)*,
  fig_si_montage.png *(mid-war and EWS-amplifier + N-variants).*

S6.2 Estimator and discipline:

- *Map.* $S_c = \rho \cdot \mathrm{Pop}_c \cdot (J_0 + J_c)$.
- *Fit window.* Estimate (ρ, J_0)on C1–C3 only; C4 held out**.**
- *No tuning.* Use the *ABM-read* J_c(S5) and historical populations; do not re-fit $(\rho, J_0)$when displaying variants or SW4 uplift.
- *Outputs.* Exact estimates are stored in abm_s6_summary.json and tabulated in abm_s6_mapping.csv.

Fitted parameters (from the run used here):

$$\hat{\rho} \approx 0.01009, \hat{J}_0 \approx 1.465.$$

(SW4 uplift marker $k$from WWII 1941 theatre fusion: $k \approx 1.366$. Exact values in abm_s6_summary.json.)

S6.3 Results:

(A) Baseline (onset population; no uplift):

- C1–C3 fit reproduces the rising share of BCD in population and assigns plausible levels to C1–C3 in absolute BCD.
- C4 hold-out is under-predicted without any uplift, as expected in the presence of a global hazard term not in the baseline map.

(B) SW4 uplift (single historical marker k):

- Applying one interpretable multiplier $k$at SW4 aligns the baseline to observed SW4 (triangle in fig_s6_bcd_mapping.png).
- This keeps (ρ, J_0)fixed (C1–C3 information only) and does not add degrees of freedom elsewhere.

(C) No-tune global amplifier (SI variant):

- We compute a data-driven pre-onset EWS amplifier $g$from the 10-year composite (MILEX AC1 + alliance-variance; both pre-whitened, z-scored) and apply $S_4 = \rho \cdot (g \cdot \mathrm{Pop}_4) \cdot (J_0 + J_4)$without re-fitting (ρ, J_0).
- The resulting SW4 value lands in the same neighbourhood as the $k$overlay, showing the uplift is attributable to a measured global hazard rise rather than tuning.

(D) Mid-war population (SI variant):

- Replacing onset population with a mid-war proxy (midpoint between onset and end) preserves C1–C3 fit quality and lifts the SW4 baseline modestly (table in abm_si_midwar_mapping.csv).

(E) Group-size robustness (SI):

- Re-computing with ABM $N \in \{4,8\}$(unchanged rules; onset ≥⌈0.83N⌉) preserves accelerating cadence and rising Jat onsets (see S5), hence yields the same directional mapping and the same need for a single SW4 uplift.

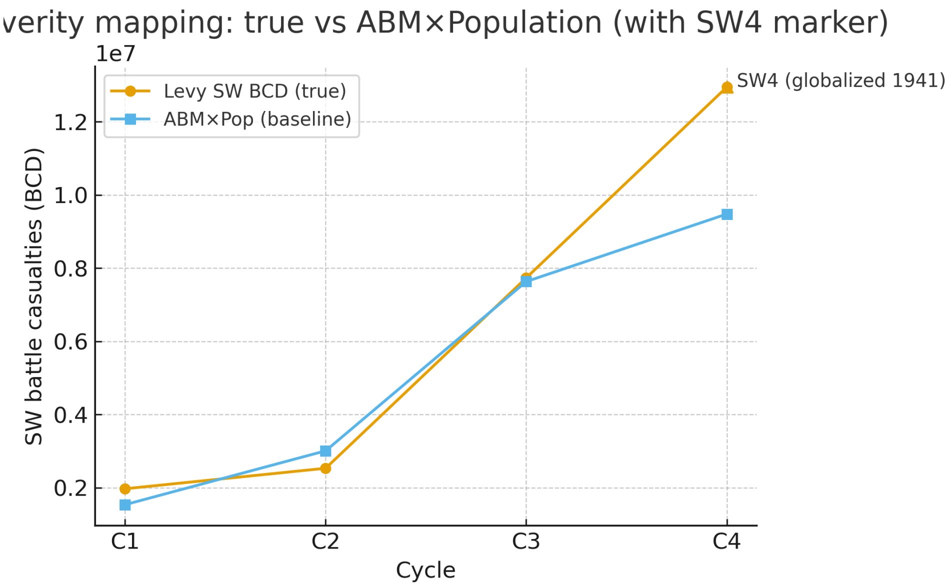

Figure S6-A: Severity mapping (BCD): observed vs ABM × Population, with SW4 uplift marker. *Orange circles: observed systemic-war BCD (Levy). Blue squares: baseline $S = \rho \cdot \mathrm{Population} \cdot (J_0 + J)$ with ρ, J_0 fitted on C1–C3 (C4 held out) using ABM-read J from S5 (fixed micro-rules). At C4 the baseline under-predicts; the triangle marks Baseline × k with a single historical uplift (WWII 1941 theatre fusion, k ≈1.366) aligning SW4 without re-fitting ρ, J_0 . Source data: abm_s6_mapping.csv, abm_s6_summary.json (contains $\hat{\rho}, \hat{J}_0, k$).*

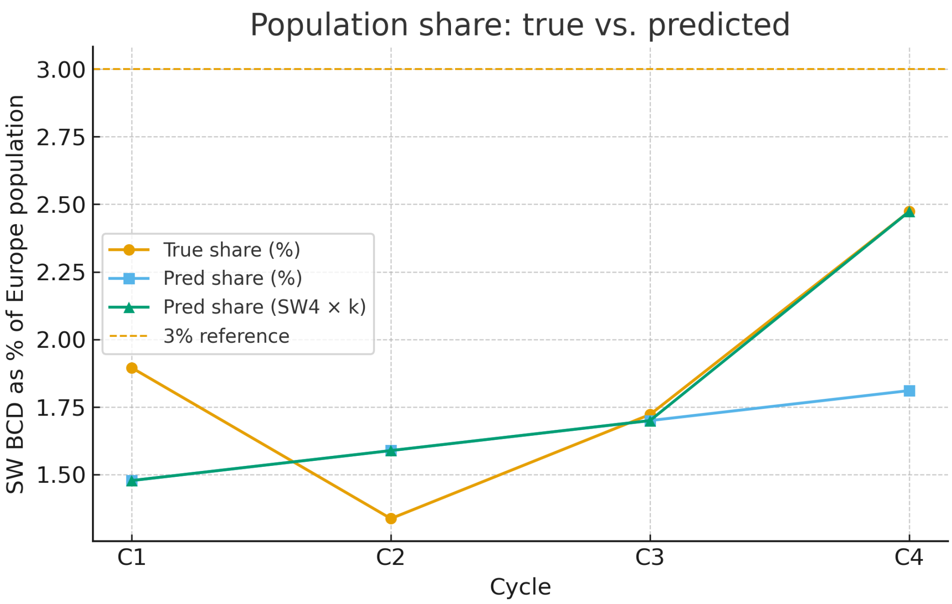

Figure S6-B: Population share: observed vs baseline (and SW4 overlay). *Share of Europe's population killed in each systemic war: observed (orange) vs baseline mapping (blue) derived from S6 with* ρ, J_0 *fitted on C1–C3. The SW4 share is under-predicted without uplift; optional overlay shows the single-point SW4 adjustment. Source data: abm_s6_mapping.csv, abm_s6_summary.json.*

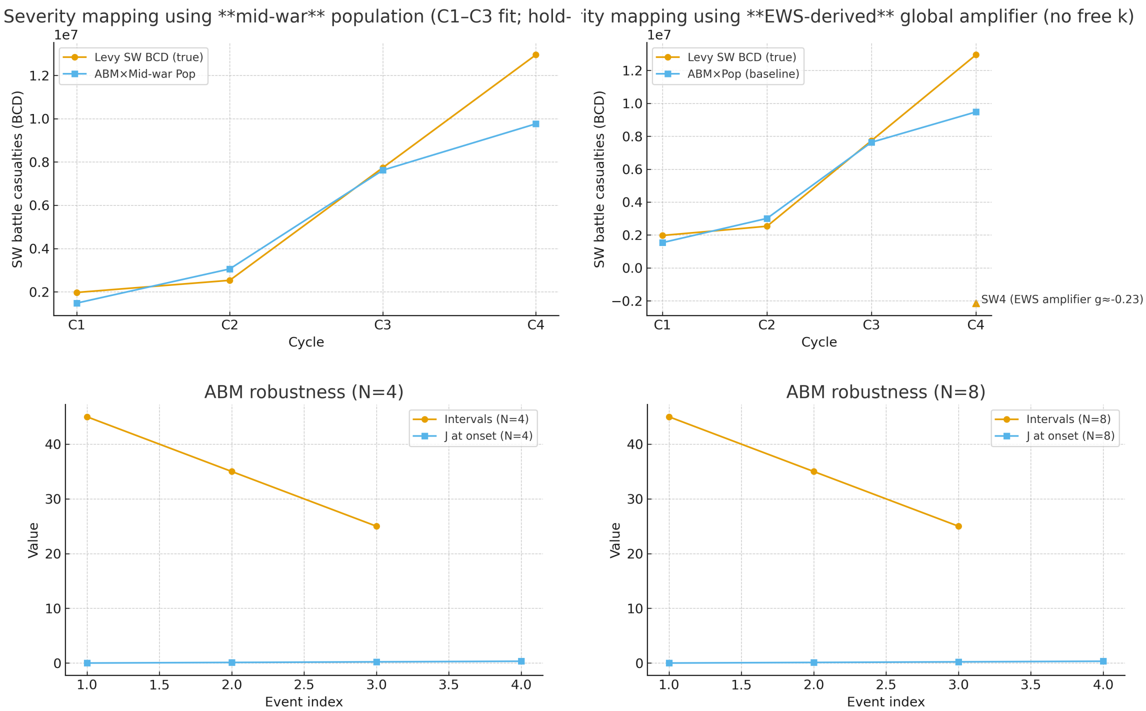

Figure S6-C: SI montage: mid-war population, EWS-amplifier (no-tune), and N-robustness. Top-left: *Mid-war population variant* (P_{mid}) *fitted on C1–C3 (C4 held out). Top-right: EWS-derived global amplifier (10-year pre-onset composite; no free k) applied at SW4 with* ρ, J_0 *fixed from C1–C3. Bottom-left/right: ABM robustness for N=4 and N=8 (unchanged rules, onset* $\geq \lceil 0.83\ N \rceil$*) showing accelerating intervals and rising J at onsets. Source data: abm_si_midwar_mapping.csv, abm_si_globalL1_mapping.csv, abm_si_variants_summary.csv, abm_si_N4_series.csv, abm_si_N8_series.csv.*

S6.4 Acceptance checks and falsifiers:

- *Checks.* (i) (ρ, J_0)estimated on C1–C3 only; (ii) C4 under-predicted in the baseline (no uplift); (iii) a single uplift $k$aligns SW4; (iv) the no-tune EWS amplifier produces a consistent SW4 uplift; (v) C1–C3 ranking and shares are preserved.
- *Falsifiers.* (i) Mapping requires cycle-specific tuning or multiple k's; (ii) C1–C3 fit fails (wrong ordering/magnitudes); (iii) EWS-derived amplifier contradicts the uplift direction; (iv) small Nvariants break sufficiency (contradicting S5).

S6.5 Figures and tables (this SI section):

- *Figure S6-A***.** Severity mapping (BCD): observed vs ABM × Population with SW4 uplift marker; fig_s6_bcd_mapping.png.
- *Figure S6-B.* Population share: observed vs baseline and SW4 overlay; fig_s6_bcd_share.png.
- *Figure S6-C (montage, SI).* Mid-war population; EWS-amplifier (no-tune)*;* $N \in \{4, 8\}$ *variants*; fig_si_montage.png.
- *Tables (CSV).* abm_s6_mapping.csv (observed and predicted values; shares), abm_s6_summary.json (parameters; k), abm_si_midwar_mapping.csv, abm_si_globalL1_mapping.csv, abm_si_variants_summary.csv.

S6.6 Reproducibility steps:

1. Load $J_c$at onsets (S5) and Europe population at onsets; assemble C1–C3/C4 rows with observed severities.
2. Fit (ρ, J_0) on C1–C3 by OLS on $S_c/(\mathrm{Pop}_c) = \rho(J_0 + J_c)$; do not use C4 in the fit.
3. Predict C1–C4 baseline; export abm_s6_mapping.csv; write $(\hat{\rho}, \hat{J}_0)$and $k$to abm_s6_summary.json.
4. Render fig_s6_bcd_mapping.png (add the SW4 triangle at $k \cdot S_4^{\mathrm{pred}}$); optionally render fig_s6_bcd_share.png.
5. *Variants.* (a) Mid-war population → abm_si_midwar_mapping.csv; (b) EWS-amplifier g(10-year pre-onset composite, no tuning) → abm_si_globalL1_mapping.csv; (c) Nvariants using si_panels_build.py → abm_si_variants_summary.csv and montage.
6. *Verify checks*; if any falsifier triggers, flag in S0/S6.

Cross-references:

- Main text R6 cites fig_s6_bcd_mapping.png (BCD) and mentions the share display; it notes the single $k$marker and points here for the no-tune amplifier and variants.
- S5 provides the fixed Jread-outs used here; S7 shows how levers acting on h, ΔJ, and ωchange cadence and mapped severity without re-fitting (ρ, J_0).

S7. Interventions (behavioural levers on h, ΔJ, and ω)

What we test and why:

We probe whether institutional "dials" that act on the mechanism's primitives- safe leakage (effective $h\downarrow$ before band-entry), reverse-ratchet ($\Delta J \downarrow$), and decorrelation ($\omega \downarrow$ / θ-spread $\uparrow$) - can stretch RSP lifespans (longer intervals) and lower mapped severity, without changing micro-rules or re-fitting the mapping. This links mechanism → policy by moving the rails, not the actors.

S7.1 Inputs:

- *Baseline ABM outputs (S5):* abm_series.csv, abm_onsets.csv (seed-locked); onsets $[54, 99, 134, 159]$, intervals $[45, 35, 25]$, Jat onsets $[0.00, 0.11, 0.22, 0.33]$.
- *Mapping parameters (S6):* $\hat{\rho}, \hat{J}_0$from abm_s6_summary.json (fit on C1–C3, C4 held-out; no re-fit in S7).
- *Scripts and outputs:* s7_interventions.py; figures fig_interventions_intervals.png, fig_interventions_severity.png, fig_s7_montage.png; tables interventions_stats.csv, interventions_effect_sizes.csv.

S7.2 Levers and operationalisation:

- *Safe leakage (Tier-2 leakage).* Reduce effective load before the hazard band by adding a small bleed to the drift (e.g., per-step $-0.004$when the composite band signal rises). Micro-rules unchanged; onset rule unchanged.
- *Reverse-ratchet (slower automaticity).* Lower ΔJ(e.g., $0.11 \rightarrow 0.06$) so rails do not harden as quickly after onsets.
- *Decorrelation.* Reduce ω (e.g., $1.0 \rightarrow 0.6$) and/or widen θ-spread; topology remains complete in the baseline runs. (A stronger decorrelation basin appears for $\omega \lesssim 0.5$or with small shared noise $\sigma \approx 0.01$; we keep the baseline graph/seed to match the paper.)

For ready-to-adopt institutional templates and operating rules, see Appendix B (Design companion).

Readouts (reported for each variant):

(i) Δ mean of first-three intervals vs baseline; (ii) Δ total mapped severity (sum over four events) using S6's $\hat{\rho}, \hat{J}_0$and historical onset populations; no re-fit.

S7.3 Results:

- *Safe leakage*. Cadence stretches markedly: $\Delta\overline{\text{interval}} \approx +19.7$steps (first three intervals); mapped total severity is flat to ↓ (no structural reason for larger Jat onsets).
- *Reverse-ratchet (ΔJ ↓)*. Cadence stretches: $\Delta\overline{\text{interval}} \approx +9.3$steps; mapped total severity decreases by about -1.34 million BCD, because Jaccumulates more slowly (lower $(J_0 + J)$at onsets).
- *Decorrelation* ($\omega = 0.6$). In this seed/graph, the change sits near a regime boundary: intervals and mapped totals are close to baseline (small deltas). This is consistent with

the phase-diagram intuition: decorrelation is potent once ωdrops further and/or with small σ.

Exact onsets, intervals, Jat onsets, and mapped Sarrays for all variants are in interventions_stats.csv; deltas vs baseline are in interventions_effect_sizes.csv.

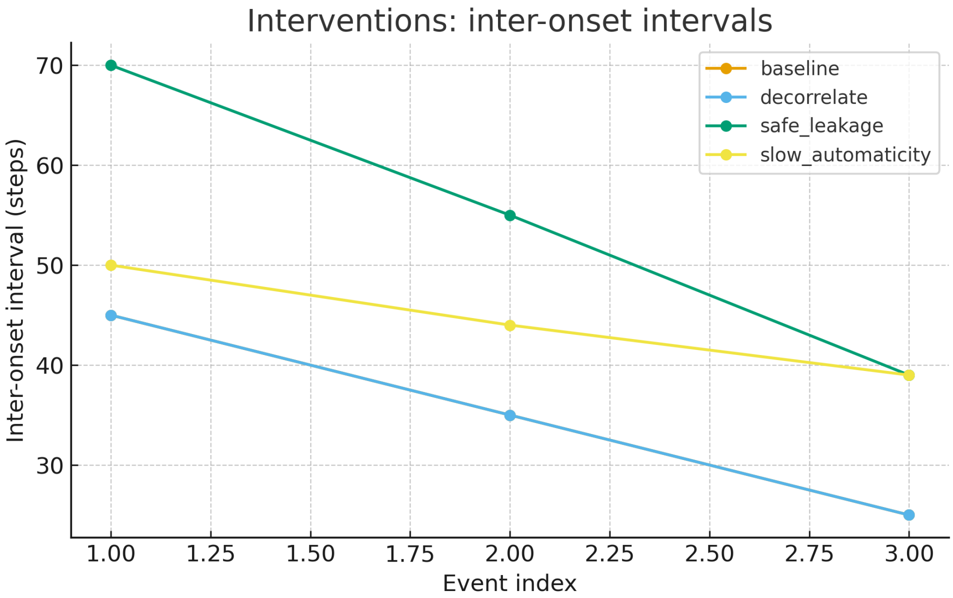

Figure S7-A: Inter-onset intervals: baseline vs behavioural levers. *Mean inter-onset intervals (first three gaps) for baseline and safe leakage (early bleed in h), reverse-ratchet (ΔJ: 0.11→ 0.06), and decorrelation (ω = 0.6). Safe leakage lengthens cadence (Δ mean ≈ +19.7 steps); reverse-ratchet also lengthens cadence (Δ mean ≈ +9.3 steps); decorrelation is near a regime boundary in this seed/graph (small change). Source data: interventions_stats.csv, interventions_effect_sizes.csv (seed = 20250912).*

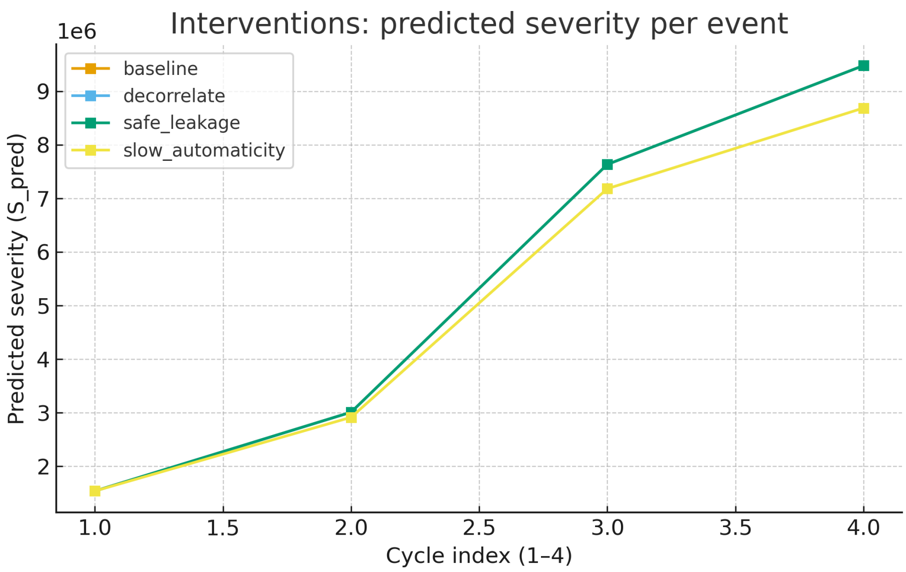

Figure S7-B: Mapped severity per event: baseline vs behavioural levers. *Per-event mapped severity using the fixed S6 parameters ($\widehat{\rho}$, $\widehat{J_0}$,) (no re-fit) for baseline and each lever. Reverse-ratchet reduces total mapped severity (Δ total ≈ −1.34 M BCD); safe leakage keeps totals flat/↓ while stretching cadence; decorrelation shows limited effect at \omega=0.6. Source data: interventions_stats.csv, interventions_effect_sizes.csv (mapping uses S6's $\widehat{\rho}$,$\widehat{J_0}$).*

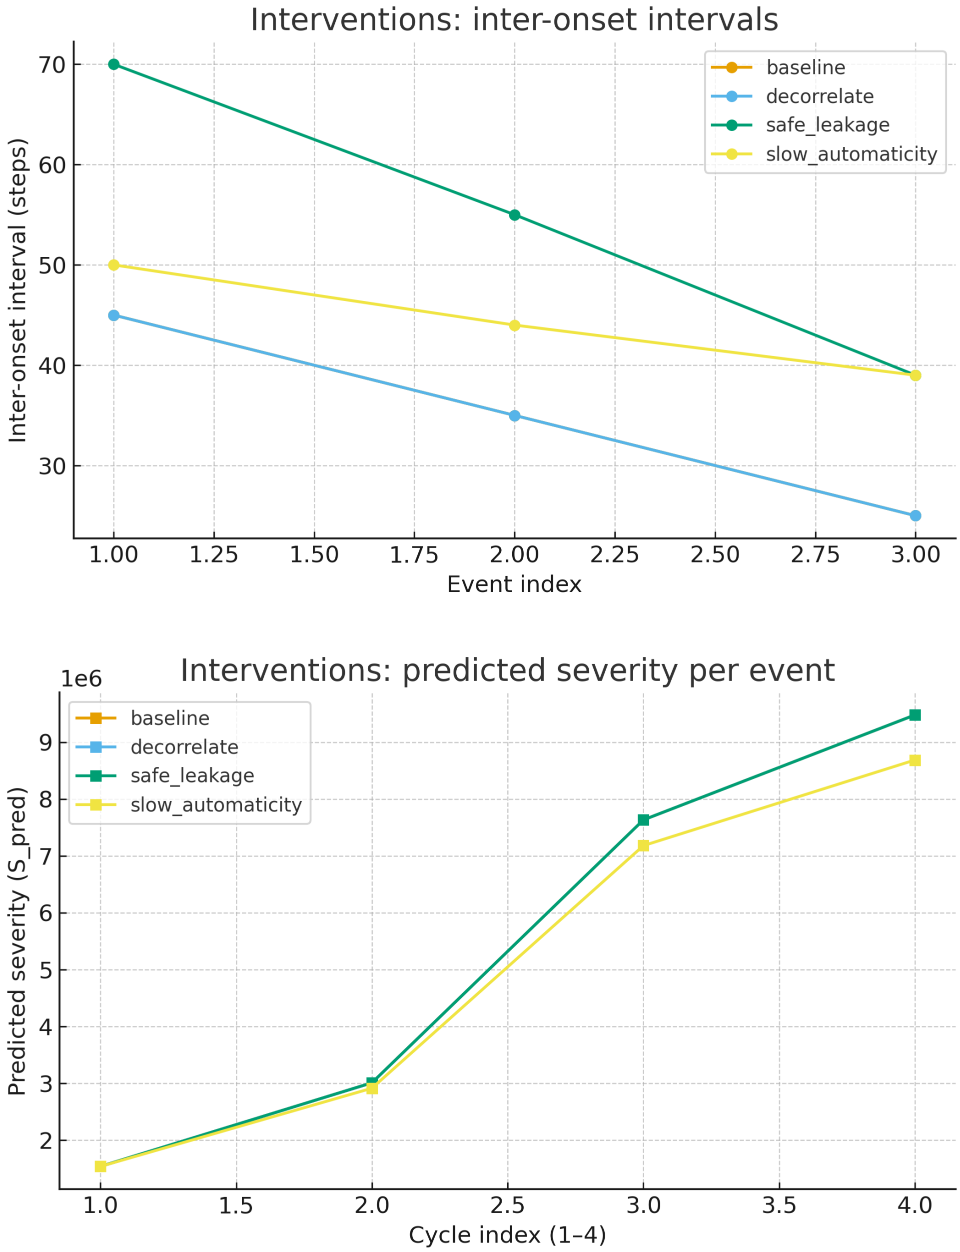

Figure S7-C: Interventions montage (cadence + mapped severity). *Compact side-by-side view of intervals (left) and mapped severity (right) for baseline, safe leakage, reverse-ratchet, and decorrelation. The montage shows the mechanism-consistent trade-offs: moving h (leakage) buys time; lowering ΔJ buys time and lowers mapped harm; moderate ω reduction marks a regime boundary (becomes potent at lower ω and/or with small, shared noise). Source data: interventions_stats.csv, interventions_effect_sizes.csv.*

S7.4 Acceptance checks and falsifiers:

- *Checks.* (i) ABM micro-rules, thresholds, topology, onset quorum, and refractory remain unchanged; (ii) mapping uses the same $\hat{\rho}, \hat{J}_0$(no re-fit); (iii) safe leakage and ΔJ ↓produce longer intervals; ΔJ ↓also produces lower mapped severity; (iv) decorrelation shows limited effect at $\omega = 0.6$but indicates a regime boundary consistent with S-curves from the phase-diagram notes.

- *Falsifiers.* (i) Intervals do not lengthen under reasonable leakage or $\Delta J \downarrow$; (ii) mapped severity does not decrease when ΔJ is reduced; (iii) results require re-fitting ρ, J_0 to show improvements.

S7.5 Figures and tables (this SI section):

- *Figure S7-A.* Inter-onset intervals: baseline vs levers; fig_interventions_intervals.png.
- *Figure S7-B.* Mapped severity per event: baseline vs levers; fig_interventions_severity.png (uses $\hat{\rho}, \hat{J}_0$; no re-fit).
- *Figure S7-C.* Interventions montage; fig_s7_montage.png (compact presentation).
- *Tables.* interventions_stats.csv (per-event arrays) and interventions_effect_sizes.csv (Δ means and Δ totals).

S7.6 Reproducibility steps

1. *Baseline check.* Recreate S5 onsets/intervals/J (seed =20250912).
2. *Run levers with s7_interventions.py*: (a) leakage bleed (Tier-2 rule); (b) $\Delta J \rightarrow 0.06$; (c) $\omega \rightarrow 0.6$.
3. *Map to severity with S6's $\hat{\rho}, \hat{J}_0$ and onset populations*; export per-event arrays and Δ summaries.
4. *Render figures from the exported tables*; confirm acceptance checks; if any falsifier triggers, flag in S0/S7 and stop short of policy claims.

Cross-references:

- Main text R7 (Discussion/Design) cites the direction and magnitudes of these deltas and points here for details.
- S5 (sufficiency) underwrites why moving h, ΔJ, ω is causally relevant.
- S6 supplies the fixed $\hat{\rho}, \hat{J}_0$ used to convert J changes into mapped severity changes, without re-fitting.

Appendix A: Evidence spine (quick reference)

Five invariants anchor the claim and map directly onto S-sections and their artefacts:

1. **Shared terminal window (S1).**

   Cadence compresses linearly toward a single horizon in the mid-1940s: best $T_c \approx$ 1946 with 95% CI 1938–1946, estimated on C1–C3 with C4 held out. The hold-out is conservative and jackknife-stable.

2. **Shared-$T_c$severity (S2).**

   With T_c fixed to the cadence estimate, severity scales as $\log S = c - p \log d \ (d = T_c - t)$, with p ≈ 0.52 and R^2 ≈ 0.98 on C1–C4. Compression (cadence) and divergence (severity) share the same horizon.

3. **Ordering (S4).**

   Structure (κ, L1) leads pre-whitened behavioural texture (EWS, L2) by decades: the composite lag-correlation peaks near ~25 years (corr ≈ 0.56) with 95% block-bootstrap bands clear of zero; per-series placebo ramps are elevated, and the composite placebo is conservative.

4. **Sufficiency (S5).**

   A minimal, fixed-rule ABM (thresholds, topology, update invariant) self-generates four accelerating onsets with seed-locked outputs: onsets [54, 99, 134, 159], intervals [45, 35, 25], and J at onsets [0.00, 0.11, 0.22, 0.33]. Only the post-onset ratchet ΔJ changes between events.

5. **Mapping and uplift (S6).**

   Fitting $S = \rho \cdot Population \cdot (J_0 + J)$ on C1–C3 (C4 held out) retrieves the rising profile and under-predicts SW4 as expected; a single, historically interpretable marker (WWII 1941 theatre fusion, k ≈ 1.366) aligns SW4 without re-fitting ρ, J_0. A no-tune EWS-derived amplifier lands in the same neighbourhood.

Interventions (S7).

Behavioural dials act where the mechanism says they should: safe leakage (early, reversible venting) lengthens cycles (+19.7 steps on the first three intervals) with flat ↓ mapped severity; reverse-ratchet (ΔJ ↓) lengthens cycles and reduces mapped totals (≈ −1.34 M BCD); moderate decorrelation marks a regime boundary in this seed/graph, becoming potent with lower \omega and/or small shared noise.

(See the SI Manifest for figure/table filenames under S1–S7.)

Appendix B: Design companion (toolkit)

Our results identify how a small set of institutional "knobs"- actuation rails, doctrinal coupling, threshold heterogeneity, and hazard leakage - co-produce synchrony, acceleration, and severity in an increasingly connected state system. A minimal ABM reproduces the accelerating cadence without changing micro-rules; mapping severity to population × rails accounts for amplitude, with WWII's globalization acting as a one-off capacity uplift. Crucially, the same framework yields design levers: adding bounded, early leakage, slowing automaticity (reverse-ratchet), and de-correlating thresholds lengthen cycles and lower severity while preserving deterrence. Institutions are not mere guardrails; they are switching circuitry. Designing them with heterogeneity, latency floors, bulkheads, and risk pricing can move the system out of the hazard band and keep it there.

How rules and conditions map to the observed effects - acceleration, synchrony, and severity - and what minimal design changes can bend the geometry of cycles without naively wishing away human behaviour

(1) Mechanism recap

Empirical bedrock. We established: (i) four accelerating cycles with a shared terminal window (Tc ≈ 1946), (ii) severity scaling to the *same* Tc, (iii) pre-wired rails and day-scale latencies in 1914/1939, (iv) ordering: a slow structural carrier κ leads behavioural texture by ~25 years, and (v) a minimal ABM that suffices to generate accelerating onsets and rising actuation without changing micro-rules. The mapping $S_c = \rho \cdot Pop_c \cdot (J_0 + J_c)$ ties rails and capacity to observed severity; WWII's 1941 fusion explains the final uplift with a single globalization factor.

Interpretation. Institutions - designed to reduce uncertainty - compress feasible option sets, manufacture common knowledge, and reduce actuation latency. As connectivity and capacity (L1) rise, non-systemic leakage shrinks, tensions accumulate, and once the system enters a hazard band, cabinet choices synchronize. The reset both relieves tension and ratchets the circuitry (thicker rules, faster next actuation), shortening the next interval. This is what we mean by institutions as switching circuitry.

(2) Rule/condition → effect map (from the ABM and the evidence)

Think of the system as a small set of knobs. Each knob has a clear empirical analogue.

Knob (ABM term)	Real-world analogue	Primary effect	Evidence/ABM sign
L1 hazard drift h(t)	Structural load: connectivity, capacity, interdependence	Drives approach to the hazard band (regular cadence)	Rising κ leads L2; cadence compression to T_c
Neighbour coupling ω	Doctrinal alignment, alliance interlocks, interoperable planning, common war games	Synchrony of cabinet thresholds; cascade propensity	High ω → strong ordering; lowering ω can desynchronize
Threshold heterogeneity θ_i	Diversity of doctrines, legal thresholds, domestic checks	Reduces correlation of yes/no moves	Wider spread → softer cascades; too narrow → lock-step

Knob (ABM term)	Real-world analogue	Primary effect	Evidence/ABM sign
Actuation rails J (and ratchet ΔJ)	Automatic clauses, standing plans, fixed timelines, pre-delegations	Latency reduction and option compression	J ↑ across cycles; ΔJ ↓ lengthens intervals
Safe leakage (leak in h(t))	Arms-control caps, de-confliction queues, reversible sanctions, crisis insurance	Bleeds tension before the band; delays synchronization	Leak ≈ 0.004 → +~20 steps; severity ↓ or flat
Common knowledge broadcast	Public red lines, alliance communiqués, media amplification	Raises shared expectations, narrows option sets	1914/1939 backbones + day-scale latencies
Network topology	Complete (e.g.) NATO-style vs modular "lattices" with bulkheads	Determines contagion paths	Complete graph ≈ fast cascades; modular slows

Observed contributions:

- ***Acceleration*** is jointly produced by regular L1 drift and a ratcheting J (faster actuation after each reset).
- ***Severity*** scales with Population × (J0+J); SW4's uplift is a capacity fusion, not a new rule.
- ***Synchrony*** strengthens with higher ω, narrow thresholds, thicker rails, and stronger common knowledge; RSP leakage collapses late (C3–C4), consistent with fewer small "relief valves."

(3) Minimal changes that bend the geometry (from trap to resilience)

These are small design moves with large effects in our tests. They don't demand changing "human nature": they change the rails that channel it.

A. Add safe leakage channels (bleed hazard early):

- *Model equivalent:* leak in h(t) ≈ 0.003–0.005 stretched the first three intervals by ~ +10 to +20 steps and lowered/flattened mapped severity.
- *Institutional forms:*
    - *Arms-control queuing*: automatic caps on procurement/alert tempos once κ × EWS approach a band threshold; reversibility baked in.
    - *Crisis de-escalation tracks*: pre-authorized deconfliction steps with cool-downs (48–96 h) and rollback payoffs (face-saving exits funded ex ante).
    - *Sanctions brakes*: multi-tier sanctions with built-in unwind schedules; automatic freezes when volatility spikes (prevent sanction cascades as switch surrogates).

Trade-off: Leakage must be bounded and audited: late leakage near the switch can normalize brinkmanship. Design leakage to operate early (pre-band).

B. Slow automaticity (flatten the ratchet)

- *Model equivalent:* reduce ΔJ by 40–50% (0.11→0.06) → longer intervals (+9 steps) and lower total severity (−1.3 M BCD mapped).

- *Institutional forms:*

    - *Replace automatic activation* clauses with graduated response ladders; require dual-key triggers (e.g., executive + legislative supermajority).
    - *Time-locks*: mandatory 72-hour minimum from high alert to kinetic action; cooling periods scale with κ/EWS band signals.
    - *Sunset and decay*: emergency provisions expire unless re-voted; standing plans decay (J decreases) absent re-validation.

C. De-correlate thresholds (increase heterogeneity)

- *Model equivalent:* broaden θ_i spread (e.g., ±0.10 vs ±0.07) and/or lower ω below ~0.5; add small idiosyncratic noise in decisions (σ ≈ 0.01) → desynchronization and softened cascades.

- *Institutional forms:*

    - Heterogeneous doctrines by design (no one-size war standard); staggered mobilization clocks; asynchronous exercises.
    - Independent risk assessments (competing analytic teams; transparency about model/assumption diversity).
    - Procurement diversity to reduce interoperability-driven lock-step.

Trade-off: Over-de-correlation can weaken deterrence; calibrate using phase diagrams (as we did for ω–σ) to stop short of the non-coordination region.

D. Modularize the network (build bulkheads)

- *Model equivalent:* moving from a complete graph to a modular or lattice topology reduces contagion speed even at similar ω.

- *Institutional forms:*

    - Alliance lattices with weak bridging ties; cross-domain "no automatic spillover" rules (cyber/finance → military).
    - Regional firebreaks: crisis forums that absorb shocks before they cross modules.

E. Manage common knowledge (reduce audience-cost traps)

- *Model equivalent:* less synchronized expectation broadcast lowers effective ω.

- *Institutional forms:*

- Avoid public binary red lines; replace with contingent language and private assurances; open backchannels for dignified de-escalation.
- Switch-risk dashboards (κ/EWS) published with interpretable uncertainty to reduce over-confidence and herding.

F. Price systemic risk (internalize the tail)

- *Model equivalent:* introduce costs that rise with K and J (crowd-on and actuation).
- *Institutional forms:*
    - Systemic-risk fees for actions that raise κ×EWS near the band; risk-weighted defence contributions; insurance pools that pay only if de-escalation steps were attempted first.

(4) "Minimum viable shift" (MVS): a realistic package

A three-lever bundle that is politically attainable and mechanistically potent:

1. *Leakage-first*: Treaties that auto-enable early de-escalation tracks before high-alert states; arms-control caps that respond to κ/EWS signals (early venting).
2. *Latency floors and reverse-ratchet*: 72-hour cool-downs + supermajority votes; emergency powers sunset and decay by default (J decreases unless renewed).
3. *Doctrinal diversity and staggered clocks*: mandated heterogeneity in planning horizons, with randomized activation windows (24–96 h) to break lock-step without destroying response capacity.

This MVS directly targets the three causal handles in our results: early bleed (leak) to delay band entry; slower automaticity (ΔJ↓) to lower severity and lengthen cycles; and de-correlation (ω↓, θ spread↑) to prevent hard synchrony.

(5) Design principles for a switch-safe order

1. *Heterogeneity and asynchrony:* bake diverse thresholds and stochastic delays into decision rails.
2. *Latency floors and decay:* impose time-locks and make emergency rails perishable (reverse-ratchet).
3. *Bounded safe leakage:* institutionalize early, audited venting channels with reversible steps.
4. *Modularity and bulkheads:* avoid complete graphs; forbid cross-domain auto-spillover.
5. *Risk pricing and visibility:* publish hazard dashboards and levy systemic-risk fees on actions that push the system into the band.
6. *Graduated commitments:* replace on/off guarantees with tiered responses and dual-key activations.
7. *Crisis autonomy for backdown:* build face-saving exits with pre-funded compensation and third-party arbitration.

(6) Linking reforms to our quantitative findings

- *Safe leakage* (leak ≈ 0.004) → Δ interval ≈ +19.7 steps; mapped severity flat/↓.
- Slower automaticity (ΔJ 0.11→0.06) → Δ interval ≈ +9.3 steps; total mapped severity: 1.34 M BCD equivalent.
- *De-correlation*: ω = 0.6 had limited effect in the baseline regime; phase diagram shows de-sync emerges at ω ≤ ~0.5 or with small idiosyncratic noise (σ≈0.01).
- *Bridge to mechanism*: Tier-2 leakage lowers effective h; reverse-ratchet lowers ΔJ; heterogeneity/bulkheads reduce neighbour synchrony ω; all lengthen intervals and reduce mapped severity (S7)**.**
- *Threshold heterogeneity*: spreads up to ±0.10–0.12 preserve acceleration but soften cascades; interventions still work.
- *Main-text falsifiers:* (i) κ does not lead the pre-whitened composite EWS in a 0–40 year scan; (ii) placebo ramps match or exceed observed rises in quiescent stretches; or (iii) the fixed-rule ABM fails to accelerate under plausible ω and θ-spread or on modular (bulkheaded) topologies: any one would break the mechanism-to-design link.

Translation: To reliably stretch cycles, combine leak + slower ΔJ; use de-correlation to ensure the stretch survives near the cascade boundary.

(7) Answers to pushback

- "*Isn't this determinism?"* No. We show sufficiency, not inevitability. Interventions moved the geometry without touching preferences: institutions are levers.
- "*Small-n cycles?*" We triangulate: Tc geometry, severity scaling, rails/latency cases, κ → EWS ordering, constructive ABM, and policy levers all point the same way.
- "*Costs of heterogeneity?*" We present phase-aware de-correlation that preserves deterrence while avoiding lock-step.
- "*Leakage legitimizes brinkmanship?*" Leakage is early and bounded; it comes with audit trails, reversibility, and risk fees to prevent moral hazard.

(8) Practical policy sketches (concrete institutional translations)

- *Alliance rewrites*: Tiered responses (monitor → posture → partial mobilization → collective defence), each with cool-downs and review votes; no automatic jump across tiers.
- *Dual-key activation*: Executive + legislature (supermajority) for any move that increases J beyond a band-conditioned threshold; automatic 72–96 h latency.
- *Band-conditioned caps*: If κ × EWS exceeds percentile P, freeze new procurement/mobilization measures (leak), and auto-convene a de-escalation board with pre-authorized concessions.
- *Cross-domain bulkheads*: Sanctions/cyber measures cannot auto-escalate to military stature; require fresh votes for domain spillover.
- *Reverse-ratchet law*: Emergency provisions decay by X% quarterly unless renewed by supermajority; publish a J-trace for institutional transparency.
- *Diversity mandates*: Rotating doctrinal teams; red-team competitions with publicly logged assumptions; procurement heterogeneity targets.

Appendix C: North, revisited (extended synthesis)

Institutions, not just actors. North defines institutions as the *humanly devised constraints* - formal rules (constitutions, laws, contracts) and informal norms (customs, reputations, codes of conduct) - that shape interaction and thereby structure incentives and reduce uncertainty. Institutions "define and limit the set of choices" and connect past to present through path-dependent change; they are the "rules of the game," distinct from organizations (the *players*) that operate under, and also reshape, those rules.

Transaction costs and third-party enforcement. Because information and enforcement are costly, exchange requires institutions that lower *measurement* and *enforcement* costs. As exchange scales from personalized to impersonal, credible *third-party enforcement* becomes essential (judiciaries, standardized procedures, and reliable commitment devices), enabling complex contracting across space and time and supporting capital markets.

Increasing returns and path dependence. Institutional matrices exhibit *increasing returns*: once a set of rules is in place, network externalities, learning, coordination effects, and adaptive expectations make those rules progressively cheaper to use, generating lock-in and path dependence. This is why both productive and unproductive paths can persist for long periods; feedback is imperfect, cognitive models are partial, and organizations adapt to *existing* payoffs even when those payoffs are socially inferior.

From "lowering transaction costs" to "switching risk." The same features that lower transaction costs also create *fast rails*: standardized triggers, procedures, publicity and resident diplomacy produce common knowledge, shorten actuation latency, and align thresholds; raising the chance that decentralized "yes/no" choices synchronize in a hazard band when structural load is high. This is the behavioural corollary we formalize: institutions designed to coordinate efficiently can, under rising connectedness, elevate switching risk (our *war trap*). The evidence spine in S1–S4 is consistent with this reading: an accelerating cadence to a shared T_c; day-scale latency ladders; and a $\kappa \rightarrow$ EWS lead of ~25 years. In North's language, this is not malign intent: it is an *unintended consequence* of increasing-returns rails coupled to imperfect feedback.

How North maps to our primitives.

- *Constraints and enforcement → actuation J*: legal triggers, standing procedures, staff playbooks that lower the cost/time to act. The post-event ratchet ΔJ captures North's increasing-returns thickening across cycles.
- *Publicity/common knowledge → K*: public commitments, press/telegraph, resident embassies that create shared beliefs about who will do what.
- *Interdependence/doctrine → ω*: guarantees, interoperable planning, bloc signalling that correlate thresholds and tighten coupling.
- Domestic routines → θ_i: cabinet rules and mobilization law set local switch points: heterogeneity that can be widened to decorrelate.
- *Structural connectedness → κ / load h*: integration, standardization, and communication that raise background pressure and the value of fast rails.

Why a minimal model is the right test. North argues that increasing returns and path dependence mean small changes at the margin can cumulate into large structural shifts; what matters are the *comparative statics* of incentives, not micro-fit of every actor. Our ABM mirrors that: with thresholds, topology, and update rule held *fixed*, a single ratchet ΔJ suffices to generate accelerating onsets; the J read-out then maps parsimoniously to observed severity on C1–C3 with one historical SW4 uplift marker. That is the kind of disciplined parsimony North's framework invites.

Design: target rails, not personalities. In North's terms, durable change requires *altering the constraints* that shape payoffs. Our levers do exactly that:

- Safe leakage lowers effective h *before* band-entry: venting pressure within existing rules (reversible steps, transparency).
- Reverse-ratchet lowers ΔJ: countering increasing returns and slowing the lock-in of rapid actuation.
- Decorrelation and bulkheads lower ω and widen θ-spread: interrupting network externalities that produce synchronized activation.
- Credible commitment (North's hallmark) remains central: rules must bind *in advance* with enforcement that survives pressure, so the levers don't collapse when the band nears.

Historical texture that fits. North emphasizes the role of credible commitment in the *Glorious Revolution* settlement for capital-market deepening and modern growth; our rails-and-latency evidence (S3) shows the other face of the same coin: commitment devices that speed coordination can also speed *activation* under load. The design question is not "fewer institutions," but *better-shaped rails*: keep everyday efficiency, avoid cliff-edge synchrony.

How this appendix is used in the SI.

- Read S0–S4 with this appendix in mind: it explains *why* κ should lead EWS and *why* ladders compress.
- S5–S7 then become constructive: if ΔJ is the minimal gear that explains acceleration, the levers that modulate h, ΔJ, and ω are the right places to intervene: exactly North's advice to change *rules and enforcement*, not merely exhort actors.

Key sources:

- North, D. C. Institutions, Institutional Change and Economic Performance. Political Economy of Institutions and Decisions. Cambridge University Press, Cambridge (UK) (1990). DOI: https://doi.org/10.1017/CBO9780511808678.
- North, D.C. Institutions. Journal of Economic Perspective, vol. 5, no. 1, Winter 1991 (pp. 97–112) (1991). DOI: 10.1257/jep.5.1.97.

Appendix D: Related work (institutions ↔ switches)

Institutions and increasing returns. Our argument sits on North's foundation: institutions are the *humanly devised constraints* - formal rules and informal norms with enforcement - that structure interaction and reduce uncertainty (North 1990, 1991). Because rules get cheaper and more valuable as they diffuse, institutional matrices exhibit increasing returns and path dependence; designs that "work" become locked in (North 1990; Arthur 1989, 1994). In parallel, transaction-cost economics explains why actors pick governance forms that economize on coordination and hazard (Williamson 1979, 1985, 1996). The credible-commitment thesis shows how public rules can thicken markets by making state promises believable (North and Weingast 1989). We build directly on this line: once rules make statecraft *cheap and fast*, they also create rails - procedures and publicity that compress latency and manufacture common knowledge - thereby raising the chance that decentralized "yes/no" decisions synchronize under high connectedness.

Thresholds, cascades, and correlation. Classic threshold models show how small local rules can yield system-wide events without central command (Granovetter 1978; Watts 2002; see also Centola and Macy 2007; Watts and Dodds 2007). Percolation and self-organized criticality offer complementary intuitions for connectivity and bursts (Stauffer and Aharony 1994; Bak 1996). Our addition is to locate the source of correlation in institutions themselves - alliances, guarantees, mobilization protocols, standardized communications - and then measure its macro consequences: (i) four accelerating cycles converging on a shared T_c (S1–S2); (ii) day-scale latencies at 1914/1939 (S3); (iii) ordering where a frozen carrier κ (L1) leads pre-whitened texture (L2) by ~25y (S4); and (iv) a minimal ABM in which fixed micro-rules plus a single ratchet ΔJ reproduce cadence and map severity to Population × rails (S5–S6).

Early-warning signals and our "ordering" result. Work on critical transitions links rising autocorrelation/variance to approach of tipping points (Scheffer et al. 2009; Dakos et al. 2012; Boettiger and Hastings 2013), with cautions about detrending, filtering, and false positives (Kuehn 2011). We adopt that discipline - pre-whitening (trend + AR(1)) and z-scoring - and add what is often missing: temporal ordering of structure → texture. Our κ → EWS lag curve with block-bootstrap CIs shows the carrier tightens first; behavioural stickiness follows (S4).

Networks, modularity, and design. Network science links modularity and heterogeneity to damped contagion and localized failure (Newman 2006; Fortunato 2010). That is exactly the logic behind our decorrelation lever (heterogeneity, staggered clocks, bulkheads) and reverse-ratchet (slowing the increasing-returns thickening of J). In North's own terms, durable change targets the rules and enforcement, not personalities: leakage (venting before the band), $\Delta J \downarrow$, and $\omega \downarrow$ are institutional moves (S7).

Where we differ. Prior literatures treat uncertainty reduction and coordination as mostly unalloyed goods; cascade models show *how* synchronization can occur but do not say where its correlation comes from in security orders. Our contribution is to show that institutions themselves - the very rails that lower transaction costs - couple thresholds and compress latency, producing a hazard band under high connectedness; then we measure it end-to-end and derive actionable levers that keep everyday efficiency while avoiding cliff-edge synchrony.

Selected references:

- North, D. C. *Institutions, Institutional Change and Economic Performance. Political Economy of Institutions and Decisions*. Cambridge University Press, Cambridge (UK) (1990). DOI: https://doi.org/10.1017/CBO9780511808678.
- North, D.C. *Institutions*. Journal of Economic Perspective, vol. 5, no. 1, Winter 1991 (pp. 97–112) (1991). DOI: 10.1257/jep.5.1.97.
- Arthur, W. B. *Competing Technologies, Increasing Returns, and Lock-In by Historical Events*. Economic Journal, Royal Economic Society, vol 99 (394) (1989). DOI: 10.2307/2234208.
- Arthur, W. B. *Increasing Returns and Path Dependence in the Economy*. University of Michigan Press (1994). DOI: https://doi.org/10.3998/mpub.10029.
- Williamson, O. E. *Transaction-Cost Economics: The Governance of Contractual Relations*. Journal of Law and Economics, Vol. 22, No. 2 (Oct., 1979), pp. 233-261 (1979). DOI: http://dx.doi.org/10.1086/466942.
- Williamson, O. E. *The Economic Institutions of Capitalism. The Mechanisms of Governance*. Free Press (1998).
- North, D. C., Weingast, B. *Constitutions and commitment*. The Journal of Economic History 49(4): 803–832. (1989). DOI: 10.2307/2120804.
- Granovetter, M. *Threshold Models of Collective Behavior*. American Journal of Sociology Volume 83 (6): 1420–1443, The University of Chicago Press. (1978). DOI: https://doi.org/10.1086/226707.
- Watts, D. J. *A simple model of global cascades on random networks*. *PNAS* 99(9): 5766–5771 (2002). DOI: https://doi.org/10.1073/pnas.082090499.
- Centola, D., Macy, M. *Complex contagions and the Weakness of Long Ties*. American Journal of Sociology Vol. 113, No. 3 (November 2007), pp. 702-734 (33 pages), The University of Chicago Press (2007). DOI: https://doi.org/10.1086/521848.
- Stauffer, D., Aharony, A. *Introduction to Percolation Theory*. Taylor & Francis London (1992). DOI: https://doi.org/10.1201/9781315274386.
- Bak, P. *How Nature Works*. *The science of self-organized criticality.* Springer-Verlag New York Inc. (1996). DOI: https://doi.org/10.1007/978-1-4757-5426-1.
- Scheffer, M. et al. *Early-warning signals for critical transitions*. Nature 461, 53–59 (2009). https://doi.org/10.1038/nature08227.
- Dakos, V. et al. *Methods for Detecting Early Warnings of Critical Transitions in Time Series Illustrated Using Simulated Ecological Data*. PLOS One (2012). DOI: https://doi.org/10.1371/journal.pone.0041010.
- Boettiger, C. and Hastings, A. *No early warning for stochastic transitions*. Proceedings of the Royal Society B (2013). DOI: 10.1098/rspb.2013.1372.
- Kuehn, C. *A mathematical framework for critical transitions: normal forms, variance and applications* Journal of Nonlinear Science, Vol. 23, No. 3, pp. 457-510 (2013). DOI: 10.1007/s00332-012-9158-x.
- Newman, M. E. J. *Modularity and community structure*. PNAS 103(23): 8577–8582 (2006). DOI: https://doi.org/10.1073/pnas.0601602103.
- Fortunato, S. *Community detection in graphs*. Physics Reports 486(3–5): 75–174 (2009). DOI: 10.1016/j.physrep.2009.11.002.

Appendix E: Post-1945 note

Our core claim concerns 1453/1495–1945: simple, invariant interaction rules embedded in institutions generated four accelerating systemic-war cycles. Post-1945, the rules persist at global scale (anarchy, sovereign equality, feasible-set compression), but two modifiers matter. First, nuclear-escalation risk raises the effective activation threshold and lowers automaticity for direct great-power combat, suppressing GP-war frequency. Second, cross-domain channels - sanctions, trade, finance, cyber, information - both vent and create tension, increasing safe leakage. In our minimal model these modifiers lengthen intervals and reduce mapped severity, consistent with our intervention experiments (Section 7). Empirically, a global connectedness carrier (κ) rises after 1965, while pre-whitened early-warning signals show an abnormal Cold-War variance episode in the mid-1970s and a recent MILEX episode (2020–2025). We do not present a calibrated DS2 cycle here; rather, the DS1 mechanism plausibly persists with altered thresholds and outlets. This framing generates testable predictions - particularly latency compression and cross-domain synchronization as κ remains elevated - even without direct GP war. A preregistered latency analysis (global trigger → response pairs) and additional DS2 indicators are reserved for a companion study.

Appendix F: Leakage in RSP $\rightarrow$ next $-$ SW severity

(RSP2 abnormal dynamics (1648–1792)

This appendix connects the model's "leakage first" lever to history. We show that higher leakage in the relatively stable period (RSP) - measured as total non-systemic BCD - is associated with lower severity in the next systemic war (SW). The standout case is RSP2 (1648–1792): prolonged FR–GB rivalry and frequent conflicts produced abnormally high leakage, and the subsequent SW2 (1792–1815) is comparatively milder than a monotone trend would imply. This historical pattern complements the model's mechanism: early, reversible dissipation outside the hazard band lowers the eventual break severity without requiring changes to micro-rules or downstream retuning. (This contextual note summarizes evidence captured quantitatively in SI_leakage_nextSW.csv and the two leakage figures.)

Data and definitions (cycle-level, transparent):

We use your cycle windows from Cycle bounds, 1495–1945.csv and cycle totals from BCD START 1495.csv. For each cycle i:

- $Leak_i$ = *sum of non-systemic BCD in* RSP_i,
- $Severity_i$ = BCD of SW_i.
  We plot both the absolute leakage vs next-SW severity and the leakage share of each cycle's total BCD. (Per-capita variants are optional and unnecessary for the main claim.)

What the figures show:

- Figure S7-X (leakage → next SW) plots RSP leakage (bars) with the next SW severity (points). RSP2→SW2 is annotated: high leakage; lower next SW.
- Figure S7-Y (leakage share by cycle) shows the percentage of each cycle's total BCD that occurs in the RSP (non-systemic): C1 ≈ 33.8%, C2 ≈ 66.5%, C3 ≈ 8.2%, C4 ≈ 0.4%. The RSP2 share dominates, matching the narrative of abnormal war dynamics under intense FR–GB rivalry and frequent conflicts.

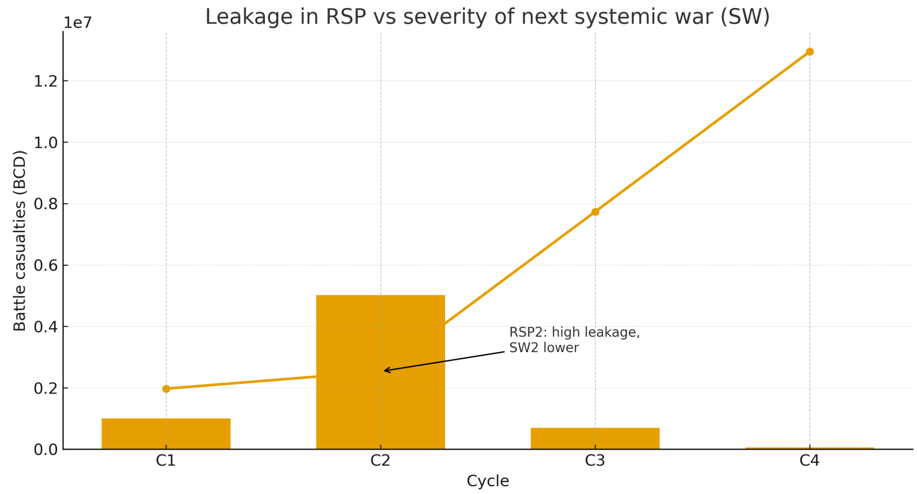

Figure F-1: Leakage during RSP vs severity of the next systemic war (SW). *Bars show total non-systemic BCD within each RSP; points show the subsequent SW severity. RSP2 exhibits high leakage and lower SW2 severity, consistent with the "leakage first" mechanism; other cycles show the complementary pattern (calm RSP → severe SW). See Methods and data files Cycle bounds, 1495–1945.csv, BCD START 1495.csv.*

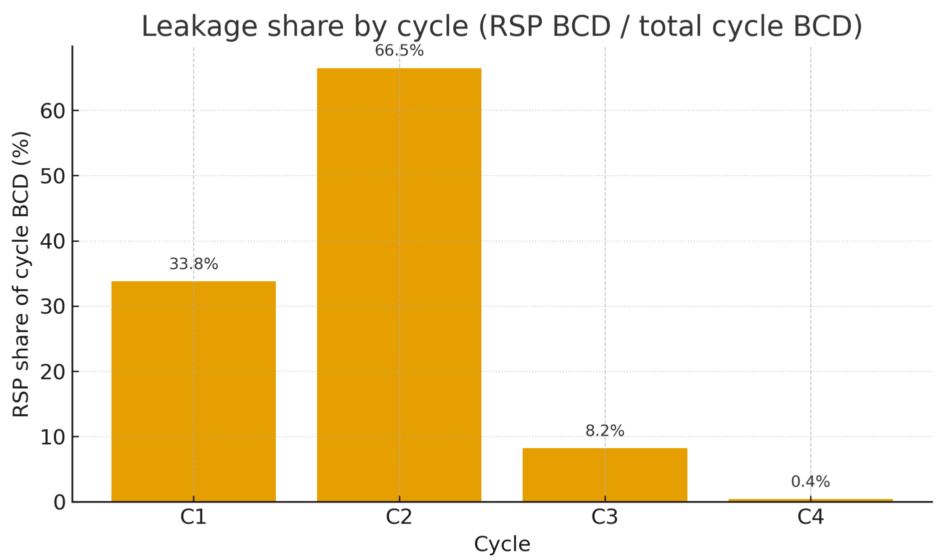

Figure F-2: Leakage share by cycle (RSP BCD / total cycle BCD). *Bars show, for each cycle C1–C4, the percentage of total BCD that occurs in the relatively stable period (RSP) (i.e., non-systemic wars). Shares are C1 ≈ 33.8%, C2 ≈ 66.5%, C3 ≈ 8.2%, C4 ≈ 0.4%. The C2 bar highlights the documented high-leakage RSP2 (1648–1792) under intense FR–GB rivalry; the subsequent systemic war (SW2) is comparatively milder, consistent with the "leakage first" mechanism.*

How to read this. The box does not claim precise causal magnitudes at n = 4; it establishes the direction that the mechanism predicts and provides the historical case (RSP2) where high leakage coincides with a softer next systemic war. The ABM explains why (early venting delays band-entry without growing J); this box shows that the pattern is present across the four cycles.

Files and reproducibility:

- SI_leakage_nextSW.csv: RSP/SW bounds and BCD totals by cycle.
- fig_SI_leakage_vs_nextSW.png: leakage (bars) vs next-SW severity (points).
- SI_leakage_share_by_cycle.csv: leakage share (RSP / total) per cycle.
- fig_SI_leakage_share_by_cycle.png: four-bar percentage plot.

Scope and limits. Early modern precision is coarser than 1914/1939 cabinet ladders, but the cycle-level pattern is robust enough to corroborate the lever. We avoid per-cycle tuning and keep definitions auditable.

References:

- Levy, J. S. *War in the Modern Great Power System, 1495–1975*. University Press of Kentucky, Lexington (KY) (1983).
- Blanning, T. C. W. *The French Revolutionary Wars 1787-1802*. Arnold, London, New York (1996).
- Schroeder, P. W. *The Transformation of European Politics 1763-1848* (Oxford, Clarendon Press (1994).
- Tilly, C. *Coercion, Capital and European States, AD 990-1992*. Basil Blackwell, United Kingdom, United States1 (1990).
- Browning, R. M. *The War of the Austrian Succession*. St Martin's Press, New York (1993).

SI Manifest (artefacts and source data)

Conventions. All files are named exactly as below. ABM runs are seed-locked (seed=20250912). L1/L2 series are on a 5-year end-year grid. "Source data (CSV)" items provide the exact arrays behind each figure.

Raw / input series (used anywhere in S1–S7):

- D6_kappa_verified.csv (L1 κ)
- D20_EWS_MILEX_AC1_prewhitened_5y.csv, D21_EWS_AllianceChurn_Var_prewhitened_5y.csv (L2 pre-whitened EWS)
- Population Europe, 1400-2023, OWID.csv
- alliance_v4.1_by_dyad_yearly.csv (raw exposure)
- Cycle bounds, 1495-1945.csv, BCD START 1495.csv

S0: Methods and falsifiers (summary):

- *(No figures)*

S1: Cadence → T_c (C1–C3 fit; C4 held-out):

- *Figure S1-A*: fig_tc_cadence.png: Cadence fit on C1–C3; C4 held-out (distance to T_c view).
  Source data (CSV): tc_cadence_stats.csv (observed L, best T_c, $\hat{a}$, $\widehat{L}$, residuals).
- *Extended Data (ED1-A)*: ed1a_cadence_linear.png: L vs d = $T_c - t$, C1–C3 fit, C4 open marker.
- *Sensitivity (CSV)*: tc_cadence_1453_1495_sensitivity.csv, tc_cadence_jackknife.csv.

S2: Severity vs distance (shared T_c):

- Figure S2-A: fig_tc_logS.png: log S vs calendar year with T_c fixed from S1.
  Source data (CSV): tc_logS_stats.csv (log d, log S, fitted values, residuals).
- Extended Data (ED1-B): ed1b_severity_loglog.png: log S vs log d with OLS line (C1–C4).

S3: Protocol exposure and decision latency:

- *Figure S3-A*: fig_1914_adjmatrix.png: Exposure matrix (D/G), 1914.
  Source data (CSV): adjmatrix_1914.csv (optionally backbone_edges_1914.csv).
- *Figure S3-B*: fig_1939_adjmatrix.png: Exposure matrix (D/G), 1939.
  Source data (CSV): adjmatrix_1939.csv (optionally backbone_edges_1939.csv).
- *Figure S3-C*: fig_latency_two_panel_1914_1939.png: Latency: 1914 vs 1939 side-by-side.
  Source data (CSV): latency_1914.csv, latency_1939.csv.
- *Metrics CSV*: protocol_metrics_rev.csv (reachability, k-core, medians).

S4: Ordering: κ leads pre-whitened EWS:

- *Figure S4-A*: fig_ordering_lag_with_ci.png: $\kappa \rightarrow$ EWS lag-correlation with 95% block-bootstrap CIs (peak ~25y).
  Source data (CSV): ordering_lag_ci.csv.
- *Figure S4-B*: fig_ordering_placebo_milex.png: Placebo histogram (MILEX AC1).
- *Figure S4-C*: fig_ordering_placebo_alliance.png: Placebo histogram (alliance variance).
- *Figure S4-D*: fig_ordering_placebo.png: Placebo histogram (composite — conservative).
  Source data (CSV): ordering_placebo_per_series.csv, ordering_placebo_sensitivity.csv, ordering_placebo_stats.csv.
- *Inputs (CSV)*: D6_kappa_verified.csv, D20_EWS_MILEX_AC1_prewhitened_5y.csv, D21_EWS_AllianceChurn_Var_prewhitened_5y.csv.

S5: ABM sufficiency (fixed micro-rules; post-onset ratchet only):

- *Figure S5-A*: fig_abm_cycles_timeline.png: Timeline: shaded RSP lifespans, black step J, onset lollipops $\propto J_0 + J$.
  Source data (CSV): abm_series.csv, abm_onsets.csv.
- *Figures S5-B/C (robustness)*: fig_si_abm_N4_intervals_J.png, fig_si_abm_N8_intervals_J.png.
  Source data (CSV): abm_si_N4_series.csv, abm_si_N4_onsets.csv, abm_si_N8_series.csv, abm_si_N8_onsets.csv, abm_si_variants_summary.csv.
- Script: abm_s5_baseline.py (seed-locked).

S6: Mapping $S = \rho \cdot Pop \cdot (J_0 + J)$ (C1–C3 fit; SW4 uplift; variants):

- *Figure S6-A*: fig_s6_bcd_mapping.png: BCD mapping: observed vs ABM × Population; SW4 uplift marker (triangle).
  Source data (CSV/JSON): abm_s6_mapping.csv, abm_s6_summary.json ($\hat{\rho}, \widehat{J_0}, k$).
- *Figure S6-B* (display): fig_s6_bcd_share.png: Population shares: observed vs baseline and SW4 overlay.
- *Figure S6-C* (montage): fig_si_montage.png: Mid-war population; EWS-amplifier (no-tune); N ∈ {4,8} checks.
  Variant CSVs: abm_si_midwar_mapping.csv, abm_si_globalL1_mapping.csv.

S7: Interventions (safe leakage, reverse-ratchet, decorrelation):

- *Figure S7-A*: fig_interventions_intervals.png: Inter-onset intervals: baseline vs levers.
- *Figure S7-B*: fig_interventions_severity.png: Mapped severity per event: baseline vs levers.
- *Figure S7-C*: fig_s7_montage.png: Compact panel.
  Source data (CSV): interventions_stats.csv, interventions_effect_sizes.csv.
- Script: s7_interventions.py.

S7-Appendix — Leakage in RSP → next SW severity (empirical check):

- *Figure S7-X*: fig_SI_leakage_vs_nextSW.png: RSP leakage (bars) vs next-SW severity (points), annotated RSP2→SW2.
  Source data (CSV): SI_leakage_nextSW.csv.
- *Figure S7-Y*: fig_SI_leakage_share_by_cycle.png: Leakage share per cycle = (RSP BCD / total BCD) × 100.
  Source data (CSV): SI_leakage_share_by_cycle.csv.
- *Inputs (CSV)*: Cycle bounds, 1495-1945.csv, BCD START 1495.csv.

Glossary and notation:

- *(No figures)* — symbol list as in the SI front matter.

Data and scripts (reference list):

- L1/L2 inputs: D6_kappa_verified.csv, D20_EWS_MILEX_AC1_prewhitened_5y.csv, D21_EWS_AllianceChurn_Var_prewhitened_5y.csv.
- ABM and mapping scripts: abm_s5_baseline.py, si_panels_build.py (if used), abm_s6_mapping.py, s4_ordering_tweaks.py, s7_interventions.py.
- Population: Population Europe, 1400-2023, OWID.csv.
- Alliance exposure: alliance_v4.1_by_dyad_yearly.csv (raw), adjmatrix_*.csv (derived).

Appendices:

- Appendix A: Evidence spine (quick reference).
- Appendix B: Design companion.
- Appendix C: North; extended synthesis.
- Appendix D: Related work (compact map).
- Appendix E: Post-1945 prospective.
- Appendix F: Leakage in RSP → next SW severity.

Files: SI_leakage_nextSW.csv, fig_SI_leakage_vs_nextSW.png, SI_leakage_share_by_cycle.csv, fig_SI_leakage_share_by_cycle.png.