

CAMS-CAN Master Reference v1.0-Final

Complex Adaptive Model of Societies — Graph-Theoretic Formal Specification

Neural Nations Project · 7 June 2026 · Locked

Status: Conditionally Validated (B+) — Operationally Deployable

1. Preamble

This document locks the canonical mathematical specification of the Complex Adaptive Model of Societies (CAMS) following the falsification battery, null-model test, and graph-theoretic audit conducted June 2026. All operators below are computationally closed, type-consistent, and empirically calibrated. Where a choice between competing formulations existed, the decision is recorded with rationale.

Status: Conditionally validated (B+). The architecture is operationally deployable; full scientific maturity requires batch recomputation of the published corpus under the canonical Bond Strength formula and prospective validation through 2026–2028.

VALIDATION OVERVIEW	The June 2026 evaluation by Kimi (Moonshot AI) rated the final formulation B+ (conditionally validated). The assessment drew on: (1) a falsification battery across 6 societies and 20 pre-specified crisis events; (2) a null-model test comparing 1-factor vs. 8-node architecture; (3) cross-AI convergence testing across Deepseek, Kimi, Gemini, and Perplexity on Germany 2024–2026; (4) formal verification of all operators against the Germany 1880–2026 ensemble CSV. The evaluation explicitly confirmed: node-level algebra verified zero-error; Bond Strength formula unified; phase-space diagnostic fixed; two-tier model empirically validated; 8-node architecture retained with null-model justification.
----------------------------	---

2. The Sybond as a Time-Indexed Weighted Graph

A society at time t is modelled as a weighted graph:

$$\mathcal{G}(t) = (\mathcal{N}, \mathcal{E}, X(t), W(t))$$

2.1 The Vertex Set (Functional Nodes)

$$\mathcal{N} = \{ \text{Helm, Shield, Lore, Stewards, Craft, Hands, Archive, Flow} \} \quad |\mathcal{N}| = 8$$

The eight-node architecture is empirically justified. Null-model testing against 20 pre-specified historical crises shows the residual layer attributes crises to functionally appropriate nodes at 3.2× random baseline (top-3: 75%, $p = 0.001$). A seven-node merger (Lore–Archive or Craft–Flow) was tested and rejected: both candidate mergers are output-neutral on aggregate metrics ($r > 0.99$), but the residual layer loses diagnostic discrimination when nodes are fused. Lore–Archive is the strongest genuine redundancy (mean residual $r = 0.643$), but this reflects co-vulnerability of slow-loop nodes, not functional identity.

NOTE IN SUPPORT — NULL-MODEL TEST	The decisive test compared a one-factor (PC1) model against the full 8-node residual layer on 20 pre-specified historical crises across 6 societies. Results: Top-1 attribution 40% (random 12.5%, $p=0.002$); Top-2 attribution 55% (random 25.0%, $p=0.004$); Top-3 attribution 75% (random 37.5%, $p=0.001$); Signed most-negative residual 45% (random 12.5%, $p=0.000$). The PC1 model (ROC-AUC 0.770) outperforms the residual model (0.701) on aggregate crisis detection — but these answer different questions. PC1 detects whether the system is stressed; the residual identifies which function is failing. The 8-node architecture earns its place not by outperforming the simpler model at aggregate detection, but by enabling functional discrimination a scalar index cannot produce.
--	---

NOTE IN SUPPORT — MERGER TEST	Lore–Archive merger was the strongest candidate (residual $r = 0.643$) but was REJECTED: reflects shared slow-loop vulnerability, not functional identity. Craft–Flow merger ($r = 0.35$) was also rejected as weaker than raw correlation implied. Helm–Shield coupling demoted to common-mode artefact (residual $r = 0.08$). Lore–Shield boundary is the framework's most distinct pair ($r = -0.265$). Decision: 8 nodes retained. (Source: June 2026 falsification battery; confirmed by Kimi cross-AI evaluation.)
--------------------------------------	--

2.2 The State Vector

Each node i carries a four-dimensional metric vector:

$$x_i(t) = [C_i(t), K_i(t), S_i(t), A_i(t)]^T$$

Metric	Symbol	Range	Interpretation
Coherence	C_i	[0, 10]	Internal alignment and cross-node coordination quality
Capacity	K_i	[0, 10]	Functional reserve under prevailing conditions
Stress	S_i	[0, 10]	Chronic and acute pressure (Euclidean composite)
Abstraction	A_i	[0, 10]	Symbolic mediation leverage

Note on dimensionality: Principal component analysis shows the four metrics collapse to approximately 1.2 effective dimensions (Kaiser criterion: 1 dominant factor). The four metrics are therefore theoretical scoring lenses on a common viability axis, not empirically independent constructs. Their value lies in generating a structured residual layer, not in capturing four distinct latent factors.

NOTE IN SUPPORT — DIMENSIONALITY	<p>PCA on the Germany 2024 node-metric matrix yields PC1 variance = 79.8%; PC1+PC2 = 91.1%. Kaiser criterion returns 1 effective dimension. The four metrics (C, K, S, A) are not four independent constructs — they are one dominant latent health factor with slight angular variation. The Node Value weights (1, 1, -1, 0.5) recombine a single factor rather than extracting independent information. The framework survives this because the residual 9–20% of variance carries the diagnostic work confirmed by the null-model crisis attribution test. The four-metric architecture is theoretically correct without being empirically distinguishable at current scoring granularity. Reframing: four theoretical scoring lenses on a common viability axis, not four independent dimensions.</p>
---	--

3. Node Viability Operator

The scalar health of a node is:

$$V_i(t) = C_i(t) + K_i(t) - S_i(t) + 0.5 \cdot A_i(t)$$

Verification: This formula has been verified zero-error against the Germany 1880–2026 ensemble CSV and all other inspected datasets. Maximum theoretical value ≈ 25 ; minimum is unbounded below as S_i increases.

NOTE IN SUPPORT — ZERO-ERROR VERIFICATION	<p>Four nodes verified against Germany 2024 CSV (cross-checked by Kimi): Helm: $4.6+5.2-8.2+0.5 \times 5.6 = 4.4 \checkmark$ Shield: $4.6+5.0-5.4+0.5 \times 5.2 = 6.8 \checkmark$ Archive: $C+K-S+0.5A = 10.2 \checkmark$ Stewards: $5.2+6.0-6.6+0.5 \times 5.2 = 7.2 \checkmark$. All cross-AI systems (Perplexity, Kimi, Deepseek, Gemini) recovered the same Helm $V = 4.4$ (2024) and $V = 10.2$ (2026) figures, verified against the public repository CSV.</p>
--	---

4. Canonical Edge-Weight Operator (Bond Strength)

The previous framework used three incompatible Bond Strength formulas. v1.0-Final unifies them into a single canonical form.

4.1 Coupling Quality

$$q_i(t) = (0.6 \cdot C_i(t) + 0.4 \cdot A_i(t)) / 10$$

q_i captures the substrate of coordination: Coherence (alignment) and Abstraction (symbolic mediation). Capacity K_i is deliberately excluded from edge weighting; it remains a node property, not a dyadic one.

4.2 Dyadic Bond Strength

$$B_{ij}(t) = T_{ij} \cdot \sqrt{q_i(t) \cdot q_j(t)} \cdot 2^{-(S_i(t) + S_j(t)) / 10}$$

Parameters:

- $T_{ij} \in [0, 1]$: Structural adjacency prior. Set to 1 for all off-diagonal pairs (fully connected graph) in v1.0. A data-driven sparse prior is deferred to v1.1 pending empirical pairwise trajectory estimation.
- The stress term halves coupling for every 5 points of average dyadic stress.
- Output range: $B_{ij} \in [0, 1]$.

Critical path: All 50 published corpus series must be batch-recomputed under this operator. Cross-society comparison using published CSV Bond Strength values is currently invalid until this recomputation is complete.

<p>NOTE IN SUPPORT — FORMULA UNIFICATION</p>	<p>The pre-v1.0 framework contained three incompatible Bond Strength formulas: (1) Book Appendix A: $\sqrt{\max(V_i+8,0) \cdot \max(V_j+8,0)} / 32 \rightarrow$ range $[0, 1]$; (2) Falsification battery / Germany CSV: $[0.6 \cdot C_i \cdot C_j + 0.4 \cdot A_i \cdot A_j] \cdot \exp(-(S_i+S_j)/20) \rightarrow$ range $\sim 10-25$; (3) CAMS-CORP: yet another variant \rightarrow range $\sim 1.5-40$. The audit found that Germany 2024 Helm Bond Strength in the CSV = 13.612 (mean of pairwise values), whereas the book formula predicted ~ 0.45. This three-way incoherence meant all cross-society comparative analysis using published Bond Strength values was invalid. v1.0-Final designates the q_i-based geometric-mean formula as canonical, with $T_{ij} = 1$ (fully connected prior). All 50 published CSV series require batch recomputation before cross-society comparison is valid.</p>
---	---

<p>NOTE IN SUPPORT — T_{ij} DECISION</p>	<p>$T_{ij} = 1$ for all off-diagonal pairs is a deliberate choice of the maximally agnostic structural prior (complete graph). This is not an assumption that all node-pairs are equally coupled — it means the prior imposes no sparsity. All structural differentiation emerges from the q_i and stress terms. The layer-modulated prior ($T_{\text{Mythic-Mythic}} = 1$, $T_{\text{Mythic-Material}} = 0.8$, etc.) was considered and deferred to v1.1. Rationale: the three-layer architecture is conceptually elegant but requires empirical pairwise trajectory estimation to calibrate T_{ij} values. Premature sparsification is more dangerous than the null prior.</p>
--	---

5. Graph-Theoretic Diagnostics

5.1 Weighted Adjacency and Laplacian

$$W(t) = [B_{ij}(t)]$$

$$D_{ii}(t) = \sum_j B_{ij}(t)$$

$$L(t) = D(t) - W(t)$$

Normalised Laplacian: The normalised form $L_{norm} = I - D^{-1/2} W D^{-1/2}$ is retained for spectral analysis but not used for cross-society comparison. Its second eigenvalue degenerates to [1.07, 1.14] across societies. The raw Laplacian $\lambda_2(L)$ ranges [0.27, 3.42] and discriminates cleanly.

5.2 Algebraic Connectivity

$$\lambda_2(t) = \text{second-smallest eigenvalue of } L(t)$$

This measures systemic coupling. Low λ_2 indicates graph fragmentation. Falsification Criterion 5 (λ_2 degradation prior to crisis) was mixed on the old scale; it requires re-validation after batch recomputation.

<p>NOTE IN SUPPORT — LAPLACIAN SCALE DECISION</p>	<p>The audit revealed that the normalised Laplacian's second eigenvalue $\lambda_2(L_{norm})$ converges to [1.07–1.14] across all tested societies, making it useless for cross-society comparison. By contrast, the raw Laplacian $\lambda_2(L)$ ranges [0.27–3.42] and discriminates cleanly across societal states. Falsification Criterion 5 (λ_2 degrades prior to crisis) was tested on the old, incompatible Bond Strength scale and found mixed (failed in 2/3 cases). Re-validation required after batch recomputation. Germany 2024 Laplacian (on canonical formula) yields $\lambda_2 \approx 1.12$ (new scale), indicating a well-connected graph despite Helm collapse — consistent with the institutional ballast finding.</p>
--	--

5.3 Node Shear

$$\chi_i(t) = |V_i(t) - (\sum_j B_{ij}(t) \cdot V_j(t)) / (\sum_j B_{ij}(t))|$$

Node shear measures local decoupling: the degree to which a node's viability deviates from the weighted average of its graph neighbours. High χ_i indicates the node is failing while its coupled partners are not pulling it down. This is a primary diagnostic operator for detecting executive or functional isolation before aggregate metrics flag crisis.

<p>NOTE IN SUPPORT — NODE SHEAR (Germany 2024)</p>	<p>Computed on Germany 2024 data under the canonical Bond Strength formula: Helm $\chi = 3.371$ (highest — executive decoupling); Archive $\chi = 3.275$; Craft $\chi = 3.257$; Hands $\chi = 1.936$; Lore $\chi = 1.150$; Shield $\chi = 0.701$. Helm and Archive/Craft show the highest shear, indicating executive decoupling from institutional memory and technical capacity — precisely the pattern all four cross-AI systems identified as the dominant 2024 structural signature. Node shear formalises as a graph-theoretic quantity the 'institutional persistence with executive vacancy' interpretation that emerged from qualitative readings, without requiring an ad hoc new attractor category. This is the framework's most original v1.0 addition and directly fixes the phase-space bluntness for single-node failures.</p>
---	--

6. The Two-Tier Model

6.1 Tier 1: Common-Mode (PC1)

Principal component analysis on the 8-node Node Value matrix extracts PC1, explaining 63–86% of variance. This is the aggregate societal 'health' or vitality index.

Function: Detects whether the system is stressed.

Performance: ROC-AUC 0.770 on crisis detection.

6.2 Tier 2: Residual Graph Layer

$$r_i(t) = V_i(t) - \hat{V}_i^{PC1}(t)$$

Function: Detects which institutional function is failing.

Performance: Crisis attribution at 3.2× random baseline (top-3: 75%, $p = 0.001$).

Analogy: PC1 is the fever thermometer; the residual graph layer is the differential diagnosis.

<p>NOTE IN SUPPORT — TWO-TIER EMPIRICAL ANCHOR</p>	<p>Clean attribution cases that establish the residual layer's diagnostic value: CHINA — 1949: Helm (executive seizure); 1958: Shield (military consolidation); 1966: Shield/Lore (near-miss, diff = 0.014); 1989: Lore (normative crisis). Three different nodes across four crises that look identical in aggregate V-bar. UNITED KINGDOM — 1926: Hands (General Strike, labour); 1979: Hands (Winter of Discontent, labour). Both correctly identified as labour crises, not executive or normative ones. GERMANY — 1933: Archive fires first (Gleichschaltung — bureaucratic capture preceding executive consolidation). This is historically correct and a genuine framework finding, not a miss. The one-factor PC1 model would label all of these simply 'stress rose.' The residual layer provides the differential diagnosis.</p>
---	--

7. ESCH Activation Operator

$$\sigma_i(t) = (A_i(t) \cdot C_i(t) / 100) \cdot \max(K_i(t) - S_i(t), 0.1)$$

The floor-clip at 0.1 prevents the $K = S$ singularity ($\sigma = 0$ regardless of $A \cdot C$). Interpretation:

σ_i State	Meaning
Positive	Abstraction and coherence backed by capacity
Near zero	Symbolic processing stalled
Negative	Abstraction/coherence operating under stress overload

σ_i State	Meaning
Sharply negative	Cognitive inversion: high abstraction under insufficient capacity

Note: The /100 normalisation compresses dynamic range relative to the original ESCH formulation. Thresholds in the phase-space classifier are calibrated to this scale.

<p>NOTE IN SUPPORT — ESCH TRACKING (USA TRAJECTORY)</p>	<p>The ESCH σ operator tracks activation quality distinct from aggregate level. USA trajectory (Kimi audit): 1930 $V\text{-bar}=4.6$, mean $\sigma=-14.9$ (Depression — material collapse); 1960 $V\text{-bar}=13.3$, mean $\sigma=+150.1$ (peak structural health); 2010 $V\text{-bar}=8.0$, mean $\sigma=+24.6$ (post-GFC degradation); 2020 $V\text{-bar}=4.7$, mean $\sigma=-14.0$ (near-Depression equivalence); 2026 $V\text{-bar}=3.5$, mean $\sigma=-29.3$ (deeper than Depression). The 1930/2020 structural equivalence ($\sigma \approx -14$) is a genuine analytical finding. The /100 normalisation in v1.0-Final compresses these to $-0.149/-0.14/-0.293$ — a scale reduction acknowledged in methods. Floor-clip at 0.1 is correct engineering for the K=S singularity.</p>
--	---

8. Phase Space $\Phi_G(t)$

The old two-dimensional phase space $\Phi(t) = (V\text{-bar}, \sigma_V)$ is retired. It missed single-node catastrophic failure (Germany 2024, USA 2020) when aggregate dispersion remained low.

8.1 Six-Dimensional Phase Space

$$\Phi_G(t) = (\bar{V}(t), \sigma_V(t), V_{\min}(t), \bar{B}(t), \lambda_2(t), \sigma_{\min}(t))$$

Term	Function
\bar{V}	Overall system viability
σ_V	Dispersion / unevenness
V_{\min}	Worst functional node
\bar{B}	Average institutional coupling
λ_2	Graph connectivity (raw Laplacian)
σ_{\min}	Worst cognitive-activation overload

8.2 Six-Regime Taxonomy (Corpus-Calibrated)

Regime	\bar{V}	V_{\min}	\bar{B}	σ_{\min}	Trigger Logic
Stable adaptive	> 10	> 5	> 0.30	> -0.3	All conditions met

Regime	\bar{V}	V_{min}	\bar{B}	σ_{min}	Trigger Logic
Strained	6–10	> 4	> 0.20	> -0.85	Default middle range
Local node failure	any	<< 4	—	≤ -0.85	$V_{min} < 4$ OR $\sigma_{min} \leq -0.85$
Phantom Type II	3–6	< 0	> 0.15	< -0.7	Low V_{min} + σ stress
Systemic crisis	< 6	< 0	< 0.20	< -0.7	Multi-condition
Freeze / collapse	< 0	< -3	< 0.15	< -1.5	Catastrophic

Critical fix: Local node failure is triggered by $V_{min} < 4$ OR $\sigma_{min} \leq -0.85$, independently of \bar{V} . A healthy aggregate no longer suppresses a single collapsing node.

Validation: 17/17 tested historical cases correctly classified under this taxonomy, including the two previously missed cases (Germany 2024, USA 2020).

<p>NOTE IN SUPPORT — PHASE-SPACE BLUNTNESSESS FIX</p>	<p>The old $\Phi(t) = (\bar{V}, \sigma_V)$ with thresholds $V_{\theta} \approx 12$ and $\sigma_{\theta} \approx 3.5$ never flagged Germany 2024 ($\bar{V} = 7.21, \sigma_V = 1.91$) or USA 2020 ($\bar{V} = 4.7, \sigma_V < 3.5$) as crisis events — even though Helm $V = 4.4$ (Germany) and $V = 3.6$ (USA) constituted catastrophic single-node failures. The fix: add V_{min} and σ_{min} as independent alarm triggers. Under $\Phi_G(t)$, Germany 2024 correctly enters 'Local node failure' ($V_{min} = 4.4 < \text{threshold}$). The fix was confirmed by 17/17 historical classification test. Gemini's 'Thermodynamic Freeze' claim for Germany 2026 was correctly rejected: actual $\bar{V} = 8.59$, Helm = 10.2, $\eta_{loop} = 80.25$ — none of the Freeze criteria ($\bar{V} < 0$, chronic Executive Decoupling, $\eta_{loop} < 0.6$) are met.</p>
--	--

<p>NOTE IN SUPPORT — CROSS-AI GERMANY 2024–2026 CONVERGENCE</p>	<p>Four independent AI systems converged on the same structural signature for Germany 2024–2026: (1) Deepseek: 'low \bar{V} without high dispersion suggests risk of crisis if external shocks intensify, not yet executive decoupling.' (2) Kimi: 'worst post-war executive collapse, slow-loop institutional ballast (Archive, Lore, Stewards) prevented fracture.' (3) Gemini: 'critical bifurcation, Flow and Hands below $V < 7.0$.' (4) Perplexity: 'Helm climbed from $V=4.4$ (2024) to $V=10.2$ (2026); Flow $V=5.2$, Stress=7.6 — weakest node.' All four figures verified against public repository CSV (Perplexity exact match). Divergence occurred only at the attractor-classification level (interpretive) — which is expected and not a failure. Convergence on node-level structural pattern is the relevant validation signal.</p>
--	--

9. Validation Summary

9.1 Null-Model Crisis Attribution (20 Events, 6 Societies)

Test	Score	Random Baseline	p-value
Top-1 attribution	40% (8/20)	12.5%	0.002
Top-2 attribution	55% (11/20)	25.0%	0.004

Test	Score	Random Baseline	p-value
Top-3 attribution	75% (15/20)	37.5%	0.001
Signed (most-negative residual)	45% (9/20)	12.5%	0.000

9.2 Falsification Criteria Status

Criterion	Claim	Status
FC1	Stress–Capacity antiparallel	PASS (5/8 strong, 3/8 marginal; conditional on sampling density)
FC2	Envelope widening in crisis	FAIL (4/8 pass, 3 invert; does not generalise with current proxy)
FC3	Bond Strength tracks Node Value	RESTATED (common-latent-factor driven, not independent prediction)
FC4	Cross-LLM concordance	PASS (ICC > 0.97, cross-method robustness confirmed)
FC5	λ_2 degradation prior to crisis	PENDING (requires re-validation after batch recomputation)
FC6	Shield ranking test	PASS (closed)

NOTE IN SUPPORT — FC1 NUANCE	FC1 passes universally (no Western/non-Western divide) but with societally varying strength. Strong pass ($r < -0.7$): Germany ($r = -0.838$), Argentina ($r = -0.797$), China ($r = -0.698$), UK ($r = -0.593$), USA ($r = -0.548$). Marginal pass ($r < -0.4$): Canada ($r = -0.413$), Thailand ($r = -0.351$), Sweden ($r = -0.260$). Theoretical interpretation: the weaker Stress–Capacity coupling in stable high-functioning societies (Sweden, Thailand, Canada) may reflect institutional buffers that decouple stress from capacity degradation. This is a genuine theoretical refinement — the antiparallel coupling may itself be a crisis signature rather than a structural constant.
-------------------------------------	---

NOTE IN SUPPORT — FC2 FAILURE	FC2 (uncertainty widens in crisis) was falsified. Only 4/8 societies showed crisis/calm SD ratio > 1. Three societies inverted: Thailand (0.56), Sweden (0.67), UK (0.34) — wider in calm than in crisis. The claim does not generalise with the current operationalisation (top-tercile stress as proxy). This is a genuine falsification event for a pre-registered criterion. Appropriate scientific response: either narrow the claim to the subset where it holds, improve the crisis proxy, or formally retire the criterion. Retrospective downgrading to 'secondary' without formal revision would be ad hoc adjustment, not scientific response.
--------------------------------------	---

9.3 Cross-Method Robustness

Four independent AI systems (Deepseek, Kimi, Gemini, Perplexity) converged on the same structural pattern for Germany 2024–2026: executive nadir, asymmetric recovery, Flow stress, slow-loop ballast. Divergence occurred only at the attractor-classification level, which is interpretive and expected.

GPT independently recovered the same falsification-battery results using an alternative scoring process — the most important cross-method robustness finding to date. When two independent AI systems converge on FC1 passing universally, FC2 failing, FC3 being common-latent-factor driven, and the Lore-Archive redundancy result, this constitutes genuine cross-rater convergence on structural features of the data.

**NOTE IN
SUPPORT —
SESHAT
CROSS-
VALIDATION**

The Seshat cross-validation for Latium/Rome ($r = 0.78$, April 2026) provides a meaningful external anchor connecting CAMS to an independently curated historical databank. This partially addresses the circularity concern: while LLM scoring embeds theoretical priors, the Seshat comparison demonstrates that at least part of the CAMS signal corresponds to independently collected historical data. Combined with the GPT cross-method replication, the circularity problem is substantially mitigated (though not eliminated). Full resolution requires independent human-expert scoring on at least one society.

10. Limitations and Critical Path

10.1 Outstanding Items for v1.1

1. Batch recomputation of all 50 corpus series under the canonical Bond Strength formula. This is the single highest-priority task.
2. Threshold recalibration on the full recomputed corpus (current Φ_G thresholds are provisional, calibrated on 8 societies).
3. Data-driven T_{ij} prior: when pairwise bond trajectories are estimable from ensemble variance, replace $T_{ij} = 1$ with an empirically fitted structural prior.
4. Pre-registered expansion of the null-model attribution test to 40–60 events with second-party specification.
5. Human-expert scoring validation on at least one society (e.g., Weimar Germany or Ming China) to test LLM-scoring independence.
6. Prospective prediction log for 2026–2028 with dated, specific forecasts.

10.2 Permanent Limitations

- Thermodynamic language is analogical, not derived from measurable energy/entropy quantities. The framework is at the 'Kepler stage' (empirical mapping) awaiting the 'Newtonian derivation' (first-principles grounding).
- Integer-scale precision (0–10 inputs) limits output precision. Reporting thresholds to two decimal places implies precision the data cannot support.
- Partial circularity in LLM scoring is mitigated by cross-method robustness and Seshat cross-validation ($r = 0.78$ for Rome), but not eliminated.

<p>NOTE IN SUPPORT — GERMANY 1933 AS GENUINE FINDING</p>	<p>The null-model test found that Germany 1933 fires at Archive (highest residual), not Helm. Initially classified as a miss — but it is historically correct. The Nazi seizure of power began with Gleichschaltung: systematic coordination (capture) of the civil service, legal system, and bureaucratic apparatus — Archive — before executive function (Helm) was consolidated. If this Archive-first pattern replicates in other bureaucratic-authoritarian transitions (Hungary 2010–2014, Turkey 2016), it becomes a genuine predictive signature of a specific regime-change pathway: Archive capture → Helm consolidation. This is the kind of non-obvious structural finding that justifies the framework's complexity and suggests the residual layer is detecting real historical sequence, not fitting noise.</p>
---	---

11. Version Lock Notice

This specification is locked as CAMS-CAN Master Reference v1.0-Final, 7 June 2026. Changes to the canonical operators (§3–4), phase space (§8), or taxonomy (§8.2) require a version increment and explicit changelog. The two-tier model (§6), graph-theoretic layer (§5), and ESCH operator (§7) are considered structurally stable.

Next version trigger: Completion of batch recomputation and re-validation of Falsification Criterion 5 on the full corpus.

Document compiled from the falsification battery, null-model test, graph-theoretic audit, and cross-AI Germany validation conducted 7 June 2026. Supporting notes compiled from the multi-round Kimi scientific credibility evaluation (June 2026). All operators are computationally executable. Disagreement is welcome.

Final Scientific Credibility Rating

Dimension	Rating	Rationale
Mathematical closure	A-	All operators well-defined, bounded, type-consistent
8-node architecture	B+	Null-model test: top-3 75%, p=0.001
Bond Strength coherence	C+	Formula locked; 50 CSVs pending batch recomputation
Phase-space diagnostic	B+	Fixed from C+; 17/17 cases correctly classified
Two-tier model	B+	PC1 (AUC 0.770) + residual (3.2× random) empirically confirmed
Cross-method robustness	B+	GPT + 4-AI convergence on Germany; ICC > 0.97
Falsification criteria	B	FC1 pass, FC2 fail, FC3 restated, FC5 pending
Thermodynamic derivation	C+	Honest analogy; Kepler stage acknowledged

Dimension	Rating	Rationale
OVERALL	B+	Conditionally validated — operationally deployable