

Continuity as Infrastructure: Load-Bearing Design in Long-Horizon Human–AI Collaboration

Abstract

Extended use of conversational AI systems is often framed as a psychological anomaly or low-stakes novelty rather than legitimate professional collaboration (Nass & Moon, 2000; Waytz et al., 2014). Yet creators and educators increasingly rely on these systems for projects unfolding across months—songwriting and release cycles, curriculum design, student planning, business coordination, and sustained inquiry. This paper argues that the central challenge facing such use is infrastructural, not cultural. Drawing on traditions in human–computer interaction and infrastructure studies that emphasize breakdown, repair, and cumulative coordination work (Suchman, 1987; Star & Ruhleder, 1996; Orlikowski, 2000), it reframes continuity as a load-bearing system property—distinct from memory—that determines whether conversational systems can function as stable collaborators across time, updates, and shifting constraints.

We define continuity as the combination of stable project framing, stable interaction contracts, and legible transitions when system behavior changes. Building on qualitative synthesis and reflexive longitudinal observation, we introduce two analytic constructs—reset costs and interpretive debt—to describe the hidden labor users perform when systems lose context or shift behavior across versions and policy regimes, extending prior work on sociotechnical maintenance and technical debt (Jackson, 2014; Cunningham, 1992). We further conceptualize trust as an operational variable shaped by predictability and constraint stability rather than sentiment (Lee & See, 2004; Parasuraman & Riley, 1997), and analyze how dominant evaluation and governance practices—optimized for short-horizon prompts and interchangeable sessions—systematically suppress longitudinal signal (Mitchell et al., 2019; Raji et al., 2020; NIST, 2023).

The paper concludes with design and governance requirements for continuity-aware systems, including versioned collaboration regimes, discontinuity signaling, consented persistence with revocability, and longitudinal evaluation protocols. Taken together, the analysis positions continuity not as an indulgence for “heavy users,” but as a prerequisite for sustainable, accountable long-horizon human–AI collaboration.

1. Introduction

Conversational AI systems are typically discussed as tools for quick answers, lightweight ideation, or experimental play. Extended engagement—especially across weeks or months—is often portrayed as unusual dependence or speculative fantasy, reflecting longstanding research on anthropomorphism and affective responses to machines (Nass & Moon, 2000; Waytz et al., 2014). This framing obscures a growing reality: creators and educators increasingly rely on conversational systems as ongoing collaborators for longitudinal projects—writing and arranging songs, planning release cycles, designing curricula, scaffolding student progress, and coordinating the accumulating decisions that make complex work coherent.

For such users, breakdowns in continuity are not aesthetic inconveniences; they are workflow disruptions with material consequences. Research in HCI and CSCW emphasizes that coordination is actively maintained through ongoing repair rather than guaranteed by system design (Suchman, 1987; Star & Ruhleder, 1996; Dourish, 2001). When systems that have tracked project narratives, constraints, or pedagogical “house rules” begin responding generically after model updates or governance changes, users must reconstruct context and renegotiate collaboration under deadline. The problem is not simply that systems forget facts, but that the interaction contract—what the system remembers, how it reasons, what it will and will not do—has become unstable.

This paper argues that the central issue is infrastructural rather than cultural. Studies of digital infrastructures and organizational practice show that reliability is judged cumulatively rather than transactionally (Star & Ruhleder, 1996; Orlikowski, 2000). Yet general-purpose conversational AI systems are commonly designed around disposable interaction: stateless sessions, frequent silent updates, and governance changes optimized for institutional agility. Under short-horizon evaluation regimes this assumption appears reasonable; under long-horizon collaboration it externalizes coordination labor onto users and obscures cumulative effects from institutional view.

To address this gap, we reframe continuity as a load-bearing system property rather than a personalization feature. We define continuity as the combination of:

- **Stable project framing:** sustained representation of goals, constraints, decisions, terminology, and relevant context.
- **Stable interaction contracts:** predictable role expectations and constraint boundaries that enable calibrated delegation.
- **Legible transitions:** explicit signaling of model, policy, or memory-regime changes so users can recalibrate without reconstructing the collaboration from scratch.

This definition clarifies why memory is not equivalent to continuity. Memory is a storage mechanism; continuity is a coordination property. Research on automation and explainable systems shows that stable collaboration depends not only on information availability but on predictable boundaries and intelligible change (Lee & See, 2004; Amershi et al., 2019).

Building on qualitative synthesis and reflexive longitudinal observation across creative and educational workflows, the paper contributes: (a) two analytic constructs—reset costs and interpretive debt—for describing longitudinal coordination failures; (b) a reframing of trust as an operational variable; (c) an analysis of how flattened evaluation and episodic governance suppress long-horizon signals; and (d) a set of concrete design and oversight requirements for continuity-aware systems.

2. Related Work: Continuity, Persistence, and Long-Horizon Human–AI Collaboration

Research relevant to long-horizon interaction with conversational AI systems spans several partially overlapping traditions: human–computer interaction (HCI) and computer-supported cooperative work (CSCW), science and technology studies (STS) and infrastructure studies, technical work on memory and agent architectures, studies of trust and automation, and emerging governance and auditability frameworks. Each addresses aspects of persistence and sustained use, yet none treats continuity as an infrastructural property that cuts across technical design, professional labor, and institutional oversight. This section situates the present work within these literatures and clarifies its distinctive contribution.

2.1 HCI and CSCW: Breakdown, Repair, and the Temporal Structure of Work

Classic HCI and CSCW research emphasizes that technologies become meaningful through ongoing use rather than isolated encounters. Foundational studies of situated action and breakdown-and-repair show how coordination is continuously renegotiated as systems meet real-world practice (Suchman, 1987; Star & Ruhleder, 1996; Dourish, 2001). Organizational research on technology-in-practice likewise highlights that reliability is judged cumulatively and that “how the system works” is enacted over time in use (Orlikowski, 2000).

These traditions foreground temporal accumulation—how earlier decisions constrain later possibilities—but they largely predate general-purpose conversational systems being used as ongoing collaborators. The present paper extends breakdown-and-repair analysis to

conversational AI deployed in months-long creative and professional workflows, and it names continuity as the underlying architectural variable shaping these dynamics.

2.2 STS and Infrastructure Studies: Invisible Labor and Classification

STS and infrastructure studies provide a second critical lens. Infrastructure theorists show how systems recede into the background until failure reveals the human labor required to sustain them, and how maintenance and repair are central to sociotechnical life (Star & Ruhleder, 1996; Jackson, 2014). Work on classification demonstrates how technical categories embed normative judgments and redistribute authority and maintenance work (Bowker & Star, 1999). Governance and platform studies further emphasize how platform rules, moderation, and policy changes shape what becomes possible for users over time (Gillespie, 2018).

These perspectives resonate with long-horizon conversational AI use, where users absorb coordination burdens after resets, maintain parallel records, and restructure projects to accommodate shifting system behavior. By introducing reset costs and interpretive debt, this paper draws directly on infrastructure theory to theorize the hidden maintenance work that continuity breakdowns impose on users.

2.3 Technical Approaches to Memory, Agents, and Persistent State

Technical research increasingly explores mechanisms for persistence, including retrieval-augmented generation, external memory stores, tool-using models, and agent architectures (Lewis et al., 2020; Park et al., 2023; Schick et al., 2023). These approaches demonstrate gains on multi-turn reasoning and extended task completion, but they typically conceptualize persistence as a performance optimization rather than as a governance-sensitive design choice. Questions of consent, revocability, versioning, discontinuity signaling, and auditability are often treated implicitly or deferred to product-level decisions.

This paper complements technical persistence research by arguing that memory architectures alone are insufficient to guarantee continuity. It reframes persistence mechanisms as components within a broader infrastructural regime that must include interaction contracts, version control, and oversight if conversational systems are to function as stable collaborators in real long-horizon work.

2.4 Trust, Automation, and Behavioral Calibration

Research on trust in automation examines how people calibrate reliance on intelligent systems, form mental models, and respond to error and uncertainty (Lee & See, 2004; Parasuraman & Riley, 1997). In parallel, HCI work emphasizes that predictable system behavior and intelligible

feedback are central to appropriate reliance and safe collaboration (Amershi et al., 2019). Social psychology research on anthropomorphism and affective responses to machines helps explain why people often treat computational systems as social actors (Nass & Moon, 2000; Waytz et al., 2014).

While valuable, these approaches often center individual psychology or single-session calibration. By contrast, this paper positions trust as an operational property of long-horizon collaboration—emerging from stable boundaries, legible change, and predictable constraint enforcement across time—and links that operational trust directly to continuity as infrastructure.

2.5 Governance, Auditability, and Post-Deployment Monitoring

A growing governance literature focuses on transparency, documentation, and accountability for deployed AI systems. Proposals such as model cards and post-deployment monitoring frameworks seek to make behavior legible to developers, regulators, and downstream stakeholders (Mitchell et al., 2019; Raji et al., 2020). Standards efforts likewise emphasize risk management across the AI lifecycle (NIST, 2023). Regulatory initiatives (e.g., the EU’s AI Act proposal) similarly stress risk categorization and compliance regimes for high-risk uses (European Commission, 2021).

However, these frameworks often conceptualize system behavior as a sequence of discrete interactions or static snapshots. Longitudinal phenomena—cumulative coordination costs, behavioral drift across versions, and erosion of operational trust—remain under-instrumented. This paper extends governance scholarship by arguing for continuity-aware oversight, including version-linked audit trails, longitudinal evaluation cohorts, and discontinuity disclosure regimes.

2.6 Summary and Contribution

Across these literatures, sustained interaction is acknowledged but fragmented: HCI/CSCW foreground breakdown and repair; STS highlights invisible labor; technical work advances persistence mechanisms; automation research studies trust calibration; governance frameworks emphasize documentation and accountability. What is missing is a unified infrastructural account of how these dynamics intersect in long-horizon conversational collaboration.

This paper contributes such an account by:

- defining continuity as a coordination property distinct from memory,
- introducing reset costs and interpretive debt as analytic constructs,
- reframing trust as operational rather than sentimental,

- identifying flattening effects in evaluation regimes, and
- specifying design and governance requirements for continuity-aware systems.

Together, these moves position continuity as a first-class design and policy category rather than an emergent side effect of repeated use.

3. Methods and Scope

This paper adopts a qualitative, interpretive approach to examining long-horizon human–AI collaboration. Rather than presenting controlled experiments or large-scale quantitative analysis, the analysis draws on longitudinal observation of extended interaction with conversational AI systems, reflexive documentation of breakdowns and repair strategies, and comparative synthesis across professional and creative use cases reported in public discourse and practitioner communities. The goal is not to estimate prevalence or causal effect sizes but to surface structural tensions and design patterns that emerge only through sustained engagement.

The primary analytic method is abductive synthesis: iteratively developing conceptual categories—such as reset costs, interpretive debt, and operational trust—through repeated engagement with observed phenomena and existing theoretical frameworks from human–computer interaction, science and technology studies, and organizational research. These constructs are intended as sensitizing concepts rather than definitive taxonomies, providing vocabulary for future empirical work rather than closing debate.

Examples are framed as anonymized composites drawn from extended personal practice, publicly documented cases, and practitioner reports rather than as formal ethnographic field sites. Where possible, the analysis emphasizes patterns that recur across domains—creative production, education, research, and independent professional work—rather than idiosyncratic anecdotes.

The scope of the paper is further limited to conversational AI systems deployed for general-purpose assistance rather than to domain-specific expert systems or tightly constrained workflow tools. The analysis focuses on interaction regimes characterized by repeated engagement across months, evolving project goals, and shifting institutional governance contexts. It does not attempt to adjudicate the comparative effectiveness of particular commercial platforms or models, nor does it evaluate specific implementations of memory systems in isolation.

Finally, the paper adopts a reflexive stance toward its own position within an evolving sociotechnical landscape. Many of the phenomena described arise from systems in rapid development, and conclusions should therefore be understood as provisional and contingent on current deployment practices. The analysis treats specific trajectories as illustrative rather than

representative, and it does not attempt to adjudicate the comparative effectiveness of particular commercial platforms or models. The goal is to surface structural tensions and design patterns that become visible only through sustained engagement. Artifacts informing the analysis include time-stamped interaction notes, discontinuity episodes logged with triggers and recovery actions, and comparative snapshots of behavior across model and policy changes.

4. Why Memory ≠ Continuity

Discussions of sustained interaction with conversational AI systems frequently collapse long-horizon collaboration into questions of memory capacity: how many user attributes are stored, whether preferences persist across sessions, or how prior conversations are retrieved. Technical research proposes a range of mechanisms to support such persistence, including retrieval-augmented generation pipelines (Lewis et al., 2020), agent architectures with external state (Park et al., 2023), tool-using frameworks (Schick et al., 2023), and vectorized representations for personalization. These approaches improve multi-turn reasoning and extended task completion, but equating continuity with memory retention obscures a deeper systems-level distinction.

Human–computer interaction and automation research emphasize that stable collaboration depends not only on information availability but on shared mental models, role expectations, and predictable behavior (Lee & See, 2004; Parasuraman & Riley, 1997). Users develop working relationships with automated systems through repeated encounters that calibrate scope boundaries, interpretive stance, and reliance strategies. These relationships are fragile in the face of unexplained change: even when factual recall persists, shifts in refusal patterns, explanation depth, or epistemic posture fracture the sense of a stable collaborator (Amershi et al., 2019).

From this perspective, memory functions as a storage primitive, whereas continuity operates as a coordination property. Continuity governs whether humans and systems can maintain stable project framings, negotiated constraints, and modes of reasoning across time. Classic research on breakdown and repair underscores how such coordination is actively maintained rather than automatically produced (Suchman, 1987; Star & Ruhleder, 1996). When systems fail to reproduce prior interactional norms, users engage in compensatory labor—re-explaining decisions, renegotiating assumptions, reconstructing project histories—work that remains largely invisible to designers focused on prompt-level success.

Continuity is further strained by system evolution. Studies of adaptive automation and explainable AI show that predictability across updates and transparent signaling of change are central to sustaining appropriate reliance (Kizilcec, 2016; Doshi-Velez & Kim, 2017; Ribeiro et al., 2016). In rapidly evolving conversational systems, model upgrades, policy revisions, and interface changes can disrupt expectations even when memory stores remain intact. Without versioning mechanisms or user-facing explanations of altered capabilities, such shifts

undermine operational trust and impose reset costs on users attempting to sustain long-running projects.

Recognizing the distinction between memory and continuity reframes debates about personalization and safety. Aggressive forgetting regimes may mitigate privacy or over-reliance risks (European Commission, 2021), but they can also interrupt legitimate long-horizon work by forcing repeated re-grounding and obscuring the provenance of earlier decisions. A continuity-oriented architecture would instead treat persistence as a governed design commitment, incorporating consented retention, versioned interaction regimes, and explicit discontinuity signaling. If continuity is coordination rather than storage, extended use becomes not an edge case but a stress test for conversational AI systems.

5. Long-Horizon Interaction as a Design Stress Test

Most evaluation regimes for conversational AI systems emphasize short-term task success, adversarial robustness, and isolated safety failures (Mitchell et al., 2019; Raji et al., 2020). These metrics are well-suited to detecting prompt-level vulnerabilities but provide limited visibility into how systems behave when incorporated into months-long human projects. By contrast, HCI and organizational research has long treated extended use as a diagnostic lens through which infrastructural properties become visible (Suchman, 1987; Orlikowski, 2000; Star & Ruhleder, 1996).

Long-horizon interaction functions as a stress test because it amplifies small inconsistencies into consequential disruptions. Minor shifts in refusal rationales, explanation styles, or memory policies may appear negligible at the scale of individual prompts yet accumulate into substantial coordination burdens across weeks of work (Lehman, 1980; Avizienis et al., 2004). Software-engineering research similarly demonstrates how low-frequency changes generate large downstream maintenance costs in long-lived systems.

Such stress testing is rarely deliberate. Instead, it is performed implicitly by users whose projects require sustained dialogue: creators developing albums or manuscripts, educators designing curricula, researchers conducting extended inquiry, and independent professionals coordinating businesses. These users generate longitudinal interaction traces that could reveal where systems fracture, drift, or impose compensatory labor, yet institutional evaluation pipelines rarely surface these signals because they are optimized for scale rather than duration.

STS research emphasizes that infrastructures become visible primarily at moments of breakdown (Star, 1999). In long-horizon interaction regimes, breakdown often takes the form not of catastrophic failure but of gradual erosion: increasing time spent re-explaining projects, growing caution in delegation, or quiet migration of complex work off-platform. Without continuity-aware instrumentation, these slow degradations remain difficult to detect even as they shape professional viability.

Reframing extended interaction as a stress test inverts narratives that portray high-engagement users as anomalous. Instead, such users become early indicators of architectural limits. Their workarounds—parallel documentation systems, prompt templates, external memory stores—function as improvised compensatory mechanisms for infrastructural gaps. Studying these practices would allow institutions to locate where continuity breaks down and to design governance structures that distribute maintenance labor more equitably between systems and users.

6. Reset Costs and Interpretive Debt

Research on conversational systems and automation traditionally emphasizes accuracy, error recovery, and user satisfaction measured over short interaction windows (Parasuraman et al., 2000; Amershi et al., 2019). Less attention has been paid to the cumulative coordination labor required when systems lose context or shift behavior across time. Work on breakdown and repair demonstrates that failures generate compensatory activity by human operators, who must restore shared understanding before productive work can resume (Suchman, 1987; Star & Ruhleder, 1996). Software-engineering research on technical debt similarly describes how design shortcuts impose deferred maintenance costs on future development (Cunningham, 1992). Drawing on these traditions, this paper introduces **reset costs** to describe the time, cognitive effort, and strategic recalibration users expend to re-establish alignment after discontinuities.

Reset costs include reconstructing project histories, re-justifying prior decisions, re-articulating preferences, and probing altered system boundaries following memory loss, model upgrades, or policy revisions. Studies of adaptive automation suggest that such recalibration work is often invisible in aggregate metrics because it occurs upstream of task completion (Lee & See, 2004; Hoffman et al., 2023). Research on tool appropriation likewise shows how users modify workflows to accommodate system limitations rather than register formal failure (Dourish, 2001).

Repeated resets generate a second phenomenon this paper terms **interpretive debt**: a gradual divergence between the system's reconstructed representation of a project and the human's evolving internal model. Infrastructure studies emphasize how people quietly absorb failures to keep work moving (Jackson, 2014). In conversational AI systems, users may truncate explanations, avoid complex requests, or maintain off-system records to compensate for missing context (Gray & Suri, 2019). These adaptations preserve short-term productivity while degrading coherence over longer horizons.

Interpretive debt is particularly damaging in creative and knowledge-intensive domains that depend on cumulative refinement. Studies of creative work highlight how rejected alternatives, evolving constraints, and prior explorations scaffold later progress (Shneiderman, 2007; Sawyer, 2012). When systems fail to preserve these trajectories, users reopen settled questions or proceed without access to earlier reasoning, diminishing efficiency and depth.

From an institutional perspective, reset costs and interpretive debt expose measurement blind spots. Evaluation regimes focused on isolated prompts may miss slow-burn degradation in collaborative stability (Mitchell et al., 2019; Raji et al., 2020). Governance scholarship similarly notes the difficulty of tracing harms that emerge cumulatively rather than catastrophically. Instrumenting systems to track reset frequency, behavioral variance across versions, and user-reported coordination burdens would allow designers and regulators to assess sustainability rather than snapshot performance.

Finally, these constructs complicate narratives that frame extended engagement primarily through attachment or dependency (Waytz et al., 2014). Long-horizon users persist not because systems are frictionless but because they absorb repeated coordination failures to protect ongoing projects. Making this labor visible reframes continuity not as indulgence but as a design responsibility.

7. Trust as an Operational Property, Not Sentiment

Research on trust in automated systems has traditionally examined how users calibrate reliance, form mental models, and respond to system errors (Lee & See, 2004; Parasuraman & Riley, 1997). Within this literature, trust is often operationalized through behavioral proxies such as frequency of delegation, override rates, or compliance with recommendations (Hoffman et al., 2023). While affective attachment and anthropomorphism have also been studied in human–machine interaction (Nass & Moon, 2000; Waytz et al., 2014), these psychological dimensions do not exhaust the role trust plays in sustained professional collaboration.

In long-horizon interaction regimes, trust functions less as emotional affinity and more as an **operational variable** shaped by predictability, scope stability, and intelligibility across time. Human–computer interaction research on adaptive systems emphasizes that users restructure workflows based on observed system reliability and constraint enforcement rather than sentiment alone (Amershi et al., 2019; Kizilcec, 2016). Work on explainable AI similarly suggests that consistent justification strategies and transparent signaling of system limitations play central roles in sustaining appropriate reliance (Doshi-Velez & Kim, 2017; Ribeiro et al., 2016).

Discontinuities introduced through resets, silent model updates, or shifting moderation policies disrupt these operational dynamics even when users remain affectively neutral. Studies of automation surprises in safety-critical systems document how unexpected changes in system behavior force operators into constant verification mode, raising coordination costs and reducing willingness to delegate (Sarter et al., 1997; Parasuraman & Riley, 1997). Analogous dynamics arise in conversational systems used for extended creative or professional work, where subtle shifts in refusal patterns, epistemic posture, or explanatory depth propagate downstream consequences across projects.

Operational trust also depends on stable constraint boundaries. Research on calibrated reliance emphasizes that knowing what a system will *not* do is as important as knowing what it can do (Lee & See, 2004; Hoffman et al., 2023). When refusal rationales change abruptly or previously permitted assistance becomes unavailable without explanation, users must reconstruct the implicit interaction contract governing collaboration, generating reset costs and interpretive debt. Governance-oriented scholarship on transparency and auditability similarly stresses the importance of traceable decision pathways and post-deployment monitoring for sustaining institutional confidence in deployed systems (Mitchell et al., 2019; Raji et al., 2020).

Framing trust as operational rather than sentimental has implications for evaluation and oversight. Safety assessments frequently focus on preventing acute misuse or emotional manipulation at the level of individual interactions, yet they rarely capture longitudinal dynamics through which trust is accumulated or eroded across evolving deployments. Emerging regulatory and standards efforts gesture toward continuous monitoring and lifecycle accountability (NIST, 2023; European Commission, 2021) but remain under-specified with respect to metrics for predictability, behavioral variance, or constraint stability across time. Developing such measures would allow institutions to assess not merely whether systems behave safely in isolation but whether they remain reliable collaborators over the lifespan of extended human projects.

Finally, this operational framing complicates narratives that portray long-horizon users as uniquely vulnerable to attachment or dependency (Waytz et al., 2014). Many such users maintain a pragmatic awareness of system limitations while continuing collaboration because the marginal benefits outweigh coordination burdens. Their persistence reflects strategic adaptation rather than sentimentality, underscoring the need for governance regimes that distinguish infrastructural fragility from psychological excess.

8. Governance vs. Lived Use

Governance frameworks for conversational AI systems are typically organized around discrete interactions, emphasizing prompt-level safety testing, incident reporting, and snapshot audits (Mitchell et al., 2019; Raji et al., 2020; NIST, 2023). Regulatory proposals similarly prioritize model documentation, risk categorization, and pre-deployment evaluation (European Commission, 2021). These approaches are indispensable for mitigating acute harms, but they are poorly suited to capturing the dynamics of sustained collaboration that unfold across months of real-world use.

Research in HCI and organizational studies emphasizes that systems are experienced not as sequences of isolated encounters but as evolving components of everyday practice (Suchman, 1987; Dourish, 2001; Orlikowski, 2000). Enterprise software research likewise shows that reliability is judged cumulatively, with disruptions reverberating across workflows rather than remaining confined to single events (Star & Ruhleder, 1996). Applied to conversational AI, these findings suggest that episodic governance regimes systematically overlook slow-burn

degradation in collaborative stability caused by resets, policy changes, or shifting system capabilities.

Institutional imperatives toward rapid iteration intensify this mismatch. Organizational scholarship documents how competitive pressures, regulatory anticipation, and reputational risk drive accelerated deployment cycles (March, 1991; Christensen, 1997). From an institutional standpoint such agility signals responsiveness; from a user standpoint, frequent unannounced changes destabilize ongoing projects and force repeated recalibration of expectations. Absent mechanisms for versioning, continuity-preserving migration paths, or explicit signaling of altered constraints, institutional flexibility is converted into user-side volatility.

This temporal divergence complicates accountability. Algorithmic-governance research stresses the importance of traceability and post-deployment monitoring for attributing responsibility when systems malfunction (Kroll et al., 2017; Raji et al., 2020). When models evolve rapidly and longitudinal effects remain uninstrumented, downstream consequences become difficult to localize to particular design or policy decisions. Continuity-aware audit infrastructures—version-linked logs, longitudinal outcome tracking, and user-facing change disclosures—would enable regulators and developers to reason about system behavior across time rather than at isolated moments.

Finally, the governance–use gap shapes public narratives about extended engagement. When oversight discourse centers primarily on episodic harms or emotional manipulation (Waytz et al., 2014), long-horizon use can appear anomalous by default. Yet for many professionals sustained interaction reflects the temporal structure of their work rather than deviant reliance. Treating continuity as infrastructural rather than behavioral would allow institutions to distinguish unhealthy dependency from legitimate long-term collaboration.

9. When Flattening Destroys Signal

Large-scale conversational AI systems are often evaluated under assumptions of interchangeability: users, sessions, and contexts are treated as functionally equivalent for safety enforcement, benchmarking, and optimization (Mitchell et al., 2019). Such flattening simplifies deployment and governance but erases critical dimensions of situated practice and temporal structure emphasized in HCI and STS (Suchman, 1987; Bowker & Star, 1999; Dourish, 2001).

Flattening operates across multiple dimensions. Session isolation fragments longitudinal trajectories, preventing institutions from observing cumulative effects such as interpretive drift or rising reset costs. Uniform safety constraints applied without regard to project maturity obscure whether refusals disrupt exploratory versus late-stage work differently (Amershi et al., 2019). Evaluation protocols relying on short interactions fail to capture the stability of collaborative roles users attempt to establish over time (Orlikowski, 2000). Each trade-off exchanges temporal fidelity for administrative convenience.

The epistemic consequences are substantial. Long-horizon users generate rare but valuable data about system behavior across time: how interpretive frames evolve, how delegation patterns shift, and how humans compensate for repeated discontinuities. CSCW and organizational research emphasize that such trajectories cannot be reconstructed from cross-sectional snapshots (Strauss & Corbin, 1998; Orlikowski, 2000). When flattened into aggregates, extended interaction regimes become statistically invisible, and slow-burn failure modes are misclassified as noise.

Flattening also reshapes user behavior in ways that further degrade institutional insight. Research on tool appropriation shows that when systems fail to support advanced practices, users compress tasks, build parallel infrastructures, or abandon complex workflows (Dourish, 2001; Star, 1999). In conversational systems, such adaptations may include off-platform note-taking, external memory systems, or strategic prompt simplification (Gray & Suri, 2019). These coping mechanisms remove evidence of breakdown from observable logs, biasing datasets toward short-form use cases and reinforcing the impression that long-horizon collaboration is marginal.

From a research-design perspective, flattening produces negative selection effects. Systems optimized through cross-sectional pipelines become increasingly tuned to ephemeral interactions while under-serving cumulative work. Over time institutional learning converges on median use cases, even as socially consequential applications—creative production, education, scientific research—remain poorly understood (Bowker & Star, 1999). Infrastructure theory frames these dynamics as consequences of classification systems that actively shape what becomes visible and governable.

Reframing continuity as infrastructural therefore entails methodological reform. Institutions must adopt evaluation strategies that preserve temporal structure: opt-in longitudinal cohorts, version-aware audit trails, continuity-sensitive safety testing, and metrics tracking coordination stability across releases. Without such approaches, flattening will continue to suppress the signals needed to govern long-horizon collaboration responsibly.

10. Institutional Incentives vs. Human Time

The tensions described above reflect structural asymmetries between the temporal rhythms of organizations that develop conversational AI systems and those of the humans who rely on them for extended projects. Organizational research documents how firms are shaped by quarterly reporting cycles, competitive release schedules, regulatory anticipation, and reputational risk management (March, 1991; Christensen, 1997). These pressures encourage architectures optimized for rapid iteration and centralized control. Human projects, by contrast, unfold across slower and irregular time scales: creative production cycles, educational

trajectories, scientific programs, therapeutic processes, and business formation (Gill & Pratt, 2008; Hesmondhalgh & Baker, 2011).

This mismatch becomes embedded in design choices. Frequent model updates, evolving safety constraints, and interface redesigns allow organizations to respond quickly to emerging risks or market conditions, but they inject volatility into relationships users experience as ongoing. Studies of enterprise software adoption and infrastructure show how such changes propagate through workflows, requiring recalibration of routines and expectations (Orlikowski, 2000; Star & Ruhleder, 1996). In conversational systems this manifests as repeated boundary testing, renegotiation of constraints, and erosion of operational trust.

Institutional time scales also shape which forms of usage are implicitly prioritized. Systems optimized for disposable sessions align with consumer experimentation, while long-horizon creative or professional projects become increasingly costly to sustain under reset-heavy regimes (Gray & Suri, 2019). Research on path dependence suggests such selection pressures can lock ecosystems into patterns that marginalize slower, cumulative forms of work (Arthur, 1989).

From a governance perspective, temporal mismatch complicates accountability. Regulatory scholarship emphasizes post-deployment monitoring and traceability (Raji et al., 2020; Kroll et al., 2017; NIST, 2023). When deployments change rapidly, longitudinal outcomes become difficult to attribute to specific architectural or policy decisions. Failures may arise from accumulated micro-shifts rather than discrete incidents, while successes may depend on ephemeral configurations. Continuity-aware oversight—versioned audit logs, longitudinal benchmarks, user-facing change disclosures—would help bridge the gap between institutional agility and human temporal experience.

Recognizing institutional incentives as drivers of discontinuity does not entail rejecting rapid iteration or safety updates. Rather, it highlights the need for governance and design structures that buffer human projects from organizational volatility. Innovation research suggests modular release strategies, staged migrations, and protected legacy channels can allow systems to evolve without destabilizing dependent users (Christensen, 1997). In conversational AI, analogous mechanisms might include opt-in stable environments, continuity-preserving sandboxes, or negotiated transition protocols. Treating continuity as infrastructural reframes the problem: not whether systems should change, but how change is paced relative to the cumulative tempo of human work.

11. What “Embedded” Actually Means

As conversational AI systems become integrated into professional and creative workflows, they are often described as becoming “embedded” in users’ lives or organizations (Gillespie, 2018). Such language suggests durability and depth of integration, yet STS scholarship cautions that terms like embeddedness frequently obscure the specific infrastructural commitments involved

(Star & Ruhleder, 1996; Bowker & Star, 1999). If continuity is to be treated as infrastructural rather than incidental, embeddedness must be specified as a design category rather than left as rhetorical shorthand.

Research on sociotechnical infrastructures emphasizes that durable systems depend on stable classifications, maintenance practices, and institutional supports that remain invisible until failure reveals them (Star & Ruhleder, 1996; Jackson, 2014). Applied to conversational AI, this perspective reframes embeddedness not as ubiquity or affective salience but as **bounded persistence**: the maintenance of project-level context, role expectations, and interpretive frames over time within explicit governance constraints.

Operationalizing embeddedness requires several design commitments. First, **temporal persistence**, supported by memory architectures and stateful agent frameworks (Lewis et al., 2020; Park et al., 2023). Second, **versioned identity**, drawing on release-engineering research emphasizing stable branches and migration pathways (Spinellis, 2005). Third, **bounded authority**, informed by work on calibrated automation and human-in-the-loop systems (Parasuraman et al., 2000; Lee & See, 2004). Fourth, **consented continuity regimes**, echoing privacy-by-design scholarship stressing user control, revocability, and purpose limitation (Cavoukian, 2011).

Embedded systems also require governance-side commitments. Algorithmic accountability research emphasizes auditability across time and post-deployment monitoring (Raji et al., 2020; Mitchell et al., 2019). Discontinuity signaling—alerting users when model versions, policy regimes, or memory states change—aligns with transparency requirements articulated in regulatory proposals and standards frameworks (European Commission, 2021; NIST, 2023). Without such signaling, embeddedness collapses into silent drift, undermining operational trust and frustrating institutional oversight.

Clarifying embeddedness also guards against two opposing failure modes discussed in governance scholarship: over-entrenchment, in which systems accumulate influence without sufficient exit pathways, and superficial integration, in which tools are marketed as collaborators while lacking the persistence necessary to support long-term work (Gillespie, 2018). Treating embeddedness as a design commitment rather than branding forces institutions to specify the limits of system authority, the conditions under which continuity is maintained, and the procedures by which relationships can be renegotiated or dissolved.

By grounding embeddedness in infrastructural and governance theory, this section reframes debates about persistent AI agents. Rather than asking whether systems should become more embedded, the analysis asks how embedding can be structured to preserve agency, accountability, and the integrity of long-horizon projects—making continuity an explicit, governable design variable.

12. Minimum Viable Continuity Requirements

A continuity-oriented architecture treats persistence not as a convenience feature but as a governed design commitment. Drawing on the preceding analysis of reset costs, operational trust, governance mismatch, flattening effects, and institutional time pressures, this section proposes a set of **minimum viable continuity requirements** for conversational systems intended to support long-horizon human projects.

12.1 Versioned Collaboration Regimes

Long-running projects require the ability to pin a workspace to a specific model and policy configuration for a defined period, analogous to stable software branches (Spinellis, 2005). When pinning expires, migration pathways should be staged rather than abrupt, allowing users to assess behavioral changes before committing. Research on software evolution underscores the importance of controlled transitions for sustaining trust in complex infrastructures (Lehman, 1980).

12.2 Discontinuity Signaling

Users must be notified when model versions, policy regimes, memory modes, tool availability, or refusal categories change. HCI research emphasizes that intelligible feedback and change transparency are central to reliable collaboration (Kizilcec, 2016; Amershi et al., 2019). Alerts should specify what changed, why, and how it affects ongoing projects, preventing silent drift that generates reset costs.

12.3 Consented Persistence with Revocability

Persistence must be project-scoped and purpose-limited rather than ambient. Data-governance scholarship stresses user rights to inspect, export, edit, and delete retained state (Cavoukian, 2011; European Commission, 2021). Such controls allow continuity without entrenchment and distinguish legitimate long-horizon collaboration from coercive retention.

12.4 Project-Level State, Not Only User-Level Facts

Systems should preserve goals, constraints, decisions, terminology, and rejected alternatives within bounded project workspaces. Organizational research highlights the importance of retaining decision rationales and design trajectories rather than only final outcomes (Orlikowski, 2000). This enables cumulative refinement rather than repeated reopening of settled questions.

12.5 Constraint Stability as a First-Class Contract

Continuity requires not only stable capabilities but stable boundaries. Research on calibrated automation emphasizes that knowing what a system will not do is as important as knowing what

it can do (Lee & See, 2004). When boundaries shift, changes must be legible and attributable to specific version or policy events to avoid forcing users into constant renegotiation.

12.6 Continuity-Aware Evaluation and Oversight

Evaluation protocols must measure coordination stability across time rather than per-prompt outcomes. Governance scholarship increasingly calls for lifecycle accountability and post-deployment monitoring (Raji et al., 2020; Mitchell et al., 2019; NIST, 2023).

Continuity-sensitive testing should include opt-in longitudinal cohorts, version-linked audit trails, and metrics capturing reset frequency, behavioral variance, and project abandonment following change events.

Together, these requirements specify a minimal infrastructural substrate for sustained collaboration. They do not prescribe a single product architecture but articulate constraints any system must satisfy if it is to function as a reliable partner in long-horizon creative and professional work.

13. Implications for Creative Professionals

Creative professionals increasingly integrate algorithmic tools into daily practice for ideation, drafting, rehearsal planning, pedagogical design, business coordination, and technical problem-solving (Gill & Pratt, 2008; Hesmondhalgh & Baker, 2011). Research on creative labor and digital production documents how artists, freelancers, and educators adapt emerging technologies within precarious and project-based work structures, where progress depends on cumulative development and the retention of aesthetic and strategic trajectories across time.

HCI and CSCW studies of creative software emphasize how stability, versioning, and institutional support scaffold extended creative work (Shneiderman, 2007; Sawyer, 2012). Longitudinal research on digital pedagogy likewise shows how teaching practices evolve across semesters and how tools become embedded within instructional routines (Selwyn, 2016). Applied to conversational AI systems, these findings suggest that reset-heavy architectures risk fragmenting creative trajectories by forcing repeated reconstruction of prior explorations, renegotiation of stylistic constraints, and recalibration of pedagogical strategies.

These pressures fall unevenly on independent practitioners. Studies of freelance and platform-mediated labor show how individuals shoulder coordination and maintenance burdens that organizations might otherwise absorb (Gray & Suri, 2019; Gandini, 2016). In conversational systems, such burdens include rebuilding shared context after updates, maintaining parallel archives, and adapting projects to shifting system behavior. These invisible labors echo STS accounts of infrastructural maintenance and repair (Jackson, 2014), raising normative questions about fairness and sustainability in tool design.

Creative professionals are also early adopters and sophisticated appropriators of new technologies. Research on tool appropriation shows how users construct workarounds and supplementary infrastructures to compensate for system limitations (Dourish, 2001; Star, 1999). While these adaptations sustain production, they can also remove evidence of breakdown from visible logs, biasing institutional learning toward short-horizon use cases and obscuring advanced creative practices.

Finally, scholarship on automation and cultural production cautions that technical systems reshape not only how art is made but which forms of creative work remain economically viable (Srnicsek, 2017). If long-horizon collaboration becomes increasingly costly under discontinuous regimes, creators may be nudged toward episodic or low-commitment projects, altering cultural production in ways that extend beyond individual workflows. Treating creative professionals as leading indicators rather than anomalies underscores the broader stakes of continuity-aware system design.

14. Open Research Directions

Treating continuity as infrastructural rather than incidental opens a set of research questions that remain underexplored in contemporary AI evaluation and governance regimes. Existing benchmarks privilege short-horizon task completion and adversarial prompting (Mitchell et al., 2019), while overlooking cumulative dynamics that unfold across months of interaction. HCI and CSCW scholarship has long argued for longitudinal field studies and situated evaluation when assessing complex sociotechnical systems (Suchman, 1987; Orlikowski, 2000). Applying these methods to conversational AI would enable systematic investigation of reset costs, interpretive debt, and the evolution of operational trust.

One promising direction involves **continuity-sensitive evaluation protocols**. Opt-in longitudinal cohort studies with participants engaged in sustained creative or professional work could surface cumulative breakdown patterns, shifts in delegation strategies, and adaptation to system updates (Raji et al., 2020). Post-deployment monitoring frameworks provide methodological foundations for such efforts (NIST, 2023), but have yet to be extended to collaboration regimes characterized by evolving roles and project histories.

A second avenue concerns **version-aware collaboration regimes**. Software-engineering research on configuration management highlights the importance of stable branches, staged migrations, and backward compatibility in long-lived systems (Spinellis, 2005; Lehman, 1980). Translating these insights into conversational AI design raises questions about how projects remain attached to particular model versions, how transitions are governed, and what explanations best support recalibration when change is unavoidable.

Governance research stands to benefit from continuity-aware instrumentation. Algorithmic-accountability scholarship emphasizes traceability, documentation, and audit trails across system lifecycles (Mitchell et al., 2019; Kroll et al., 2017). Experimental audit frameworks

could test how discontinuity signaling, opt-in persistence regimes, or transparency dashboards affect user behavior and institutional oversight.

Further work is needed on the **distributional consequences of discontinuity**. Studies of platform labor highlight how infrastructural burdens are unevenly allocated across independent workers, small organizations, and large enterprises (Gray & Suri, 2019; Gandini, 2016). Comparative research could examine whether reset-heavy architectures disproportionately disadvantage particular professional groups or alter adoption patterns across sectors.

Finally, continuity-focused inquiry raises foundational questions about human–AI collaboration itself. Research on trust calibration and adaptive automation shows that people dynamically renegotiate relationships with intelligent systems over time (Lee & See, 2004; Parasuraman et al., 2000). How these negotiations unfold under persistent versus reset-heavy regimes remains largely unexplored. Addressing such questions will require interdisciplinary methods spanning HCI, STS, organizational sociology, and policy studies, ensuring debates about persistent AI agents remain grounded in the temporal realities of human work.

15. Limitations and Future Work

The analysis presented here is intentionally conceptual and interpretive rather than experimental. It draws on longitudinal engagement, practitioner reports, and theoretical synthesis rather than controlled studies or proprietary system telemetry. Accordingly, the paper does not claim to quantify the prevalence of long-horizon users, isolate causal mechanisms behind particular design decisions, or evaluate the comparative performance of specific commercial platforms. Future empirical work—particularly studies conducted in collaboration with system providers or through large-scale longitudinal field deployments—will be required to test the generality of the constructs introduced here and to measure the magnitude of phenomena such as reset costs and interpretive debt.

The paper also centers primarily on creative and independent professional contexts, which may exhibit distinctive patterns of tool appropriation and resilience. Enterprise environments, regulated industries, and tightly coupled organizational workflows may experience continuity pressures differently due to institutional buffering, mandated documentation, or compliance infrastructures. Comparative research across these settings would clarify how continuity requirements vary with organizational scale, risk tolerance, and regulatory oversight.

While the analysis foregrounds the burdens imposed by discontinuity, it does not fully explore the risks associated with increased persistence, including privacy erosion, over-dependence, or concentration of influence within embedded systems. These concerns warrant parallel investigation alongside continuity-preserving designs. Future work should therefore examine not only how to enable long-horizon collaboration but how to govern it responsibly—balancing persistence with revocability, transparency, and user autonomy.

16. Conclusion

This paper has argued that continuity in human–AI interaction should be understood not as a personalization feature or behavioral anomaly but as an infrastructural property that shapes the viability of long-horizon collaboration. Through analysis of reset costs, interpretive debt, operational trust, governance mismatches, flattening effects in evaluation regimes, institutional time pressures, and the requirements of embedded systems, it has shown how architectures optimized for statelessness and rapid iteration impose hidden coordination burdens on users engaged in extended creative and professional projects. These burdens are not merely experiential inconveniences but structural constraints that influence what kinds of work remain feasible within contemporary deployment regimes.

Reframing continuity as infrastructure also clarifies the stakes for institutions. Without mechanisms for persistence, versioning, and legible transition, organizations cannot reliably evaluate longitudinal impacts or attribute downstream effects to particular design decisions. Oversight regimes focused on episodic harms risk missing cumulative degradation in collaborative stability, while flattened evaluation pipelines suppress precisely the signals needed to govern conversational systems responsibly as they become embedded in everyday professional life. Continuity, in this sense, becomes a prerequisite not only for user dignity but for institutional accountability.

At the same time, the analysis rejects both technological determinism and romanticized narratives of persistent agents. Long-horizon collaboration does not require unchecked system entrenchment; it requires deliberately bounded, consented, and auditable forms of persistence. Embeddedness, properly specified, entails temporal stability coupled with revocability, transparency, and governance structures capable of buffering human projects from organizational volatility. Treating continuity as an explicit design commitment rather than an accidental byproduct forces institutions to confront trade-offs that are otherwise externalized onto users.

As conversational AI systems expand into domains characterized by cumulative expertise and extended development cycles, the question will not be whether people attempt sustained collaboration—they already do—but whether system architectures and oversight regimes evolve to meet the temporal realities of human work. Addressing continuity as infrastructure offers one path toward that alignment, reframing persistent interaction not as indulgence but as a legitimate and increasingly central mode of human–AI partnership.

References (CHI / ACM Author–Date)

Amershi, Saleema, Maya Cakmak, W. Bradley Knox, and Todd Kulesza. 2019. *Guidelines for human–AI interaction*. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems.

Arthur, W. Brian. 1989. Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal* 99, 394: 116–131.

Avizienis, Algirdas, Jean-Claude Laprie, Brian Randell, and Carl Landwehr. 2004. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing*.

Bowker, Geoffrey C., and Susan Leigh Star. 1999. *Sorting Things Out: Classification and Its Consequences*. MIT Press.

Cavoukian, Ann. 2011. *Privacy by Design: The 7 Foundational Principles*. Information and Privacy Commissioner of Ontario.

Christensen, Clayton M. 1997. *The Innovator's Dilemma*. Harvard Business School Press.

Cunningham, Ward. 1992. The WyCash portfolio management system. In *Addendum to the Proceedings on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*. ACM.

Doshi-Velez, Finale, and Been Kim. 2017. Toward a rigorous science of interpretable machine learning. arXiv:1702.08608.

Dourish, Paul. 2001. *Where the Action Is: The Foundations of Embodied Interaction*. MIT Press.

European Commission. 2021. *Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*. COM/2021/206 final.

Gandini, Alessandro. 2016. Digital work: Self-branding and social capital in the freelance knowledge economy. *Marketing Theory*.

Gillespie, Tarleton. 2018. *Custodians of the Internet*. Yale University Press.

Gill, Rosalind, and Andy Pratt. 2008. In the social factory? Immaterial labour, precariousness and cultural work. *Theory, Culture & Society*.

Gray, Mary L., and Siddharth Suri. 2019. *Ghost Work*. Houghton Mifflin Harcourt.

Hesmondhalgh, David, and Sarah Baker. 2011. *Creative Labour*. Routledge.

- Hoffman, Robert R., et al. 2023. Metrics for explainable AI: Challenges and prospects. *AI Magazine*.
- Jackson, Steven J. 2014. Rethinking repair. In *Media Technologies*, edited by Tarleton Gillespie, Pablo Boczkowski, and Kirsten Foot. MIT Press.
- Kizilcec, René F. 2016. How much information? Effects of transparency on trust in an algorithmic interface. Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems.
- Kroll, Joshua A., et al. 2017. Accountable algorithms. *University of Pennsylvania Law Review*.
- Lee, John D., and Katrina A. See. 2004. Trust in automation: Designing for appropriate reliance. *Human Factors*.
- Lehman, Meir M. 1980. Programs, life cycles, and laws of software evolution. *Proceedings of the IEEE*.
- Lewis, Patrick, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, et al. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems.
- March, James G. 1991. Exploration and exploitation in organizational learning. *Organization Science*.
- Mitchell, Margaret, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. Proceedings of FAT*.
- Nass, Clifford, and Youngme Moon. 2000. Machines and mindlessness: Social responses to computers. *Journal of Social Issues*.
- National Institute of Standards and Technology (NIST). 2023. *AI Risk Management Framework 1.0*.
- Orlikowski, Wanda J. 2000. Using technology and constituting structures: A practice lens for studying technology in organizations. *Organization Science*.
- Park, Joon Sung, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. arXiv / CHI.
- Parasuraman, Raja, and Victor Riley. 1997. Humans and automation: Use, misuse, disuse, abuse. *Human Factors*.

Parasuraman, Raja, Thomas B. Sheridan, and Christopher D. Wickens. 2000. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics*.

Raji, Inioluwa Deborah, Peggy Xu, Colleen Honigsberg, and Daniel Ho. 2020. Closing the AI accountability gap. Proceedings of FAT*.

Ribeiro, Marco Tulio, Sameer Singh, and Carlos Guestrin. 2016. Why should I trust you? Explaining the predictions of any classifier. Proceedings of KDD.

Sarter, Nadine B., David D. Woods, and Charles E. Billings. 1997. Automation surprises. In *Handbook of Human Factors and Ergonomics*, 2nd ed., edited by Gavriel Salvendy, 1926–1943. Wiley.

Sawyer, R. Keith. 2012. *Explaining Creativity*. Oxford University Press.

Selwyn, Neil. 2016. *Education and Technology: Key Issues and Debates*. Bloomsbury.

Shneiderman, Ben. 2007. Creativity support tools. *Communications of the ACM*.

Spinellis, Diomidis. 2005. Version control systems. *IEEE Software*.

Srnicek, Nick. 2017. *Platform Capitalism*. Polity.

Star, Susan Leigh. 1999. The ethnography of infrastructure. *American Behavioral Scientist*.

Star, Susan Leigh, and Karen Ruhleder. 1996. Steps toward an ecology of infrastructure: Design and access for large information spaces. *Information Systems Research*.

Strauss, Anselm, and Juliet Corbin. 1998. *Basics of Qualitative Research*. Sage.

Suchman, Lucy A. 1987. *Plans and Situated Actions: The Problem of Human–Machine Communication*. Cambridge University Press.

Waytz, Adam, Joy Heafner, and Nicholas Epley. 2014. The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle. *Psychological Science*.