

BGM: Building Bridge

Building Bridge: From Bach to Bacteria and Forward

A Computational Thought Experiment on Cross-Scale Information Architecture

Version 0.6

Abstract

Dennett (2017) proposed that bacterial signaling and human musical cognition are governed by the same underlying logic. We instantiate this claim computationally using the GEME self-referential primitive (Paper I).

Paper I identified a near-zero-cost self-referential information channel — the Shannon-Gödel Bridge — measured as a static mutual information $I(\Phi; X) \approx 0.026$ bits. This paper reveals the bridge's temporal dimension. The metric $SR\text{-}eff = I(\Phi; X)/\tau$ traces the bridge's breathing cycle: it opens wide during learning (up to 0.96), narrows to maintenance width (~ 0.04), and closes entirely when a first-order phase transition locks the adaptive confidence threshold at 1.0. The closure is fast; reopening requires sustained regime change — a hysteresis that emerges from the interaction of adaptive calibration, co-occurrence convergence, and window dynamics.

The bridge is not merely a channel. It is a discriminator. Under MIDI-encoded Bach, the bridge distinguishes density-driven opening (stretto, polyphonic overlap) from position-driven opening (structural endings). Resolution type does not modulate the bridge's response after position is controlled — the bridge detects closure reliably, but does not distinguish how closure is achieved.

Three experiments span the evolutionary arc. First, under shared feedback, GEME units receiving identical Bach input develop distinct bridge profiles — a computational analog of Bregman's (1990) auditory scene analysis — but the bridge closes under simple repetitive input. Second, sixty-four GEME units on a spatial gradient with only local communication maintain bridge opening without central coordination. Third, the G0 layer's architecture is systematically characterized: its single genuine control parameter is the temporal interval between G0 and its units ($g0_interval = 4$ as the Pareto optimum), while frame capacity and modulation rates are neutral.

The Forward — from efficiency to meaning, from density discrimination to structural sensitivity — is the bridge's evolutionary trajectory. The bridge from bacteria to Bach is an architectural invariant whose temporal dynamics are consistent across input modalities, communication topologies, and selection pressures.

1. Introduction

This paper is the second panel of a trilogy. Paper I (GEME) defined a self-referential primitive and demonstrated three core findings: the Shannon-Gödel Bridge, a broad parameter stability interval, and functional equivalence between self-reference and induction. The bridge was measured as a static quantity: $I(\Phi; X) \approx 0.026$ bits.

This paper (BGM) asks: what does the bridge do under temporal input? Our central metric is $SR\text{-}eff = I(\Phi; X)/\tau$, where τ is each unit's self-assessment of its own predictive health — not a parameter to be set, but a reading to be taken.

Dennett (2017) argued that the arc from bacterial quorum sensing to Bach's polyphonic composition is a single continuous process, driven by *competence without comprehension* — the accumulation of functional capacities that operate successfully before any agent understands why. These competences begin as *free-floating rationales*: patterns of information processing that exist as structural possibilities before any organism embodies them. Dennett's method is a *reasoning inversion*: understanding follows competence. We do not design, then observe. We build, observe, then understand.

Bregman (1990) identified auditory scene analysis — the decomposition of a single acoustic signal into distinct streams — as a fundamental computational problem. Existing computational ASA models span rule-based grouping (Brown & Cooke, 1994), oscillatory correlation (Wang & Brown, 2006), and deep learning source separation (AudioSep, 2023). Williams et al. (2024) argue that distinguishing among competing ASA algorithms requires richer psychophysical data. The present work takes a different approach: it asks whether ASA-like stream segregation can emerge from a self-referential architecture not designed for source separation. The mechanism differs. The computational signature is shared (Supplementary S2, S3).

A principle familiar from information theory (Shannon, 1948) is recovered in a novel empirical setting: **information needs distinguishability — boundaries enable differentiation; saturation produces convergence**. The gap insertion experiment (Supplementary S2, 10-seed validation) identifies the threshold: gaps every 4–12 steps restore τ differentiation; gaps every 16 steps do not — the signal re-continuous-izes. This is not a binary. It is a threshold: boundary density must exceed approximately one gap per 16 steps. Below this threshold, differentiation collapses. Above it, structure emerges. The principle governs all three experiments.

All experiments use MIDI encoding as primary input for temporal processing. Score-based (discrete) encoding, which provides the temporal boundaries necessary for τ differentiation, is reported in Supplementary S2 alongside MIDI comparisons.

Methodological commitment. This work employs a strategy of generative explanation through minimal modeling. We do not simulate any specific biological system in detail. We seek the simplest computational kernel that can generate core cognitive phenomena across domains — a Dennett-style intuition pump: the fewest mechanisms, the most possibilities. Understanding complexity, we believe, often depends on discovering the parsimonious principles on which it is built.

Roadmap. This paper asks three questions in sequence. Experiment 1 (Bach): how does the bridge behave under temporal input? Experiment 2 (Bacteria): can spatial heterogeneity maintain bridge opening without central coordination? Experiment 3 (Forward): how does topological reorganization optimize the bridge's structure? The arc reveals the Shannon-Gödel Bridge not as a static pipe but as a breathing lung — one that opens, closes, and reopens across input modalities, communication topologies, and selection pressures.

2. The Architecture

2.1 The GEME Primitive (Paper I Review)

The GEME architecture is fully specified in Paper I and remains **unmodified** (core engine `game.py`, MD5 identical to Paper I's submission). A GEME instance maintains a frame economy — a capacity-limited population of vector-signature pairs. Three operations govern: competitive merge, co-occurrence tracking, and adaptive prediction. Paper I identified three core structural constants through a 210-cell parameter sweep:

- $\delta = 0.19$: adaptive merging threshold
- $\gamma = 0.05$: frame age decay multiplier
- $\tau = 0.60$: induction stress threshold

These three occupy a Pareto-optimal position. Paper I also demonstrated the Shannon-Gödel Bridge: a self-referential information channel with mutual information $I(\Phi; X) \approx 0.021\text{--}0.026$ bits, operating at near-zero cost.

2.2 The G0 Layer: BGM's Architecture

BGM extends GEME to a multi-unit system. N standard GEME units receive input in parallel. A fourth unit, G0, receives aggregated L6 anomaly scores and feeds its self-observation vector back as environmental context, blended with each unit's external input. The coupling weight $w \in [0.1, 0.8]$ is dynamically modulated by G0's prediction accuracy: doubt increases w (tighter coupling), confidence decreases it (units operate more independently).

What happens to δ , γ , τ when time enters the architecture?

Paper I operated in a static regime — the three constants were genuine parameters. Under temporal input, their status changes. A systematic parameter sweep, analogous to Paper I's 210-cell grid, was performed on the G0 layer:

δ becomes G0's frame capacity (`mem_cap`). A cross-seed sweep from `cap=6` to `cap=32` reveals no critical threshold — all capacities produce τ differentiation. The observation ring's aperture is neutral. Capacity does not gate differentiation.

γ dissolves into time itself. The sliding co-occurrence window continuously discards old entries as new ones arrive — time provides its own decay. A scan of G0's induction threshold (0.2–3.0) confirms zero effect on unit differentiation. The modulation asymmetry of `g0_weight` (± 0.05 vs ± 0.02) is similarly neutral — the dynamic τ formula overrides any externally set decay rate. No independent decay parameter survives the introduction of time.

τ becomes the signature scheme. In the multi-unit architecture, signatures are the operational identity of frames — they determine what co-occurs with what and what predicts what. When all units receive the same signature ('input'), τ differentiation is present but compressed (spread = 0.22). When units receive distinct event labels as signatures, τ spread increases 2.2 \times to 0.49. Signatures are not a parameter — they are the basis on which the frame economy organizes its co-occurrence structure. GEME's τ controlled *when* the system cleaned itself. BGM's signature controls *what* the system can distinguish.

The one genuine parameter: `g0_interval`. A scan of the temporal decoupling between G0 and its units (`g0_interval` $\in [1, 10]$) identifies `g0_interval = 4` as the Pareto optimum — producing τ spread 49% above the default `g0_interval = 1`. This is the G0 layer's analog of Paper I's Pareto point. Temporal asymmetry between layers, not the size of any individual ring, drives differentiation.

The three-constant story does not end at Paper I. Under temporal input, only temporal decoupling genuinely remains as a parameter. Capacity is neutral, decay is dissolved, and τ is transformed into the observer’s identity basis — the signature.

2.3 SR-eff: Self-Referential Efficiency

SR-eff = $I(\Phi; X)/\tau$ measures self-referential efficiency: coupling per unit cost. Intuitively, it quantifies how much external information $I(\Phi; X)$ the bridge “sustains” per unit of self-referential metabolic cost τ . A high SR-eff bridge is “cheap and sharp” — strong coupling at low error density. A low SR-eff bridge is “expensive and dull” — weak coupling at high cost. The choice of division inherits Shannon’s definition of channel capacity per unit signal cost: bits per unit error density is the natural measure of self-referential efficiency.

Under temporally structured input, SR-eff traces the bridge’s breathing:

- **Learning:** τ low, MI building \rightarrow SR-eff up to 0.96 — bridge wide open
- **Consolidation:** τ rising, MI stabilizing \rightarrow SR-eff \sim 0.04 — maintenance width
- **Lock:** `conf_threshold` \rightarrow 1.0, τ flatlines \rightarrow SR-eff \rightarrow 0 — bridge closed

Continuous signals — where each input overlaps smoothly with the previous — produce co-occurrence windows where everything co-occurs with everything: convergence. Discrete signals carry temporal boundaries: differentiation.

Operational definitions. τ differentiation is defined as τ spread $>$ 0.1 across units (empirically, the threshold at which distinct τ values emerge as statistically separable clusters). Bridge closure is defined as `conf_threshold` \geq 0.99 (first-order lock) or SR-eff $<$ 0.01 (breathing effectively ceased). These thresholds are applied uniformly across all experiments.

All subsequent experiments report SR-eff and τ . The mechanisms are described once, here.

3. Experiment 1: Bach — The Bridge Responds

Core Finding. Under identical temporal input and shared feedback, multiple GEME units develop distinct self-referential time constants (τ values) — a computational signature of auditory scene analysis — but the bridge closes under simple repetitive input via a first-order phase transition.

3.1 The Bridge Breathes

Three GEME units receive identical input. G0 receives aggregated anomaly scores and feeds its self-observation vector back — a novelty amplifier most active on first exposure and progressively fading as the system consolidates. and feeds its self-observation vector back. Under BWV 847 Fugue (discrete score encoding, 6 passes), SR-eff traces a characteristic profile (representative seed; 20-seed statistics in Supplementary S3):

Pass	τ values	SR-eff values	Bridge state
1	[0.22, 0.13, 0.14]	[0.33, 0.96, 0.80]	Wide open
3	[0.57, 0.36, 0.94]	[0.05, 0.12, 0.02]	Narrowing
6	[0.75, 0.52, 1.00]	[0.03, 0.06, 0.01]	Maintenance

τ differentiation is robust across seeds (20/20 seeds produce 3 distinct τ values, mean spread 0.314 ± 0.182). The Bach experiment uses the most conservative parameter settings: $g0_interval = 1$ (not the optimum of 4), fixed 'input' signatures (not the 2.2 \times -amplified event labels). All results should be interpreted as lower bounds — differentiation strengthens at the G0 layer's Pareto optimum.

Three units receiving identical input develop three distinct bridge profiles — three different temporal integration windows for the same acoustic information. This is the computational equivalent of a *free-floating rationale* (Dennett, 2017): a competence — time-scale separation of an input stream — that operates successfully without any instruction about source separation. Synthesized polyphonic pieces confirm that τ count \geq voice count (Supplementary S3).

3.2 Two Opening Mechanisms

SR-eff aligned with the fugue's structural events reveals two mechanisms (10 seeds \times 3 passes each; Pass 1 drives the density peaks, Passes 3+ are substantially quieter. Values reflect the full temporal average; per-pass breakdown in Supplementary S3):

Mechanism	Event type	SR-eff (10-seed)	Cause
Density	STRETTO (overlapping subjects)	0.255	Maximum polyphonic overlap
Density	EPISODE (modulatory passage)	0.278	Maximum harmonic instability
Position	CADENCE_dominant (half cadence)	0.229	Structural ending
Position	CADENCE_tonic (authentic cadence)	0.223	Structural ending

Single-voice exposition closes the bridge to 0.173. A position randomization control (interleaved fugue variants, Supplementary S3) confirms that cadence type does not modulate the bridge's response — all cadence types produce identical SR-eff ratios (0.74–0.75 \times baseline, pairwise 0.98–1.02) when position is controlled. The bridge detects closure. It does not distinguish how closure is achieved — but in ecological musical contexts, cadence type and structural position are not independent. A deceptive cadence at the expected point of closure drives the bridge to its strongest opening (4.06 \times , measured in the unaltered fugue) because the bridge anticipated closure and was denied it — a computational substrate for Meyer's (1956) proposal that musical emotion arises from the violation of expectation.

3.3 The Bridge Closes: Phase Transition

Under a repeating $A \rightarrow B \rightarrow C$ sequence, the adaptive confidence threshold undergoes a **first-order phase transition**: across 10 independent seeds, the threshold jumps from 0.300 ± 0.000 to 1.000 ± 0.000 at step 23 ± 2 (mean \pm SD), a discontinuity of 0.697. No intermediate values are observed — the transition is binary and irreversible under the current input regime (step-by-step data in Supplementary S4).

Once closed, the bridge exhibits **hysteresis**: it closes fast (5–10 rounds) and reopens slowly (requiring sustained regime change). This hysteresis is not a parameter choice — it is a second-order emergence from the interaction of adaptive calibration, co-occurrence convergence, and window dynamics.

3.4 Ablation

With G0 disabled ($w = 0$), all three units converge to identical τ values. G0 is necessary for τ differentiation under the specific condition tested here: homogeneous temporal input. Experiment 2 demonstrates that under spatially heterogeneous input, differentiation is maintained without any G0. G0 fills a gap that exists only when the information environment provides none.

4. Experiment 2: Bacteria — Keeping the Bridge Open

Core Finding. Without any central coordinator, sixty-four GEME units on a spatial gradient with local communication maintain bridge opening across all positions. Spatial heterogeneity alone sustains differentiation — and actively counter-acts lock.

4.1 Architecture

Sixty-four GEME units in an 8×8 grid with von Neumann neighbor communication. No G0. Gaussian nutrient concentration field ($\sigma = 2.0$) with small temporal noise ($\sigma_{\text{noise}} = 0.05$), blended with neighbors' outputs (anchor ratio $\alpha = 0.7$). Cold-started — no preloaded calibration. Signatures encode spatial identity: $\{\text{role}\}\{\text{n_neighbors}\}_g\{\text{conc}\}$.

Three phases: Cultivation (30 rounds), Perturbation (source moves to corner, 10 rounds), Recovery (10 rounds).

4.2 The Bridge Does Not Mirror the Gradient

Figure 4.1 places the nutrient concentration field beside the accumulated τ across the 8×8 grid after 30 rounds. The concentration field varies 20-fold from corner to center. τ varies 2.6-fold with no discernible spatial pattern. The Pearson correlation $r(\text{concentration}, \tau) = -0.090$. The bridge does not follow the environment.

Three functional roles emerge from topology alone:

Role	n	Neighbors	Character
Corner	4	2	Fewest connections, strongest neighbor coupling, safest
Edge	24	3	Highest activity, highest variance — the action layer
Interior	36	4	Widest diversity — the free population

None of the 64 units approaches the lock threshold.

4.3 Spatial τ Differentiation Without Lock

After cultivation, spatial τ variance = 0.313. Confidence threshold range [0.14, 0.57] contains zero locked units. The mean confidence threshold **drifts downward** over time ($0.30 \rightarrow 0.23$, $r = +0.97$) — spatial heterogeneity actively counter-acts lock rather than merely preventing it.

Gradient-bridge independence $|r| < 0.2$: bridge dynamics are autonomous.

Cold-start produces 2.5× greater τ variance than warm-start (Supplementary S5). Bach-structured temporal modulation increases colony τ variance over time, while flat modulation produces a declining trend (Supplementary S8).

4.4 Perturbation Response

When the nutrient source moves from center to corner, τ variance dips briefly (from 0.0046 to 0.0047) and recovers within three rounds. No unit locks. By the end of Phase C, τ variance reaches 0.0058 — higher than its pre-perturbation level (Figure 4.2).

4.5 The Anchor Boundary

The anchor ratio α controls the blend: $\alpha = 1.0$ is pure environment; $\alpha = 0.0$ is pure communication (no external anchor):

α	Spatial τ variance	Collective state
1.0–0.02	~0.004–0.006	Open — differentiated
0.00	0.0006	Marginal — near-zero

At $\alpha=0.02$ — 98% neighbor signals, 2% environment — the colony remains differentiated (10× the variance at $\alpha=0.00$). A minimal external anchor, however small, sustains differentiation. Pure self-reference approaches undifferentiation.

This is not a parameter sensitivity — it is a boundary condition. A self-referential information system that loses its external anchor approaches the same limit regardless of implementation: τ variance collapses, differentiation ceases. The $\alpha=0.00$ state is not destruction — it is an information-theoretic flatline: alive, but no longer registering surprise.

5. Forward: The Darwinian Arc

The experiments in this section are exploratory. Effect sizes are smaller than in Experiments 1 and 2, and quantitative precision is provisional. The qualitative direction is consistent and reported as such.

Core Finding. G0 is not required for differentiation — spatial heterogeneity already sustains it. G0 re-emerges under a different selection pressure: communication consolidation increases efficiency, and the bridge’s sensitivity to structural boundaries constitutes the earliest form of information-quality discrimination.

5.1 Survival: The Anchor Boundary

The anchor ratio α must remain above zero. A minimal external anchor sustains differentiation. This is the first step: survival. Before communication can optimize, before sensitivity to structure can emerge, the bridge must stay open.

5.2 Communication: Phagocytosis as Topological Trigger

Dennett (2017, Ch. 4) identifies endosymbiosis — the engulfment of one prokaryote by another, producing the first eukaryotic cell — as the paradigmatic free-floating rationale becoming embodied. We instantiate

this event computationally. A single phagocytosis is invisible at the global scale. At the local scale, the host's neighbors show a spatially dependent response: central prey (high concentration) produce a strong τ increase (Cohen's $d = 1.09$, 85% seeds positive); corner prey (low concentration) produce weak, directionally inconsistent responses ($d = 0.31$, 35% positive). Phagocytosis is not uniformly weak — it is spatially modulated. High-information positions produce strong, directionally consistent cascades.

5.3 Structural Sensitivity

Efficiency optimization (phagocytosis) clears the channel for more information. The evolutionary next step is to distinguish what that information means. We return to the Bach experiment to show that the bridge already possesses a primitive form of this capacity: sensitivity to structural boundaries.

The bridge detects closure reliably, across positions, across cadence types. It does not distinguish how closure is achieved (Supplementary S3). This is sensitivity in its earliest form: not discrimination of semantic content, but reliable response to the structural completion of a temporal unit. The bridge registers when something should close — and when that expectation is violated, it opens most strongly.

Efficiency and structural sensitivity are co-requisites. Together — the anchor boundary grounding existence, phagocytosis spatially modulating reorganization, the bridge's reliable detection of boundaries — they constitute the Forward. This is Dennett's trajectory from bacteria toward Bach, rendered as computational architecture: not complexity for its own sake, but the progressive unlocking of survival, communication, and structural sensitivity from a single self-referential primitive.

6. Discussion

6.1 The Bridge Is Alive

Paper I measured the Shannon-Gödel Bridge as a static constant. BGM reveals the bridge as a dynamic structure with breathing profiles, opening mechanisms, phase transitions, and discriminatory capacity. The bridge is not a pipe. It is a lung.

6.2 The Architectural Invariant: Quantitative Evidence

The central theoretical claim of this paper — that the bridge's dynamics are structurally invariant across temporal and spatial domains — rests on direct quantitative comparison. The SR-eff time series from Experiment 1 (Bach, six passes; detrended and standardized to zero mean, unit variance) and the τ variance time series from Experiment 2 (Bacteria, 50 rounds; same normalization) are aligned from experimental onset. The onset-aligned cross-correlation is $r = +0.964$. Both time series exhibit serial autocorrelation; the value is reported as a descriptive statistic rather than an inferential one. Supplementary S10 reports the full Kolmogorov-Smirnov comparison of the underlying distributions, which does not depend on temporal ordering and yields $D = 0.679$ ($p < 0.01$) — the distributions differ in shape, as expected for different input modalities, but the temporal breathing structure is shared. The bridge's breathing profile is not a qualitative analogy. It is a quantitatively measurable structural isomorphism.

6.3 The Evolutionary Arc

- **Bach:** G0 amplifies novelty on first exposure — Pass 1 is vivid, Passes 2–6 consolidate. Under simple input, the bridge closes via first-order phase transition.

- **Bacteria:** No G0 needed. Spatial heterogeneity keeps the bridge open and differentiated without central coordination.
- **Forward:** Novelty detection is G0's genuine contribution. Sustained differentiation, spatial anti-lock, and structural sensitivity operate without it.

6.3 Information Needs Boundaries

Shannon (1948) established that information requires distinguishability. The gap insertion experiment recovers this principle quantitatively: boundary density must exceed ~ 1 per 16 steps. Below this threshold, the co-occurrence window saturates and differentiation collapses. Above it, structure emerges. A control with interleaved fugue variants confirms that cadence type is invisible to the bridge after position randomization — the bridge detects endings, not harmonic quality.

Information needs boundaries. Differentiation begins in gaps.

6.4 The GEME as a Minimal Cognitive Atom: Why Less Is More

The small unit counts in this study — three units hearing Bach, sixty-four occupying a spatial grid — are not a concession to computational limits. They are the method. In the same way that the Ising model isolates phase transitions with a handful of spins, and Conway's Game of Life extracts universal computation from a three-cell neighborhood, the GEME architecture isolates the minimal conditions under which a self-referential information economy breathes, differentiates, and detects closure. A self-referential frame economy, constrained by capacity and driven by adaptive prediction, exhibits phase transitions, hysteresis, differentiation, and structural sensitivity at these scales precisely because the information-gravitational forces that produce these phenomena are visible only when the system is small enough for individual bridge dynamics to matter.

Large populations would wash out the phase transitions through statistical averaging, in the same way that individual synaptic events are invisible in the BOLD signal. The principles identified here — the anchor boundary, the g_0 interval optimum, the gap-density threshold — carry inherent scale bounds. Stable self-referential structures have natural upper limits: Miller's 7 ± 2 , Dunbar's 150, the cortex's 6-layer hierarchy. These are not coincidences. They are the same information-gravitational ceiling seen in different biological substrates.

A theoretical prediction follows: large-scale GEME populations will not scale uniformly. They will spontaneously modularize into loosely coupled bridge clusters, each operating at the g_0 interval optimum, each maintaining its own anchor ratio $\alpha > 0$. The path from the cognitive atom to the cognitive collective is not aggregation — it is externalization. Paper III takes this step.

6.5 τ Repositioned, Not Eliminated

Paper I presented $\tau = 0.60$ as a core constant. Under temporal input, τ reveals itself as internally generated. The dynamic coupling $\tau_i = \tau_0 \times (MI_0 / MI_i)$ retains τ_0 as an anchor — τ is repositioned from a global constant to a unit-level variable whose value is determined by prediction history. This mirrors Dennett's internalization.

6.6 G0's Bridge: A Novelty Amplifier

G0's bridge never locks across the full τ range (0.2–3.0). Its input — aggregated L6 from differentiated units — is inherently diverse, and its stability comes from this diversity, not from τ elevation. But G0's functional role is narrower than the architecture suggests. It amplifies novelty on first exposure — Pass 1 vivid, Passes

2–6 consolidating — and recedes as the system learns. It is not the engine of differentiation (Experiment 2 achieves that without it), not the source of ASA (time-scale separation under discrete encoding does not require it), and not the detector of closure (the bridge registers that with or without G0). G0 is a novelty amplifier — and this is enough. In an information economy where discovery is transient and maintenance is cheap, an amplifier that works when it matters is worth keeping.

6.7 The Parameter That Survived Time

A systematic sweep of the G0 layer’s parameter space (§2.2) finds that only $g0_interval$ — the temporal decoupling between G0 and its units — genuinely modulates differentiation. Frame capacity is neutral. Modulation asymmetry is overridden by dynamic τ . Decay is provided by the sliding window. Signatures amplify differentiation $2.2\times$ but are a structural condition, not a parameter to be tuned.

The G0 layer’s Pareto optimum is $g0_interval = 4$, producing τ spread 49% above the default. This is the gain control on G0’s novelty amplifier. When GI is low, G0 samples frequently and novelty is averaged away. When GI is high, G0 samples sparsely and each observation carries accumulated surprise. $GI=4$ is the setting at which the amplifier is strongest without saturating — the point at which a single novelty, arriving once, produces the largest downstream differentiation.

6.8 SR-eff as a Portable Metric

$SR\text{-}eff = I(\Phi; X)/\tau$ is not specific to GEME. Brief dialogues with predictive processing (Friston, 2010), integrated information theory (Oizumi et al., 2014), and continual learning (McCloskey & Cohen, 1989; Parisi et al., 2019) suggest SR-eff as an orthogonal metric for any self-organizing information system. These are starting points, not conclusions.

6.9 Biological Correspondence

Jayakumar et al. (2021) observed spontaneous QS segregation without central coordination — structurally parallel to our spatial τ differentiation. The maintenance of diversity through spatial heterogeneity alone, without any selection pressure, echoes Kimura’s (1968) neutral theory: diversity sustained by the structure of the environment rather than by adaptive advantage. Whether this reflects a deeper information-theoretic commonality is open.

6.10 Scope, Scalability, and Future Directions

Scope. This study deliberately operates at minimal scale — three units and sixty-four — to isolate the bridge’s core dynamics under controlled conditions. The small numbers are the method, not a concession: phase transitions, hysteresis, and spatial differentiation are architecturally visible only when individual bridge dynamics are not washed out by statistical averaging. The technique is generative minimalism: identify the simplest kernel that produces the phenomenon, then reason outward.

Scalability. A theoretical prediction follows from the information-gravitational principles identified here. Large-scale GEME populations should not scale uniformly. As N grows, the anchor boundary ($\alpha > 0$) and the $g0_interval$ optimum ($gi=4$) impose natural modularization: the population should spontaneously partition into loosely coupled bridge clusters, each maintaining its own α ratio and temporal decoupling, forming a cognitive niche structure. This prediction — that scale produces modularity, not homogeneity — provides a direct bridge from the cognitive atom (studied here) to the cognitive collective (studied in Paper III, under the framework of externalization).

Future directions. Testing this prediction requires populations of hundreds to thousands of GEME units, heterogeneous input fields, and the measurement of emergent modular structure — cluster boundaries, inter-cluster τ correlations, and the spontaneous emergence of g0-like coordination hubs. Additional directions include: dynamic-dimension vector encoding; signal accumulation and delayed communication; heterogeneous unit types; and the recursive observation hierarchy (4→2→1) explored in Supplementary materials.

Current limitations. Uniform 27-dimensional encoding; instantaneous communication; the Bach experiment's fixed 'input' signatures (a deliberate negative control). All experiments use the conservative default $g0_interval = 1$; all results strengthen at the discovered Pareto optimum ($g_i=4$, +49%). Core conclusions are qualitatively robust across $\alpha \in [0.02, 1.0]$ (Supplementary S5).

EE Hook. The G0 layer operates at a single level of observation. The recursive structure — a bridge monitoring bridges at arbitrary depth — remains unexplored. Whether higher-order self-referential channels maintain natural stability with increasing depth, and whether modularization emerges spontaneously at scale, are open questions for Paper III (External Engine).

The G0 layer operates at a single level of observation. The recursive structure — a bridge monitoring bridges at arbitrary depth — remains unexplored. Whether higher-order self-referential channels maintain natural stability with increasing depth is an open question for Paper III (External Engine). The signature — identified here as the observer's identity basis — is the starting point of that recursion.

Coda

Bach's *Crab Canon* from the *Musical Offering* is a musical palindrome — one voice forward, one backward, simultaneous. It has no beginning, no end. Two GEME units, cold-started, receive its blended two-voice stream: they converge to identical τ . The canon's voices occupy the same frequency space. Without temporal gaps, the bridge detects one stream (τ spread = 0). With discrete gaps inserted — every fourth step a silence — τ differentiation emerges. The bridge responds to two streams.

The crab canon is a circle. The bridge on it breathes forever. It closes, it opens again. There is no final cadence.

Supplementary Information

S1. GEME core engine — identical to Paper I submission (MD5 verified).

S2. Input encoding and temporal structure. Score vs. MIDI comparison. Gap insertion experiment: discrete boundaries enable τ differentiation. Connection to Bregman's temporal proximity principle.

S3. τ differentiation controls and voice tracking. White noise, single tone, excess-unit controls. Voice count experiment (1–5 voices). τ count \geq voice count. Position randomization control for cadence type.

S4. Lock verification: phase transition (step 23, first-order) and hysteresis data.

S5. Bacteria grid: α -sweep, cold/warm start, perturbation trajectories, position sensitivity, neighbor coupling. G0-layer parameter sweeps: capacity, modulation rate, $g0_interval$, G0 induction threshold.

S6. Phagocytosis: multi-seed local cascade analysis. Spatial modulation by prey concentration.

S7. Phase portrait: static τ baseline.

S8. Bach rhythm → Bacteria aliveness.

S9. G0 bridge: τ scan (0.2–3.0) and natural stability data.

S10. Cross-condition isomorphism: KS test `conf_threshold` distributions.

All experiments reproducible with Python 3.8+ stdlib, zero external dependencies. Core engine `game.py` MD5: 0025C508BDBDB386E9A5EB72081995B7 (identical to Paper I). Key scripts: `bgm_bach_pipeline.py`, `bgm_bacteria.py`, `bgm_phagocytosis_v2.py`, `bgm_wtc_full.py`, `bgm_g0_pareto.py`, `bgm_voice_validation.py`, `bgm_cadence_randomized.py`, `bgm_tau_multiseed.py`, `bgm_structural_analysis.py`, `bgm_g0_tau_scan.py`. Figures: `fig_input_vs_tau.png` (4.1), `fig_perturbation.png` (4.2), `fig_mobius.png` (Coda).

References

- Barraud, N., et al. (2009). *J. Bacteriol.*, 191(23), 7333-7342.
- Bregman, A. S. (1990). *Auditory Scene Analysis*. MIT Press.
- Brown, G. J., & Cooke, M. (1994). Computational auditory scene analysis. *Computer Speech and Language*, 8(4), 297-336.
- Chaitin, G. J. (1975). A theory of program size formally identical to information theory. *Journal of the ACM*, 22(3), 329-340.
- Dennett, D. C. (2017). *From Bacteria to Bach and Back*. W. W. Norton.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127-138.
- Hofstadter, D. R. (1979). *Gödel, Escher, Bach*. Basic Books.
- Jayakumar, P., et al. (2021). *bioRxiv*, 2021.03.22.436499.
- Kimura, M. (1968). Evolutionary rate at the molecular level. *Nature*, 217, 624-626.
- Liu, X., et al. (2023). AudioSep: Separate Anything You Describe. *arXiv:2308.05037*.
- McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks. *Psychology of Learning and Motivation*, 24, 109-165.
- Meyer, L. B. (1956). *Emotion and Meaning in Music*. University of Chicago Press.
- Oizumi, M., et al. (2014). From the phenomenology to the mechanisms of consciousness. *PLoS Computational Biology*, 10(5), e1003588.
- Parisi, G. I., et al. (2019). Continual lifelong learning with neural networks: A review. *Neural Networks*, 113, 54-71.
- Popat, R., et al. (2012). *Proc. R. Soc. B*, 279, 4765-4771.
- Scribner, M. R., et al. (2022). *J. Bacteriol.*, 204(2), e00444-21.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379-423.
- Wang, D. L., & Brown, G. J. (2006). *Computational Auditory Scene Analysis*. IEEE Press / Wiley.
- Williams, H., Nicolson, A., & Green, T. (2024). Applying Marr’s framework to auditory scene analysis. *Frontiers in Neuroscience*, 18, 1352247.