ERES Institute for New Age Cybernetics

Bella Vista, Arkansas, USA · ORCID 0000-0001-9946-3221 · CC BY 4.0 / CCAL v2.1

Certifying Sustainability Without Overclaiming

A weakest-link KPI ladder with design / simulation / empirical tiers and claim-status discipline

Instrument: ERES-CERTKPI-2026-001 (SCKL), v0.2 — methodological overview **Status:** Design stage (GOOD). No KPI is measured on real data in this paper. The paper describes an architecture and invites its testing; it does not report results.

---

1. The problem this architecture is built against

A certification instrument fails in a characteristic way: the specification becomes elaborate, the mathematics becomes sophisticated, and the elaboration is then mistaken for evidence that the certified property holds in the world. A framework can reproduce behaviour generated by its own assumptions without those assumptions describing anything external. The design goal of SCKL is to make that failure mode structurally hard — to build an instrument that reports honestly what it *is* (a procedure) and refuses to report what it is *not* (a result) until the result exists.

Nothing below asserts that any system has been certified. The contribution is the machinery that would prevent a premature certification, presented so others can use, break, or improve it.

2. The maturity ladder: GOOD → SOUND → BEST

Certification status is not binary. SCKL uses three tiers, each with an explicit and different evidentiary bar:

- **GOOD — design consistency.** The instrument is fully specified, internally coherent, and dimensionally sound. This is a statement about the *design*, not about any system.
- **SOUND — instantiated and simulation-checked.** The operators are instantiated and the procedure runs on synthetic data, with behaviour consistent under simulation. Crucially, *simulation consistency is not empirical validation*: a model reproducing its own assumptions demonstrates coherence, not correspondence to the world.
- **BEST — real-data survival.** The instrument's claims survive testing on real, externally-sourced data, **against a pre-registered null and a simpler established baseline**. Both conditions are required; "real-data tested" without a null and a baseline is confirmation, not falsification.

The ladder's whole point is that an instrument is not promoted because its formal structure grew. Promotion is earned by evidence at the next bar. At the time of writing the instrument sits at GOOD, and the honest position is to keep it there.

3. The weakest-link certification rule

System-level status is the **minimum** over load-bearing KPIs, not an average:

```
Level(system) = min over K in {load-bearing KPIs} of Level(K).
```

An average lets a strong component compensate for a structurally unvalidated one; a certification architecture must not permit that. If the empirical-floor KPI is only GOOD, the terminal metric cannot be reported as BEST however sophisticated its computation, because the composite is pinned by its weakest load-bearing link. This single rule does most of the work of preventing overclaiming: it makes the least-validated necessary component the ceiling for the whole.

4. Claim-status discipline: STATED / DERIVED / CONSTRUCTED

Every load-bearing claim in the instrument carries an explicit provenance tag:

- **STATED** — asserted as a modelling choice or definition.
- **DERIVED** — follows by proof or computation from stated premises.
- **CONSTRUCTED** — an engineered artifact of the framework, not a claim about the world.

The tagging is not cosmetic. It localises exactly where an eventual empirical failure would land, and it forbids the quiet drift by which a CONSTRUCTED convenience (a chosen Jacobian, a symbolic design number) is later cited as if it were a DERIVED invariant or an empirical fact. The governing sentence, held across the corpus, is: **a shared mathematical form is not a shared empirical mechanism.** The same equation appearing in two applications transfers no evidence between them.

5. The empirical floor and the independence problem

The instrument's most load-bearing KPI is an **empirical floor**: a confound-independent, self-distinctive signature that must be established on real data before any downstream measure affected by a shared measurement substrate can be interpreted. The reason is an identifiability problem, not a statistical nicety. Two formally distinct channels `A(X)` and `B(X)` can both be functions of the same latent variable `Z`; their apparent multiplicity then establishes no independent evidential content.

This is why nominal evaluator count is treated as an **upper bound only** on independence:

```
n_eff <= k,      with n_eff << k possible under a shared substrate.
```

Agreement is not independence. A framework in which evaluators score artifacts produced by the same framework faces circular confirmation, and covariance correction is a *mitigation*, not a proof of independence. The architecture therefore holds that a genuinely strong certification eventually requires at least one evaluation pathway whose data-generating process is demonstrably external to the framework itself. The floor gates the rest: downstream measures may be interpreted only once the floor is empirically passed.

6. Reproducibility discipline: freeze the reference before testing

Where the instrument interprets a reference object that could itself drift, one face of that object is **frozen before testing** as `E_ref`. Without a fixed target, a reproducibility coefficient is uninterpretable because the thing being reproduced can move. Freezing `E_ref` in advance is a prerequisite for reproducibility, not an editorial nicety, and it belongs to the same family of pre-commitments as the pre-registered null and baseline of the BEST tier.

7. The terminal metric and its foundation-only attribution

The instrument's terminal criterion is a spectral-stability margin, `m = -max Re lambda(J)`, of a coupled system's Jacobian. The associated stability mathematics — the spectral abscissa, the stability criterion, and the reliability↔survival correspondence — is **foundational prior work by an external mathematician, cited on a foundation-only basis.** The continuous/discrete boundary of this criterion, and the construction of the survival functional from the propagator, are treated in separate work in preparation and are not relied upon by this paper.

The attribution is bounded by four conditions, retained here verbatim in spirit and stated so no reader can mistake foundation for endorsement. The cited source: is not the author of this instrument; has not experimentally validated this instrument; is cited as mathematical foundation, which **does not constitute endorsement**; and bears no responsibility for the instrument's empirical claims, which remain solely the instrument authors'. The foundation-only citation is subject to the cited source's standing right of reply and to correction or withdrawal at their request.

8. What is deliberately not claimed

To be exact about honest status: no KPI has been measured on real data; the empirical floor is open; illustrative Jacobians are illustrative and are not evidence that any physical system possesses that Jacobian; symbolic design numbers are design choices, not derived invariants; and simulation consistency, when reached, will be SOUND and not BEST. The instrument reports a procedure. It will report results when results exist, in a separate publication, against pre-registered nulls and baselines.

9. Why publish a design-stage instrument at all

Because the strongest form of the independence requirement in §5 is external testing, and external testing requires that the instrument and its pre-registered gates be public *before* the results are in. Publishing the architecture at GOOD is not an announcement of success; it is an invitation to falsify — and an invitation issued before the outcome is known is the credible kind.

Licensed CC BY 4.0 / CCAL v2.1. Foundational mathematical attribution is foundation-only and subject to the cited source's right of reply and consent (DVC standing-contestability). All third parties referenced in the parent instruments remain VEILED.