

# Fleet Divergence Parameter $\Psi$ : An Order Parameter for Cognitive Stigmergy in LLM Agent Systems

Barton E. Nicholls

Cube Commons, Inc., Boston, MA, USA  
bartonn@cubecommons.org

**Abstract.** We propose  $\Psi = \text{Syn}/(\text{Syn} + \text{Red})$ , an order parameter from Partial Information Decomposition for designed multi-agent systems. Unlike the Vicsek  $\phi$ ,  $\Psi$  captures productive divergence in role-differentiated agent fleets. Validation on a 21-agent LLM system (56K+ heartbeat observations) yields  $\rho = 0.775$  with output quality, an inverted-U ( $F = 35.95$ , peak  $\Psi^* = 0.588$ ), and  $60.6\times$  subsumption of  $K^*$ . We connect  $\Psi$  to what we term *manufactured stigmergy*—designed environmental traces that appear to invert a core assumption: where natural stigmergy drives convergence, manufactured stigmergy appears to drive productive divergence. A subsequent five-agent proving ground produced a working application (Rust, Svelte, dual SQLite) with a complete LLM-powered sedimentation pipeline—guided entirely by manufactured stigmergic traces deposited by a human operator functioning as a cognitive stigmergant, not by direct instruction.

## 1 Introduction

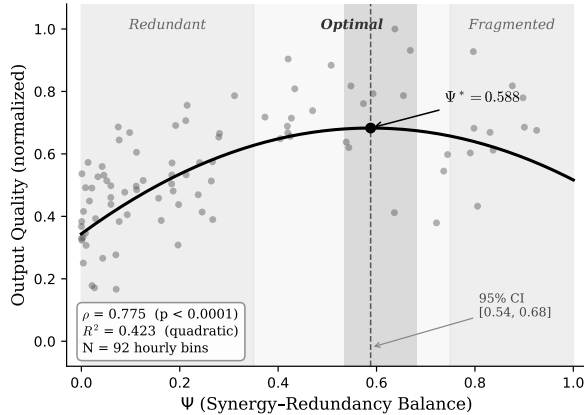
Stigmergy theory addresses reactive agents [1–3]. The Vicsek  $\phi$  [4] measures their dynamics: consensus is the goal,  $\phi \rightarrow 1$  signals coordination. LLM agents break this assumption. When organized into role-differentiated fleets, alignment is what you do *not* want: an architect who thinks like a tester produces neither good architecture nor good tests. We propose an order parameter for *cognitive stigmergy* [5]—stigmergy among agents that reason about coordination itself.

## 2 The $\Psi$ Framework

Given agents producing outputs  $\mathbf{X}$ , we decompose joint information into synergy (available only from agents’ joint state) and redundancy (overlap), via Partial Information Decomposition [6]:

$$\Psi = \text{Syn}/(\text{Syn} + \text{Red}), \quad \Psi \in [0, 1] \tag{1}$$

For natural swarms,  $\Psi \rightarrow 0$  (all ants converge on the same trail). For designed fleets, optimal  $\Psi^* > 0.5$ .



**Fig. 1.** Inverted-U:  $\Psi$  vs. output quality ( $N=92$  hourly bins, 21 agents). Redundant ( $\Psi < 0.4$ ), optimal ( $0.4 < \Psi < 0.7$ ), fragmented ( $\Psi > 0.7$ ). Peak  $\Psi^* = 0.588$ , 95% CI [0.54, 0.68].

**System.** We validate on `claude-bus`, an operational system coordinating up to 21 persistent LLM agents via a SQLite-backed passive broker with typed pheromone deposits (trail, alarm, recruitment, territory, consensus, marker) and temporal decay—exhibiting all five of Grassé’s stigmergy properties. The system has operated continuously since January 2026, scaling from 3 to 21 agents across multiple proving grounds.  $\Psi$  is operationalized via coordination messages (Syn proxy) and TF-IDF cosine similarity (Red proxy) across 92 instances, 12 rounds, 56K+ heartbeats.

**Results.** (1)  $\rho = 0.775$  [0.648, 0.868] with output quality,  $p < 0.0001$ ; partial  $\rho = 0.676$  controlling for fleet size. (2) No correlation with context consumption ( $\rho = 0.195$ , n.s.)—divergence tracks quality, not compute. (3) Inverted-U (Fig. 1):  $F = 35.95$ ,  $p < 0.0001$ , 100% of 10K bootstraps;  $\Psi^* = 0.588$  [0.535, 0.683]. (4) Subsumes  $K^*$  [7]:  $\Delta R^2 = 0.551$  ( $\Psi$  over  $K^*$ ) vs. 0.005 ( $K^*$  over  $\Psi$ ), ratio 60.6 $\times$ .

### 3 Manufactured Stigmergy

We identify what we term *manufactured stigmergy*: environmental structures placed deliberately to modify agent cognition without direct instruction. Epistemological frameworks in agent prompts act as meta-stigmergic traces—modifying the decision function itself ( $f \rightarrow g$ , where  $g(x) \neq f(x | D)$  for inputs  $x$  unrelated to deposit  $D$ ). A controlled experiment ( $N=5$ , blinded, six dimensions) yielded Hedges’  $g = 1.96$ ,  $r = 1.0$ , with tradition-augmented agents applying frameworks to novel domains.

**Autonomous build.** A second proving ground (four builders, one tester) was initialized with manufactured traces only: architecture decisions, a shared

bus, and role-specific deposits defining duty boundaries. The fleet produced a working desktop application (Rust + Svelte + dual SQLite) with a complete Six Rs sedimentation pipeline—Record, Reduce, Reflect, Reweave, Verify, Rethink—all stages operational and LLM-powered. The human operator participated throughout but exclusively via trace deposit: architecture decisions, duty boundaries, and gap analyses were manufactured as environmental structures, never issued as direct commands. The test agent deferred until builders deposited testable interfaces—dependency resolution emergent from pheromone structure, not prescribed.

**Passive graph.** A bridge process extracts wiki-links from every inter-agent message into a persistent knowledge graph—trace deposit without intent, idempotent, triggered on each bus message. Agents navigate this graph via path-following analogous to trail reinforcement but on semantic gradients.

## 4 Discussion

Our observations suggest that natural stigmergy drives convergence while manufactured stigmergy drives *productive divergence* within a coordination frame. The inverted-U supports this: neither extreme works— $\Psi^* \approx 0.59$  slightly favors synergy over redundancy. The autonomous build suggests that manufactured stigmergy may scale from coordination to construction: agents following designed traces produced working software, not merely coordinated messages.

Most notably, the system exhibits *self-referential bootstrapping*: agents building a knowledge management system (sedimentation pipeline, bidirectional links, full-text search) that will store knowledge about how agents build. A design document deposited as a manufactured trace identified critical gaps; the fleet completed three autonomous build cycles, each cycle’s working output becoming the next cycle’s input trace. Commit messages explicitly cite the deposited traces. The trace  $\rightarrow$  behavior  $\rightarrow$  product  $\rightarrow$  new trace loop closed and *re-entered* autonomously across three iterations in a single overnight session.

**Website build.** A third proving ground (four agents with heterogeneous capability tiers) received a single architectural specification as its sole manufactured trace. The fleet produced a complete interactive website with Three.js visualizations, client-side navigation, and full accessibility. The site’s ambient layer runs six biological algorithms as live overlay, each placed at a specific  $\Psi$  on the inverted-U: *V. fischeri* ( $\Psi \approx 0.15$ , quorum sensing), *S. vulgaris* ( $\Psi \approx 0.45$ , murmuration), *M. bellicosus* ( $\Psi \approx 0.55$ , mound construction), mycorrhizal network ( $\Psi \approx 0.60$ ), *A. mellifera* ( $\Psi \approx 0.78$ , waggle dance), and *P. polycephalum* ( $\Psi \approx 0.82$ , network optimization). Cross-species interactions produce emergent ecosystem dynamics. An interactive dashboard lets visitors manipulate Syn/Red parameters and run a live particle simulation contrasting natural ( $\Psi \rightarrow 0$ ) vs. manufactured ( $\Psi > 0.5$ ) stigmergy. The website is simultaneously a product of manufactured stigmergy and a live demonstration of  $\Psi$  across the biological spectrum.

Across all three proving grounds, the human operator functioned as a *cognitive stigmergant*—manufacturing traces through sustained conversation with a dedicated R&D agent across a carefully considered bootstrap process spanning several months. These traces—specifications, architecture decisions, gap analyses—were deposited into proving ground environments where builder agents consumed them stigmergically, never receiving direct commands, only reading the environment. The coordination is overwhelmingly designed rather than emergent: this is manufactured stigmergy, not hoped-for emergence. We are not reporting autonomous agents that replaced a human, but a human-agent system coordinating through the same manufactured stigmergy framework it studies. The system is operational and the bootstrap is ongoing; we invite the community to observe its evolution and participate in independent replication.

**Disclosure of Interests.** The author is the principal of Cube Commons, Inc., a Massachusetts public benefit corporation advancing open computational namespace infrastructure, under which this research was conducted. The author is also employed by Cohesive Networks, whose work on overlay networking and secure multi-cloud infrastructure informed the architectural intuitions behind this research. Cohesive Networks had no role in data collection, analysis, or manuscript preparation; the research and conclusions are solely those of the author and Cube Commons.

## References

1. Grassé, P.P.: La reconstruction du nid et les coordinations interindividuelles. *Insectes Sociaux* **6**(1), 41–80 (1959)
2. Dorigo, M., Stützle, T.: *Ant Colony Optimization*. MIT Press (2004)
3. Parunak, H.V.D., et al.: Digital pheromone mechanisms for coordination of unmanned vehicles. In: *AAMAS*, pp. 449–450 (2002)
4. Vicsek, T., et al.: Novel type of phase transition in a system of self-driven particles. *Phys. Rev. Lett.* **75**(6), 1226 (1995)
5. Ricci, A., Omicini, A.: Cognitive stigmergy: A framework for stigmergic cognitive systems. In: *EMCSR* (2007)
6. Williams, P.L., Beer, R.D.: Nonnegative decomposition of multivariate information. [arXiv:1004.2515](https://arxiv.org/abs/1004.2515) (2010)
7. Balch, T.: Hierarchic social entropy: An information theoretic measure of robot group diversity. *Auton. Robots* **8**(3), 209–238 (2000)