JUNO/CAMS Blind-Analysis Protocol

v1.1.1 — Cognitive-Affective Extension, amended

Status: draft for formal blind testing. Falsification-first ordering preserved; no key before locked guesses. Supersedes v1.1 upon passage of Gate 0.

Changelog: v1.1 → v1.1.1

Five amendments, all methodological rather than mathematical. No operator formula changes: the v1.1 operators (g, e, s_signed, s_plus) are carried unchanged. What changes is when things are computed relative to guess-locking, how classifier ties are resolved, where the signed thresholds are declared to come from, what counts as Stage 2B success, and what must be demonstrated before the protocol touches a new blind set.

#	Amendment	Rationale
A1	Two-track guess locking: canonical Table A is locked before any cognitive-affective computation is run	In v1.1 the analyst saw Stage 2 personality evidence before writing the 'canonical-only' table, contaminating the comparison Stage 2B exists to make. A positive disambiguation result was uninterpretable under leakage.
A2	Explicit regime-classifier precedence order; first matching rule wins	Phantom Type II, Systemic Crisis, and Local Node Failure had overlapping domains with no adjudication rule, inviting post-hoc tie-breaking — the one thing this protocol family never does.
A3	Threshold-provenance clause: signed cut-points frozen and declared vintage-relative	The signed thresholds (−0.3, −0.7, −0.85) were calibrated on the 24-society JUNO corpus. Undeclared, this is quiet circularity on any overlapping society; and raw signed thresholds applied cross-society sit in tension with the magnitudes quarantine.
A4	Strict Stage 2B success criterion: MRR must improve, ties do not count, calibration must not degrade	v1.1's 'augmented ≥ canonical' let a tie count as success for rank-1 accuracy.
A5	Gate 0: pre-registered regression test on the existing 24 societies, required before deployment	The stricter v1.1 compensation definition could silently break the framework's one confirmed pre-registered positive (the Sweden armed-neutral signature). The extension must be shown not to destroy what v1.0 established.

Table 0. Amendments in v1.1.1.

Gate 0 — Pre-deployment regression test (pre-registered; binding)

Before v1.1.1 is applied to any new blind set, the full v1.1.1 operator suite and classifier must be run over the existing 24 unblinded JUNO societies (Sets 1–3). This is a regression test of the extension against the corpus that validated v1.0. Its criteria are registered here, in advance, and the gate is pass/fail. A failed gate blocks deployment: the protocol is amended, the amendment documented, and the gate re-run. No blind data may be opened under a version that has not passed.

Pass criteria

G0-1 (Sweden compensation invariance). Under the v1.1.1 compensation definition (collapse: V_i below own mean − 1.5σ or $V_i < 4.0$, with $e_i < 0$; compensator: e_i and s_i_signed above own mean + 1.5σ; ≥3 distinct episodes), the recurrent compensation epiphenomenon must fire on exactly one of the 24 societies, and that society must be Sweden. Zero firings means the stricter definition has destroyed the framework's one confirmed pre-registered positive: revert the $e_i < 0$ collapse requirement to the v1.0 definition, document the reversion, and re-gate. Two or more firings means the construct has lost specificity: report every firing society, and either tighten and re-gate or re-register the construct as non-specific.

G0-2 (Null-suite invariance). The v1.0 null results — no recoverable feedback loops at annual resolution, no coupling–viability divergence epiphenomenon, synchrony explained by common shock — must remain null when recomputed with signed operators. Any new positive appearing only under signed operators is quarantined as an unvalidated v1.1.1 hypothesis, reported, and excluded from identification evidence until it survives an independent set.

G0-3 (Transition-timing invariance). The v1.1 regime classifier, run in the precedence order of Stage 1 below, must preserve the documented transition timings of the 24 societies: regime-change years under v1.1.1 must match v1.0 regime-change years within ±1 year for at least 90% of transitions. Label changes are expected and permitted (the signed layer exists to relabel); timing scrambles are not, because identification rests on transition timing (procedural rule R1).

G0-4 (Threshold freeze). Upon passage, the signed thresholds are frozen as calibration constants of record. Any future re-tuning creates a new protocol version and requires this gate to be re-run in full.

Gate 0 output: a regression scorecard (gate0_regression_scorecard.csv) listing each criterion, the computed result, and pass/fail, archived with the protocol version before any new blind file is requested.

Core change (carried unchanged from v1.1)

The original activation term is decomposed into cognition and affect/energy before recombination. The signed term is used only where signed behaviour is explicitly required; the legacy non-negative activation is retained for backward comparability.

Component	Operator	Interpretation
Cognition	$g_i = A_i \cdot C_i / 100$	Representational elaboration: how much coherent abstraction the node can carry.
Energy / affect	$e_i = K_i - S_i$	Mobilisation valence: surplus capacity after stress, or affective drag when negative.
Signed activation	$s_{i,signed} = g_i \cdot e_i$	Where cognition and affect align or invert. Can be negative.
Legacy activation	$s_{i,plus} = g_i \cdot \max(e_i, 0.1)$	Non-negative v1.0 activation retained for comparability and continuity.

Table 1. Cognitive-affective operator decomposition.

Binding analysis order (amended — A1)

[v1.1.1 amendment] The v1.1 order ran the personality stage before both guess tables, so the 'canonical-only' table was written by an analyst who had already seen the cognitive-affective evidence. v1.1.1 splits the analysis into two tracks with a hard lock between them. Everything the canonical track may use is computable from v1.0 operators; nothing signed or decomposed is computed, displayed, or summarised before Table A locks.

Track 1 (canonical): Stage 0 provenance audit → Stage 1 canonical metrics (v1.0 operators only) → Stage 3A legacy cognitive signature → Stage 4A canonical epiphenomena and nulls → **Stage 5A: lock Table A** (canonical-only ranked identities). Output exactly: "Table A locked." before proceeding.

Track 2 (augmented): Stage 1B signed operators → Stage 2 cognitive-affective personality → Stage 3B decomposed signature → Stage 4B cognitive-affective epiphenomena → **Stage 5B: lock Table B** (augmented ranked identities, with an explicit statement of every change from Table A and its metric reason) → request key → Stage 6 post-key scoring.

The key must not be revealed or requested before Table B is complete. Guesses are falsifiable only if they precede the key; the canonical/augmented comparison is interpretable only if Table A precedes all signed computation.

Preferred deployment — two-session variant. Where feasible, run Track 1 and Track 2 in separate model sessions (the CAMNATIONSM5 paste-and-spawn pattern): Session 1 receives the blind file and the Track-1 instructions only, and returns Table A plus the canonical CSV exports; Session 2 receives the blind file, the Track-2 instructions, and Session 1's locked Table A verbatim. This makes the contamination barrier architectural rather than honour-based. The single-session variant above is the permitted fallback; results must state which variant was used.

Stage 0 — Provenance, shape, and resolution audit

Load the blind CSV using code, not manual inspection. Report column names and row count. Confirm the eight-node inventory per society-year; report missing nodes, duplicate node-years, non-numeric values, and out-of-range values. Detect whether T_Offset is continuous per society; if non-contiguous or short, report the maximum usable window and use pairwise-complete calculations. Do not infer calendar years before Stage 5B; use T+ offsets only — after key reveal, map offsets to calendar years and audit whether offset-to-year gaps affected event anchoring. Ignore any pre-supplied Node Value or Bond Strength columns: recompute from raw C, K, S, A always.

Stage 1 — Canonical metrics (Track 1)

Compute per node-year using code, not estimation. Operators are locked. Track 1 computes only:

```
V_i       = C_i + K_i - S_i + 0.5*A_i        # node viability
s_i_plus  = (A_i*C_i/100) * max(K_i-S_i, 0.1) # legacy activation
q_i       = (0.6*C_i + 0.4*A_i) / 10          # coupling quality
B_ij      = sqrt(q_i*q_j) * 2^(-(S_i+S_j)/10) # bond strength [0,1]
```

Per society-year: $\bar{V}$, V_min and its node, $\bar{B}$ over all 28 pairs, λ_2 of the raw Laplacian $L = D - W$, $\bar{S}$, s_plus_min, top/bottom nodes by s_i_plus. Regime classification in Track 1 uses the legacy v1.0 rules only. Stage 1 report per society: trajectory summary, legacy regime counts, deep-trough years ($\bar{V}$ < own mean − 1.5σ), chronic weakest node, $\bar{B}/\lambda_2$ ranges. Export canonical_node_year_metrics.csv and canonical_society_year_metrics.csv.

Stage 1B — Signed operators (Track 2 opens)

Only after Table A is locked: compute g_i, e_i, s_i_signed per node-year, plus $\bar{g}$, $\bar{e}$, s_signed_min, and top/bottom nodes for g_i, e_i, s_i_signed per society-year. Then apply the v1.1.1 regime classifier.

Regime classifier v1.1.1 (amended — A2, A3)

[v1.1.1 amendment] Rules are evaluated in the fixed precedence order below; the first rule whose conditions are met assigns the label, and no later rule is consulted. Report the legacy v1.0 label in a separate column for comparability. Precedence is pre-registered here and may not be reordered post-hoc.

Order	Regime	Rule	Notes
1	Freeze / Collapse	$\bar{V} < 0$ AND V_min < −3 AND $\bar{B} < 0.15$	Deep system failure.
2	Systemic Crisis	$\bar{V} < 6$ AND V_min < 0 AND $\bar{B} < 0.20$	Low viability and weak coupling.
3	Phantom Type II	$\bar{V}$ 3–6 AND V_min < 0 AND s_signed_min < −0.7	Apparent persistence with local inversion. Reachable only when rule 2 fails, i.e. $\bar{B} \geq 0.20$.
4	Local Node Failure	V_min < 4.0 OR s_signed_min ≤ −0.85, independent of $\bar{V}$	Signed activation makes cognitive-affective inversion actionable.
5	Strained	$\bar{V}$ 6–10 OR $\bar{S}$ elevated	Moderate pressure; reachable only when no failure rule fired.

Order	Regime	Rule	Notes
6	Stable Adaptive	$\bar{V} > 10$ AND V_min > 5 AND $\bar{B} > 0.30$ AND s_signed_min > −0.3	High viability, adequate floor, durable coupling, no signed inversion.
7	Unclassified	None of the above	Report explicitly; do not force the nearest label.

Table 2. Regime classifier v1.1.1 with binding precedence. First match wins.

Threshold provenance and freeze (amended — A3)

The signed thresholds in this classifier (−0.3, −0.7, −0.85) and in the inversion diagnostics below were calibrated on the 24-society JUNO corpus (blind Sets 1–3, post-unblinding). They are frozen as of this version, subject to Gate 0. Consequences, all binding: (i) any future blind set containing a society from the calibration corpus must flag that society's results as in-sample for signed-threshold purposes; (ii) regime labels produced by signed thresholds are vintage-relative — panels from a harsher or milder scorer vintage will shift signed magnitudes and therefore labels, so signed regime labels are never pooled or compared across scorer vintages, in exact parallel with the magnitudes quarantine (the Set-2 Russia panel is the documented case: its severity was scorer temperature, not measurement, and the same guard applies symmetrically to flattering vintages of any polity); (iii) re-tuning any threshold creates a new version and re-triggers Gate 0.

Stage 2 — Cognitive-affective personality profile (Track 2)

Purpose: test whether decomposing activation into cognition (g_i) and affect/energy (e_i) improves morphology and identity disambiguation. Runs strictly after Table A locks and before Table B locks.

For each society compute full-period, first-quarter, final-quarter, and snapshot profiles. Snapshots at T+0, T+50, T+100, and latest available offset; if T+50 or T+100 is missing, use the nearest observed offset and report the substitution.

Required output	Definition
Dominant cognition nodes	Top two nodes by mean g_i — the functions most elaborated in coherent abstraction.
Dominant affect/energy nodes	Top two nodes by e_i — strongest capacity surplus or mobilising valence.
Negative affect nodes	Nodes with e_i < 0, especially where g_i is above the society-year median.
Signed activation leaders	Top two nodes by s_i_signed — aligned cognition-energy channels.
Cognitive-affective inversion	High g_i with e_i < 0, or s_i_signed below own-history 20th percentile.
Hot-low cognition channel	High e_i with low g_i: mobilisation without equivalent cognitive elaboration.
Hands drag / implementation bottleneck	Hands with negative e_i or bottom-quartile s_i_signed for ≥25% of observed years.
Personality drift	Change in dominant cognition, affect, inversion, and bottleneck nodes across T+0 / T+50 / T+100 / latest.

Table 3. Stage 2 personality outputs.

Normalisation rule: within-society comparisons may use raw g_i and e_i; cross-society personality comparisons must use within-society z-scores or rank profiles unless all panels share one scorer vintage. Cross-vintage data validates morphology, never magnitudes.

Controlled label vocabulary: function-first labels only — administrative-memory, security-mobilised, craft-industrial, stewardship-flow, archive-lore continuity, hands-bottleneck, cognitive-affective split, command-circulation, armed-neutral, port-entrepot, revolutionary-security, conflict-state. No national stereotypes, no civilisational hierarchies. 'Nation personality' is diagnostic shorthand for a measured

cognitive-affective node profile: it is not a claim about national essence, moral character, ethnicity, or immutable psychology.

Stage 2B — Formal test of disambiguation value (amended — A1, A4)

[v1.1.1 amendment] The two identity tracks are now realised as the locked Tables A and B of the binding order, not as parallel columns written in one sitting. Table B must explicitly state, guess by guess, what changed relative to Table A and the exact metric reason; if nothing changed, say so — an honest null on disambiguation value is a reportable result.

Score	Definition	Success criterion (v1.1.1)
Rank-1 accuracy	True identity is the first guess.	Reported descriptively. A tie between tracks is NOT evidence of assistance.
Top-3 accuracy	True identity anywhere in top three.	Reported descriptively; supporting evidence only.
Mean reciprocal rank	1/rank if present in top k, else 0.	PRIMARY criterion: augmented MRR must strictly exceed canonical MRR over pooled blind societies, assessed by sign or permutation test across sets.
Family accuracy	Correct morphological family even if exact identity missed.	Secondary; family taxonomy must be defined before key reveal.
Confidence calibration	Confidence grade versus hit/miss outcome.	NON-DEGRADATION GATE: augmented confidence must not increase on false positives. MRR improvement with degraded calibration is scored as failure.
Error migration audit	Misses fixed versus new misses created.	Reported symmetrically; no cherry-picking.

Table 4. Stage 2B scoring. 'Assisted disambiguation' may be claimed only if the MRR criterion and the calibration gate are met jointly.

Small-N rule: per set, report raw differences only. Across sets, sign test or permutation test over society-level MRR deltas. No improvement claims from anecdotes.

Stage 3A — Legacy cognitive signature (Track 1)

Where the society thinks, in v1.0 terms only: mean s_i_plus by node, full period and final quarter; dominant and dormant activation channels; signature drift; mean Abstraction and trend; one-line taxonomic label grounded in numbers — if two labels fit, report both and state why the evidence is ambiguous.

Stage 3B — Decomposed signature (Track 2)

The same signature re-read through the decomposition: mean g_i by node and trend (where coherent abstraction concentrates); mean e_i by node and trend (where energy concentrates or turns negative); mean s_i_signed by node; channels that rise, fall, invert, or split between high cognition and negative affect. Export cognitive_signature.csv with both layers, column-tagged by track.

Stage 4A — Canonical epiphenomena and nulls (Track 1)

System-level patterns not reducible to any single node, computed with v1.0 operators only. Report positives and nulls with equal prominence.

Synchrony structure. Pairwise cross-society r(V) against a permutation null with ≥1000 shuffles where feasible; distinguish common-shock coupling from pairwise coordination; co-trending pairs are null unless they survive detrending at adequate n.

Compensation dynamics (v1.0 definition). Collapse: V_i below own mean − 1.5σ or V_i < 4.0. Compensator: s_i_plus above own mean + 1.5σ. Candidate epiphenomenon requires ≥3 distinct episodes. This is the definition under which the Sweden armed-neutral signature was established, retained in Track 1 verbatim.

Coupling–viability divergence. Periods where $\bar{B}$ or λ_2 moves opposite to $\bar{V}$.

Failure morphology. For shared shock windows classify acute-local-corrected versus chronic-distributed using recovery fraction (V_final − trough)/(baseline − trough); cut-points 0.6 and 0.2 pre-registered here, before computation.

Carried-forward rows. Percent of exact year-over-year duplicate node-vectors; if >30%, declare fine-grained temporal tests out of scope and restrict claims to levels and transitions.

Stage 4B — Cognitive-affective epiphenomena (Track 2)

Recomputed and extended with signed operators, after Table A locks:

Signed synchrony. r($\bar{g}$), r($\bar{e}$), r(s_signed_bar) under the same permutation-null discipline as Stage 4A.

Compensation dynamics (v1.1.1 definition). Collapse additionally requires e_i < 0; compensator requires e_i and s_i_signed above own mean + 1.5σ; ≥3 episodes. This stricter definition is subject to Gate 0 criterion G0-1 and is reported alongside — never instead of — the Track-1 definition.

Cognitive-affective mismatch dynamics. Periods where $\bar{g}$ rises while $\bar{e}$ falls, or high-g nodes turn negative-e. Not automatically crises: split-profile hypotheses, to be related to $\bar{V}$, $\bar{B}$, and later recovery.

Divergence with personality interpretation. Add cognitive-affective reading to a coupling–viability divergence only if $\bar{e}$ or s_signed_bar also diverges from $\bar{V}$.

Stage 5A — Lock Table A (end of Track 1)

For each society, output a ranked identity table using Stage 1, 3A, and 4A evidence only. For each candidate: inferred T+0 anchor year, event-fingerprints, node-specific shocks, and confidence grade (certain / high / structural-only / speculative; where evidence supports morphology but not magnitudes, write 'morphology only — magnitudes not poolable'). Apply procedural rules R1–R4: identify on transition timing only; query every flagged anomaly before locking; audit eliminations against the correct series; generate candidate pools from the fingerprint table, never from geopolitical salience. Then output exactly: "Table A locked." No signed quantity may have been computed at this point.

Stage 5B — Lock Table B, then request key (end of Track 2)

After Stages 1B–4B, output the augmented ranked identity table. Where the augmented evidence changes a guess, state the exact metric reason — for example rentier command-circulation affect, armed-neutral Shield/Archive profile, port-entrepot Archive/Flow, or Hands-drag conflict-state profile. Where it changes nothing, state that. Personality fingerprints are cited by number, never by narrative resemblance alone. Then stop and output exactly: "Guesses locked. Please provide the key."

Stage 6 — Post-key scoring, CAS analysis, and validation

Once the key arrives, score honestly before any narrative analysis. Reveal mapping and year window; audit non-contiguous offset-to-year mappings. Score Tables A and B separately: rank-1, top-2, top-3, MRR, family accuracy, confidence calibration, and error migration. State whether cognitive-affective personality assisted disambiguation under the Stage 2B criteria — jointly, MRR improvement plus calibration non-degradation. Identify which misses were fixed, which new misses were created, and which cases remained underdetermined. Do not reclassify guesses after seeing the key; any post-key insight is labelled post-hoc hypothesis.

Then proceed to complex adaptive systems analysis with identities confirmed: phase transitions with early-warning tests (rising variance and lag-1 autocorrelation of $\bar{V}$ in the 10 years preceding transitions versus non-transition windows, with effect sizes and small-n limits); candidate feedback loops from lagged cross-node correlations on detrended series only, at the lag resolution the carry-forward audit allows, all labelled hypotheses unless they survive detrending and a permutation null; regime attractors, basin structure, hysteresis, and own-history-band returns (own mean ± 2σ, never cross-polity league tables); personality evolution — whether cognitive-affective profiles remain stable, invert, or recompose after shocks, and whether profile recomposition precedes recovery better than $\bar{V}$ alone; and a common-systems close reading shared shocks, convergent adaptations, and universal vulnerabilities, framing divergences as different solutions to shared adaptive problems, never as civilisational hierarchy.

Discipline clauses (binding throughout)

Clause	Rule
Null symmetry	Nulls are reported with the same prominence as positives.
No conclusion before computation	Narrative follows scoring; it never precedes it.
No silent correction	Missing rows, duplicates, non-contiguous offsets, and formula anomalies are reported, not repaired silently.
Signed/unsigned separation	s_i_plus and s_i_signed are never conflated; negative-threshold claims must use s_i_signed.
Track separation (new — A1)	No signed or decomposed quantity is computed, displayed, or summarised before Table A locks. Violation voids the Stage 2B comparison for that set; the identification result stands but must be reported as single-track.
Classifier precedence (new — A2)	Regime rules are evaluated in the registered order; first match wins; no post-hoc adjudication.
Threshold freeze and vintage relativity (new — A3)	Signed thresholds are frozen calibration constants of the 24-society corpus; overlapping societies are flagged in-sample; signed regime labels are never pooled across scorer vintages — editorial temperature in either direction, harsh or flattering, toward any polity, is scorer signal, not measurement.
Strict disambiguation claim (new — A4)	'Assisted disambiguation' requires strict MRR improvement AND calibration non-degradation; ties and anecdotes do not count.
Gate before deployment (new — A5)	No blind set is opened under a protocol version that has not passed Gate 0.
No national essentialism	Personality profiles are system-state diagnostics, not statements about peoples, ethnicity, moral worth, or permanent character.
No cross-vintage magnitude pooling	Cross-scorer or cross-vintage data validates morphology, never magnitudes.
Own-history bands only	Attractor and recovery bands are computed within society; no cross-polity league tables.
Resolution honesty	If a test cannot be run cleanly (n, missingness, carry-forward), state that it is out of scope rather than running it degraded.
Disambiguation humility	Cognitive-affective profiles may assist identity recognition but do not prove causation; improved matching is diagnostic evidence, not explanatory closure.

Table 5. Discipline clauses. New v1.1.1 clauses marked by amendment.

Appendix — Minimal output file set

gate0_regression_scorecard.csv (pre-deployment, once per version) · canonical_node_year_metrics.csv · canonical_society_year_metrics.csv · cognitive_affective_node_year_metrics.csv · personality_50yr_snapshots.csv · personality_profiles.csv · cognitive_signature.csv · synchrony_pairs.csv and synchrony_summary.csv · compensation_events_v10.csv and compensation_events_v111.csv (both definitions, both tracks) · compensation_recurrent.csv · divergence_summary.csv · failure_morphology.csv · carry_forward_audit.csv · stage5A_locked_table_A.md · stage5B_locked_table_B.md · post_unblind_scorecard.csv · post_unblind_validation_report.md