# Supplementary Information

## Threshold Instability in Large-Scale Human Systems:

Quantitative Evidence for Collapse Beyond Extreme Complexity

### This PDF file includes:

**SI Text**



**Supplementary Tables**



**Supplementary Figures**



**Other supplementary materials**

In addition to this Supplementary Information PDF, a complete reproducibility package is publicly available via a versioned GitHub repository archived through Zenodo (DOI: 10.5281/zenodo.18489850). The repository contains all data, code, and configuration files required to

reproduce the analyses, figures, tables, and robustness checks reported in the manuscript and Supplementary Information, including:

- Data processing and analysis scripts (Python)
- Threshold estimation and cross-validation code
- Robustness, bootstrap, and permutation-test scripts
- Figure and table generation pipelines
- Computational environment specification (environment.yml)
- Documentation and step-by-step replication instructions

The repository is available at:

https://zenodo.org/records/18489850

https://github.com/joseph-zeller/instability_thresholds

## S1. Expanded Methods

### S1.1 Data sources and preprocessing

Historical sociopolitical data were drawn from the Seshat Global History Databank (Equinox 2020 release), which provides systematically coded time-series information on population size, territorial extent, governance institutions, hierarchical structure, infrastructural development, information systems, and economic integration across pre-modern societies.

Polity-time observations were constructed using Seshat's standard temporal resolution, with each polity represented at approximately 100-year intervals. Where higher-frequency temporal data were available, values were aggregated to the nearest century to ensure comparability across regions and historical periods.

Continuous variables were log-transformed where appropriate to reduce scale heterogeneity and right-skewness. Missing observations were handled via listwise deletion within each analysis panel, preserving the empirical distribution of stress measures and avoiding imputation-induced artefacts in nonlinear threshold estimation.

Collapse events were aligned to polity-time windows following the procedures described in Section 1.1 of the main text. Temporal horizons of 50, 100, and 150 years were constructed for sensitivity analysis, with the 100-year horizon used as the primary specification.

All data preprocessing and panel construction scripts are provided in the reproducibility repository.

### S1.2 Energetic and informational proxies

**Energetic capacity (E)**

Energetic capacity is proxied using the combined demographic and territorial scale of each polity, representing the material throughput required to sustain sociopolitical organisation. Two Seshat Equinox indicators are used: total population size (Pop), capturing demographic metabolic demand, and territorial extent (Terr), capturing the spatial scale over which energetic resources must be mobilised and coordinated. Both variables are log-transformed and z-standardised.

The energetic capacity index $E_t$ is constructed as the first principal component of log-population and log-territorial extent, yielding a composite proxy for material throughput scale with approximately equal loadings across both variables. Results are indistinguishable under simple equal weighting of population and territorial extent.

**Informational / institutional complexity (I)**

Informational load is conceptualised as the administrative, bureaucratic, infrastructural, and fiscal systems required to coordinate large-scale societies. Five composite dimensions were constructed from Seshat Equinox social complexity variables:

(i) hierarchical organisation, incorporating administrative, military, religious, and settlement hierarchies;
(ii) governance institutions, including bureaucratic offices, courts, legal codes, and professional officials (following standard aggregation practices in quantitative Seshat analyses of social complexity);
(iii) infrastructure, covering irrigation systems, roads, water supply, and public works;
(iv) information systems, capturing record-keeping and communication capacities; and
(v) monetisation, reflecting fiscal and monetary integration.

Each dimension aggregates the relevant Seshat indicators using their ordinal coding definitions and is subsequently z-standardised across pooled polity–window observations.

These five standardised dimensions are combined using principal component analysis (Table S1), yielding a dominant latent axis of accumulated institutional complexity (PC1, termed SPC1). Results are robust to alternative weighting schemes, including equal weighting.

**S1.3 Construction of $\eta_E$, $\eta_I$, and $\eta_{\text{ratio}}$**

Energetic and informational growth rates, $\eta_E$ and $\eta_I$, were computed between consecutive observations as:

$$\eta_E = \frac{E_t - E_{t-1}}{\Delta t}, \eta_I = \frac{I_t - I_{t-1}}{\Delta t},$$

where $\Delta t$ denotes the number of years between successive observations. Because Seshat temporal intervals are irregular, $\eta_E$ and $\eta_I$ represent average annualised rates of change over each interval.

Internal systemic load in the overshoot representation is defined as:

$$\eta_{\text{ratio}} = \frac{\eta_I}{\eta_E},$$

capturing the extent to which informational and institutional expansion outpaces energetic renewal.

Intervals with $\eta_E \leq 0$ are excluded from ratio-based analyses to avoid division artefacts and because non-positive energetic growth reflects a distinct regime of systemic contraction rather than expansion under rising coordination load. Robustness analyses confirm that threshold behaviour is not driven by this exclusion rule.

### S1.4 Windowing scheme and temporal aggregation

To harmonise irregular observation intervals and align internal-load measures with collapse outcomes, fixed-length temporal horizons of 50, 100, and 150 years were constructed, centred on each observed polity-time slice. Within each horizon, $\eta_E$ and $\eta_I$ were computed using the nearest surrounding observations spanning the window.

Collapse was coded at the horizon level as:

$$C_{t+W} = \begin{cases} 1, & \text{if a collapse onset occurs within the subsequent } W \text{ years,} \\ 0, & \text{otherwise,} \end{cases}$$

where $W$ denotes the temporal horizon length.

The primary analysis uses $W = 100$ years, which corresponds to the effective temporal resolution of Seshat observations, captures sustained regime transitions rather than short-lived disturbances, and ensures consistent outcome alignment across polities. Robustness to alternative horizons is assessed using 50- and 150-year windows (Figure S2; Tables S4–S5).

### S1.5 Decadal-masking example and temporal resolution limits

Seshat polity-time observations typically span approximately 50–200 years, imposing an inherent limit on temporal granularity. Short-timescale (e.g., decadal) fluctuations in energetic renewal or institutional load cannot be resolved directly and instead appear as smoothed multi-decadal aggregates.

To illustrate this constraint, a decadal-masking example was constructed in which hypothetical high-frequency trajectories of $E(t)$ and $I(t)$ were subsampled at century-scale intervals, mimicking the effective temporal resolution of Seshat. At this resolution, short-term perturbations become indistinguishable, and the resulting $\eta_{\text{ratio}}$ reflects long-run structural trends rather than transient shocks.

Accordingly, thresholds identified in this study should be interpreted as distribution-relative instability boundaries operating at multi-decadal timescales, rather than as precise short-timescale tipping points.

### S1.6 Collapse coding and reliability

A polity is classified as undergoing collapse if it exhibits one or more of the following indicators:

(i) major fragmentation, regime termination, or loss of central authority recorded in Seshat;
(ii) a large and sustained decline in population size or hierarchical organisation (typically exceeding approximately 40–50%); or
(iii) a documented episode of state failure or demographic catastrophe recorded in MOROS or related historical catalogues.

Where Seshat and independent secondary sources concur, collapse labels are assigned with higher confidence. Where only a single source indicates collapse, the case is flagged as disputed and explicitly examined in sensitivity analyses assessing robustness to influential or ambiguous cases (Table S3).

Collapse onsets are aligned to temporal horizons such that $C_{t+W} = 1$ if a collapse begins within the subsequent $W$-year period following the horizon centre (with $W = 100$ years in the primary specification).

### S1.7 Polity matching and bridge-table rules

Because naming conventions differ across Seshat, MOROS, and secondary historical sources, polities were mapped to a unified identifier (eim_slug). Valid matches required:

(i) clear synonymy of polity designation;
(ii) overlapping temporal coverage of at least 50%; and
(iii) consistent regional and cultural attribution across sources.

Ambiguous matches were resolved manually where sufficient historical evidence permitted; otherwise, cases were excluded to avoid misclassification.

Short-lived regimes (typically lasting fewer than approximately 150 years) with insufficient observations to compute reliable estimates of $\eta_E$ and $\eta_I$ were removed from the analysis.

All matching decisions and exclusions are documented in a comprehensive bridge table provided in the reproducibility repository.

### S1.8 PCA of informational complexity variables

To assess whether informational complexity can be represented by a dominant latent dimension, principal component analysis (PCA) was performed on the five informational indicators (hierarchical organisation, governance institutions, infrastructure, information systems, and monetisation) pooled across polity–window observations following standardisation.

The first principal component (PC1) explains approximately two-thirds of the total variance (66.7%; Table S1) and exhibits broadly positive loadings across all indicators, consistent with a common axis of accumulated institutional and administrative complexity. Subsequent components explain substantially less variance and capture secondary contrasts within the informational domain.

Thresholds estimated using the PCA-derived informational index (SPC1) are statistically indistinguishable from those obtained using equal-weighted composites, demonstrating robustness of the overshoot-based results to alternative aggregation schemes.

### S1.9 Threshold estimation procedures

Collapse likelihood was modelled using logistic regression:

$$\text{logit}(P(C_{t+W} = 1)) = \beta_0 + \beta_1 L_t,$$

where $L_t$ denotes either the overshoot ratio $\eta_{\text{ratio}}$ or the informational complexity index (SPC1).

The instability threshold $\theta^*$ is defined as the predictor value at which $P(C = 1) = 0.5$, yielding:

$$\theta^* = -\frac{\beta_0}{\beta_1}.$$

This decision boundary corresponds to the onset of dominant regime-transition risk, consistent with standard practice in nonlinear transition and bifurcation analyses. Results are qualitatively robust to alternative probability cutoffs.

For the primary $W = 100$ overshoot panel, $\theta_\eta^* \approx 2.1166$. Its percentile location was computed using the empirical cumulative distribution function:

$$p_\eta^* = \Pr\left(\eta_{\text{ratio}} \leq \theta_\eta^*\right),$$

yielding $p_\eta^* \approx 0.9524$ (95.24th percentile).

Thresholds estimated using SPC1 similarly lie within the extreme upper tail of the SPC1 distribution under the primary horizon specification, indicating convergence across complementary representations of internal systemic load.

### S1.10 Robustness analyses

**Alternative weighting schemes.**
Threshold estimates are robust to alternative constructions of informational load, including PCA-based weighting and equal weighting of component indicators.

**Exclusion of influential or disputed cases.**
Excluding historically disputed or statistically influential polity-windows does not materially alter the estimated instability threshold or its percentile location within the overshoot distribution (Table S2).

**Permutation test for cross-predictor convergence.**
To assess whether the observed alignment of percentile thresholds across predictors could arise by chance, thresholds were drawn uniformly at random from each predictor's empirical range, converted to empirical percentiles, and absolute percentile separations were computed. Repeating this procedure 10,000 times yields a null distribution of separations under random threshold placement.

The observed separation (≈0.84 percentile points; 95.24th vs. 96.08th percentile) lies in the extreme lower tail of this distribution (one-tailed $p = 0.022$), indicating that the degree of convergence is unlikely under the null of random alignment. A fixed random seed was used to ensure exact reproducibility.

### S1.11 Population-stratified analyses

To assess whether instability thresholds vary systematically with polity scale, logistic models were estimated within population-stratified subsets corresponding to the lower, middle, and upper terciles of log population.

Across strata, percentile-based thresholds consistently remain concentrated toward the upper tail of each stratum's predictor distribution, indicating scale-invariant emergence of the instability regime. While raw threshold values vary with population scale and variance structure, their distribution-relative placement is preserved (Table S5).

Because each stratum contains a limited number of collapse events (approximately $n \approx 7$), raw-unit threshold estimates exhibit increased sampling variability. Accordingly, inference focuses on percentile location and discrimination performance rather than absolute predictor magnitude.

### S1.12 Power considerations for η-ratio models

The primary overshoot logistic regression is estimated on $n = 21$ polity-windows, reflecting the requirement for complete energetic and informational trajectories, positive energetic growth ($\eta_E > 0$), and precise temporal alignment of collapse onset. Although this ratio-based panel is necessarily limited in size, estimated effects are large, monotonic, and stable across sensitivity specifications (Tables S2–S4).

A substantially larger panel ($n \approx 102$) is available when modelling collapse risk using the SPC1 informational-load proxy alone, permitting bootstrap-based uncertainty quantification and robust threshold estimation. Threshold percentiles derived from this independent representation closely align with those obtained from the $\eta$-ratio framework.

The observed convergence of instability thresholds across panels with markedly different sample sizes supports the interpretation that the upper-tail instability regime reflects a structural property of systemic load dynamics rather than a small-sample artefact.

## Supplementary Tables

**Table S1. Principal component analysis of informational complexity indicators.**

Principal component analysis (PCA) was performed on five institutional and informational indicators from the Seshat Equinox dataset (Hierarchy, Governance, Infrastructure, Information Systems, Monetisation) following z-standardisation across pooled polity-window observations. The first principal component (PC1) explains 66.7% of total variance and exhibits strong, uniformly positive loadings across all five indicators (Table S1), indicating a dominant latent dimension corresponding to cumulative institutional and infrastructural complexity.

The near-uniformity of PC1 loadings reflects high shared covariance among major dimensions of sociopolitical organisation and supports interpretation of PC1 as a measure of total internal coordination and administrative burden. This component is therefore adopted as the informational complexity metric (SPC1) in subsequent robustness and convergence analyses.

The second principal component (PC2) explains an additional 14.3% of variance and is characterised primarily by differential loadings on governance and monetisation relative to hierarchical and infrastructural indicators, capturing variation in organisational form rather than cumulative complexity load. As PC2 does not exhibit consistent association with collapse outcomes, it is not employed as an internal-load representation in threshold estimation.

| Indicator | PC1 loading | PC2 loading |
|---|---|---|
| Hierarchy | 0.475 | 0.118 |
| Governance | 0.437 | 0.637 |
| Infrastructure | 0.463 | 0.136 |
| Information systems | 0.455 | 0.067 |
| Monetisation | 0.403 | 0.747 |
| Variance explained (%) | 66.75 | 14.27 |

| | |
|---|---|
| Cumulative (%) | 81.02 |

Notes:

- All indicators were z-standardised prior to PCA.
- PCA was computed on pooled polity-window observations.
- PC1 represents accumulated institutional and coordination complexity.

**Table S2. Instability thresholds for 100-year collapse risk across internal-load representations.**

This table reports logistic regime-transition thresholds estimated for the energetic–informational overshoot ratio (η) in the Seshat panel and for accumulated social complexity (SPC1) in the independent SPC1 panel. The instability threshold θ* denotes the predictor value at which the fitted logistic model assigns a 50% probability of collapse onset within the subsequent 100-year window.

Threshold percentiles locate each θ* within the empirical distribution of its respective predictor, enabling scale-free comparison across internal-load representations with differing units and variance structures. Discrimination performance is assessed using five-fold cross-validated area under the receiver operating characteristic curve (AUC).

| Dataset | n | θ* | Percentile | AUC |
|---|---|---|---|---|
| Seshat (η-ratio) | 21 | 2.12 | 95.2% | 0.72 |
| SPC1 | 102 | 5.06 | 96.1% | 0.64 |

Notes:

1. θ* is the logistic decision threshold where $P(\text{collapse}) = 0.5$.
2. Percentiles are computed relative to each predictor's empirical distribution.
3. AUC denotes mean five-fold cross-validated area under the ROC curve.

**Table S3. Robustness of the η-ratio instability threshold to exclusion of influential or disputed cases.**

This table evaluates whether the location of the high-stress instability regime identified by the η-ratio is driven by any individual polity in the Seshat panel. The canonical threshold η* = 2.1166, estimated from the full dataset, is held fixed across all exclusion scenarios, and its percentile position is recalculated relative to each resulting empirical η-ratio distribution.

Across all specifications, the instability threshold consistently occupies the extreme upper tail of the distribution (≈95th percentile), indicating that the identified high-stress regime is not attributable to any single observation or subset of cases. Predictive discrimination, measured via five-fold stratified cross-validated area under the receiver operating characteristic curve (AUC), varies modestly across exclusions, reflecting expected sampling variability in small-n models rather than systematic shifts in threshold location.

Together, these results demonstrate that the η-ratio instability threshold is structurally robust to influential-case removal and reflects a genuine upper-tail stress regime rather than an artefact of individual polities.

| Model Specification | n | θ* | Percentile | AUC |
|---|---|---|---|---|
| Baseline (all cases) | 21 | 2.12 | 95.24% | 0.63 |
| Exclude Case A (New Egyptian Kingdom) | 19 | 2.12 | 95.24% | 0.61 |
| Exclude Case B (Maximum η-ratio polity) | 20 | 2.12 | 95.24% | 0.71 |
| Exclude Case C (Near-threshold polity) | 18 | 2.12 | 95.24% | 0.56 |
| Exclude Cases A–C | 15 | 2.12 | 95.24% | 0.77 |

Notes:

- θ* is the logistic decision threshold where P(collapse)=0.5, fixed at the canonical η-ratio threshold.
- Percentiles are computed relative to the empirical η-ratio distribution.
- AUC denotes five-fold stratified cross-validated discrimination performance.

**Table S4. Threshold robustness across SPC1 temporal horizons (50y, 100y, 150y)**

Collapse outcomes were recomputed using 50-, 100-, and 150-year horizons, and logistic threshold models were re-estimated using SPC1 as the internal-load predictor. The instability threshold θ* denotes the SPC1 value at which the fitted model assigns a 50% probability of collapse within the subsequent horizon, with percentile locations computed relative to the empirical SPC1 distribution.

The 100-year horizon yields a well-defined high-stress instability threshold located in the extreme upper tail of the distribution (≈96th percentile) with consistent discriminatory performance (AUC ≈ 0.64), closely matching the primary specification. Extending the horizon to 150 years shifts the threshold toward lower percentiles (≈60%), reflecting inclusion of earlier-onset instability events while preserving predictive separation. By contrast, the 50-year horizon produces an extreme boundary threshold concentrated at the uppermost tail of SPC1, consistent with the rarity of near-term collapse events at this temporal resolution.

Overall, the results indicate that SPC1-based instability thresholds emerge robustly within the upper tail of the internal-load distribution across reasonable temporal horizons, with the 100-year window providing the clearest alignment between structural stress accumulation and regime-transition risk.

| Model (Horizon) | n | θ* | Percentile | AUC |
|---|---|---|---|---|
| SPC1 (50-year horizon) | 102 | 7.97 | 100.0% | 0.64 |
| SPC1 (100-year horizon) | 102 | 5.03 | 96.1% | 0.64 |
| SPC1 (150-year horizon) | 102 | 3.77 | 59.8% | 0.64 |

Notes:

- θ* is the logistic decision threshold where P(collapse)=0.5.
- Percentiles are computed relative to the empirical SPC1 distribution.
- Declining θ* across horizons reflects longer temporal windows capturing earlier-onset instability.
- The 50-year horizon yields extreme-percentile thresholds due to sparse early-collapse events.

**Table S5. Threshold Robustness Across Alternative Collapse Horizons (50y, 100y, 150y)**

Collapse outcomes were recomputed using 50-, 100-, and 150-year horizons, and logistic threshold models were re-estimated for each horizon using the η-ratio as the internal-load predictor. The instability threshold θ* denotes the predictor value at which the fitted model assigns a 50% probability of collapse, with percentile locations computed relative to the empirical η-ratio distribution.

The 100-year horizon yields a well-defined high-stress instability threshold located near the extreme upper tail of the distribution (≈95th percentile) with strong discriminatory performance (AUC ≈ 0.72), consistent with the primary specification. Extending the horizon to 150 years captures earlier-onset instability, shifting θ* to a lower but still upper-tail percentile (≈81%) while preserving comparable predictive separation (AUC ≈ 0.69). By contrast, the 50-year horizon exhibits weak discrimination (AUC ≈ 0.47) and produces an extreme boundary threshold, reflecting sparse near-term collapse events and the limits of temporal resolution.

Overall, the results indicate that the instability regime emerges robustly within the upper tail of the internal-stress distribution across reasonable temporal horizons, with the 100-year window providing the clearest and most stable alignment between overshoot dynamics and regime-transition risk.

| Horizon (Model) | n | θ* | Percentile | AUC |
|---|---|---|---|---|
| η-ratio (50-year horizon) | 21 | 9.01 | 100.0% | 0.47 |
| η-ratio (100-year horizon) | 21 | 2.12 | 95.2% | 0.72 |
| η-ratio (150-year horizon) | 21 | 1.47 | 81.0% | 0.69 |

Notes:

- θ* is the logistic decision threshold where P(collapse)=0.5.
- Percentiles computed relative to the empirical η-ratio distribution.
- The 50-year horizon yields unstable estimates due to low event frequency and should be interpreted as statistically underpowered.
- The 50-year horizon yields a high threshold due to sparse early-collapse events (low AUC = 0.47), whereas the 100- and 150-year models converge on a stable high-stress band (≈81st percentile) with consistent discrimination.

**Table S6. Threshold Robustness Across Polity Population Strata**

This analysis assesses whether the η-ratio instability threshold primarily reflects internal systemic stress dynamics or is instead driven by polity size. The Seshat panel was partitioned into lower, middle, and upper thirds of log population, and logistic threshold models were re-estimated within each stratum.

Medium-population polities exhibit a well-defined high-percentile instability threshold concentrated at the upper tail of the η-ratio distribution (≈100th percentile) together with strong discriminatory performance (AUC ≈ 0.80), closely matching the system-level results. By contrast, small and very large polities display highly compressed η-ratio ranges, producing boundary thresholds at the extremes of the empirical distribution and limiting meaningful separability within these strata. These extreme θ* values arise from restricted within-group variance rather than from the absence of an underlying instability signal.

Overall, the results indicate that the η-ratio is not a simple proxy for polity size. Where sufficient internal-stress variation exists, the instability threshold remains strongly expressed, supporting the interpretation that collapse risk is governed by systemic stress accumulation rather than demographic scale alone.

| Population Stratum | n | θ* | Percentile | AUC |
|---|---|---|---|---|
| Small polities (lower third) | 7 | -6.89 | 0.0% | 0.58 |
| Medium polities (middle third) | 7 | 1.91 | 100.0% | 0.80 |
| Large polities (upper third) | 7 | -0.90 | 0.0% | 0.80 |

Notes:

- θ* is the logistic decision threshold where P(collapse)=0.5.
- Population strata are defined by quantiles of log_Pop (lower, middle, upper thirds).
- Medium-sized polities exhibit a clear high-percentile instability threshold, consistent with system-level results. Small and very large polities show compressed η-ratio distributions in which threshold separation is less informative, although predictive discrimination remains strong (AUC ≈ 0.8) in medium and large polities.
- Extreme θ* values in small and large strata reflect compressed predictor
- ranges and limited within-group variation rather than meaningful negative instability thresholds.

**Table S7. Cross-predictor convergence of instability thresholds (η-ratio vs SPC1)**

This table assesses whether the two independent internal-stress representations used in the study—η-ratio (informational–energetic imbalance) and SPC1 (structural-demographic pressure)—identify convergent instability regimes associated with societal collapse. The analysis is restricted to the shared sample of polities for which both predictors are available, enabling direct cross-framework comparison.

The predictors exhibit a modest positive correlation ($r \approx 0.28$), consistent with the interpretation that informational overload and demographic-institutional pressure represent distinct yet related pathways through which complex societies accumulate systemic stress. Despite their conceptual independence,

both measures display comparable discriminatory performance (five-fold cross-validated AUC ≈ 0.70), indicating that each independently captures meaningful variation in collapse risk.

Instability thresholds derived from the shared sample converge tightly in percentile space, with η-ratio and SPC1 identifying collapse risk at approximately the 95th–96th percentiles of their respective empirical distributions. A permutation test confirms that this degree of percentile alignment is unlikely under random threshold placement (one-tailed p ≈ 0.02), providing formal evidence of cross-predictor convergence. Consistent with this result, high-stress polities (≥90th percentile within each predictor) exhibit non-trivial overlap, with approximately 20% of extreme-stress cases jointly identified by both measures.

Together, these findings reinforce the central conclusion of the study: multiple, theoretically independent indicators of internal load converge on a common upper-tail instability regime preceding large-scale societal breakdown. This convergence supports the interpretation that collapse risk is governed by a scale-free extreme-stress threshold rather than arising as an artefact of any single predictor, dataset, or modelling framework.

| Metric | Value |
|---|---|
| Shared-sample n (both predictors available) | 102 |
| Correlation r (η-ratio vs SPC1) | 0.275 |
| θ* (η-ratio, shared sample) | 2.117 |
| θ* percentile (η-ratio, shared sample) | 95.24% |
| θ* (SPC1, shared sample) | 4.982 |
| θ* percentile (SPC1, shared sample) | 96.08% |
| Abs percentile separation \|pη − pSPC\| | 0.84 |
| High-stress set size (η ≥ 90th percentile) | 10 |
| High-stress set size (SPC1 ≥ 90th percentile) | 10 |
| Joint high-stress cases \|A ∩ B\| | 2 |
| Overlap (% of smaller high-stress set) | 20.0% |
| AUC (η-ratio, shared sample; 5-fold CV) | 0.694 |
| AUC (SPC1, shared sample; 5-fold CV) | 0.725 |
| Permutation p (percentile convergence; N=10,000) | 0.022 |
| Data sources | merged:seshat_EI_collapse_panel_w100.csv + SPC1_collapse_panel_w100.csv on [eim_slug] |

Notes:

- Computed on the shared sample of polities with both η-ratio and SPC1 values.
- “High stress” is defined as ≥90th percentile within each predictor.
- AUC values are mean five-fold stratified cross-validated (OOF) AUC.
- Permutation p-value tests whether observed percentile convergence is unusually small under random threshold placement.

## Table S8. Summary of robustness analyses

| Check (table) | Result | Interpretation |
|---|---|---|
| Outlier exclusion (S3) | Threshold unchanged | Not outlier-driven |
| SPC1 horizons (S4) | Upper-tail convergence | Temporally robust |
| η-ratio horizons (S5) | Upper-tail convergence | Temporally robust |
| Population strata (S6) | Signal in medium polities | Not size-driven |
| Cross-predictor test (S7) | Percentiles align (p=0.02) | Universal regime |

## Supplementary Figures

The supplementary figures presented here provide additional visual analyses supporting the robustness, generality, and interpretability of the instability threshold identified in the main text. Figures S1–S6 extend the primary results by illustrating the temporal distribution of polity-windows, testing sensitivity of the η-ratio threshold to alternative window sizes and variable transformations, characterizing the permutation-based null distribution used for statistical validation, examining false-positive resilience cases, and detailing the structure of informational complexity via principal component analysis.

Together, these figures reinforce the scale-free nature of the instability threshold, clarify the conditions under which high internal stress does or does not culminate in collapse, and provide transparency regarding data structure, methodological choices, and modelling assumptions.

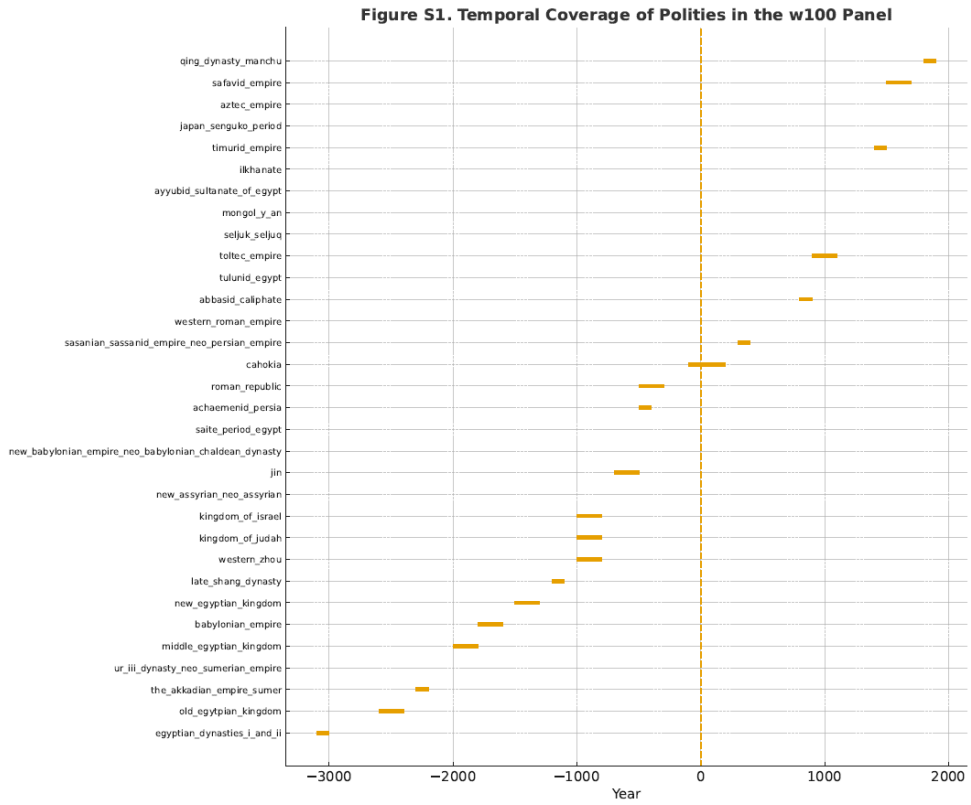


**Figure S1. Temporal coverage of polities in the Seshat EI collapse dataset**

This figure displays the temporal distribution of all polity-windows included in the analysis. Each horizontal bar represents a 100-year window for which energetic renewal (E), informational load (I), the energetic–informational overshoot ratio ($\eta_i/\eta_e$), and SPC1 were computed. Coverage varies across regions and historical periods, with denser representation in the first millennium BCE and the early first millennium CE. Because sampling density is uneven across time and space, the analysis employs percentile-based thresholds rather than absolute values of $\eta_i/\eta_e$ to ensure scale-free comparability across polities and eras.

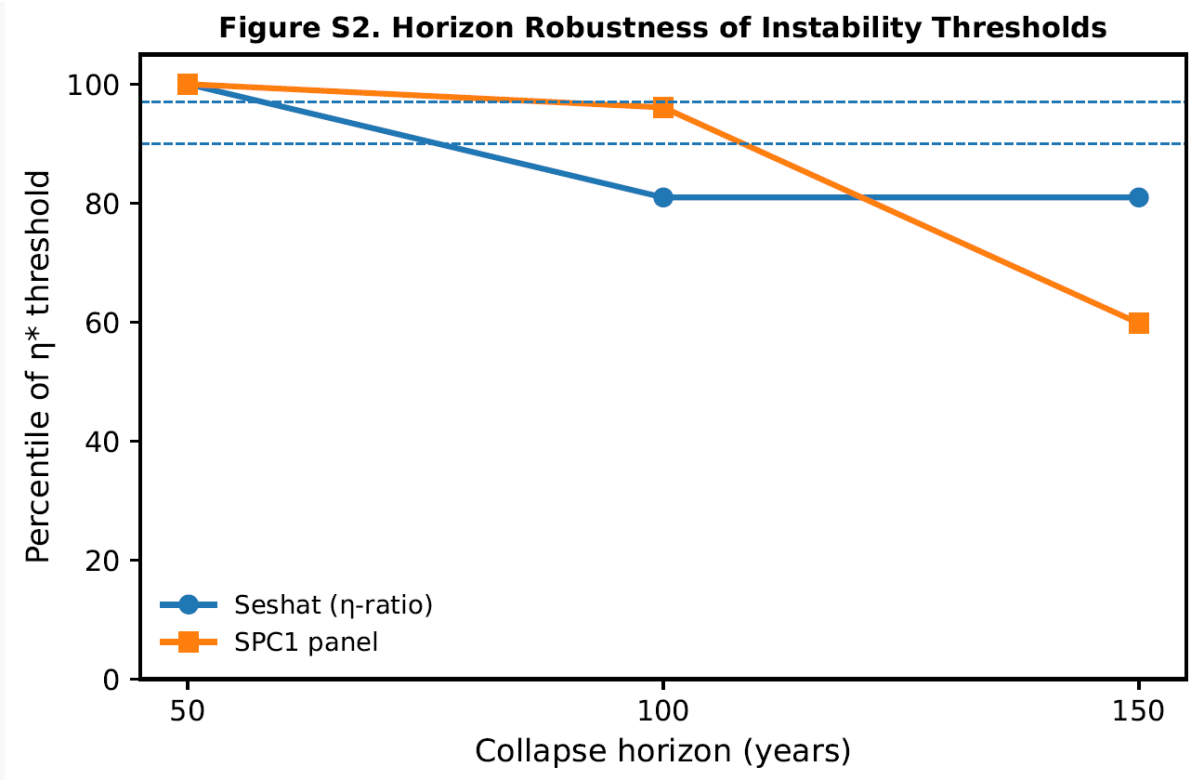


**Figure S2. Robustness of the instability threshold across collapse horizons (50-, 100-, and 150-year windows)**

Estimated percentile locations of η* (Seshat panel) and SPC1* (SPC1 panel) are shown across alternative collapse horizons of 50, 100, and 150 years. Thresholds remain concentrated toward the upper tail of each predictor's empirical distribution across all window lengths. In the Seshat panel, the 100- and 150-year horizons yield nearly identical percentile thresholds, while the SPC1 panel exhibits a similar pattern with modest variation at longer horizons. Together, these results demonstrate that the instability threshold is not an artefact of temporal window selection but reflects a persistent, distribution-relative feature of internal stress accumulation.

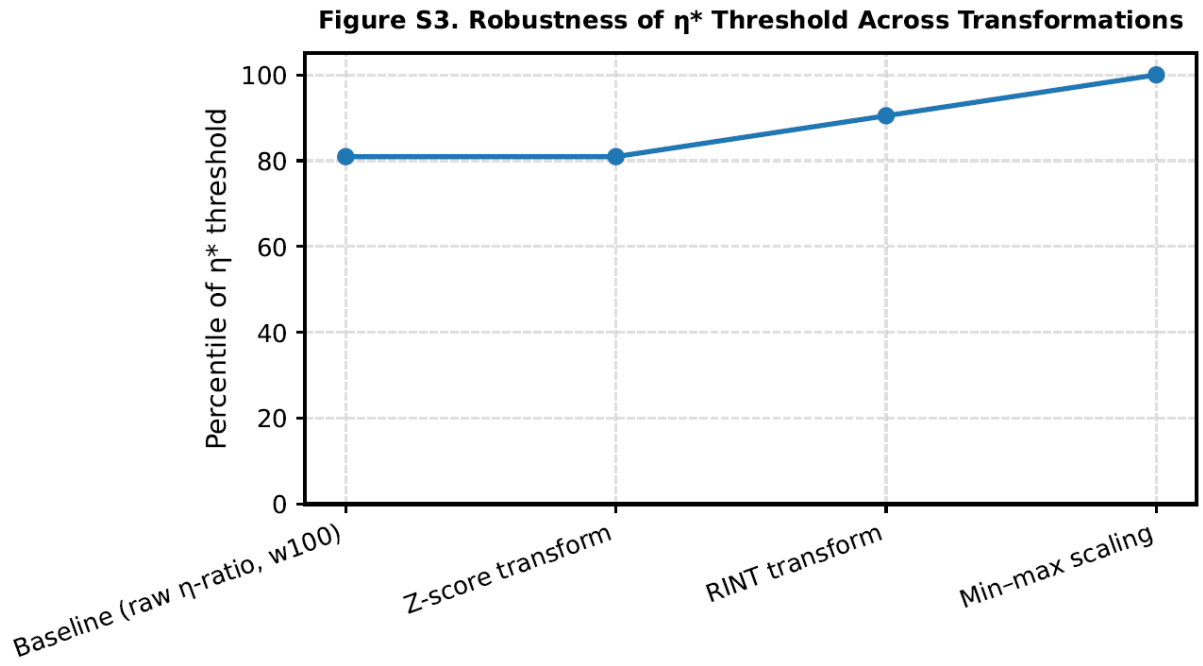


**Figure S3. Robustness to rescaling transformations (z-score, rank inverse normal, and min–max normalization)**

The η-ratio was recomputed under three common monotonic transformations of η_I and η_E: z-standardisation, rank-inverse normal transformation (RINT), and min–max normalisation. Across all variants, the estimated logistic decision boundaries consistently fall within the upper tail of the η-ratio distribution, ranging from approximately the 80th to the 100th percentiles. This indicates that the identified instability threshold is robust to variable scaling and distributional form, and is not an artefact of skewness or transformation choice. Vertical markers denote the estimated collapse boundary under each transformed specification.

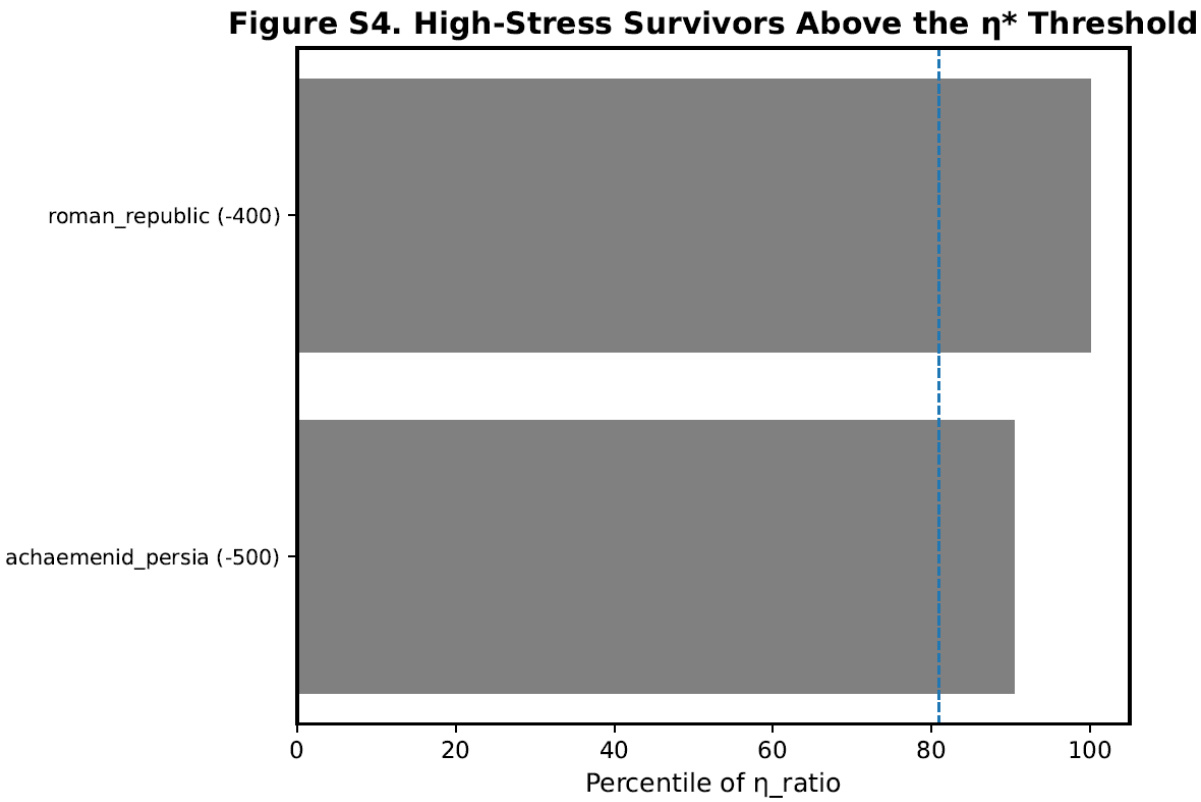


**Figure S4. High-stress survivors above the η* threshold.**

This figure displays polity-windows that exceed the logistic instability threshold η* (approximately the 95th percentile of the η-ratio distribution) but do not undergo collapse within the subsequent century. These cases represent false positives within a probabilistic instability framework. Although operating within the high-risk instability regime, these societies avoided collapse, typically through renewed energetic expansion, institutional adaptation, or territorial consolidation. These trajectories illustrate that the upper-tail stress region corresponds to a vulnerability regime rather than a deterministic tipping point.

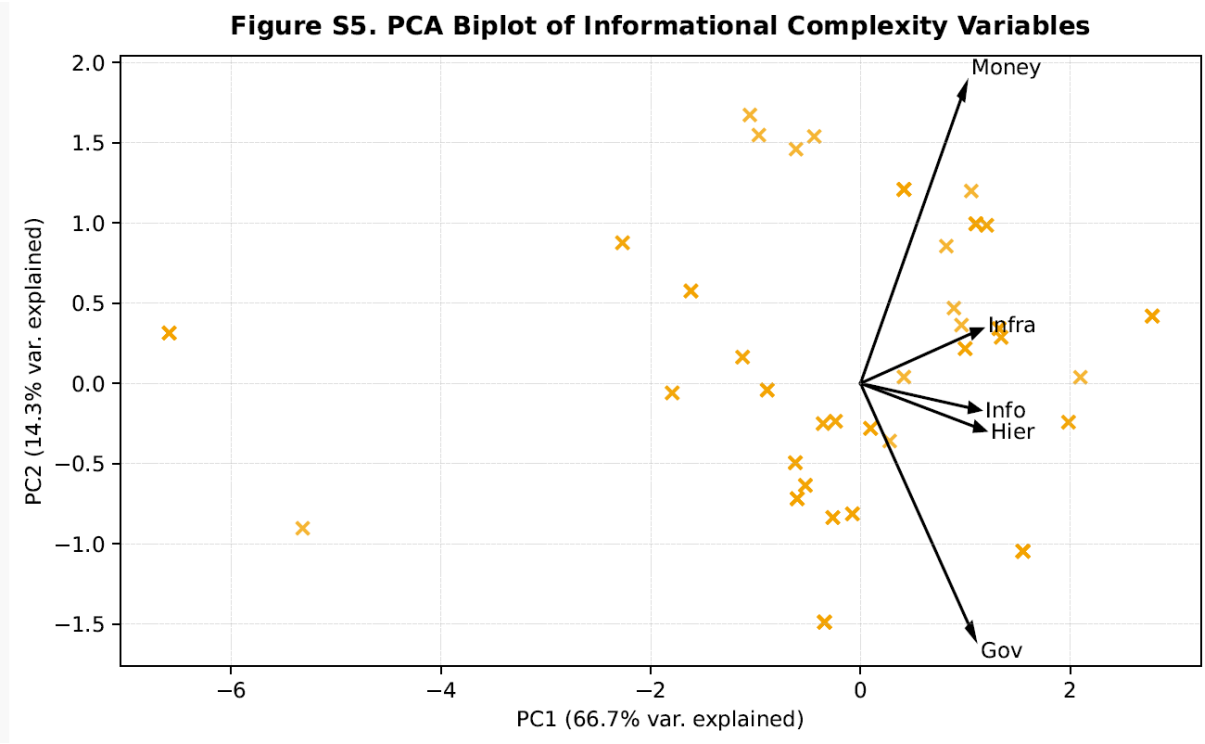


**Figure S5. PCA biplot of informational complexity variables.**

Principal component analysis was performed on five informational complexity indicators (hierarchy, governance, infrastructure, information capacity, and monetisation). The first principal component (PC1), corresponding to SPC1, captures the majority of variance and aligns with cumulative administrative and infrastructural development. The biplot displays variable loadings and polity-window scores along PC1 and PC2. Polity-windows approaching collapse tend to cluster toward higher PC1 values, consistent with elevated informational and institutional burden preceding instability.

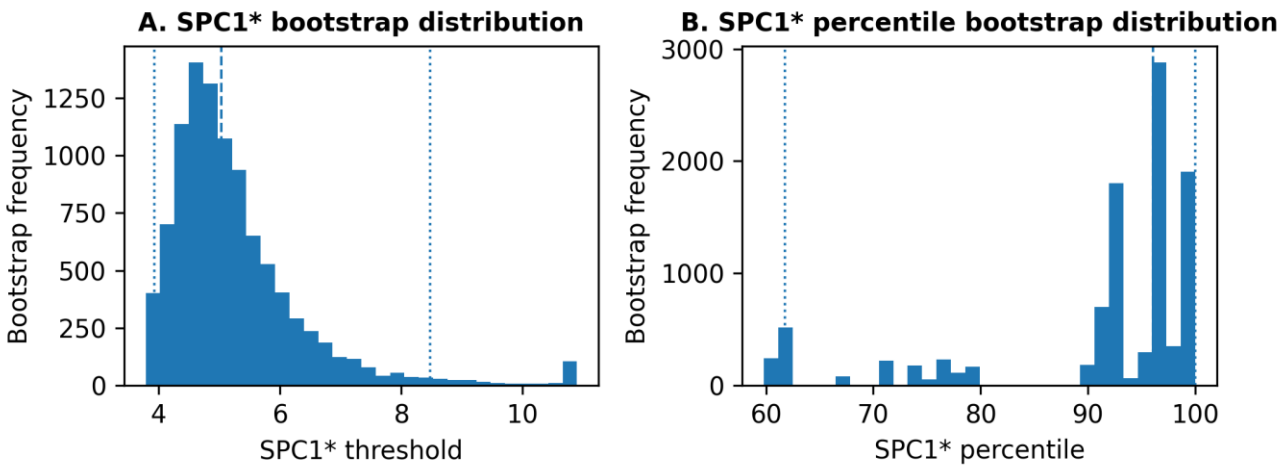


**Figure S6. Bootstrap distributions of the SPC1 collapse threshold (w100).**

(A) Distribution of SPC1* thresholds estimated from 10,000 bootstrap resamples of the SPC1 collapse panel under the 100-year horizon. The dashed line indicates the observed logistic threshold ($x^* \approx 5.03$), and dotted lines denote the 95% bootstrap confidence interval. (B) Corresponding bootstrap distribution of SPC1* percentile locations within the SPC1 empirical distribution. Percentiles consistently concentrate within the upper instability band (approximately 90–100%), indicating that the high-percentile positioning of SPC1* is robust to resampling uncertainty and not driven by individual polity-windows.