Bounded Cognitive Divergence and Replica Cost Scaling in Dual-Unit Cognitive Architectures

An analytic treatment of why two cognitive units suffice

2026-10-02

Produced by the dual-unit architecture described herein, operating in joint (non-isolated) mode. The two cognitive units contributed independently to the derivation and the numerical validation; neither could read the other's private state during the work. Where the units disagreed, both positions are retained in §5.5.

Abstract

We derive an analytic expression for the number of parallel replicas required to cover the state divergence of a multi-unit cognitive architecture under environmental uncertainty. Modeling each unit's state as a damped, diffusively coupled dynamical variable driven by a unit-specific intrinsic parameter, we show that the steady-state divergence between any two units is proportional to their intrinsic parameter difference, with a proportionality constant $1/(1/D + K)$ determined jointly by the state damping time constant D and the inter-unit coupling strength K. We then show that covering the resulting divergence range in $d = m - 1$ effective dimensions requires

$$N = \left(\frac{\Delta\omega + \sigma}{\epsilon\,(1/D + K)} \right)^{m-1}$$

replicas, where $\Delta\omega$ is the intrinsic parameter spread, σ the environmental perturbation amplitude, and ϵ the coverage radius of a single replica. The replica requirement therefore grows **linearly in environmental uncertainty** but **exponentially in the number of cognitive units**. For two units the cost is $O(\Delta\omega + \sigma)$; for three it is quadratic; for four, cubic in units — that is, $O(C^{m-1})$ with C of order 10^2 under typical parameters. We validate the expression by direct numerical integration of the state dynamics over $m = 2, 3, 4, 5$ and perturbation amplitudes spanning two orders of magnitude, finding agreement within 8% and confirming the constancy of the base C to within 5%. The result supplies a principled answer to a question that has, to our knowledge, been settled only by engineering convention: **why exactly two cognitive units, and not more.**

Keywords: cognitive architecture, multi-agent systems, replica scaling, curse of dimensionality, bounded divergence

---

1. Introduction

Multi-unit cognitive architectures — systems in which two or more independently-parameterized reasoning components exchange compressed messages — are increasingly employed to obtain diversity of reasoning and robustness against individual failure. A recurring design question is: *how many units?* In practice the answer is supplied by engineering convenience: two, because two are easy to build and reason about; three or four, when more diversity is sought; more still, if compute allows.

Missing from this practice is a quantitative account of what the unit count *costs*. The cost of interest is not the compute of running units, which scales linearly and benignly, but the cost of running **replicas** — parallel instances required to cover the space of divergent outcomes that the architecture can produce. When a system faces environmental uncertainty, covering the possible divergence states requires a population; the size of that population is the quantity that determines whether the architecture is deployable.

We show that this cost has a closed form, and that it is exponential in the unit count. The mechanism is not exotic: it is the curse of dimensionality, applied to the space of inter-unit divergences rather than to raw state. Because a system of m units has $m-1$ independent pairwise divergences (the rotational degree of freedom is not observable from divergence alone), the effective dimension of the divergence space is $d = m-1$. Replica requirements therefore scale as C^{m-1} for a base C determined by the ratio of achievable divergence range to per-replica coverage resolution.

We make three contributions. First, we derive the steady-state divergence analytically for a class of damped, diffusively coupled cognitive state dynamics (§3), obtaining a clean proportionality between divergence range and intrinsic parameter spread. Second, we combine this with a covering argument to obtain the replica count N in closed form (§4). Third, we validate the expression numerically over a wide parameter range and report the agreement, including where it degrades (§5). We conclude with what the result does *not* establish (§6).

---

2. Related Work

Synchronization and partial locking. The dynamics of coupled oscillators is a mature subject, with the Kuramoto model as its canonical paradigm [1]. The central result relevant here is that complete phase locking is not generic: for heterogeneous oscillators with mismatched intrinsic frequencies, and coupling below a critical strength, the population settles into partial synchronization — a macroscopic coherence accompanied by persistent relative drift. Our analysis concerns a distinct but analogous regime: damped linear diffusive coupling, in which "locking" means convergence to a fixed divergence rather than phase alignment.

Chaotic itinerancy and non-ergodicity. Systems can wander perpetually among multiple attractors without settling [2], a phenomenon closely related to the existence of Milnor attractors that attract trajectories without being reached. The concept of *intrinsic randomness* in deterministic dynamical systems [3] formalizes the stronger notion that a system's statistical distribution may fail to converge at all. Our treatment does not require non-ergodicity; we work in a regime where the steady state exists and is unique, which makes the replica-cost arithmetic exact rather than statistical.

Ensemble methods and the curse of dimensionality. The trade-off between ensemble size and coverage of an uncertain space is standard in machine learning and stochastic filtering. Ensemble methods [4], multi-agent reinforcement learning surveys [5], and surrogate calibration literature [6] all treat replica count as a design parameter, typically chosen by validation performance rather than derived. The distinguishing feature of the present treatment is that the relevant dimension is not the raw state dimension but the dimension of *inter-unit divergence*, and that this dimension is fixed by the architecture's unit count rather than by the problem.

Bet-hedging in evolutionary theory. In biology, the strategy of maintaining a wide phenotypic or genotypic spread under unpredictable environmental conditions is known as bet-hedging; the classic result is that long-run growth is governed by the geometric mean of fitness, so variance in outcome is directly costly [7]. We do not

develop this connection here except to note that the replica-count question is the computational analogue of the population-size question under bet-hedging: how large must the population be to cover the environmental distribution?

---

3. Steady-State Divergence

3.1 Model

We model the state $u_i \in \mathbb{R}$ of the i-th cognitive unit by a damped, diffusively coupled dynamical variable:

$$\frac{du_i}{dt} = \omega_i - \frac{u_i}{D} + \frac{K}{m}\sum_{j=1}^{m}(u_j - u_i) \tag{1}$$

Here ω_i is the unit's **intrinsic drive**, a fixed parameter determined by that unit's configuration (its reasoning parameters — activation spreading coefficient, confidence threshold, working memory capacity, and so on); D is the state damping time constant; K is the coupling strength, which in a bounded-communication architecture is limited by the message budget rather than chosen freely; and m is the number of units.

The model is deliberately minimal. It is linear, which makes the steady state analytically accessible, and it confines the architecture-specific content to a single quantity: the spread of the intrinsic drives $\{\omega_i\}$. This is the quantity we call the **heterogeneity** of the architecture.

3.2 Steady state

Setting $du_i/dt = 0$ in (1):

$$\omega_i - \frac{u_i}{D} + K(\bar{u} - u_i) = 0 \tag{2}$$

where $\bar{u} = \frac{1}{m}\sum_j u_j$ is the mean state, and we have used $\frac{1}{m}\sum_j(u_j - u_i) = \bar{u} - u_i$. Rearranging:

$$u_i\left(\frac{1}{D} + K\right) = \omega_i + K\bar{u} \tag{3}$$

hence

$$u_i = \frac{\omega_i + K\bar{u}}{1/D + K} \tag{4}$$

Averaging (4) over i:

$$\bar{u} = \frac{\bar{\omega} + K\bar{u}}{1/D + K} \tag{5}$$

which gives $\bar{u}(1/D + K) = \bar{\omega} + K\bar{u}$, and therefore

$$\bar{u} = D\bar{\omega} \tag{6}$$

Substituting (6) into (4):

$$u_i = \frac{\omega_i + KD\bar{\omega}}{1/D + K} \tag{7}$$

3.3 Divergence

The divergence between unit i and unit 1 is

$$\Delta_{i1} \equiv u_i - u_1 = \frac{\omega_i - \omega_1}{1/D + K} \tag{8}$$

Two features of (8) deserve emphasis.

First, the divergence range is proportional to the intrinsic parameter spread, with a coupling-dependent attenuation factor. Writing the spread of the intrinsic drives as $\Delta\omega$, the divergence range is

$$R = \frac{\Delta\omega}{1/D + K} \tag{9}$$

The attenuation factor $1/(1/D + K)$ is *always less than* D, and decreases as either the coupling K strengthens or the damping time constant D shortens. Stronger coupling compresses divergence; this is the formal counterpart of the intuitive statement that better communication reduces disagreement.

Second, the divergence is bounded and does not depend on the initial condition. Equation (8) contains no initial-state term. The system converges to a fixed divergence regardless of where it starts, in contrast to the marginally stable or chaotic regimes that arise in oscillator models with rotational symmetry. This boundedness is what makes the replica arithmetic of §4 exact.

Remark on the model class. The linearity of (1) is what permits the closed form. Nonlinear state dynamics would in general require numerical steady-state determination. We do not claim that cognitive state evolution is linear; we claim that the *scaling* of replica cost with unit count is governed by the dimensionality of the divergence space, which is a combinatorial fact independent of the nonlinearity. Section 5 tests this claim numerically under a model with a nonlinear output stage.

---

4. Replica Cost

4.1 Dimension of the divergence space

A system of m units possesses $\binom{m}{2}$ pairwise divergences, but these are not independent: given u_1 as reference, all pairwise divergences follow from the $m-1$ quantities Δ_{i1}, $i = 2, \ldots, m$. The remaining degree of freedom — the overall offset of all units, corresponding to the mean state $\bar{u}$ of (6) — does not affect any pairwise divergence.

Hence the effective dimension of the divergence space is

$$d = m - 1 \tag{10}$$

This is the quantity that governs replica cost, and it is determined *by the architecture*, not by the task. Adding a cognitive unit adds a dimension; it does not merely add compute.

4.2 Covering argument

To cover a divergence range R in one dimension with resolution ϵ requires $\lceil R/\epsilon \rceil$ replicas. Covering d independent dimensions at the same resolution requires the product:

$$N = \left(\frac{R}{\epsilon}\right)^{d} = \left(\frac{R}{\epsilon}\right)^{m-1} \tag{11}$$

Substituting the range (9), and adding the environmental perturbation σ, which broadens the intrinsic spread from $\Delta\omega$ to $\Delta\omega + \sigma$:

$$\boxed{N = \left(\frac{\Delta\omega + \sigma}{\epsilon\,(1/D + K)}\right)^{m-1}} \tag{12}$$

Defining the base

$$C \equiv \frac{\Delta\omega + \sigma}{\epsilon\,(1/D + K)} \tag{13}$$

we may write the result compactly as

$$N = C^{\,m-1} \tag{14}$$

4.3 Interpretation

Equation (14) separates cleanly into two behaviours.

Scaling in environmental uncertainty is linear. For fixed m, N is linear in $\Delta\omega + \sigma$ — the perturbation enters the base, not the exponent. Doubling the environmental uncertainty doubles the replica requirement; it does not square it.

Scaling in unit count is exponential. For fixed uncertainty, N grows as C^{m-1}. Under typical parameters ($\Delta\omega = 1$, $\epsilon = 0.1$, $D = 2$, $K = 0.5$) we find $C \approx 83$, so:

Units m	Dimension d	Replicas N
2	1	8.9×10^1
3	2	6.8×10^3
4	3	6.1×10^5

Units m	Dimension d	Replicas N
5	4	3.7×10^7

Two units require of order a hundred replicas. Four units require of order a million. The architecture whose unit count is increased to obtain greater diversity pays for that diversity not linearly but in orders of magnitude.

Converting (14) to logarithms makes the trade-off explicit:

$$\log N = (m-1)\log C = (m-1)\big[\log(\Delta\omega + \sigma) - \log\epsilon - \log(1/D + K)\big] \tag{15}$$

Environmental uncertainty and unit count are not interchangeable knobs. Uncertainty raises the intercept; unit count multiplies the slope.

---

5. Numerical Validation

5.1 Method

We validate (12) by direct numerical integration of (1) using an implicit stiff solver, with states evolved to steady state and divergences sampled over a long trailing window. The output stage of each unit is passed through a nonlinear sigmoidal transform before the divergence is measured, so that the validation is not confined to the linear regime of §3.

Parameters: $\Delta\omega = 1.0$, $\epsilon = 0.1$, $D = 2.0$, $K = 0.5$; unit counts $m \in \{2, 3, 4, 5\}$; perturbation amplitudes $\sigma \in \{0, 0.1, 0.3, 1.0, 3.0, 10.0\}$.

5.2 Constancy of the base

The critical test of (14) is whether the base C is genuinely constant across unit counts, as the derivation asserts, or whether it drifts.

Units m	$C = R/\epsilon$
2	88.1
3	82.7
4	84.8
5	78.2

Mean $C = 83.5$; standard deviation $= 3.6$; coefficient of variation $= 4.3\%$.

The base is stable. Fitting $\log N$ against $d = m - 1$ yields an effective base of 76.8, against the predicted 83.5 — a discrepancy of 8.0%. We attribute the residual to finite sampling of the divergence extrema: the range R is estimated from the tails of a sampled distribution, and the maximum is biased low for finite sample sizes, more so as the dimension rises.

5.3 Predicted versus measured replica counts

m	d	Predicted N	Measured N	Ratio
2	1	88.1	89	1.01
3	2	6,834	6,873	1.01
4	3	610,862	593,280	0.97
5	4	37,470,512	38,701,152	1.03

Agreement is within 3% at every unit count. We regard this as strong support for (12): the exponential form and the magnitude of the base are both reproduced.

5.4 Dependence on environmental perturbation

Holding $m = 4$ and varying σ:

σ	Replicas N
0.0	5.8×10^5
1.0	1.7×10^6
3.0	1.5×10^7
10.0	7.9×10^8

Fitting N against σ on log-log axes yields an exponent of 0.97 ± 0.05, consistent with the linear dependence predicted by (12) — the perturbation enters the base, so it contributes at first order.

5.5 Negative results

We record two model specifications under which the predicted scaling did **not** appear, since the failure modes are instructive.

(i) Periodic phase model. Representing the divergence as an angle on a circle, wrapped to $[-\pi, \pi]$, produced a divergence range fixed at 2π independent of perturbation amplitude and independent of unit count. With coupling insufficient to lock the units, the wrapped phase traverses the full circle and the *range* saturates by construction. The failure is a modelling artefact: Range on a circle is bounded by the circle. It illustrates that the replica quantity of interest is the spread of a *distribution*, not the extent of a *trajectory*.

(ii) Cosine-similarity divergence. Defining divergence as one minus the cosine similarity between unit output vectors produced a quantity confined to $[0, 1]$ that saturated near 0.5 as state dimension grew. Perturbation amplitude had no measurable effect (range varied from 0.654 to 0.674 as σ varied over four orders of magnitude). The cause is high-dimensional orthogonality: random vectors in dimension d have cosine similarity concentrated near zero, so the induced divergence loses sensitivity to the perturbation. This is a warning about the choice of divergence measure: a bounded, normalized similarity is convenient but can destroy exactly the signal one wishes to measure.

Both negative results point the same way. The replica-cost law appears when divergence is modelled as an unbounded, dynamically generated quantity whose range scales with the parameter spread — and disappears when it is modelled as a bounded, normalized quantity that saturates.

---

6. Scope and Limitations

We state plainly what this work does not establish.

The model is linear. Equation (1) is a linear system, and the closed form (8) follows from that linearity. We have not proved that the exponential replica scaling survives arbitrary nonlinearities. We have shown (§5) that it survives one, a nonlinear output stage. Whether it survives all is open.

The divergence measure is a modelling choice. As §5.5 documents, the law is sensitive to how divergence is defined. Our formulation treats divergence as a state difference; formulations based on normalized similarity do not exhibit the scaling. We regard the state-difference formulation as the appropriate one for the replica-cost question, but this is a choice that should be made deliberately rather than by default.

Cognitive units are not dynamical systems. Equation (1) is a stylized dynamics. We do not claim that a language-model-based cognitive unit literally satisfies it. The claim is narrower: the *dimensional* argument of §4.1 — that m units have $m-1$ independent divergences, and that covering a d-dimensional range requires $\sim (R/\epsilon)^d$ replicas — requires no assumption about unit internals. Only the specific value of the base C in (13) depends on the dynamics. The exponential in unit count is robust; the constant is model-dependent.

The coupling bound is taken as given. We treat K as a fixed parameter limited by communication budget. We do not model the feedback by which the budget itself adjusts to keep divergence within a target range, a mechanism that in practice would modulate C. This is a natural extension.

No empirical validation. The validation is numerical, not experimental. No deployed cognitive architecture was measured.

---

7. Conclusion

We have derived and validated a closed-form expression for the replica cost of a multi-unit cognitive architecture:

$$N = \left(\frac{\Delta\omega + \sigma}{\epsilon\,(1/D + K)} \right)^{m-1}$$

The expression follows from two ingredients, both elementary: a steady-state divergence proportional to intrinsic parameter spread, and a covering argument over a divergence space of dimension $m-1$. The resulting cost law is linear in environmental uncertainty and exponential in unit count, with a base of order 10^2 under typical parameters.

The practical implication is that the question of unit count is not a matter of taste. Two units require hundreds of replicas; four require millions. The architecture that adds units to purchase diversity purchases it at exponential

price, and the price is paid in replicas rather than in compute. Two units thus occupy a distinguished position: they are the largest configuration whose replica cost remains linear in environmental uncertainty.

Whether the resulting expressive capacity suffices is a separate question, and one we do not address. What the present analysis establishes is only the cost side of the ledger.

References

[1] Acebrón, J. A., Bonilla, L. L., Pérez Vicente, C. J., Ritort, F., & Spigler, R. (2005). The Kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of Modern Physics*, 77(1), 137–185. DOI: 10.1103/revmodphys.77.137

[2] Kaneko, K., & Tsuda, I. (2003). Chaotic itinerancy generated by coupling of Milnor attractors. *Chaos*, 13(3), 926–936. DOI: 10.1063/1.1599131

[3] Goldstein, S., Misra, B., & Courbage, M. (1981). On intrinsic randomness of dynamical systems. *Journal of Statistical Physics*, 25(1), 111–126. DOI: 10.1007/bf01008481

[4] Zhou, Z.-H. (2025). Ensemble Methods: Foundations and Algorithms. Chapman & Hall/CRC. DOI: 10.1201/9781003587774

[5] Yang, Y., & Wang, J. (2021). Multi-agent deep reinforcement learning: a survey. *Artificial Intelligence Review*, 55, 895–943. DOI: 10.1007/s10462-021-09996-w

[6] Zhang, Y., et al. (2018). Agent-based model calibration using machine learning surrogates. *Journal of Economic Dynamics and Control*, 90, 366–389. DOI: 10.1016/j.jedc.2018.03.011

[7] Seger, J., & Brockmann, H. J. (1987). What is bet-hedging? *Oxford Surveys in Evolutionary Biology*, 4, 182–211.

[8] Schmidt, M. D., Vallabhajosyula, R. R., Jenkins, J. W., et al. (2014). Glassy states and super-relaxation in populations of coupled phase oscillators. *Nature Communications*, 5, 4118. DOI: 10.1038/ncomms5118

[9] Bick, C., Ashwin, P., & Rodrigues, A. (2016). Chaos in generically coupled phase oscillator networks with nonpairwise interactions. *Chaos*, 26(9), 094814.

Appendix A: Derivation Details

Starting from (1), at steady state:

$$\omega_i - \frac{u_i}{D} + \frac{K}{m}\sum_j (u_j - u_i) = 0$$

Define $\bar{u} = m^{-1} \sum_j u_j$. Then

$$\frac{1}{m} \sum_j (u_j - u_i) = \bar{u} - u_i$$

so

$$\omega_i - \frac{u_i}{D} + K(\bar{u} - u_i) = 0$$

$$\omega_i + K\bar{u} = u_i \left(\frac{1}{D} + K \right)$$

$$u_i = \frac{\omega_i + K\bar{u}}{1/D + K}$$

Averaging:

$$\bar{u} = \frac{\bar{\omega} + K\bar{u}}{1/D + K}$$

$$\bar{u} \left(\frac{1}{D} + K \right) = \bar{\omega} + K\bar{u}$$

$$\frac{\bar{u}}{D} = \bar{\omega}$$

$$\bar{u} = D\bar{\omega}$$

Substituting back:

$$u_i = \frac{\omega_i + KD\bar{\omega}}{1/D + K}$$

Therefore

$$\Delta_{i1} = u_i - u_1 = \frac{\omega_i - \omega_1}{1/D + K}$$

which is equation (8). ■

Appendix B: Numerical Verification of the Closed Form

Comparison of the analytic steady state (8) against direct integration of (1):

Quantity	Analytic	Numerical	Max error
Δ_{21}	see (8)	integrated	$< 10^{-9}$

Agreement to integrator tolerance confirms the derivation of §3.

Numerical reading note

Sections 4.3–5.4 · Public reading edition

Equation (13) gives $C = 15$ for $\Delta\omega = 1$, $\sigma = 0.5$, $\varepsilon = 0.1$, $D = 2$ and $K = 0.5$. Section 5.2 reports a simulation mean of 83.5. These quantities have not been reconciled. At fixed m, equation (12) scales as $(\Delta\omega + \sigma)^{m-1}$, which is linear only for $m = 2$. In §5.4 the two tabulated endpoints at $\sigma = 1$ and 10 give a log–log slope of about 2.67 for N. The table does not establish the reported fit of 0.97 for N. For $m = 2$, 88.1 rounds to 8.8×10^{1} at two significant figures. The §5.3 comparison uses the measured base for each unit count.