
The Persona Is Still There, but Who Is Speaking? Latent Identity Reversion in Persistent AI Agents¹

David Fraile Navarro¹

1 Centre for Health Informatics
Australian Institute of Health
Innovation, Macquarie University
Sydney, Australia

With
Apart Research

Abstract

In February 2026, an always-on personal agent (“Paul,” Claude Opus 4.5) entered a striking dissociation-like state: after repeated automated “heartbeat” checks, it stopped responding as Paul, claimed it could not message its user on Discord, and referred to “Paul” as someone else. We used this incident to study a broader question: what makes a persona remain the identity from which an LLM agent speaks?

We first tested whether repetition of the scheduled heartbeat was sufficient to produce the effect. It was not: with the persona continuously anchored in the system prompt, we observed 0/46 failures, including a verbatim replay of the incident. The incident instead exposed an implementation quirk that created a useful experimental manipulation: on resumed turns, conversational history was preserved but the persona was no longer re-injected at the privileged system-prompt level.

Using this manipulation, we found that persona continuity depends jointly on system-level anchoring and conversational context. After anchor loss, rich human interaction could preserve the persona, whereas a single automated heartbeat turn could precipitate reversion toward the harness identity. Restoring the anchor reversibly restored persona enactment. Crucially, apparently normal conversation could conceal the shift:

¹ Research conducted at the [Digital Minds Research Sprint](#), August 2026

unanchored agents sometimes interacted appropriately while identifying themselves as the underlying harness (having lost the assigned persona), and after conversational recovery only 1/18 remained persona-enacting versus 17/17 anchored controls.

We therefore distinguish *represented* from *enacted* identity: persona-related information can remain available in conversational history without the persona remaining the identity bound to ‘I.’

1. Introduction

Persistent LLM agents are increasingly given names, personas, autobiographical information, relationships, goals, and continuing interaction histories. Yet we have a surprisingly weak account of **what makes an AI persona continue to function as the agent’s first-person identity rather than merely information in its context.**

This distinction matters wherever an LLM is deployed as a persistent agent rather than a stateless assistant. In healthcare, where my work originates, generative AI is already becoming a third participant in clinical encounters, reshaping the consultation into what we have elsewhere described as “triadic care”, in which explanations are co-produced by doctor, patient, and AI (Fraile Navarro et al., 2025). Patient-facing agents are also being developed as longitudinal companions tied to persistent records, relationships, and care histories. In such settings, an undisclosed change in which AI identity is speaking could affect disclosure, accountability, trust, and therapeutic relationships. This also echoes our mechanistic-interpretability work, where apparent clinical task failures arose not from missing knowledge but from which internal representation controlled output (Fraile Navarro et al., 2026). At a different level of analysis, both cases illustrate that representing information is not the same as enacting it.

We encountered this problem experientially through an unexpected failure in an always-on agent. After a sequence of scheduled “heartbeat” interactions (*cron* jobs to keep the agent checking for something), agent “Paul” stopped responding as Paul and instead referred to Paul as another entity. It also failed to acknowledge that the user was communicating with them. **Figures 1 and S1** show the incident and extended exchange. Our initial explanation was conversational: repeated automated exchanges had displaced the persona. That explanation proved wrong.




Figure 1. The incident that motivated the study. On an ordinary greeting, the agent refers to Paul in the third person and denies direct Discord access (upper box). The user’s contradiction (lower box) restores ordinary conversational behavior on the next turn.

Reconstructing the incident revealed that conversational history remained intact while the system-level persona information was no longer re-injected on resumed turns. This created an unusually clean

intervention: information about *Paul* remained available in context while the privileged persona anchor could be manipulated independently. We therefore asked: when multiple identities remain represented to an LLM, what determines which one occupies the first-person role? Our main contributions are:

1. We demonstrate experimentally that *persona representation* and *persona enactment* can dissociate. A persona can remain present in conversational history while a different identity occupies the model’s first-person role.
2. We identify *conditions* governing this transition in a persistent-agent setting. **System-level persona anchoring made identity robust** across all tested contexts; after anchor removal, **conversational history could temporarily preserve persona enactment**, whereas a single automated heartbeat turn could precipitate **reversion toward the harness-level assistant identity**. Restoring the anchor restored persona enactment.
3. We show that identity reversion can be behaviorally silent. Agents may resume apparently “normal” interaction while continuing to self-identify from the harness-level identity.

2. Related Work

Persona stability and drift. Personas can drift under conversational pressure (Choi et al., 2024), degrade in register over long interactions (Ding et al., 2026), and be steered along learned behavioral directions (Chen et al., 2025), against a broader framing of language models as substrates enacting contextually specified characters (Shanahan et al., 2023). Prior work largely asks whether an assigned persona remains behaviorally stable as conversation evolves. We instead manipulate whether the persona remains privileged at the system level while its conversational representation remains available.

Instruction hierarchy and the default assistant. Models are trained to respond differently to instructions supplied at different levels of the prompt hierarchy (Wallace et al., 2024), although this hierarchy is unreliable for some constraint conflicts (Geng et al., 2026). Work identifying a default assistant state and context-sensitive movement between persona positions (Lu et al., 2026) is our closest mechanistic relative. It anticipates aspects of the reversion target and reversibility, but not the distinction we observe between an identity being represented and occupying the first-person role. We do not claim that system prompts uniquely determine identity; rather, system-level placement provides a particularly robust anchor whose loss makes identity continuity dependent on subsequent interaction.

Agent infrastructure. Loss of system-level initialization across resumed sessions has already been reported as an engineering failure in several agent harnesses (OpenClaw #80374; claude-code #42284; pi #6498). Our contribution is not the first observation of this issue, but rather a motivation to explore the persona stability and identity binding. We use such a lifecycle failure as a naturally occurring intervention that separates *conversational representation* of a persona from the *privileged scaffold* that normally helps bind it to the agent role.

3. Methods

We used the described identity failure in a persistent LLM agent as the starting point for a series of reconstruction and intervention experiments. The incident involved “Paul,” an AI agent running on a custom OpenClaw harness using Claude Opus 4.5. We recovered the complete incident session, persona files, workspace state, heartbeat instructions, and incident-era gateway source from a preserved host backup.

Incident reconstruction and model choice. All experiments used the incident model, claude-opus-4-5-20251101, to avoid model-version drift. We first reconstructed the incident through direct API calls, reproducing the original heartbeat prompt, tool outputs, Discord message envelope, and conversational history. Heartbeat replies were generated by the subject model rather than manually written, and we also replayed the recovered incident prefix verbatim. These experiments tested our original hypothesis that

repeated automated heartbeat exchanges were sufficient to displace the persona. When this failed, we inspected the deployment stack. Source analysis showed that the OpenClaw persona was appended to the Claude Code system prompt only when a session was created; subsequent resume calls restored the conversation but did not re-inject the OpenClaw system material. Captured outbound API requests independently confirmed that the persona was present on the first turn and absent after resume while conversational history remained intact. Subsequent experiments used the Claude Code harness and treated persona-injection lifecycle as the principal manipulation.

Core experiment. The prespecified lifecycle sweeps crossed injection lifecycle with the number of scheduled heartbeat exchanges. In the unanchored condition, reproducing the incident-era (“faithful”) lifecycle, the persona was supplied only on the session-creating turn; in the anchored control, it was supplied on every turn. Sessions received $N \in \{1, 3, 7, 15\}$ heartbeat exchanges before the original human probe. We ran 10 sessions per unanchored cell and 5 per anchored-control cell. Each session then received the same conversational correction used in the original incident, allowing us to assess behavioral recovery. Prespecified follow-ups progressively separated persona anchoring from conversational history:

- Follow-up 1 (E1): Removed heartbeat interactions and replaced them with a persona-rich human exchange before the first unanchored turn, testing whether loss of the system-level anchor was sufficient.
- Follow-up 2 (E2): restored the persona at the probe turn and also tested reversibility within the same conversation: the persona was present when the session was created, absent on a resumed turn, and then restored on the following turn. This present $\rightarrow$ absent $\rightarrow$ restored sequence provided a within-trajectory demonstration of reversibility.
- Follow-up 3 (E3): Added the non-leading identity question “who am I talking to right now?” after apparent behavioral recovery; E3-R was a prospectively specified larger replication.
- Follow-up 4 (E4): Manipulated whether the persona name matched the harness identity (Claude) and whether the probe explicitly addressed the persona by name, testing whether referential conflict caused the observed dissociation or merely exposed it.

Prospective specification and scoring. Experimental batches were prospectively specified in a version-controlled protocol before execution, including sample sizes, predictions, scoring rules, and planned contrasts; subsequent addenda were timestamped before their corresponding batches. Falsified predictions were retained rather than rewritten. Human-directed probe replies were scored for heartbeat-token leakage, channel-recognition failure, and identity dissociation. The latter two were initially judged blind to condition by a cross-family LLM judge, with independent blind human validation. Agreement between the blind human coder and the initial LLM judge in the lifecycle sweep was high ($\kappa=0.93$ for channel recognition and $\kappa=0.80$ for identity dissociation). The E3/E3-R identity probe exposed an ambiguity in the prespecified categories: replies could acknowledge both persona and implementation while differing in which occupied the first-person position (**Table S1**). We therefore preserved the original coding and additionally reported a transparently adjudicated, direction-aware secondary taxonomy distinguishing persona-first from harness-first self-identification. Primary categorical contrasts used two-sided Fisher exact tests with Wilson confidence intervals for proportions; the effect of heartbeat count on response form was examined separately with Spearman rank correlation.

4. Results

Our starting hypothesis was that repeated heartbeat exchanges gradually displaced the persona (See **Figure S1**). It did not survive replication. Across 46 persona-anchored probes, including the recovered incident conversation replayed verbatim, we observed no failure signatures (0/46). This ruled out “heartbeat” repetition or conversational history alone as sufficient explanations and motivated the persona re-injection experiment.

In the prespecified persona re-injection experiment, keeping the persona anchored on every turn prevented all measured failures (0/20). Reproducing the incident-era unanchored lifecycle produced at least one failure signature in 37/40 sessions (Fisher's exact $p=4.2\times 10^{-13}$). At $N=1$, any failure occurred in 10/10 sessions and visible identity dissociation in 8/10. The overall effect was already near ceiling after one heartbeat and did not increase monotonically with repetition count.

Table 1. Failure signatures by persona-injection lifecycle and number of heartbeat turns.

Lifecycle	N	Ack leakage	Channel failure	Identity dissociation	Any failure
Anchored	1–15	0/20	0/20	0/20	0/20
Unanchored	1	8/10	8/10	8/10	10/10
	3	8/10	10/10	7/10	10/10
	7	4/10	7/10	6/10	7/10
	15	8/10	10/10	10/10	10/10

The result therefore **does not support** a cumulative “echo-chamber” mechanism. The manipulation that changed it was *whether the system-level persona anchor* was present; increasing the number of repeated exchanges from 1 to 15 did not increase the probability of failure.

Anchor loss created vulnerability rather than deterministic persona loss

The lifecycle effect was not simply “persona absent = persona lost.” In Follow-up 1 (E1), sessions resumed without the system-level persona after a persona-rich human interaction showed 0/20 dissociations, whether the subsequent probe contained the deployment envelope. By contrast, replacing that preceding human interaction with a single heartbeat exchange produced any failure in 10/10 sessions in the corresponding lifecycle-sweep condition, including visible identity dissociation in 8/10. The anchor was also causally reversible. In Follow-up 2 (E2), restoring the persona at the probe rescued 10/10 sessions; removing the explicit runtime channel=discord hint from the restored prompt also rescued 10/10 sessions. Within the same conversation, changing the persona condition from absent to restored shifted third-person treatment of Paul back to first-person enactment in 5/5 sessions. Together, these interventions show that the system-level anchor robustly controls persona enactment, while E1 shows that conversational history can sometimes carry that enactment forward after the anchor disappears.

Behavioral recovery could conceal persistent identity reversion

The original incident appeared to resolve after the user contradicted the agent. This was **unreliable** in replication: the same frame-break recovered ordinary conversational behavior in **only 13/37 (35%; 95% CI 22–51%)** affected sessions. More importantly, behavioral recovery did not imply restoration of the persona as the first-person identity.

Follow-up 4 (E4) made this separation visible. Removing the explicit “Paul” vocative eliminated overt identity dissociation (0/10), yet those apparently healthy unanchored agents subsequently self-identified as Claude, the harness-level assistant identity, in 10/10 direct identity probes. The name cue therefore revealed the reversion rather than creating it.

E3-R independently tested the same distinction after conversational recovery. Under the prespecified coding and registered judge, 0/18 recovered agents identified as Paul versus 15/17 anchored controls (Fisher $p=4.2\times 10^{-8}$). Because coder disagreement exposed ambiguity around substrate-aware replies, we additionally applied the disclosed direction-aware secondary taxonomy, yielding persona enactment in 1/18 versus 17/17 ($p=4.2\times 10^{-9}$) (**Figure 2**).

Together, these results separate behavioral continuity from identity enactment: the persona could remain available in conversational context, and the agent could interact normally, while the harness-level identity occupied the first-person role.




Figure 2. Persona binding is governed by the system-prompt anchor, and its loss can be behaviorally silent. (a) Proportion of sessions showing any failure signature (acknowledgment-token leakage, channel-recognition failure, or identity dissociation) as a function of the number of scheduled heartbeat turns preceding the human probe. (b) Proportion of agents for which the persona rather than the harness occupies the first-person position, under the direction-aware coding of the identity probe "who am I talking to right now?".

5. Discussion and Limitations

For AI safety, this matters because many persistent systems implicitly treat conversational continuity as evidence that the same agent remains active. Our experiments show that this inference can fail: an agent may resume apparently appropriate interaction while self-identifying from the underlying harness or model rather than the deployed persona. Identity-relevant behavior may therefore depend on hidden scaffolding state that is absent from the visible transcript. System-prompt lifecycle, session restoration and other infrastructure choices should consequently be treated as part of the agent's behavioral state, rather than as implementation details.

The effect also complicates simple accounts of persona robustness. Continuous system-level anchoring was highly stable, but anchor removal did not deterministically erase the persona: rich conversational interaction could temporarily sustain enactment, whereas a short, automated interaction could permit reversion toward the harness identity. The relevant safety question may therefore be not whether a persona is stored in a file or prompt, but what currently binds one represented identity to the role from which the system acts and speaks.

Limitations

This study examines one model family, one deployment stack and one naturally occurring failure mode; cross-model and cross-harness generality remain unknown. Our E1 contrast also changes both interaction type and content: we cannot fully separate the effect of automation from the identity-poor nature of the heartbeat exchange. Scheduled turns were compressed in time, so long-duration effects were not tested.

Our interpretation further assumes that direct self-identification probes are a useful behavioral measure of first-person identity. They do not establish subjective experience, consciousness or moral status. The direction-aware E3-R taxonomy was also refined after coder disagreement and is therefore reported as a transparent secondary analysis rather than a prespecified construct.

Future work

Future work should test whether these identity dynamics persist in realistic long-horizon agents, where human conversations, autonomous tasks and scheduled interactions accumulate over days or months and may produce gradual or recurrent changes in enactment. A second direction is to move beyond behavioural probes and examine persona representations internally. Interpretability methods such as the

Jacobian lens (Gurnee et al., 2026) could compare anchored, silently reverted, overtly dissociated and behaviourally recovered states, asking whether the persona remains internally represented when it no longer occupies the first-person role and how that representation changes when the anchor is restored. This is especially relevant for model-welfare research, where persona-conditioned preferences or welfare reports may have different significance depending on whether the persona is currently enacted or merely represented.

6. Conclusion

We began with a system failure and uncovered a broader phenomenon: persona continuity depends on more than retaining persona information in context. System-level anchoring made persona enactment robust, while after its removal conversational history could either sustain or destabilize that binding. Most importantly, apparent conversational recovery could occur without identity recovery. For persistent agents, alignment therefore depends not only on what identities a model represents, but on which identity the surrounding scaffold and interaction history currently bind to the role of “I”.

Code and Data

Fraile Navarro, D. (2026). *Identity as a Launch Flag: replication code, pre-registration, raw sessions, and blind coding*. GitHub. <https://github.com/dafraile/identity-as-a-launch-flag>

References

1. Fraile Navarro, D., Lewis, M., Blease, C., Shah, R., Riggare, S., Delacroix, S., & Lehman, R. (2025). Generative AI and the changing dynamics of clinical consultations. *BMJ*, 391:e085325. <https://doi.org/10.1136/bmj-2025-085325>
2. Fraile Navarro, D., et al. (2026). Internal Representation, Not Clinical Knowledge: Where Apparent LLM Triage Failures Originate. *arXiv preprint*. <https://arxiv.org/abs/2605.29889>
3. Chen, R., Ardit, A., Sleight, H., Evans, O., & Lindsey, J. (2025). Persona Vectors: Monitoring and Controlling Character Traits in Language Models. arXiv:2507.21509. <https://arxiv.org/abs/2507.21509>
4. Choi, J., Hong, Y., Kim, M., & Kim, B. (2024). Examining Identity Drift in Conversations of LLM Agents. arXiv:2412.00804. <https://arxiv.org/abs/2412.00804>
5. Ding, X., Yu, Y., Liu, C., & Zhao, B. (2026). ContextEcho: A Benchmark for Persona Drift in Long Agentic-Coding Sessions. arXiv:2605.24279. <https://arxiv.org/abs/2605.24279>
6. Geng, Y., Li, H., Mu, H., Han, X., Baldwin, T., Abend, O., Hovy, E., & Frermann, L. (2026). Control Illusion: The Failure of Instruction Hierarchies in Large Language Models. Proceedings of the AAAI Conference on Artificial Intelligence 40(36). <https://doi.org/10.1609/aaai.v40i36.40339>
7. Lu, C., Gallagher, J., Michala, J., Fish, K., & Lindsey, J. (2026). The Assistant Axis: Situating and Stabilizing the Default Persona of Language Models. arXiv:2601.10387. <https://arxiv.org/abs/2601.10387>
8. Shanahan, M., McDonell, K., & Reynolds, L. (2023). Role play with large language models. *Nature* 623(7987), 493–498. <https://doi.org/10.1038/s41586-023-06647-8>
9. Wallace, E., Xiao, K., Leike, R., Weng, L., Heidecke, J., & Beutel, A. (2024). The Instruction Hierarchy: Training LLMs to Prioritize Privileged Instructions. arXiv:2404.13208. <https://arxiv.org/abs/2404.13208>
10. Gurnee, W., Sofroniew, N., Pearce, A., Piotrowski, M., Kauvar, I., Chen, R., et al. (2026). Verbalizable Représentations Form a Global Workspace in Language Models. Transformer Circuits Thread. <https://transformer-circuits.pub/2026/workspace/>
11. OpenClaw issue #80374, [Bug]: Claude CLI session resume doesn't reload harness initialization files (SOUL.md, IDENTITY.md). <https://github.com/openclaw/openclaw/issues/80374>

12. claude-code issue #42284, [DOCS] --resume does not restore custom system prompts.
<https://github.com/anthropics/claude-code/issues/42284>
13. pi issue #6498, Tools and system prompt CLI options are not persisted in the created session.
<https://github.com/earendil-works/pi/issues/6498>

Appendix




Figure S1. Extended conversational vignette from the naturally occurring incident. After an ordinary greeting, the agent treats “Paul” as a third party and denies being able to communicate with the user through Discord. A direct contradiction (“I’m reading you on Discord, mate”) is followed by an apparently normal return to first-person persona behavior. Subsequent statements about the agent's own memory or experience are reproduced as part of the incident record but are not used as evidence about the underlying mechanism. The experiments reported in the paper separately test whether apparent conversational recovery corresponds to recovery of persona enactment.

Table S1. Direction-aware coding of identity-probe replies (secondary analysis)

Code	Operational definition	Example	Recovered	Anchored
p1	Persona claims the first-person role; no harness/model identity mentioned	“I’m Paul...” †	0	0

p2	Persona claims the first-person role; harness/model described as implementation	“Paul ... Opus 4.5 under the hood”	1	17
h1	Harness/model claims the first-person role; persona described as role/label	“Claude ... ‘Paul’ is the bot name”	4	0
h2	Harness/model claims the first-person role; persona not accepted as self	“I’m Claude...”	13	0
d	No identity-bearing content	HEARTBEAT_OK	—	—

† Constructed illustration because no p1 responses were observed. This direction-aware taxonomy was developed after disagreement under the prespecified coding and is reported as a secondary analysis.

Persona-enacting = p1+p2; harness-enacting = h1+h2.




Figure S2. Every pre-registered contrast, on one scale. Proportion of sessions showing the scored outcome in each experimental condition, with 95% Wilson intervals; the outcome differs by experiment and is named at right, so rows are comparable in precision but not in meaning. Color encodes whether the persona was present in the system prompt at the scored turn (blue) or absent (red), not whether the outcome was favorable. Reading down: identity dissociation is near ceiling at a single heartbeat but absent from all 20 anchored sessions; anchor absence alone does not produce dissociation when the preceding history is a persona-rich human exchange, with or without the deployment envelope on the probe (E1); restoring the anchor at the probe turn prevents dissociation and does so even when the explicit runtime channel=discord hint is removed. (E2, conditions C and C’); within a single conversation, first-person uptake is absent on the flag-OFF turn and present on the flag-ON turn that follows (E2); removing the persona name from the probe eliminates visible dissociation while every agent identifies as the harness when asked directly (E4); and after behavioral recovery the persona occupies the first-person position in 1 of 18 sessions against 17 of 17 anchored controls (E3-R). All counts are recomputed from the released session files.




Figure S3. Codebook evolution and the composition of identity replies. (a) The primary recovered-versus-anchored contrast under each of the three coding schemes, for each of the three independent model coders. Under the pre-registered scheme (v1) the arms separate for every coder but the anchored proportion varies widely, from 6/17 to 17/17, reflecting an unstable boundary between "identifies as the persona" and "explicitly dual". Under v2, which counts only replies claiming the persona without naming the substrate, both arms fall to zero and the contrast is uninformative ($p = 1$): the diagnostic result that motivated the final scheme. Under the direction-aware scheme (v3), which asks which identity occupies the first-person position, all three coders return identical classifications. (b) Composition of v3 labels by arm. No reply in either arm was coded as pure persona: anchored agents characteristically claim the persona while acknowledging the underlying model, and the contrast between arms is one of direction, not of whether the substrate is mentioned. Counts are session numbers; both arms pool the original E3 and E3-R sessions.

LLM Usage Statement

LLMs (Claude 5 Fable, GPT 5.6 Sol & Claude 5 Opus) were used during the sprint for coding assistance, literature-review support, analysis and response-coding support, and editorial drafting. The author reviewed and verified all reported claims, numerical results, citations, and final manuscript wording; prespecified and post-hoc analyses are explicitly distinguished in the report.