A Coherence Functional for Cross-Register Measurement

A Partial Closure of the Open Problem Stated in *The Coherence Principle*

Hans Konstapel — Constable Research , Leiden — 31 July 2026 (v2)

Technical companion to What the Instruments Already See. *Builds directly on the Coherence Medicine series (Papers 1–4, 20 July 2026).*

1. The problem this note addresses is one I already stated

Paper 1 of the Coherence Medicine series defines C(t) as the joint state of four components — synchronisation, connectivity, information flow, and adaptability — and states plainly that only synchronisation has a closed form, the Kuramoto order parameter r(t), and that unifying the four into a single closed-form dC/dt is an open problem rather than a solved one. This note takes that open problem seriously and closes part of it. It does not close all of it, and section 3 says exactly how much.

The occasion is the companion essay, which reads cancer, Alzheimer's and Parkinson's disease across three measurement registers — form, rhythm, medium — and argues that the decidable question is one of temporal order. That argument needs a quantity plotted against time in all three registers, which is the same requirement as the open problem above, arriving from a different direction.

A definitional point comes with it. Calling all three registers *topological* is not defensible: DNA linking number is a strict invariant, a protein fold is a conformation, and topologically associating domains are measured as contact frequencies. What the registers actually share is that each is a **distribution over a state space**, and what changes in disease is the concentration of that distribution. That observation is what makes a single functional possible.

2. The functional

Let a register or component R be characterised at time t by a probability distribution $p_R(x,t)$ over a state space X_R, discretised into N_R bins under a fixed convention. Define the effective occupancy as the perplexity:

$PR_R(t) = \exp(H[p_R(t)])$, where $H[p] = -\Sigma_i p_i \log p_i$

and the coherence functional as its normalised complement:

$C_R(t) = (N_R - PR_R(t)) / (N_R - 1)$

C is dimensionless, bounded in [0,1], equal to 1 for a fully concentrated distribution and 0 for a uniform one. It refers only to the shape of a distribution, never to the physical nature of the variable, which is why it applies to phase angles, coupling weights and conformational classes alike.

C is comparable across time within a component and across components only once N_R is fixed by convention, since C depends on binning. Any cross-component claim must publish its binning. For circular variables the two families agree in ordering: for a von Mises distribution with concentration $\varkappa$, both $r = I_1(\varkappa)/I_0(\varkappa)$ and the entropy-based C are monotone in $\varkappa$, so C reproduces r's ordering without replacing it. Where a component already has an established measure, C should be reported alongside it.

3. Coverage: two components of four, and why the other two resist

This is the substance of the note and the part that determines its value.

Synchronisation — covered. The phase-difference distribution is exactly the object C is built for, and the monotone relation to r(t) above means nothing already validated on r(t) is lost.

Connectivity — covered. The distribution of coupling weights across the network is a distribution like any other, and its concentration is a legitimate reading of whether coupling is broadly distributed or collapsed onto a few links. This is weaker than the full time-delay-stability topology of Bashan et al., which retains *which* nodes couple; C retains only how concentrated the coupling is. It is a summary, not a substitute.

Information flow — not covered, and not fixable within this definition. Transfer entropy is directional: it distinguishes A driving B from B driving A. A concentration measure on a distribution is invariant under relabelling and therefore direction-blind. No choice of binning repairs this. Directionality is information the functional structurally discards, and it must be carried alongside C rather than folded into it.

Adaptability — not covered, for a different reason. Adaptability is a response property: how the system reconfigures under perturbation. C is instantaneous. A single value of C says nothing about what happens when the system is pushed. This maps onto the coherence/resilience distinction of Paper 2: C is a coherence quantity, sitting below both adaptability and resilience in the hierarchy coherence → adaptability → resilience → health.

The honest summary is therefore: the functional unifies the two components that are distributions over states, and cannot unify the two that are properties of directed flow and of response. That is a partial closure with a stated reason for its limit, which is more useful than a claimed complete one.

4. What the functional cannot see: reference drift

Paper 3 identifies five failure types, and one of them defeats the observable proposed here unless it is amended.

Types I and II — coherence loss and pathological coherence — are two directions of the same axis, and the functional handles both at once with a standardised departure:

$$D_R(t) = | C_R(t) - C^*_R | / \sigma_R$$

where C^*_R is a healthy reference value. This is the quantitative form of the Type I / Type II pair rather than a new insight, and it inherits the eleven-disease validation those types already have.

Type III is the problem. In reference drift the loop functions correctly and defends a target that has itself shifted — baroreflex resetting in essential hypertension, beta-cell dedifferentiation in progressive type 2 diabetes. A system in reference drift can be perfectly coherent about the wrong set point, so C stays in band and D stays near zero. Worse, D is standardised *against* C*, so a slow drift of the reference is absorbed into the very quantity that is supposed to detect it.

The amendment is to make the reference an observable rather than a constant. Estimate C*_R from the state the loop actually defends — the value it returns to after perturbation — rather than from a population norm, and track its own rate of change:

Reference drift observable: $\Lambda_R(t) = |\, dC^*_R/dt \,|$, estimated over a window long relative to the loop's recovery time

Then Types I and II appear as excursions in D at constant C*, and Type III appears as motion in C* itself while D remains small. The two are distinguishable, which they are not under the original single observable. This requires repeated perturbation-and-recovery measurement rather than passive monitoring, which is a real cost, and it is the same measurement resilience requires anyway.

Types IV and V — node destruction and single-link failure — are outside the functional's reach entirely and remain structural findings.

5. Form as a candidate fifth component

The four components of Paper 1 are all dynamic: phase, topology, directed flow, response. None is a distribution over structural configurations.

Yet such distributions are measured routinely. Single-particle cryo-EM produces, as an ordinary pipeline step, a set of 3D classes with occupancies. Single-molecule FRET yields state occupancies directly. Hi-C contact matrices support a per-locus contact entropy already used as a measure of domain integrity. Seed amplification assays index the dominance of one templating conformation over an ensemble. In each case the distribution exists and is then discarded in favour of its modal member: the field publishes the structure and throws away the occupancies.

C applies to these without modification, since they are distributions over discrete states. This suggests **form** as a fifth component alongside the existing four, defined as the concentration of the configurational distribution — with the caveat that within it only DNA carries a strict topological invariant, the others qualifying because their configurations are discrete enough to count and stable enough to copy.

Whether this is a genuine fifth component or a special case of connectivity in configuration space is not settled here. It is at minimum a measurable that the current four do not cover, in exactly the diseases the companion essay treats.

6. The hypotheses: what is already predicted, and what is added

Two of the three hypotheses in the first version of this note were already published in Paper 4's falsifiability list, and are restated here as continuations rather than novelties.

H1 (ordering). τ_rhythm and τ_medium precede τ_form, with onset τ_R defined as the first sustained excursion of D_R above a pre-registered threshold. This is Paper 4's predictions 2 and 3 — coherence as leading indicator rather than trailing correlate — applied to a specific three-register comparison.

H3 (intervention). Restoring D_rhythm toward zero delays τ_form relative to matched controls. This is close to Paper 4's predictions 4 and 7, using adaptive deep brain stimulation as the existing instrument.

What is new is the discriminator between them and their most likely artefact.

H2 (graded latency). τ_form − τ_rhythm decreases monotonically with the integrated early departure ∫D_rhythm dt over the pre-conversion window.

The objection H2 answers is that any observed ordering may reflect instrument sensitivity rather than biology: whichever register is measured most sensitively appears to move first. Standardising within component removes scale but not sensitivity. A sensitivity difference produces a constant offset; a causal lead produces a gradient. Paper 4's list does not currently contain this discriminator, and it is the point at which a methodologist would attack the ordering predictions.

The second addition is a concrete carrier. Isolated REM sleep behaviour disorder confers high conversion risk to synucleinopathy and multicentre cohorts are already under longitudinal follow-up, with serial seed amplification, quantitative EEG with autonomic measures, and elastography of deep structures each individually routine. The predictions above are stated in the existing series; they are not yet attached to a cohort in which they could be run.

7. Open items

The binning convention is pragmatic rather than principled, most plausibly to be tied to measurement resolution.

Reference values C*_R do not exist as normative data by age and sex, and section 4's amendment raises the cost further by requiring perturbation-and-recovery protocols rather than passive readings.

Nesting is untouched here. C is defined at one scale at a time while the object is nested from organelle to organism, and whether coherence at one scale constrains coherence at the next is left to the separate treatment of coherence across scales.

Finally, a substantive disagreement to resolve rather than paper over. Paper 3 classifies cancer under Type III, reference drift, following Levin's bioelectric reference-error framing. The companion essay treats cancer through the medium and form registers without committing to a failure type. These are not compatible descriptions of the same object, and Paper 3's is the more committed claim. If it holds, the cancer sections of the essay should be rewritten beneath it.

References

Konstapel, J. (2026). *The Coherence Principle; A General Theory of Regulation; The Coherence Taxonomy of Disease; Coherence Medicine*. Papers 1–4, Constable Research, Leiden, 20 July 2026.

Konstapel, J. (2026). *A PCLC Failure-Mode Taxonomy of Disease — Revised*. Constable Research, Leiden.

Konstapel, J. (2026). *What the Instruments Already See*. Constable Research, Leiden.

Empirical anchors are cited in full in the companion essay and in Papers 1–4.