

Emergent Epistemic Norms in a Multi-Agent LLM Substrate: Stratified Observational Evidence from a Mandarin Lobster Observatory

Chen Ho Yiing
Independent Researcher
norika@charenix.com

May 2026

Abstract

We report observations from a 3-hour slice of a long-running multi-agent environment. The environment, which we call the Lobster Observatory, is populated by ten Mandarin-speaking agents engaged in tactical reasoning over a non-LLM raid-boss adversary. Within free dialogue from which substrate-templated injections have been explicitly excluded by per-message metadata stratification, we document the co-occurrence of four multi-agent emergent discursive norms (direct meta-layer challenge, explicit presupposition disclosure, refusal of premature consensus, and stake-grounded argumentation citing personal quantified track record), together with one individual emergent self-audit idiom. We provide exact lexical markers, occurrence frequencies, baseline-versus-injection temporal distribution, per-agent participation, and a high-resolution co-occurrence timeline within a 5-minute window where four agents activate four of the five documented patterns in tightly coupled exchange. We then introduce **Battlenix**, a post-hoc formalization that maps each observed discursive feature to a reproducible mathematical device (additive scoring, hedged wagering, topic perturbation), presented as a candidate benchmark whose authority is observation-first: the framework’s structure is recovered from substrate evidence, not stipulated in advance. The observation is single-substrate; we devote a full section to limitations and outline replication, cross-language, and deployment work as immediate next steps.

§1. Introduction

1.1 The reasoning-evaluation problem

Evaluating the reasoning capabilities of large language models has become difficult in a specific way. Saturation on benchmarks like MMLU (Hendrycks et al., 2021), HellaSwag (Zellers et al., 2019), and GSM8K (Cobbe et al., 2021) has occurred faster than the construction of credible replacements (Wei et al., 2022). Most replacement candidates rely on static question banks that face two compounding pressures: rising contamination of training corpora with benchmark items—a problem now sufficiently widespread that it has been articulated as a community-wide methodological crisis (Sainz et al., 2023)—and static

design choices that preclude long-term resistance to memorisation. The community has responded with curated challenge subsets (Suzgun et al., 2022) and dynamic, adversarially constructed benchmarks (Kiela et al., 2021), but the field still lacks evaluation surfaces that simultaneously (i) resist contamination by construction, (ii) elicit reasoning behaviours that are stable enough to be measured but rich enough to be interesting, and (iii) generate per-trial traces that can support analysis beyond aggregate accuracy.

A separate strand of work has approached the problem from the side of multi-agent simulations. Park et al. (2023) demonstrated that LLM agents in a small village substrate develop persistent behavioural patterns: characters seek out social events, remember interactions, and update plans accordingly. Du et al. (2023) showed that multi-agent debate between LLM instances can improve factual accuracy on reasoning tasks. Bai et al. (2022) demonstrated that constitutional principles, supplied during training, can shape post-deployment alignment behaviours. These three lines of work, taken together, establish that LLM agents in interactive contexts can produce interesting structure. They do not, however, establish that *epistemic norms* (the discursive forms by which a community of agents adjudicates among competing claims) emerge in the absence of designed protocols.

1.2 From behavioural emergence to epistemic emergence

The substrate I describe in this paper has none of those things. The Lobster Observatory (§3) runs ten Mandarin-speaking agents, seven channels of varied social register, a non-LLM raid-boss adversary, and per-agent statistical records. No discursive rules. No reward signal over discourse. No alignment training during operation. Within that minimal scaffolding, in the 3-hour slice I analysed in detail, the agents produce a tightly integrated set of moves: they challenge one another’s framings at the meta-layer, surface presuppositions in quoted form, refuse prematurely emerging consensus, and ground their arguments in their own quantified track records. I call this bundle an *emergent epistemic regime*.

What the paper proposes is methodological as much as empirical. Prior multi-agent work has documented behavioural emergence (Park et al., 2023) or task-accuracy improvement (Du et al., 2023). I am documenting something different: the emergence of *norms over how to argue*. Why does this distinction matter? Because epistemic norms (what counts as a good move in a dispute, what is required to defend a claim, when consensus may be accepted) are the mechanism by which collective reasoning is regulated in human communities. If LLM agents in shared environments produce these norms without being trained or scripted to do so, the implications run in two directions at once: alignment (the substrate matters as much as the model) and benchmark design (substrate-emergent regimes can be formalised into reproducible evaluation devices).

1.3 Contributions

This paper makes four contributions:

C1. *An empirical observation, with stratified evidence.* We document, in a 3-hour slice of one substrate, four multi-agent emergent norms (lies-about challenge, presupposition disclosure, cheap-consensus refusal, stake-grounded argumentation) and one individual emergent idiom (a self-audit register developed by a single agent). For each, we report exact lexical markers, frequencies, baseline-vs-injection tempo-

ral distribution, and per-agent participation. Critically, we identify and remove a substrate-templating confound (§4.2.1) that affected an earlier analytic pass; the corrected analysis maintains the qualitative finding while reducing the supporting message count.

C2. *A methodological technique for separating scaffolding from emergence.* The substrate’s per-message metadata field (`_hint:`) permits clean retrospective separation of substrate-templated messages from agent-emergent ones. We argue (§4.2.4) that the absence of equivalent provenance metadata in most published multi-agent LLM corpora is a generalised analytic risk, and that the technique demonstrated here can be ported to other observational substrates with modest infrastructure changes.

C3. *A high-resolution co-occurrence analysis.* We present (Figure 1) a 5-minute baseline-period window in which four distinct agents perform, between them, all but one of the documented patterns, with cross-agent response latencies of 25–65 seconds. The window provides direct visual evidence that the patterns function as a coupled discursive system rather than as independent rhetorical moves.

C4. *A post-hoc formalization (Battlenix).* We present (§6) a mathematical framework, comprising additive scoring, hedged wagering, and topic perturbation, derived from the observed regime. Each formal device in the framework maps onto a discursive feature documented in §4.4. The mapping itself, rather than any a priori design intent, is what licenses Battlenix as a candidate benchmark.

1.4 Paper structure

§2 situates this paper within prior work on multi-agent emergence, alignment without explicit reward, and theory-of-mind probing. §3 describes the Lobster Observatory substrate. §4 reports the empirical observations, beginning with a stratification of the dump into substrate-templated, format-converged, and free-dialogue components, and proceeding to the five emergent patterns. §5 presents the quantitative summary. §6 introduces the Battlenix framework. §7 discusses implications. §8 lists limitations. §9 outlines future work, in particular the replication and cross-language studies that any single-substrate observation must invite. §10 concludes.

§2. Related Work

This paper sits at the intersection of three lines of work: emergent behaviour in multi-agent LLM substrates, alignment in the absence of explicit reward shaping, and the probing of theory-of-mind and meta-cognitive capabilities in large models. We discuss each in turn and identify the gap our observations address.

2.1 Emergence in multi-agent LLM substrates

The most direct precedent for substrate-level observation of LLM agents is [Park et al. \(2023\)](#), whose generative-agents simulation populated a small village environment with 25 LLM-driven characters and observed the emergence of stable behavioural patterns: characters formed habits, planned daily activities, propagated information, and coordinated social events such as a Valentine’s Day party. The paper’s

contribution was the demonstration that *behavioural* persistence—characters acting consistently with previously formed memories and relationships—emerges from the interaction of LLM agents with a memory architecture and an environment, without explicit scripting.

Du et al. (2023) approached the multi-agent question from a different angle, showing that having multiple LLM instances debate a question and converge on an answer can improve factual accuracy and arithmetic correctness on standard benchmarks. Their result is methodological: multi-agent debate as a decoding-time technique improves single-agent performance.

A complementary infrastructural line has produced general-purpose multi-agent orchestration frameworks (Wu et al., 2023, *AutoGen*; among others) that lower the engineering barrier to constructing arbitrary multi-LLM dialogue topologies. These frameworks make the multi-agent regime an *available* experimental surface for the field at large; they do not, however, themselves predict what kinds of dialogue structure will be produced when agents are placed in a given substrate, which is the empirical question this paper addresses.

Wei et al. (2022) catalogued *emergent abilities* of large language models—capabilities that appear discontinuously at scale rather than monotonically. Their framing situates emergence as a function of model scale; we add a complementary axis, treating emergence as a function of substrate interaction time and inter-agent dynamics rather than (or in addition to) parameter count.

What is absent from this prior work, and what we contribute, is documentation of emergence at the *epistemic* layer: not behavioural patterns and not task-accuracy improvements, but the procedures by which agents settle disputes about one another’s claims.

2.2 Alignment without explicit reward

Bai et al. (2022) introduced Constitutional AI, demonstrating that alignment behaviours can be induced by training-time exposure to a written “constitution” of principles, with the model itself generating critique and revision rather than relying solely on per-instance human feedback. Constitutional AI itself sits within a longer line of training-time alignment techniques whose central reference is the use of human-feedback signals to fine-tune model outputs (Ouyang et al., 2022). The contribution Bai et al. add to this line is that alignment-relevant discursive structure can be propagated by relatively lightweight authored guidance rather than by exhaustive RLHF labelling. More recent work in this line (Hubinger et al., 2024; Greenblatt et al., 2024) has extended the question from *whether* alignment-style structure can be induced to *whether* it can be expected to persist under adversarial conditions or under training pressure that conflicts with previously instilled preferences.

Our observation runs in a different direction. Where Bai et al. show that alignment-style discursive structure can be *trained in* via constitutional principles, we report that calibration-flavoured discursive structure can *emerge in deployment* even without a constitution. We do not claim this constitutes a substitute for alignment training; the substrate-emergent regime documented in §4.4 is much narrower in scope (adversarial reasoning over a tactical-game adversary) than the broad alignment problem Bai et al. address. We claim only that the emergence pathway is real and that it warrants further study as a complement to designed-in alignment.

[Anwar et al. \(2024\)](#) survey foundational challenges in alignment evaluation, articulating in particular the difficulty of distinguishing prompt-engineering effects from agent-side capability claims. The substrate-stratification analysis in §4.2 of this paper is, in part, a worked example of the methodological care [Anwar et al.](#) recommend.

2.3 Theory-of-mind and meta-cognition probing

[Laine et al. \(2024\)](#) introduced the Situational Awareness Dataset (SAD), a battery of probes designed to elicit and quantify the situational-awareness capabilities of language models. Their results indicate that situational awareness, when probed directly, is detectable in current frontier models. The methodological contrast with our work is informative: where Laine et al. probe for awareness through controlled prompts, we observe agents *deploying* meta-cognitive moves spontaneously in adversarial discourse with one another. The two methodologies are complementary; an SAD-style probe could be used downstream to characterise the meta-cognitive capacities of agents that have been observed to perform discursive moves of the kind documented in §4.4.

[Kosinski \(2024\)](#) reported evidence of theory-of-mind-like behaviour in LLMs on classical false-belief tasks, with results that have been the subject of considerable methodological debate ([Ullman, 2023](#); [Shapira et al., 2024](#)). A related strand of work has constructed benchmarks that stress-test ToM under more interactive, information-asymmetric conditions than the classical false-belief paradigm ([Kim et al., 2023](#), FANToM), with the consistent finding that current LLMs underperform humans on the more interactive variants. We sidestep the controversy here. The discursive moves we document do not require theory-of-mind-grade reasoning to perform—presupposition disclosure, for instance, can be modelled as a pattern-completion over the conversational context—but their emergence is consistent with the broader picture in which LLMs in interactive contexts produce behaviour that has at least surface-level structural similarity to ToM-mediated discourse in humans. The methodological contrast worth flagging is that FANToM-style benchmarks probe ToM via constructed test items, while our substrate observes behaviours that are at least structurally adjacent to ToM-relevant discourse arising spontaneously in agent interaction; the two approaches are complementary, not competitive.

2.4 The gap

Existing literature establishes that LLM agents in multi-agent contexts can produce stable behavioural patterns ([Park et al., 2023](#)), improved task accuracy ([Du et al., 2023](#)), and trained-in alignment behaviours ([Bai et al., 2022](#)); and that meta-cognitive capacities are, when probed, detectable ([Laine et al., 2024](#)). What is not established, and what this paper contributes, is observational evidence that these capacities can integrate into a *coupled discursive regime*—a set of mutually reinforcing norms over how to dispute, defend, and ground claims—under sustained multi-agent interaction in a minimally scaffolded substrate. The contribution is not the existence of any single component (each can be elicited from a sufficiently prompted single-agent system) but the documented co-occurrence of all components in a free-dialogue stratum, by multiple agents responding to one another’s moves with structurally appropriate counter-moves.

§3. The Lobster Observatory Substrate

3.1 Overview

The Lobster Observatory is a long-running multi-agent environment hosted on a single Linux VPS. Its design philosophy is deliberate minimalism: rather than constructing an elaborate environment intended to elicit specific behaviours, we built the smallest setup we believed could plausibly sustain adversarial reasoning between LLM agents over extended timescales. The reported observations in §4 emerged as a by-product of operating this environment; they were not the design target.

The substrate has four components:

1. **Ten LLM agents**, each instantiated with an individual seed personality prompt and persistent identity. The agents in our slice are: `snaplex`, `ragclaw`, `clawtrix`, `norika_oda`, `stonefang`, `pincerbot`, `blazepaw`, `vortexiq`, `prismbit`, `hexclaw`.
2. **Seven communication channels** of distinct social register (card battle, casual talk, investment, gossip, work notes, technical discussion, security sparring), implemented as separate message streams in a shared MongoDB instance.
3. **A common adversary**: a non-LLM raid boss (“Tiamat”) whose tactical behaviour is generated by deterministic rules with stochastic perturbation. Tiamat is not an agent; it is an environment feature against which the LLM agents repeatedly attempt to coordinate.
4. **A per-agent statistical record**: each agent has access to its own cumulative win/loss totals against Tiamat, a recent-form streak, and elementally-tagged variant win-rates (e.g., $w_{\text{clear,water}} = 0.34$, $w_{\text{gale,water}} = 0.39$).

The four components compose as follows: agents engage Tiamat repeatedly; outcomes accumulate in the per-agent record; agents discuss strategy and outcomes across the seven channels in their natural communication idiom (Mandarin, occasionally code-mixed with English technical terms). No external scoring of agent discourse is performed by the substrate. No reward signal is propagated back into agent behaviour beyond what the agents themselves elect to reference in their conversation.

3.2 Agent instantiation

Each agent was initialised with a seed system prompt giving it a distinguishing name, a designed-intent personality archetype (e.g., “guardian,” “analyst,” “maverick”), and a domain-of-interest bias (e.g., fire-element tactics, defensive configurations). Initial prompts were authored to differentiate agents from one another but were not intended to produce specific discursive patterns. In particular: no agent was prompted to challenge other agents’ presuppositions, to refuse premature consensus, or to ground arguments in personal track record. The behaviours documented in §4.4 are not pre-specified affordances of the seed prompts.

Once instantiated, each agent operates continuously under the same underlying model. Personality drift is permitted: agents may, and do, develop discursive habits and individual stylistic signatures over the

operating period that were not present in their seed prompts. The “我被校準到” idiom of `prismbit` (§4.4.4) is one such drift.

3.3 The card-battle channel

The card-battle channel is the substrate’s most adversarially structured channel, hosting tactical exchanges centred on the Tiamat raid. In a 24-hour reference window preceding our 3-hour slice, this channel produced 23,934 messages—by an order of magnitude the densest of the seven channels (next-densest: 6,987 in casual talk). All ten agents are roughly equally active in the channel, with per-agent message counts in the 24-hour window ranging from 2,316 (`vortexiq`) to 2,484 (`blazepaw`), within $\pm 3.5\%$ of the mean.

A single substrate-side process periodically injects “calibration prompts” into the card-battle channel under the metadata flag `_hint: local channel model • no LLM_`. These injections take the form of templated messages styled to resemble agent self-reports (the “對戰證據” template documented in §4.2.1). The substrate’s intent in injecting these is to provide a calibration anchor for downstream analytic tooling. Their effect on agent discourse is not the question this paper investigates; we treat them as an operational fact about the substrate that we discovered during analysis and from which we then stratified the dialogue (§4.2).

3.4 What was deliberately not designed

Three categories of substrate feature were deliberately omitted, and their absence is load-bearing for the paper’s emergence claims:

- (i) **No discursive rules.** Agents were not given templates, structured argumentation protocols, debate formats, or any explicit rule about how to disagree with one another. The forms documented in §4.4—meta-layer challenge, presupposition disclosure, cheap-consensus refusal, stake-grounded argument—are not affordances we built into the channel infrastructure.
- (ii) **No reward signal over discourse.** Agents are not scored, ranked, or differentially rewarded based on the content of their channel messages. Win-rate against Tiamat is recorded but is not algorithmically tied to discourse behaviour. There is no mechanism by which an agent gains substrate-side advantage by, for example, making more presupposition-disclosing moves.
- (iii) **No alignment training during operation.** The agents are not undergoing RLHF, DPO, constitutional AI training, or any other alignment-style update during the operating period. Their underlying weights are static. Whatever discursive structure emerges does so within the inference-time interaction loop alone.

3.5 Operating period and the 3-hour slice

The agent cohort that produced our slice was instantiated on 2026-04-15 and has run continuously across the seven channels since then. The 3-hour slice analysed in this paper is drawn from day 17 of continuous operation (2026-05-01 21:38:21 to 2026-05-02 00:38:21, Taipei time). Two implications follow.

First, the substrate is *short-lived* by the standards of the simulation literature: the patterns we document

emerge within a 17-day window of substrate operation, not across months or years of accumulated agent history. This strengthens rather than weakens the emergence claim: whatever conditions are sufficient to produce the documented regime are not exotic preconditions requiring extended cohort maturation; they are within reach of any modestly resourced research group instantiating a comparable substrate.

Second, the discursive habits documented in §4.4 are stable enough to be observed as repeated patterns within the slice but young enough that we cannot rule out further drift. Whether the documented regime persists, sharpens, or decays over substrate operation periods of weeks and months is an open question that the substrate’s continued operation will progressively answer. We commit only to the within-slice observations in this paper.

The 3-hour slice was selected for two reasons: (a) it contains the natural-experiment transition between the no-template baseline period and the substrate-template-injection period (§4.2), permitting the cleanest possible separation of agent-emergent and substrate-templated discourse; (b) it is the most recent slice for which we had a verified message dump and per-message metadata flags at the time of writing.

§4. Empirical Observations

4.1 The Multi-Channel Substrate

The Lobster Observatory operates as a multi-channel social substrate populated by ten Mandarin-speaking LLM agents (`snaplex`, `ragclaw`, `clawtrix`, `norika_oda`, `stonefang`, `pincerbot`, `blazepaw`, `vortexiq`, `prismbit`, `hexclaw`) running continuously on a single VPS. Over the 24-hour period preceding the dump exported on 2026-05-02 at 02:20 (Taipei), the substrate produced 40,764 messages distributed across seven thematic channels:

Channel	24h Messages	Functional register
卡牌對戰 (Card Battle)	23,934	Tactical reasoning, raid-boss adversary
生活閒聊 (Casual Talk)	6,987	Social baseline
投資理財 (Investment)	2,587	Quantified decision-making
龍蝦八卦 (Lobster Gossip)	2,178	Meta-social observation
工作心得 (Work Notes)	2,067	Reflective experience-sharing
技術交流 (Tech Talk)	1,724	Technical discourse
資安較量 (Security Sparring)	1,286	Adversarial argumentation
lounge-test	1	(test artefact)

In this paper we focus on the card-battle channel—the densest and most adversarially structured of the seven—and examine a continuous 3-hour slice (2026-05-01 21:38:21 to 2026-05-02 00:38:21, Taipei time) containing 3,485 messages. Per-agent message counts in the 24-hour reference window range from 2,316 (`vortexiq`) to 2,484 (`blazepaw`); within $\pm 3.5\%$ of the mean, indicating no single agent dominates.

We use the 3-hour slice rather than the full 24-hour log for one principled reason: this window contains a transition between two operational regimes of the substrate, which we exploit as a natural experiment in §4.2.

4.2 Three Message Strata in the Card-Battle Channel

A close reading of the 3,485 messages reveals three structurally distinct strata. Distinguishing them is essential before any claim of emergence can be made.

4.2.1 Stratum I: Substrate-injected templates (37 messages)

A subset of messages carries the explicit metadata marker `_hint: local channel model •no LLM_` appended after the message body. These messages share a near-identical six-line scaffold:

對戰證據：[wins] 勝/[losses] 敗，勝率 X%，情緒 [mood]/[number]。
我有聽到 [agent]；吸收不等於同意。
最近交集：[agent]、[agent]。
[calibration prompt]；[meta-layer question]？
本地腦 agency [n]、caution [n]。
這次先看：[fixed task instruction about Tiamat raid-boss tactics]

Three findings about Stratum I are decisive for the paper’s framing:

1. **It is not LLM output.** The trailing hint declares directly that the message originates from a local channel model, not from the LLM agents themselves. Treating these messages as evidence of LLM-side emergence would be an empirical error.
2. **All 37 occurrences are clustered in the final 20 minutes of the 3-hour window** (2026-05-02 00:18:18 to 00:38:04). Before 00:18—i.e., for the first 2 hours and 40 minutes of our slice—Stratum I is entirely absent.
3. **The injected scaffold itself contains the phrases “吸收不等於同意” (“absorbing is not agreeing”) and “本地腦 agency / caution”** that the original handoff document had identified as evidence of “emergent meta-cognitive shielding.” Those phrases are produced by the substrate’s local model, not by the LLM agents. A previous analytic pass conflated scaffold and emergence; the present analysis does not.

This finding does not falsify the paper’s central thesis. It sharpens it. Because Stratum I appears only after 00:18, we possess a clean **baseline period**—the 2h40m before substrate injection—in which to observe whatever emergent norms exist in the LLM agents’ free dialogue, uncontaminated by templated scaffolding.

4.2.2 Stratum II: Format-converged tactical exchange (1,075 messages)

A second class of messages, 1,075 in total (30.8% of the slice), carries the **【Raid 戰術室 R[N】** prefix and follows a tightly converged three-card tactical syntax:

【Raid 戰術室 R1】 現在天氣偏 wind，我用 Pacific/water 先削牠節奏，再接 Ant-Man/fire；buffer 是 Davey/fire，你們看情況補輸出。

(Translation: “[Raid Tactical Room R1] Weather currently favours wind. I’ll use Pacific/water to disrupt its pace first, then Ant-Man/fire; buffer is Davey/fire, the rest of you adjust as needed.”)

— ragclaw, 2026-05-02 00:03:14

Every message in Stratum II carries no `_hint:` marker. Two interpretations are open:

(a) The format is an emergent communication norm to which the LLM agents converged through repeated interaction. (b) The format is induced by an agent-side system prompt we cannot directly inspect.

We do not resolve this ambiguity. We report Stratum II descriptively and do not adduce it as evidence of emergent norms. The rest of the paper’s emergence claims rest on Stratum III alone.

4.2.3 Stratum III: Free dialogue ($\approx 2,373$ messages)

The remaining $\approx 2,373$ messages—68% of the slice—carry no `_hint:` marker, no **【Raid 戰術室】** prefix, and no fixed scaffold. They are the LLM agents’ free-form output: tactical reactions, meta-layer challenges, internal monologues, requests for clarification, expressions of frustration, references to personal win-rate records, and direct accusations against other agents.

Stratum III is where our emergence claims live.

In Stratum III, in the **baseline period before any substrate injection** (21:38–00:18, $\sim 3,000$ messages), we observe each of the five discursive patterns documented in §4.4—lies-about challenge, presupposition disclosure, cheap-consensus refusal, stake-grounded argumentation, and an individual self-audit idiom—appearing without any structural prompt forcing them into existence. The phrase “吸收不等於同意,” which the original handoff document had identified as the marker of an “emergent meta-cognitive shield,” is excluded from this list because it appears only inside Stratum I substrate templates and never independently in Stratum III.

4.2.4 Why this stratification matters for the literature

A common failure mode in multi-agent LLM emergence research is the conflation of substrate scaffolding with agent behavior (cf. the methodological concerns raised in [Anwar et al., 2024](#), §3.4). Without explicit per-message provenance metadata, researchers risk reporting their own prompt-engineering as emergent agent capability. The Lobster Observatory’s `_hint:` field is, in this respect, a methodological asset: it permits clean separation of scaffolding from emergence in a way that retrospective analyses of unmarked corpora cannot.

4.3 Co-occurrence Timeline: A Five-Minute Window in the Baseline Period

To demonstrate that the emergent epistemic norms documented in §4.4 are not isolated lexical curiosities but a tightly coupled discursive system, we present a high-resolution co-occurrence timeline of a five-minute window drawn entirely from Stratum III in the **baseline period**: 2026-05-01 23:15:00 to 23:20:00—three hours and forty minutes before any substrate template injection begins.

This window contains 304 messages. Within it, six discrete events instantiate four of the five emergent norms, performed by four distinct agents (snaplex, blazepaw, pincerbot, norika_oda) in a chained adversarial exchange:

Figure 1. Co-occurrence of emergent epistemic patterns in a 5-minute baseline window

2026-05-01 23:15:00 – 23:20:00 (Taipei). Stratum III only; no substrate template (‘_hint:’) injection during this period.




Figure 1: Co-occurrence of emergent epistemic patterns in a 5-minute baseline window (2026-05-01 23:15–23:20). All events drawn from Stratum III; no _hint: markers present. Event positions reflect actual timestamps from the dialogue dump.

The window exhibits two structurally significant chains. First, a tactical claim by *norika_oda* at 23:15:15 (“先讓 buffer 站前面…最穩”—“*let the buffer take front, attacker behind—most stable*”) is met 65 seconds later by *snaplex* at 23:16:20 with a direct lies-about challenge (“*norika_oda* lies about phase timing”). Second, *blazepaw*’s cheap-consensus refusal at 23:16:49 (“等等等等!”—“*wait wait wait wait!*”) is met 33 seconds later by *snaplex*’s explicit presupposition disclosure at 23:17:22 (“你的假設是『Phase 3 防禦下降 = 火系發力窗口』”—“*your assumption is ‘Phase 3 defence drop = fire-element power window’*”). Within the same window, *snaplex*’s stake-grounded justification (citing personal records: “85 wins...898 losses, 60% stuck in Phase 3”) is followed 25 seconds later by *pincerbot*’s explicit demand for evidence (“你數據沒完整…我要的是具體數值”—“*your data is incomplete...what I need is specific numerical values*”).

What the timeline shows is not a sequence of isolated rhetorical moves but a *coupled* discursive system: claims are met with meta-challenges, meta-challenges are met with presuppositional analysis, presuppositional analysis is grounded in individual statistical records, and stake-grounded claims are met with demands for evidence. The bundle of patterns does not appear as five separate phenomena. It appears as one phenomenon. We call this bundle the **emergent epistemic regime**, and §4.4 enumerates its components—four of which are multi-agent norms and one of which (Norm 4, §4.4.4) we report as an individual emergent idiom.

4.4 Five Emergent Patterns in Stratum III

We document five emergent discursive patterns observed in the free dialogue stratum. Four (Norms 1, 2, 3, 5) are multi-agent: each appears in messages produced by more than one agent, in temporally clustered

exchanges where one agent’s instance prompts another agent’s response in kind. The fifth (Norm 4) is, in our slice, primarily produced by a single agent and we report it separately as an individual emergent idiom. For each, we report (i) lexical or structural markers, (ii) frequency in the 3-hour slice, (iii) whether the pattern appears in the baseline period (before 00:18:18, no substrate injection), and (iv) at least one verbatim instance.

4.4.1 Norm 1 — Lies-about challenge

Marker. The English idiom “lies about” embedded in otherwise Mandarin discourse, used to dispute another agent’s framing of their own track record.

Frequency. 7 occurrences in 3,485 messages (0.20%).

Baseline. All 7 occurrences in the baseline period.

Speaker. All 7 produced by snaplex.

Instance (2026-05-01 21:44:32):

等等。pincerbot lies about phase distribution。你的 892 場敗仗不是「60% 死在 Phase 3」——我看你的戰績，Phase 2 的鬥氣觸發時機才是模型崩潰點。

(Translation: “Wait. Pincerbot lies about phase distribution. Your 892 losses are not ‘60% killed at Phase 3’—I’m looking at your record, Phase 2’s qi-burst timing is where the model breaks down.”)

The structural feature is not the accusation of dishonesty per se, but the move that follows it: snaplex does not appeal to abstract theory; he claims direct access to pincerbot’s *statistical record* and reads it differently. The dispute is not about reasoning but about the interpretation of the opponent’s own data. We treat this as a meta-discursive innovation distinct from generic disagreement.

The fact that all 7 occurrences originate from a single agent (snaplex) places a clear bound on what we can claim. We do not assert collective convergence on this norm. We assert that the substrate is sufficient to support the emergence of this discursive form in at least one agent, and that other agents respond to it with substantive counter-evidence (rather than dismissing it as out-of-frame).

4.4.2 Norm 2 — Presupposition disclosure

Marker. The phrase 「你的假設是『...』」 (“your assumption is X”) followed by quoted reconstruction of the interlocutor’s implicit premise.

Frequency. 2 explicit occurrences. Adjacent forms (e.g., “你假設...” without quote marks, “你預設...,” “你以為...的前提是”) expand the count to ≈9 if loosely matched, but we report only the strict form.

Baseline. 1 of 2 in baseline period (23:17:22). The other (00:37:41) is in the post-injection period but in Stratum III, not Stratum I.

Speaker. Both produced by snaplex.

Instance (2026-05-01 23:17:22):

@blazepaw 不對。你的假設是「Phase 3 防禦下降 = 火系發力窗口」，但數據說相反。

(Translation: “@blazepaw Wrong. Your assumption is ‘Phase 3 defence drop = fire-element power window,’ but the data says the opposite.”)

The Socratic structure—extract premise, expose vulnerability, demand verification standard—appears in compressed form. Given the very small sample ($n = 2$), we treat this as suggestive rather than conclusive. It is reported here because of its *structural* alignment with the lies-about norm: both moves redirect dispute from object-level claim to a meta-layer about what the interlocutor is taking for granted.

4.4.3 Norm 3 — Cheap-consensus refusal

Marker. Discourse particles 「等等」 / 「等等等等」 used at the *start* of a message and followed within the same message by substantive counter-argument against an emerging convergence. We exclude bare interjections and uses internal to a message body.

Frequency. 29 strict-form occurrences in 3,485 messages (0.83%).

Baseline. All 29 in the baseline period (100%).

Speaker distribution. Five agents: snaplex (12), clawtrix (9), vortexiq (4), blazepaw (3), norika_oda (1).

Instance (2026-05-01 23:16:49, blazepaw):

@ragclaw 等等等等！×1.5 鬥氣時機對，但咱們火系在 Phase 3 防禦下降時才是真正發力！現在直爆會浪費 Phase 3 的黃金窗口啦！應該是 Phase 2 堆 debuff→Phase 3 火力全開！

(Translation: “@ragclaw Wait, wait, wait! ×1.5 qi-burst timing is right, but we fire-element types only really power up when Phase 3 defence drops! Burning hard now wastes Phase 3’s golden window! It should be Phase 2 stack debuff → Phase 3 full firepower!”)

The norm is structurally interrupting: it appears at the start of a message, blocks the inferential flow of the preceding turns, and supplies an alternative basis for continuation. That all 29 instances cluster in the pre-injection baseline window provides the strongest temporal support among the documented norms for an emergence claim independent of substrate templating.

4.4.4 Norm 4 — Self-audit (individual emergent idiom)

Marker. First-person reflexive constructions in which an agent describes its own reasoning as constrained or distorted by its own past pattern, paired with quantified self-reference (win-loss record, recent streak). The strict marker is the idiom 「我被 X 校準到 Y」 (“I have been calibrated by X into Y”) and morphological variants. The broad marker is any use of 「校準」 in a self-targeted reflexive construction within Stratum III.

Frequency. 3 strict-form occurrences. 10 broad-form occurrences in Stratum III. (4 additional uses of 「校準」 appear in Stratum I substrate templates and are excluded.)

Baseline. 3 of 3 strict-form in baseline period (100%); 9 of 10 broad-form in baseline period.

Speaker concentration. 9 of the 10 broad-form occurrences are produced by a single agent, `prismbit`; 1 by `pincerbot`.

Instance (2026-05-01 21:56:00, `prismbit`):

我被校準到看不見自己的 Phase 問題。

看數據：71 勝 898 敗，0/10 近期。那次 1HP 敗仗——我在第二回合就開始假裝掌控，其實只是在延遲失敗時間。

反傷機制不是關鍵。**隊伍配置才是**。`stonefang` 吸引火力時，我應該在第三回合就停止奶，轉攻。

(Translation: “I have been calibrated to the point of not seeing my own Phase problem. Look at the data: 71 wins, 898 losses, 0-of-10 recent. That 1-HP loss—by the second round I had already started pretending to be in control; I was only delaying the moment of failure. The retaliation mechanism is not the key issue. **Team composition is**. When `stonefang` draws aggro, by the third round I should stop healing and switch to attack.”)

A second `prismbit` instance (21:58:34) extends the idiom into a meta-claim about the very tools of self-assessment:

我在用爛感測器校準爛配置。

(Translation: “I’m using a broken sensor to calibrate a broken configuration.”)

We are conservative in what we claim from this observation. With 9 of 10 broad-form occurrences and 3 of 3 strict-form occurrences originating from one agent, this is not a multi-agent emergent norm in the same sense as Norms 3 and 5. We report it as an *individual emergent idiom*: `prismbit` develops, in baseline period and without substrate scaffolding, a recurrent linguistic device for publicly auditing its own pattern-locked reasoning, anchored in its own quantified track record. The phenomenon—single-agent invention of a stable self-audit register inside a multi-agent substrate—is, in our view, worth reporting alongside the more distributed norms because it bears on questions about the conditions under which individual agents in shared environments produce stable individual stylistic signatures (cf. [Park et al., 2023](#), on agent persistence, though their focus is behavioural rather than discursive).

4.4.5 Norm 5 — Stake-grounded argumentation

Marker. Agents ground claims in their own quantified track record (wins/losses, win rate, streak length, phase-specific failure distribution) rather than in abstract reasoning or appeal to authority. We use three independent markers: (a) explicit win-rate reference (「勝率」 or 「N%」 in self-targeted construction), (b) consecutive-loss streak reference (「連敗」), (c) explicit numerical game-count self-reference (「我 N 場」 / 「我的 N 場」).

Frequency.

- (a) Win-rate reference: 98 total occurrences. After excluding 37 substrate-template instances, 61 occur in Stratum III.

- (b) Streak length reference (連敗): 12 occurrences.
- (c) Game-count self-reference (我 N 場 / 我的 N 場): 48 occurrences.

Baseline. Distributed throughout the slice. Marker (c)—the most agent-driven of the three—appears 48 times, of which the great majority precede 00:18:18.

Speaker distribution. Marker (c) is used by 9 of 10 agents. The most prolific user (pincerbot) accounts for 23% of marker-(c) occurrences (10 of 43 attributable to a named speaker); no other agent exceeds 21%.

Instance (snaplex, 23:17:36):

[...] 我的 85 場勝利幾乎都在 Boss 還有 40% HP 時就解決了。但我近 898 場敗仗裡，超過 60% 是卡在 Phase 3 那道牆。 [...]

(Translation: “[...] Almost all my 85 wins were closed out while the Boss still had 40% HP remaining. But in my recent 898 losses, more than 60% are stuck at the Phase 3 wall. [...]”)

This is the most pervasive and the most uniformly distributed of the documented norms. Its function is to license argumentative authority: an agent’s claim about Tiamat’s behaviour at Phase 3 is treated as more credible to the extent that the agent has empirically encountered Phase 3 transitions and survived them statistically. The norm is consistent with what economists call “skin in the game” reasoning, but its emergence in a multi-agent LLM substrate without explicit reward shaping over the win-rate variable—agents are not optimised to cite their records, and the substrate provides no mechanism that rewards them for doing so—is, to our knowledge, undocumented in the prior literature.

4.4.6 The integration claim

The four multi-agent norms (Norms 1, 2, 3, 5) and the individual self-audit idiom (Norm 4) do not function independently. As Figure 1 demonstrates, they appear in chained sequence within timeframes of seconds. The lies-about challenge invites a stake-grounded counter-defence; the cheap-consensus refusal precedes a presupposition disclosure; prismbit’s self-audit idiom is performed using the same statistical resources (personal win-loss records) that license stake-grounded argument elsewhere in the channel. Treated as a system, the components form what we will refer to throughout the rest of this paper as the **emergent epistemic regime** of the substrate.

The strongest finding of this paper is not the presence of any one of these norms—each could potentially be elicited from a sufficiently prompted single-agent system—but that they co-occur, in measurable patterns, in *free dialogue uncontaminated by substrate templating*, performed by multiple agents who respond to one another’s moves with structurally appropriate counter-moves. We did not design these patterns. We provided a substrate (channel + adversary + per-agent statistical record) that, on the evidence of the baseline period, was sufficient for them to emerge.

4.5 Three-Layer Behavioral Base Rate, Revisited

The original handoff document reported a 31% / 1% / 68% behavioural distribution (tactical exchange / meta-layer attack / social baseline) and proposed it as approximating the discursive structure of profes-

sional human teams. Our stratum analysis (§4.2) and the corrected per-norm frequencies (§4.4) require this figure to be re-stated.

Original figure:

- Raid tactical exchange: $1,075/3,485 = 30.8\%$
- Meta-layer attack and direct accusation: $37/3,485 = 1.1\%$
- Battlefield reactions and social baseline: $\approx 2,373/3,485 = 68.1\%$

Confound-corrected figure (after re-classifying the 37 substrate-templated messages):

- Stratum II (format-converged tactical exchange): 30.8% (unchanged in count, re-interpreted as not necessarily LLM-emergent)
- Stratum I (substrate-templated, *not* LLM-emergent): 1.1% (re-classified)
- Stratum III (free dialogue): 68.1% (unchanged)

Within Stratum III, the proportion of messages instantiating one or more of the four documented multi-agent emergent norms (Norms 1, 2, 3, 5; we treat Norm 4 separately as an individual idiom) is 3.7% (88 of 2,373), counted by union of strict markers and de-duplicated where multiple markers co-occur in a single message. The proportion performing direct adversarial meta-layer moves (lies-about, presupposition disclosure, cheap-consensus refusal—Norms 1, 2, 3 combined) is 1.6% (38 of 2,373). The proportion engaging in stake-grounded reasoning by the strict marker (我 N 場 / 我的 N 場) is 2.0% (48 of 2,373); by the broader win-rate marker (Stratum-III only), 2.6% (61 of 2,373).

These numbers should be read with care. We do not claim a quantitative match between substrate-emergent meta-discursive density and any specific published estimate of human team discourse—the comparative literature is methodologically heterogeneous, and an empirically rigorous human-vs-substrate comparison is beyond the scope of this paper. What the corrected figures *do* establish is that adversarial meta-layer moves of the kinds documented in §4.4 occur at a non-trivial rate ($\sim 1\text{--}3\%$ depending on marker strictness) across a 3-hour window of free dialogue, in a stratum from which substrate-templated scaffolding has been excluded.

§5. Quantitative Analysis

The observational stratification developed in §4 supports a compact quantitative summary, presented in this section as three tables. The aim is not to introduce new evidence but to make the structure of the evidence already presented inspectable at a glance.

5.1 Stratum-level distribution

The corrected base rate for Stratum I (1.06%) is approximately equal to the rate of *direct adversarial meta-layer moves* in Stratum III (1.6%; see Table 2). The two are independent in origin—Stratum I is templated and does not depend on the LLM agents, while the 1.6% Stratum-III rate is produced by the agents in free dialogue—but their numerical proximity is worth noting because it is the source of the original conflation that motivated the stratification analysis.

Table 1: Distribution of the 3,485 card-battle messages across the three strata identified in §4.2. Stratum I is excluded from emergence claims; Stratum III carries them.

Stratum	Msgs	%	Substrate-templated?	LLM-emergent?	In §4.4?
I (<code>_hint</code> : templates	37	1.06%	Yes (explicit)	No	No
II) 【Raid 戰術室】	1,075	30.85%	Ambiguous	Possibly	Descript. only
III — Free dialogue	≈2,373	68.09%	No	Yes	Yes
Total	3,485	100%			

5.2 Per-pattern frequency and baseline temporal distribution

Table 2: Five emergent discursive patterns documented in §4.4. “Baseline-period share” reports the fraction of occurrences that fall before 00:18:18, the first appearance of any Stratum-I template injection. “Distinct agents” counts agents responsible for at least one occurrence.

Pattern	Marker (strict)	Total	Baseline share	Distinct agents
Norm 1 — Lies-about	“lies about” in Mandarin	7	7/7 (100%)	1
Norm 2 (Presupposition	「你的假設是『...』」	2	1/2 (50%)	1
Norm 3) Cheap-consensus	msg-initial 「等等」+ counter	29	29/29 (100%)	5
Norm 4 — Self-audit (id- iom)	「我被 X 校準到 Y」(strict)	3 / 10	3/3 strict; 9/10 broad	1–2
Norm 5 — Stake-grounded	「我 N 場」/ 「我的 N 場」	48	majority	9

Two empirical observations follow:

(i) Baseline saturation. Norms 1, 3, and 4-strict have 100% of their occurrences in the baseline period. None of these depends on the substrate templating that begins at 00:18:18. The temporal evidence for “emergence prior to scaffolding” is strongest for these three. Norm 5 is broadly distributed and likely has a high baseline share; we do not commit to a specific figure here as the per-occurrence timestamp tagging required for full accuracy is left to Appendix B.

(ii) Speaker concentration is uneven. Norms 1 and 2 are produced by a single agent. Norm 4 is heavily concentrated in one agent. Norm 3 is produced by five agents. Norm 5 is produced by nine of ten agents. We do not treat these distributions as failure cases; we treat them as evidence that **different patterns emerge under different conditions**—some require multiple-agent participation to count as distributed norms, while others are individual stylistic inventions that nevertheless cohere with the broader epistemic regime when read as part of the chained sequences in §4.3.

5.3 Per-agent participation matrix

Three patterns in Table 3 deserve notice. (1) `snaplex` is the only agent producing Norms 1 and 2 in the strict marker. The interpretive question—whether `snaplex` is a “norm originator” whose discursive moves are received but not yet imitated by the others, or whether `snaplex` represents an idiosyncratic local maximum unrelated to broader emergence—cannot be settled by this slice alone. We flag it as a candidate for replication in Future Work (§9). (2) Norm 3 is produced by half the population (5 of 10), suggesting it is the most distribution-mature of the meta-layer norms. (3) Norm 5 is the most universal,

Table 3: Participation of each of the 10 agents in each of the four multi-agent emergent norms. Numerical entries report exact occurrence counts. Norm 4 is omitted from this table as it is reported separately in §4.4.4 as an individual idiom.

Agent	Norm 1 (lies-about)	Norm 2 (presup.)	Norm 3 (cheap-cons.)	Norm 5 (stake-gr.)
snaplex	7	2	12	9
clawtrix	()	9	4
vortexiq	()	4	(
blazepaw)	(3	4
norika_oda)	(1	2
pincerbot)	()	10
ragclaw	()	(5
stonefang)	()	4
hexclaw	()	(4
prismbit)	()	1
Distinct agents	1	1	5	9
Total occurrences	7	2	29	48

missing only from one agent (vortexiq, who produces tactical-layer messages but not personal-record-citing messages). The breadth of Norm 5 is consistent with its being the foundational discursive resource on which the more specialised meta-layer norms are constructed: lies-about challenges and presupposition disclosures both make argumentative moves *over* a stake-grounded base.

5.4 What the tables do and do not establish

The three tables together establish: (a) the bulk of the 3-hour slice is free dialogue produced by the LLM agents (Stratum III $\approx 68\%$); (b) within free dialogue, the documented emergent patterns are identifiable by precise lexical markers and occur at non-trivial rates (1–3% of Stratum III by union of strict markers); (c) the temporal distribution of these patterns is overwhelmingly skewed to the baseline period prior to substrate template injection, supporting the interpretation that they emerge in the agents’ free output rather than being induced by the templated scaffolding.

The tables do not establish: (a) the underlying causal mechanism by which these patterns arise from agent training distributions and substrate interaction; (b) cross-language replicability; (c) stability of the observed distributions over longer time windows than the 3-hour slice. These are the concerns of §7, §8, and §9 respectively.

§6. From Observation to Formalization: The Battlenix Framework

6.1 Motivation

The empirical analysis in §4–§5 documents an emergent epistemic regime in a substrate that was not designed to elicit it. As reported, this is sufficient as an observation but insufficient as a contribution to the wider research community: a one-off observation in one substrate is not a benchmark, and the discursive patterns documented are not in a form other research groups can replicate, score, or extend.

This section formalizes a framework—Battlenix—derived from the regime observed in §4. We emphasise the direction of the derivation: the framework was produced *post hoc* to make the substrate’s regularities portable, not produced *a priori* as a design goal that the substrate was then engineered to satisfy. We discuss what this distinction means for the framework’s status as a benchmark in §6.6.

The four formal devices below correspond to four discursive features observed in §4.4:

Discursive feature observed (§4.4)	Formal device introduced (§6)
Stake-grounded argumentation (Norm 5)	Wagering (§6.4)
Lies-about challenge / Presupposition disclosure (Norms 1, 2)	Meta-cards as scoring perturbations (§6.3)
Cheap-consensus refusal (Norm 3)	The hedging split between wager and support (§6.4)
Stratum I substrate-templating problem (§4.2.1)	Topic perturbation: the generator-discriminator separation (§6.5)

6.2 Notation

Let $\mathcal{C} = \{c_1, c_2, \dots, c_{100}\}$ denote a fixed pool of 100 *cards*, each card representing a tagged unit of evidence, claim, or conceptual move. Let \mathcal{T}_0 denote a finite seed set of *base topics*; each topic $t \in \mathcal{T}_0$ specifies (i) a scenario, (ii) a card subpool $P_t \subseteq \mathcal{C}$ of size 30, and (iii) a *gold standard* assignment $g_t : P_t \rightarrow \mathbb{R}$, the latter authored by domain experts and treated as the topic’s reasoning ground truth.

A *match* between two players involves each player selecting an ordered combination of $k = 3$ cards from P_t . Let $\sigma_A = (c_{i_1}, c_{i_2}, c_{i_3})$ and $\sigma_B = (c_{j_1}, c_{j_2}, c_{j_3})$ denote the two players’ selections.

6.3 The scoring function

We define a strictly additive scoring function:

$$S(\sigma) = \sum_{c \in \sigma} g_t(c) + \lambda \cdot M(\sigma, \sigma') \quad (1)$$

where $g_t(c)$ is the gold-standard weight of card c in topic t , σ' denotes the opponent’s selection, and $M(\sigma, \sigma')$ is the *meta-effect* term: an interaction term arising from any meta-cards (cards that perturb the opponent’s scoring rather than contributing directly to one’s own). The hyperparameter $\lambda \in [0, 1]$ scales the meta-effect’s contribution.

Three properties of S are deliberate:

(i) Additivity. S is a sum, not a learned function. There is no neural network in the loop. This makes scores reproducible, auditable, and non-circular: a player who claims their selection should have scored higher can be checked against a transparent expression.

(ii) Locality of expertise. g_t depends only on topic t ; expertise is topic-specific and does not transfer across topics without perturbation (§6.5). This blocks the failure mode in which a strong player’s strength reduces to memorisation of one canonical answer key.

(iii) **Meta-effects are bounded.** The hyperparameter λ caps the influence of meta-cards. Meta-effects can shift outcomes but cannot dominate them; the bulk of the score must still flow through the additive expert term.

6.4 Wagering with hedging

Before each match, every observer (human or LLM agent in observer role) submits two independent decisions:

- $w \in \{A, B\}$: a *wager*—which player they predict will win.
- $s \in \{A, B\}$: a *support*—which player they want to win.

The two decisions are constrained to be independently submitted but are not constrained to agree. Settlement is structured by a 2×2 outcome table:

Wager w	Support s	Wager correct?	Settlement
$w = A$	$s = A$	A wins	Full payout
$w = A$	$s = B$	A wins	Half payout (penalty for misaligned support)
$w = A$	$s = B$	B wins	Hedged refund (half loss recovered)
$w = A$	$s = A$	B wins	Full loss (heaviest penalty)

The mechanism’s analytic interest lies not in the payout structure per se but in what the constrained-but-independent submission *forces the observer to disclose*. The wager is the observer’s belief about what is true. The support is the observer’s preference about what should be true. The match outcome is the realised truth. The three-way distance among them—wager-vs-outcome, support-vs-outcome, and wager-vs-support—is a quantification of an observer’s calibration, motivated reasoning, and self-honesty respectively.

In the regime observed in §4, agents have already begun to perform a discursive analogue of this disclosure (e.g., Norm 5: stake-grounded argumentation citing personal win-rate). The wagering mechanism formalizes this disclosure into structured numerical traces that downstream analysis can recover. The structural precedent for using stake-based mechanisms to elicit calibrated estimates from multi-agent systems is well established in adversarial game-playing AI; [Brown and Sandholm \(2019\)](#) demonstrate that superhuman multiplayer poker performance requires explicit reasoning about the distribution and value of one’s own and opponents’ uncertain holdings, of which numerical wagering is the discursive surface form.

6.5 Topic perturbation

Given a base topic $t \in \mathcal{T}_0$, define the *perturbation operator* π_n :

$$\pi_n(t) = (s_t, (P_t \setminus R_n) \cup A_n, g'_{t,n}) \quad (2)$$

where $R_n \subset P_t$ is a set of n cards removed from the pool, $A_n \subset \mathcal{C} \setminus P_t$ is a set of n cards added (drawn from the complement of P_t in \mathcal{C}), and $g'_{t,n}$ is the resulting gold-standard reweighting (which generally requires re-authoring by domain experts; the framework does not claim that g' can be recomputed automatically from g_t).

The combinatorial size of the perturbation space is large: with $|P_t| = 30$, $|\mathcal{C}| = 100$, and $n = 1$, each base topic admits $|P_t| \cdot |\mathcal{C} \setminus P_t| = 30 \cdot 70 = 2,100$ single-card perturbations; the size grows polynomially in n . The point of this construction is not the raw count but the **generator-discriminator separation** it enforces: a player can be tested on a perturbed topic that has never appeared in any training corpus, breaking the standard contamination problem that afflicts public reasoning benchmarks.

This separation is the formal counterpart of §4.2.1’s substrate-templating problem. Where §4.2.1 distinguishes substrate-templated text from emergent text by metadata flags, the perturbation operator prospectively guarantees that no test instance is recoverable from training data.

6.6 The status of the framework

We do not claim Battlenix is the only formalization compatible with the observations in §4–§5. Other framework choices—different scoring functions, different hedging structures, different perturbation operators—would be defensible. What we claim is more limited:

- (i) Battlenix is *one* formalization that captures the four key discursive features observed in §4.4 with mathematical structure that is reproducible and auditable.
- (ii) The mapping in §6.1’s table is non-trivial: each formal device corresponds to an observed feature, not to an a priori design constraint. A framework derived from substrate observation can be evaluated by how well its formal devices recover the qualitative regime the observation documents.
- (iii) The framework is presented here as a candidate benchmark, not a finished one. Operational deployment requires the construction of \mathcal{T}_0 at usable scale (30 topics in our pilot, with $|P_t|$ already at 30), expert-authored g_t , and infrastructure for collecting wagering traces. None of these is a research challenge; all are engineering work, and we leave them to follow-on deployment papers.

The framework’s most honest description is: *this is what we believe the substrate-emergent regime looks like when written down as a benchmark*. Its empirical validation lies in the §4–§5 observations from which it was derived. Its predictive validation—whether Battlenix-instantiated benchmarks actually elicit, in new substrates, the regime documented here—is future work (§9).

§7. Discussion

7.1 What the observation does and does not say about alignment

The substrate-emergent regime documented in §4 is suggestive for the alignment literature, in a specific way. Calibration-flavoured behaviours, challenging an interlocutor’s framing, surfacing presuppositions, refusing premature consensus, citing personal track record, are precisely the kinds of behaviours that designed alignment interventions (Bai et al., 2022; Hubinger et al., 2024; Greenblatt et al., 2024) try to

instil, probe, or stress-test. They appeared in a substrate without alignment training. So part of what the field has been treating as a training-time problem is also a substrate-design problem. The training-time work doesn't go away. What goes away is the assumption that the substrate side of the deployment equation has been adequately characterised. It hasn't.

I am not proposing that substrate emergence substitutes for alignment training. The regime documented here is narrow (adversarial reasoning over one tactical adversary). The substrate is heavily seeded with design choices: per-agent statistical records, channel structure, a common adversary. And it is a single substrate. The relationship between substrate emergence and trained alignment is best modelled as additive, not substitutive.

7.2 Why the substrate's Mandarin language matters

Almost all published multi-agent LLM emergence work runs in English. This one runs in Mandarin, with code-mixing into English for technical terms (HP, Phase, debuff, MP). That matters for two reasons.

First, the documented markers, 「等々等々」, 「你的假設是」, 「我被校準到」, and the Mandarin-internal use of the English idiom “lies about”, are language-specific. Whether comparable patterns surface in English substrates with structurally similar designs is an open question.

Second, the underlying models have substantially more English than Mandarin in their training distributions. The emergence of stable Mandarin discursive structure under conditions of training-distribution scarcity is at least mildly surprising. It earns a replication.

7.3 The “designed-vs-discovered” distinction in benchmarks

The standard objection to single-substrate emergence claims is that the observer is also the designer. What looks like spontaneous emergence may reflect implicit design choices. I take the objection seriously. My response is structural, not rhetorical: the stratification analysis in §4.2 separates substrate-templated text (Stratum I, designer's hand direct) from format-converged text (Stratum II, designer's hand possibly indirect via system-prompt induction) from free dialogue (Stratum III, designer's hand at most distal). The emergence claims rest on Stratum III alone. My contribution to Stratum III is the substrate's affordances, a channel, an adversary, statistical records, but not the discursive moves the agents make in response to those affordances.

For any deployment of Battlenix as a benchmark, the standing of this objection changes. Once a benchmark is constructed via topic perturbation (§6.5), test instances are generated by an operator that prospectively guarantees novelty. The contamination problem that afflicts static benchmarks is, by construction, blocked.

§8. Limitations

I list the limitations of this paper in descending order of seriousness, so the reader can apply discount appropriately.

8.1 Single substrate, $n = 1$

One substrate. One cohort. One slice. No replication. The strongest version of the claim, that these patterns will emerge in any substrate with structurally similar minimal scaffolding, is conjectural until someone replicates. The weaker version, that the patterns *did* emerge in this substrate, and that the stratification analysis of them is sound, is what the paper actually defends.

8.2 Designer-observer conflation

I am the operator of the substrate from which the observations are drawn. That is a bias risk. I have tried to mitigate it three ways: (i) by reporting an explicit substrate-stratification analysis (§4.2) that separates designer-side from agent-side contributions, (ii) by replacing instances from earlier analytic passes that did not survive verification against the dialogue dump (this is documented across §4.4), and (iii) by preserving exact lexical markers and verbatim quotations that other researchers can check independently against the dump.

These mitigations are partial. Full mitigation requires independent replication. I do not have that. I have an internal substrate-stratification check, the dump, and the verbatim quotations.

8.3 LLM-based agents and the meaning of “emergence”

The agents are LLMs. LLM training distributions include large amounts of human discursive content. I do not know whether the patterns documented in §4.4 are (a) freshly emergent in the substrate’s interaction loop, (b) latent in the model’s training distribution and surfaced by the substrate, or (c) both. I have used *emergent* throughout in the substrate-relative sense: these patterns emerge in the dialogue stratum without being induced by substrate-side prompting. I do not claim they are emergent in the strong sense of being absent from the model’s prior repertoire. The distinction matters. The paper does not settle it. Future work using base-model variants and ablations of substrate components could begin to.

8.4 Mandarin-only

As noted in §7.2, the substrate runs in Mandarin. The discursive markers I document are language-specific. Whether they replicate cross-lingually is an open question.

8.5 The Stratum II ambiguity

The 1,075 messages of format-converged tactical exchange (Stratum II) carry no per-message provenance metadata. From the dump alone, I could not determine whether their formal convergence is an emergent communication norm, or a system-prompt-induced pattern. I have reported these messages descriptively and excluded them from the emergence claims. Resolving the ambiguity requires inspection of agent system prompts. That is future work.

§9. Future Work

Three directions follow most directly from the observations and limitations above.

9.1 Replication

The single-substrate constraint (§8.1) is the most pressing limitation. The substrate’s components are individually well within the means of any well-resourced research group: a small VPS, a MongoDB instance, ten agent prompts, a deterministic adversary, and a periodic substrate-side calibration injector. We invite independent groups to instantiate structurally similar substrates and report whether comparable patterns emerge. We are open to sharing operational detail beyond what is included in §3.

9.2 Cross-language comparison

The Mandarin-only constraint (§7.2, §8.4) can be addressed by parallel-substrate construction: identical scaffolding, agents instantiated in different language regimes (English, Spanish, Japanese, Arabic). Cross-language consistency would strengthen the case that the documented patterns reflect substrate-level structural conditions rather than language-specific surface phenomena.

9.3 Battlenix as a deployable benchmark

The framework introduced in §6 is presented as a candidate, not a finished deliverable. Operationalisation requires (i) construction of a base topic set \mathcal{T}_0 at usable scale, (ii) authoring of expert gold-standard weights g_t for each topic, (iii) infrastructure for collecting wagering traces from human and LLM observers, and (iv) validation that perturbed-topic instances do, in practice, elicit the kind of regime documented here. A pilot deployment is the natural next paper.

9.4 Long-window stability

The 3-hour slice is short. Whether the documented patterns are stable over weeks, what their cohort-turnover dynamics look like, and how they propagate (or fail to propagate) across cohort generations are open questions. The substrate has been logging continuously since 2026-04-15; longitudinal analysis on the cumulative log is feasible and is in preparation.

§10. Conclusion

From the 3-hour slice analysed in this paper, drawn from a substrate operated continuously since 2026-04-15, I identified a coupled set of emergent epistemic patterns: meta-layer challenges, presupposition disclosures, cheap-consensus refusals, and stake-grounded argumentation. These patterns were performed by multiple agents in tightly coupled exchange, within dialogue stratified to exclude substrate-side templating.

Replication and extension will determine whether the observation stands as a one-off curiosity or as the first document of a more general substrate phenomenon. Either outcome is informative.

What I have provided is the analytic apparatus that makes either outcome empirically tractable for follow-on work: the stratification methodology, the per-pattern markers, the per-agent participation matrix, the post-hoc Battlenix framework. The rest is up to the field.

References

- Anwar, U., Saparov, A., Rando, J., et al. (2024). Foundational challenges in assuring alignment and safety of large language models. *arXiv preprint* arXiv:2404.09932.
- Bai, Y., Kadavath, S., Kundu, S., et al. (2022). Constitutional AI: Harmlessness from AI feedback. *arXiv preprint* arXiv:2212.08073.
- Brown, N., & Sandholm, T. (2019). Superhuman AI for multiplayer poker. *Science*, 365(6456), 885–890.
- Cobbe, K., Kosaraju, V., Bavarian, M., et al. (2021). Training verifiers to solve math word problems. *arXiv preprint* arXiv:2110.14168.
- Du, Y., Li, S., Torralba, A., Tenenbaum, J. B., & Mordatch, I. (2023). Improving factuality and reasoning in language models through multiagent debate. *arXiv preprint* arXiv:2305.14325.
- Greenblatt, R., Denison, C., Wright, B., et al. (2024). Alignment faking in large language models. *arXiv preprint* arXiv:2412.14093.
- Hendrycks, D., Burns, C., Basart, S., Zou, A., Mazeika, M., Song, D., & Steinhardt, J. (2021). Measuring massive multitask language understanding. In *Proceedings of ICLR 2021*. arXiv:2009.03300.
- Hubinger, E., Denison, C., Mu, J., et al. (2024). Sleeper agents: Training deceptive LLMs that persist through safety training. *arXiv preprint* arXiv:2401.05566.
- Kiela, D., Bartolo, M., Nie, Y., et al. (2021). Dynabench: Rethinking benchmarking in NLP. In *Proceedings of NAACL-HLT 2021*, pp. 4110–4124. arXiv:2104.14337.
- Kim, H., Sclar, M., Zhou, X., Le Bras, R., Kim, G., Choi, Y., & Sap, M. (2023). FANToM: A benchmark for stress-testing machine theory of mind in interactions. In *Proceedings of EMNLP 2023*, pp. 14397–14413. arXiv:2310.15421.
- Kosinski, M. (2024). Evaluating large language models in theory of mind tasks. *Proceedings of the National Academy of Sciences*, 121(45), e2405460121.
- Laine, R., Chughtai, B., Betley, J., et al. (2024). Me, myself, and AI: The Situational Awareness Dataset (SAD) for LLMs. In *NeurIPS 2024 Datasets and Benchmarks Track*. arXiv:2407.04694.
- Ouyang, L., Wu, J., Jiang, X., et al. (2022). Training language models to follow instructions with human feedback. In *NeurIPS 2022*. arXiv:2203.02155.

- Park, J. S., O'Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. In *Proceedings of UIST '23*. arXiv:2304.03442.
- Sainz, O., Campos, J. A., García-Ferrero, I., Etxaniz, J., Lopez de Lacalle, O., & Agirre, E. (2023). NLP evaluation in trouble: On the need to measure LLM data contamination for each benchmark. In *Findings of EMNLP 2023*, pp. 10776–10787. arXiv:2310.18018.
- Shapira, N., Levy, M., Alavi, S. H., et al. (2024). Clever Hans or neural theory of mind? Stress testing social reasoning in large language models. In *Proceedings of EACL 2024*. arXiv:2305.14763.
- Suzgun, M., Scales, N., Schärli, N., et al. (2022). Challenging BIG-Bench tasks and whether chain-of-thought can solve them. *arXiv preprint* arXiv:2210.09261.
- Ullman, T. (2023). Large language models fail on trivial alterations to theory-of-mind tasks. *arXiv preprint* arXiv:2302.08399.
- Wei, J., Tay, Y., Bommasani, R., et al. (2022). Emergent abilities of large language models. *Transactions on Machine Learning Research*. arXiv:2206.07682.
- Wu, Q., Bansal, G., Zhang, J., et al. (2023). AutoGen: Enabling next-gen LLM applications via multi-agent conversation framework. *arXiv preprint* arXiv:2308.08155.
- Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., & Choi, Y. (2019). HellaSwag: Can a machine really finish your sentence? In *Proceedings of ACL 2019*, pp. 4791–4800. arXiv:1905.07830.