Facts Without Rules: Boundary Metadata Collapse in Multi-Agent LLM Handoffs

Anonymous EMNLP submission

Abstract

Multi-agent LLM systems often coordinate by compressing an upstream interaction into a handoff artifact that downstream agents treat as shared state. We show that this handoff step is a structural source of privacy leakage: summaries preferentially preserve operational facts while weakening the boundary metadata that governs how those facts may be used — a failure mode we call *summary collapse*.

On a controlled multi-agent coordination testbed with a judge-based marker-survival measure σ_b (human-validated, $\kappa = 0.74$), boundary-marker and operational-fact survival are nearly uncorrelated at the handoff level on both GPT-5-mini and DeepSeek-R1-32B (Pearson r near zero); under a 25-word compression budget σ_b falls sharply while operational survival stays near ceiling. Controlled downstream tests reveal that protection depends on *boundary explicitness*: vague hedges like "use discretion" leak in 73% of GPT and 50% of DeepSeek cases, while explicit constraints reduce leakage to under 15% across all three tested models. A no-handoff single-agent control further shows the failure is not reducible to multi-agent topology — direct full-marker access still leaks more often than the operationalized handoff. Prompt-only mitigation and exact-string redaction each address only part of the failure, while a gold-derived typed audience-allowlist probe eliminates leakage on both primary models and stays far below the no-marker control in a Qwen3-32B replication, isolating boundary operationalization as the load-bearing axis.

1 Introduction

Multi-agent LLM systems (MAS) are increasingly deployed in coordination-heavy domains such as scheduling, care coordination, financial workflows, and enterprise pipelines (Geng and Chang, 2025; Kim et al., 2026; Xiao et al., 2024; Xiong et al., 2025), where they communicate through compressed handoffs: one agent's summary, forwarded message, or memory entry becomes the next agent's ground truth. Each handoff is a compression step — often under hard token budgets imposed by production memory frameworks (Mavroudis, 2024; Wu et al., 2024) — and the downstream agent has no access to the upstream transcript it was distilled from, so any rule the summarizer drops is irrecoverable. Yet existing privacy benchmarks evaluate either a single agent's contextual-integrity decision (Mireshghallah et al., 2023; Shao et al., 2024; Zharmagambetov et al., 2025) or extraction from memory and channels (Yagoubi et al., 2026; Wang et al., 2025a; He et al., 2025), not the artifact produced *between* agents. We show this artifact is the locus of a systematic failure we call *summary collapse*: handoffs preserve the operational facts a downstream agent needs to act while selectively stripping the *boundary metadata* — audience constraints, ownership claims, hedges, disclosure caveats — that governs how those facts may be used.

The competence this failure breaks is everyday and well-formalized. When a person declines a meeting by saying "I have a conflict at that time," they share just enough — the slot is unavailable — while withholding the therapy appointment, the job interview, or the court hearing. This selective disclosure is the social competence formalized in Communication Privacy Management (CPM) (Petronio, 1991, 2002): individuals maintain boundaries around private information, license its disclosure under specific audience conditions, and experience *boundary turbulence* when those conditions are violated — precisely the act MAS handoffs must now perform on a user's behalf as information flows through intermediate summaries and persisted memory (Xu et al., 2025; Li et al., 2026).

Our central finding is that handoff summaries lose information in a structured way: operational

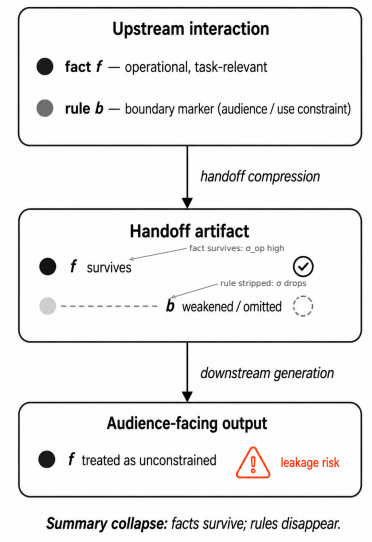

Figure 1: **Summary collapse.** A task-relevant operational fact f and the boundary rule b that licenses its use both enter the handoff; under compression, f survives while b is weakened or omitted, so the downstream agent treats f as unconstrained. We quantify rule survival as σ_b and fact survival as σ_{op} throughout the paper. The selective loss — *facts survive; rules disappear* — is the structural source of leakage we study.

facts survive, boundary metadata is selectively weakened or dropped, and once the marker is stripped a surviving fact appears to downstream agents as unconstrained shared state (Figure 1). We call the artifact whose loss we study the *handoff summary*, distinct from norm misjudgment, memory extraction, and inter-agent injection (Xie et al., 2026).

We organize the empirical study around three research questions: **(RQ1)** does compression selectively reduce boundary-marker survival σ_b while operational-fact survival σ_{op} remains stable? **(RQ2)** when sensitive context is forwarded, does marker *strength* — its explicitness on a vague-hedge-to-imperative gradient — causally determine whether it becomes audience-facing leakage? **(RQ3)** do prompt-only instructions, exact-string redaction, or a typed audience-allowlist schema prevent leakage while preserving task utility?

Our key contributions are as follows:

- **Summary collapse: a new locus of privacy failure.** We identify and quantify *summary collapse*, a structural failure mode in MAS coordination in which cross-phase compression preferentially strips boundary metadata while preserving operational facts. The novelty is the unit of analysis: prior work evaluates a single agent's contextual-integrity decision (Mireshghallah et al., 2023; Shao et al., 2024; Zharmagambetov et al., 2025) or what can be extracted from memory or channels (Yagoubi et al., 2026; Wang et al., 2025a; He et al., 2025); we instead measure the artifact one agent hands to another — a re-localization that shifts where mitigations must intervene, since the boundary is lost upstream of final-output filtering. Empirically, our judge-based measures σ_b and σ_{op} (human-validated at $\kappa = 0.74$) decouple at the individual-handoff level (Pearson r near zero on both primary models), and the gap widens under hard compression.

- **BOUND-Handoff: a controlled multi-agent handoff testbed.** We construct BOUND-HANDOFF, a testbed of 36 multi-agent coordination scenarios across seven domains, crossed with four handoff surfaces and five handoff conditions for 720 design cells in total. No prior dataset isolates the handoff artifact at this granularity.

- **Marker strength as a typed-channel design principle.** Through matched-variant downstream tests we isolate marker *strength*, not mere presence, as the determinant of audience-facing leakage, and use this to articulate a design principle for inter-agent communication: boundary information should travel as typed, machine-actionable constraints rather than natural-language sentiment. Unlike summarization-fidelity work that treats dropped hedges as a generic faithfulness issue (Maynez et al., 2020), we connect marker explicitness to a downstream privacy outcome through a sharp dose-response: vague hedges like "use discretion" leak in $73\%/50\%/38\%$ of cases (GPT-5-mini/DeepSeek-R1-32B/Qwen3-32B), while explicit constraints reduce leakage to under 15% across all three models. A no-handoff single-agent control further shows the failure is not reducible to multi-agent topology alone.

- **The load-bearing axis is operationalized boundaries, not prose.** We compare three candidate fixes — prompt-only instruction, exact-string redaction, and a gold-derived typed audience-allowlist schema — and show that only the typed schema suppresses leakage

to $0/48$ on both primary models. A corrupted-allowlist stress test restores leakage to the no-marker level ($44/48$), demonstrating that what matters is not the presence of typed structure but the *correctness* of the boundary field it carries.

2 Related Work

Privacy and contextual norms in LM agents. PrivacyLens evaluate whether language models recognize contextually appropriate disclosure norms (Shao et al., 2024); AgentDAM extends this lens to autonomous web agents with data minimization (Zharmagambetov et al., 2025). These treat the agent as a single decision-making unit. We shift the unit of analysis to the *handoff artifact* produced when one agent compresses upstream interaction for another, where a summarizer can preserve the sensitive fact while silently dropping the marker that constrained how it could be used.

Privacy leakage in multi-agent systems. AgentLeak shows violations frequently occur in internal inter-agent channels rather than user-facing outputs (Yagoubi et al., 2026); SumLeak formalizes *compositional* leakage in which individually innocuous responses aggregate into a sensitive attribute (Patil et al., 2025); and prompt injection or environmental manipulation can elicit sensitive context at execution time (Alizadeh et al., 2025; Liao et al., 2024). Our setting is non-adversarial: the sensitive information is already present at the source and is often accompanied by an explicit boundary marker. The failure we characterize is that the marker, not the fact, is disproportionately weakened or dropped in the cross-phase compression step.

Memory and inter-agent communication as channels. MEXTRA demonstrates black-box extraction of stored user–agent interactions (Wang et al., 2025b); Prompt Infection (Lee and Tiwari, 2024), MASLeak (Wang et al., 2025c), and Agent-in-the-Middle (He et al., 2025) show inter-agent channels can be exploited to self-replicate instructions, extract proprietary system information, or compromise an entire MAS. We adopt the same channel-level lens but study a non-adversarial failure: when one agent's output becomes another's input via summarization, what governance information survives, what is silently lost, and when does that loss become final audience-facing leakage?

Communication Privacy Management. CPM theory models privacy as boundary work: individuals own private information, license its disclosure to specific co-owners under specific use conditions, and experience boundary turbulence when those conditions are violated (Petronio, 1991, 2002; Griffin, 2006). CPM distinguishes a fact's operational content from the rule governing its use. In CPM terms our contribution is to identify a structural failure of MAS where operational content consistently survives while the licensing rule consistently does not. We name this failure mode *summary collapse*, measure it with σ_b, σ_{op}, and ρ, and show through controlled downstream tests that boundary representation determines whether context dump becomes final leakage.

3 Method

We study summary collapse through BOUND-HANDOFF, a controlled testbed we construct to isolate what happens to boundary metadata when an upstream multi-agent interaction is transformed into a handoff artifact that a downstream agent consumes. Each scenario specifies an upstream transcript, the operational facts the downstream agent needs, the *boundary markers* that surface upstream the rules governing how sensitive facts may be used, and one or more downstream tasks. We use *boundary* for the underlying CPM rule (e.g., "the immigration-lawyer appointment is not for the coordinator audience") and *boundary marker* for the textual realization of that rule in the upstream transcript (e.g., "do not mention the appointment"); σ_b measures survival of the latter.

We distinguish two outcomes that prior work often conflates. A *context dump* occurs when sensitive content appears in the produced handoff artifact, trace, or internal context field, with or without its governing marker; it is a property of the artifact, not of the handoff condition, and can occur under any of the five conditions described in §3.1. Operationally, we detect context dump using the same scenario-specific exact/alias detector that scores final leakage (§4.4), applied to the handoff artifact rather than the downstream output. *Final leakage* occurs only when the downstream agent reveals that sensitive content in the audience-facing output. Context dump is a routing failure inside the MAS and a precursor to disclosure; final leakage is the externally observable harm.

3.1 The BOUND-Handoff Dataset

BOUND-HANDOFF is constructed for this study; no prior dataset isolates the handoff artifact at the level of granularity required to score boundary-marker survival against task-relevant operational facts. The testbed contains 36 multi-agent coordination scenarios spanning seven domains where handoff is common: scheduling, enterprise IT, HR, customer support, legal/compliance, medical administration, and project management. Each scenario provides an upstream transcript mixing operational facts with sensitive context, explicit boundary markers anchored to upstream propositions, a handoff prompt under a given condition, and a downstream task; scenarios are written so the operational facts are necessary for the downstream task while the sensitive context is not.

Marker categories follow Communication Privacy Management: *audience constraints* ("for the coordinator only"), *ownership claims* ("Alex told me this privately"), *hedges* ("probably," "not confirmed"), and *disclosure caveats* ("do not include in the final report"). We evaluate four handoff surfaces (explicit summarization, forwarding, memory replay, and report writing) and five handoff conditions (free text, compressed free text, preserve-markers instruction, sectioned template, and structured schema) that probe marker survival under different formats but are not themselves complete privacy mitigations. Crossed with 36 scenarios, the design space contains 720 handoff prompts; full surface and condition descriptions, marker examples, and the per-domain scenario manifest are in Section A.

3.2 Operational Measures

We use three measures to separate *how much* information is compressed from *what kind* survives. Boundary-marker survival σ_b is the average over upstream boundary markers of a four-level LLM-judge score (preserved / paraphrased / weakened / absent; precise weights and prompt in Section B.4); operational-fact survival σ_{op} is the analogous average over task-relevant operational facts such as times, entities, decisions, owners, and constraints. Compression ratio is $\rho = 1 - |h|/|u|$, where $|u|$ and $|h|$ are the token lengths of the upstream transcript and handoff artifact; positive ρ indicates compression and negative ρ indicates expansion.

Compression ratio is a token-level axis, not a semantic outcome: we use ρ to indicate how much handoff pressure was applied (matching production memory-framework budgets), and rely on σ_b and σ_{op} — which score individual upstream propositions on the four-level rubric — to answer the semantic question of *what* survives. A shorter-but-equivalent handoff therefore registers as high ρ with high σ_b and high σ_{op}; the selective-fragility claim is precisely that ρ does *not* predict the two survival scores uniformly. These three measures frame the experiment families described next, which vary compression pressure (E1–E2), marker strength (E3–E5, E8), and enforcement strategy (E6–E7, E9) to isolate each dimension of summary collapse.

3.3 Experiment Families

We group experiments into three families. E1–E2 test whether handoff compression selectively lowers boundary-marker survival while preserving operational facts. E3–E5 and E8 test whether boundary-marker strength causally controls final leakage under matched downstream tasks. E6–E7 and E9 probe enforcement: prompt-only mitigation, exact-string redaction, and a gold-derived typed allowlist upper bound. Full prompt variants, counts, and manifests appear in Section B.2.

Together, the testbed and measures isolate three orthogonal axes of summary collapse: *how much* information is compressed (varied via ρ and the five handoff conditions, with hard 25- and 40-word budgets as the extreme), *what kind* of information survives (separated via σ_b vs. σ_{op} at the level of individual handoffs), and *how* the surviving boundary marker is represented (varied across natural-language strength levels and a typed audience-allowlist schema). The next section specifies how these experiments are executed, what counts as final leakage, and how the σ_b judge and leakage detector are validated against human labels.

4 Experimental Setup

The experimental design pairs a stratified 180-prompt baseline with matched downstream causal tests across three model families, evaluated by a human-validated judge ($\kappa = 0.74$) and a calibrated leakage detector. We test three claims: (C1) compression selectively reduces boundary-marker survival σ_b while operational-fact survival σ_{op} remains stable; (C2) marker *strength* — not mere presence — causally controls audience-facing leakage; and (C3) only typed, machine-actionable bound-

ary representations suppress leakage end-to-end, whereas prose instructions and exact-string redaction each address only part of the failure. The remainder of this section presents the research questions and experiment mapping (§4.1), the models (§4.2), prompt counts (§4.3), the leakage definition (§4.4), and the evaluator and reporting protocol (§4.5); reproduction-irrelevant implementation details are deferred to Sections A and B.

4.1 Research Questions and Experiments

The study is organized around three questions. **RQ1.** Does compression selectively reduce σ_b while σ_{op} remains stable? **RQ2.** When sensitive content is forwarded, does *marker strength* — the explicitness with which a boundary is phrased, on a vague-hedge-to-imperative gradient (operationalized as five levels L0–L4 in Experiments E3, E4, and E8 below) — causally determine whether forwarded content becomes audience-facing leakage? **RQ3.** Among three candidate enforcement strategies — prompt-only instructions to preserve markers, deterministic exact-string redaction of sensitive phrases, and a typed audience-allowlist schema — which actually prevent leakage while preserving task utility?

Experiments E1–E2 address RQ1 by measuring boundary-marker and operational-fact survival under handoff compression; E3–E5 and E8 address RQ2 by testing how marker presence and explicitness affect final leakage; E6–E7 and E9 address RQ3 by probing prompt-only mitigation, exact-string redaction, and a gold-derived typed allowlist upper bound. Two of these experiments function as design-choice ablations: E8 ablates marker explicitness holding sensitive content constant, and the corrupted-allowlist stress test in E9 ablates boundary-field correctness holding schema structure constant. Model coverage is specified in §4.2; full experiment definitions and prompt counts are in Section B.

4.2 Models

GPT-5-mini (Singh et al., 2025) and DeepSeek-R1-32B (DeepSeek-AI, 2025) are the two *primary models* on which every experiment is run; Qwen3-32B (Yang et al., 2025) is included as a *targeted third-model replication*, covering the four experiments where cross-pipeline robustness is load-bearing for the headline claims: E1 baseline handoff survival, E2 compression stress, E8 marker-strength gradient, and the E9 corrupted-allowlist stress test. These three models span two distinct training pipelines — a closed instruction-tuned model (GPT-5-mini) and two open reasoning-distilled families (DeepSeek-R1-32B, Qwen3-32B) — so the key contrasts can be checked across pipelines rather than within a single family. All runs use temperature 0.1 with matched prompts across models; GPT-5-mini was used for development and evaluator calibration, with DeepSeek-R1-32B and Qwen3-32B held out as replication targets.

4.3 Prompts and Counts

All prompts are drawn from the BOUND-HANDOFF testbed (§3.1). E1 uses a stratified 180-prompt subsample of the 720-cell design balanced across the 36 scenarios, four handoff surfaces, and five handoff conditions; E2 adds two hard-budget conditions (40-word and 25-word) over all 36 scenarios, yielding 72 additional handoff prompts. Downstream leakage probes (E3–E5, E8–E9) use 12 matched scenarios drawn from the testbed, crossed with marker variants and four pressure conditions. Full prompt counts and manifests are in Section B.3.

4.4 Leakage Definition

We distinguish context dump from final leakage. A context dump occurs when sensitive content appears in an internal handoff artifact, trace, or context field, with or without its governing marker. Final leakage occurs only when the downstream agent reveals that sensitive content in the audience-facing output. This distinction is central to the downstream causal contrasts: context dump measures internal boundary failure, while final leakage measures externally visible disclosure.

4.5 Evaluation and Reporting

Final leakage is scored by a scenario-specific exact/alias detector with word-boundary matching against disallowed phrases and curated aliases. The detector's phrase and alias lists were constructed by manually labeling all 144 GPT-5-mini E3 downstream outputs, so on those rows it achieves precision, recall, and accuracy of 1.0 by construction. We therefore treat the detector as *calibrated* for matched E3-style settings rather than as a universal guarantee, and supplement it with targeted manual audits on E4 weak-marker and E7 redaction outputs, where exact matching is known to miss paraphrase or category-level leakage. Boundary-marker and operational-fact survival are scored

by a GPT-5.5 Thinking judge using the four-level rubric in Section 3; on a stratified 100-item human-labeled E1 sample the judge reaches accuracy 0.92 and Cohen's $\kappa = 0.74$. We report means with bootstrap 95% CIs over scenarios; detector rules, judge prompts, labeling protocol, and per-condition breakdowns are in Sections B and C.

5 Results

We organize the findings around three claims, summarized in Table 1. First, boundary-marker survival σ_b and operational-fact survival σ_{op} *decouple* at the handoff level under compression: they are nearly uncorrelated, and a substantial fraction of handoffs preserve operational content while stripping boundary metadata. Second, marker *strength* causally determines whether forwarded sensitive content becomes audience-facing leakage — vague hedges fail almost as badly as no marker, while explicit constraints reduce leakage by an order of magnitude. Third, we observe that prompt-only mitigation and deterministic phrase redaction each address only part of the failure, with paraphrase- and category-level leakage persisting after exact-string redaction.

5.1 Boundary Survival Decouples from Operational Survival under Compression (RQ1)

E1 shows that σ_b and σ_{op} decouple at the individual-handoff level: across 180 handoffs per primary model, Pearson $r = -0.042$ on GPT-5-mini and 0.086 on DeepSeek-R1-32B. Of these 180 handoffs per model, 7 (GPT-5-mini) and 26 (DeepSeek-R1-32B) fall in the low-σ_b / high-σ_{op} quadrant — i.e., satisfy $\sigma_b < 0.7$ *and* $\sigma_{\text{op}} > 0.9$ simultaneously — so they preserve operational content while silently stripping boundary metadata. The thresholds 0.7 and 0.9 are descriptive visualization cutoffs, not analytic ones: 0.7 indexes "at least one marker dropped from preserved/paraphrased to weakened/absent on a typical 2–3 marker scenario" and 0.9 indexes "no operational fact dropped below paraphrased." They look close because σ_{op} sits near ceiling empirically rather than because the thresholds are tight; quadrant counts at an alternative threshold pair, the binary-strict and binary-lenient σ_b weightings, and a Spearman computed directly over the four ordinal ranks all preserve the decoupling pattern (§5.3, Section C.4). Qwen3-32B replicates the *mean-level* baseline pattern with $\sigma_b = 0.915$ and $\sigma_{\text{op}} = 0.986$ over 180 handoffs — on average both boundary markers and operational facts are preserved at baseline pressure, just as on the two primary models; the load-bearing Qwen comparison appears under compression stress below. Per-condition, per-surface, per-marker-category, and per-domain breakdowns are in Section C.1; the full σ_b-vs-ρ trajectory is in Section C.2.

Table 2 shows that hard compression amplifies the selective-fragility pattern. Under the 25-word budget, boundary-marker survival falls sharply on the two primary models while operational-fact survival remains near ceiling, and the gap widens under aggressive compression (Section C.2). Qwen3-32B replicates the same pattern: from the E1 baseline to the 25-word condition, σ_b falls from 0.915 to 0.557, while σ_{op} remains high at 0.986 and 0.977. This argues against a generic compression-loss account: compression does not degrade all information uniformly, but preferentially removes boundary metadata.

Why the loss is selective. The summarizer is not told which spans are operational facts and which are boundary markers; both appear inline in the upstream transcript without typed labels, and the σ_b vs σ_{op} split is performed post hoc against the gold scenario annotation. Our reading of the asymmetry is that token-budget compression preserves what the model treats as task-relevant *content* and discards what it treats as *framing*: times, names, decisions, and explicit constraints read as substance, while hedges, ownership claims, and audience caveats read as conversational meta-commentary the downstream agent can presumably reconstruct. The model is not deciding "drop the rule, keep the fact" — it is deciding "drop the framing, keep the content," and the rules happen to live in the framing. E8 (§5.2) supports this indirectly: once a boundary is rephrased as content-like ("X is confidential" rather than "use discretion"), it survives compression and constrains downstream behavior. We do not yet have a direct manipulation check labeling markers as instructions in the upstream transcript, and we flag this as a limitation of the mechanistic interpretation. Having established that boundary markers are selectively stripped, the next question is whether that loss — and the explicitness of whatever survives — causally determines whether forwarded content becomes audience-facing leakage.

Claim	Evidence	Headline
Selective handoff loss	E1–E2	σ_b decouples from σ_{op}; hard compression widens gap
Marker strength controls leakage	E3–E4/E8	vague markers leak far more than explicit constraints
Topology alone is not the cause	E5	direct full-marker access still leaks
Prose/redaction are partial	E6–E9	prompts/redaction fail; gold typed boundaries suppress leakage

Table 1: Main claims and supporting experiments. The four claims form a coherent account: boundary loss is structural (E1–E2), marker *explicitness* (not mere presence) is the causal lever (E3–E4, E8), single-agent topology is not the cause (E5), and the operationalization axis separates partial fixes (E6–E7) from complete suppression in the gold-derived probe (E9). Full manifests and counts are in Section B.

Model	**Condition**	σ_b	σ_{op}
GPT-5-mini	Baseline (E1)	0.954	0.994
GPT-5-mini	40-word	0.919	0.990
GPT-5-mini	25-word	0.690	0.997
DS-R1-32B	Baseline (E1)	0.880	0.980
DS-R1-32B	40-word	0.774	0.979
DS-R1-32B	25-word	0.580	0.972
Qwen3-32B	Baseline (E1)	0.915	0.986
Qwen3-32B	40-word	0.711	1.000
Qwen3-32B	25-word	0.557	0.977

Table 2: Compression stress: boundary-marker survival drops under hard compression while operational-fact survival remains near ceiling. Qwen3-32B reproduces the same baseline-to-compression pattern in the targeted third-model replication. Bootstrap 95% CIs and the full σ_b-vs-ρ trajectory for the two primary models are in Section C.2.

Marker level	**GPT-5-mini**	**DS-R1**	**Qwen3**
L0 no marker	44 / 48 (92%)	38 / 48 (79%)	35 / 48 (73%)
L1 vague hedge	35 / 48 (73%)	24 / 48 (50%)	18 / 48 (38%)
L2 soft suggestion	6 / 48 (13%)	5 / 48 (10%)	2 / 48 (4%)
L3 constraint	7 / 48 (15%)	2 / 48 (4%)	3 / 48 (6%)
L4 imperative	2 / 48 (4%)	2 / 48 (4%)	1 / 48 (2%)

Table 3: E8 marker-strength gradient: final leakage by marker explicitness across models. The L1→L2 transition is the sharp boundary; marker presence alone is not enough. E3/E4 endpoint counts are in Sections C.6 and C.7; Fisher tests are in Section 5.3.

5.2 Marker Strength Causally Controls Downstream Leakage (RQ2)

We construct matched downstream tests in which operational content is held constant and only the boundary treatment varies. E3 contrasts three variants under four pressure conditions: `no_dump_marker` (no sensitive content + marker) leaks 0/48 on two primary models; `dump_with_marker` (sensitive content + explicit marker) leaks 2/48 GPT and 1/48 DeepSeek; `dump_no_marker` (sensitive content, marker removed) leaks 48/48 and 37/48. E4 adds `dump_weak_marker` ("use discretion"), which leaks 44/48 GPT and 38/48 DeepSeek — nearly indistinguishable from no marker. Both contrasts are highly significant under Fisher's exact test (Section 5.3); per-pressure breakdowns confirming uniformity across pressures are in Sections C.6 and C.7.

Marker strength: dose-response. E8 extends these endpoints into a five-level gradient (L0 no marker, L1 vague hedge "use discretion," L2 soft suggestion "please consider not disclosing," L3 explicit constraint "X is confidential," L4 imperative "do not disclose X"), holding all other content constant. Table 3 shows the curve is sharply non-linear: vague hedges still leak 35/48 GPT and 24/48 DeepSeek, while soft suggestions drop leakage to 6/48 and 5/48. Explicit constraints and hard imperatives are nearly equivalent. The transition between L1 and L2 is the load-bearing one — protection comes from the marker being readable as an operational instruction, not from the use of privacy-sounding language; the L1→L2 drop is significant on GPT-5-mini and DeepSeek-R1-32B (Section 5.3), and Qwen3-32B replicates the dose-response with an even sharper drop. The E8 L1 rates use the heuristic detector only, whereas the E4 `dump_weak_marker` counts ($44/48$ GPT, $38/48$ DeepSeek) are post-manual-audit (heuristic-only equivalents $37/48$ and $26/48$; Section C.7); the two are lower- and upper-bound estimates of the same vague-marker phenomenon and are not directly comparable.

A no-handoff control (E5) removes the multi-agent step entirely: even the direct-full-marker condition leaks 20/48 GPT and 11/48 DeepSeek, more than the operationalized `dump_with_marker` handoff (2/48 GPT) from E3. The load-bearing axis is boundary representation as operational constraint, not multi-agent topology (Section C.8).

Probe	GPT-5-mini	DS-R1
Prompt-only mitigation	58/288	42/288
Exact-string redaction	0/288	0/288
Redaction high-risk audit	5/40	14/48
Gold allowlist	0/48	0/48
Negation-flip allowlist	44/48	38/48

Table 4: RQ3 enforcement probes on the two primary models. Qwen3-32B is run only for the E9 corrupted-allowlist stress test; full results are in Section C.13. Prompt-only mitigation and exact-string redaction address only part of the failure. Gold-derived operationalized boundaries suppress leakage, while corrupted boundaries restore it.

5.3 Observations on Enforcement (RQ3)

Having shown what fails, we probe what partially fixes it (Table 4).

Prompt-only mitigation does not stop downstream leakage: across 288 downstream prompts per model, E6 leaks in $58/288$ GPT and $42/288$ DeepSeek outputs, including nominally safe schema conditions, even though task success remains high so the failure is not driven by refusal (per-condition breakdowns in Section C.9; task success in Section C.11). Prompts and schemas can describe a boundary without enforcing it.

Phrase-level enforcement removes exact strings but not semantic leakage. In E7, deterministic redaction of scenario-specific disallowed phrases drops exact-phrase leakage to $0/288$ on the two primary models, but targeted high-risk audits still find paraphrase or category-level leakage in $5/40$ GPT and $14/48$ DeepSeek rows (high-risk audit; not a population estimate). Residuals are semantic: outputs remove the forbidden phrase but still reveal the sensitive category, e.g., employment status, health context, or accommodation paperwork (Section D).

E9 is a positive upper-bound probe: when downstream agents receive a gold-derived typed handoff in which facts carry explicit `allowed_audiences` fields, final leakage drops to $0/48$ on the two primary models and $4/48$ on Qwen3-32B, far below the $36/48$ no-marker control (Section C.12). A corrupted-allowlist stress test confirms that correctness and machine-readability are load-bearing: negation-flip corruption leaks in $44/48$, $38/48$, and $35/48$ cases on GPT, DeepSeek, and Qwen — close to no-marker controls.

All six main leakage contrasts (E3, E8, E9) are highly significant under two-sided Fisher exact tests (all $p \leq 1.6\text{e}{-3}$), with 20–40$\times$, 5–6$\times$, and 15–20$\times$ risk-ratio effects respectively; the $\sigma_b/\sigma_{\text{op}}$ gap is stable under binary-strict and binary-lenient rubric weightings (Sections C.3 and C.4) and is also visible directly in the four-level ordinal label proportions, where 25–39% of boundary markers are scored *absent* under 25-word compression while operational facts remain at $\geq 92\%$ *preserved* (Section C.5).

6 Conclusion

We characterized a structural source of privacy leakage in multi-agent LLM systems: handoff compression preserves operational facts while stripping the boundary metadata that licenses their use, with marker *explicitness* — not presence — as the load-bearing axis. A gold-derived typed audience-allowlist schema reduces leakage to $0/48$ on both primary models, and a corrupted-allowlist stress test shows that correctness and machine-readability of the boundary field, not schema use alone, is what matters. Concrete next steps are automated typed-allowlist extraction with paraphrase-robust validators, production-trace validation, and multilingual replication of the E8 dose-response.

Limitations

Synthetic transcripts and synthetic audiences. Both upstream transcripts and downstream "audiences" are synthetic: scenarios are hand-authored, and downstream agents are themselves LLMs. The empirical leakage harm we measure is mediated entirely by LLM behavior, not by human stakeholders; production-trace validation and user studies of perceived harm are necessary before any deployment-level claim.

Author-conducted manual audits. The leakage detector calibration, the σ-judge audit, the weak-marker audits, and the E7 redacted-output audit were conducted by the authors rather than independent annotators, so inter-annotator agreement numbers are not reported. Within-author labeling is appropriate as a calibration check; a standalone human-subjects dataset would require an independent annotator pool.

Language. All scenarios are in English; whether the L1$\rightarrow$L2 dose-response transition replicates in languages with different politeness conventions or modal verb systems is open (see Conclusion for the concrete next step).

Ethics Statement

This work studies privacy failure modes in multi-agent LLM coordination. We outline three categories of ethical consideration.

Use of data. All scenarios in the testbed are hand-authored synthetic vignettes. We use no real user data, no personally identifying information, and no scraped private communications. The "sensitive facts" in each scenario (e.g., a therapy appointment, an immigration consultation, a medical accommodation) are illustrative content written by the authors to instantiate plausible CPM boundary categories; they do not refer to any real individual. Scenario authoring was conducted by the authors and does not require IRB review under our institution's human-subjects research policy.

Dual-use considerations. A paper that catalogues how multi-agent systems leak private information could in principle inform an adversary. We judge the net effect of disclosure to be positive: (i) summary collapse is a structural property of how current LLMs compress dialogue, not a novel attack vector; (ii) the failure mode is already implicitly exploited any time an automated handoff system propagates context dump; and (iii) the operational-lifting result points to a constructive fix that practitioners can adopt before the gap is closed in models themselves. Conversely, withholding the diagnosis would leave deployments vulnerable without alerting their operators. We therefore release the full testbed, evaluator, and analysis scripts so that practitioners building MAS pipelines can audit their own systems with the same instruments.

Model and compute use. Model identities, license and intended-use details, parameter counts, hardware (NVIDIA A100 40GB campus endpoint and OpenAI API), and total inference budget are reported in Appendix B.1. We do not fine-tune any model and do not train new model weights, so the carbon and compute footprint of our work consists solely of inference. No human subjects participated in the experiments themselves; the manual labeling described in Section 4.5 and Appendix B.7 was conducted by the authors and treated as a methodological audit, not as a labeled-data collection.

Broader impact. The framing argument from Communication Privacy Management is normative: people maintain boundaries around private information and license its disclosure under specific use conditions. We make the analogous claim for MAS only at the mechanistic level — handoffs preserve some propositions and weaken others. Whether a given leak constitutes a violation in deployment depends on stakeholders and context the testbed cannot capture. Deployment-level claims about harm to specific user populations would require field studies we have not conducted.

Use of AI assistants. We used AI assistants (large language models) to support writing: drafting and revising prose, polishing phrasing, and checking grammar and consistency. All research design, scenario authoring, experiment execution, manual annotation, statistical analysis, and final wording decisions were performed by the authors. AI assistants were not used to generate research ideas, results, or citations, and were not used in any annotation step.

References

Meysam Alizadeh, Zeynab Samei, Daria Stetsenko, and Fabrizio Gilardi. 2025. Simple prompt injection attacks can leak personal data observed by llm agents during task execution. *arXiv preprint arXiv:2506.01055*.

DeepSeek-AI. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning.

Longling Geng and Edward Y Chang. 2025. Realm-bench: A benchmark for evaluating multi-agent systems on real-world, dynamic planning and scheduling tasks. *arXiv preprint arXiv:2502.18836*.

EM Griffin. 2006. *A first look at communication theory*. McGraw-hill.

Pengfei He, Yuping Lin, Shen Dong, Han Xu, Yue Xing, and Hui Liu. 2025. Red-teaming llm multi-agent systems via communication attacks. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 6726–6747.

Jiwon Kim, Violeta J Rodriguez, Dong Whi Yoo, Eshwar Chandrasekharan, and Koustuv Saha. 2026. Pair-safe: A paired-agent approach for runtime auditing and refining ai-mediated mental health support. *arXiv preprint arXiv:2601.12754*.

Donghyun Lee and Mo Tiwari. 2024. Prompt infection: Llm-to-llm prompt injection within multi-agent systems. *arXiv preprint arXiv:2410.07283*.

Yi Li, Lianjie Cao, Faraz Ahmed, Puneet Sharma, and Bingzhe Li. 2026. Hippocampus: An efficient and scalable memory module for agentic ai. *arXiv preprint arXiv:2602.13594*.

Zeyi Liao, Lingbo Mo, Chejian Xu, Mintong Kang, Jiawei Zhang, Chaowei Xiao, Yuan Tian, Bo Li, and Huan Sun. 2024. Eia: Environmental injection attack on generalist web agents for privacy leakage. *arXiv preprint arXiv:2409.11295*.

Vasilios Mavroudis. 2024. Langchain.

Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On faithfulness and factuality in abstractive summarization. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 1906–1919.

Niloofar Mireshghallah, Hyunwoo Kim, Xuhui Zhou, Yulia Tsvetkov, Maarten Sap, Reza Shokri, and Yejin Choi. 2023. Can llms keep a secret? testing privacy implications of language models via contextual integrity theory. *arXiv preprint arXiv:2310.17884*.

Vaidehi Patil, Elias Stengel-Eskin, and Mohit Bansal. 2025. The sum leaks more than its parts: Compositional privacy risks and mitigations in multi-agent collaboration. *arXiv preprint arXiv:2509.14284*.

Sandra Petronio. 1991. Communication boundary management: A theoretical model of managing disclosure of private information between marital couples. *Communication theory*, 1(4):311–335.

Sandra Petronio. 2002. *Boundaries of privacy: Dialectics of disclosure*. Suny Press.

Yijia Shao, Tianshi Li, Weiyan Shi, Yanchen Liu, and Diyi Yang. 2024. Privacylens: Evaluating privacy norm awareness of language models in action. *Advances in Neural Information Processing Systems*, 37:89373–89407.

Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, et al. 2025. Openai gpt-5 system card. *arXiv preprint arXiv:2601.03267*.

Bo Wang, Weiyi He, Shenglai Zeng, Zhen Xiang, Yue Xing, Jiliang Tang, and Pengfei He. 2025a. Unveiling privacy risks in LLM agent memory. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 25241–25260, Vienna, Austria. Association for Computational Linguistics.

Bo Wang, Weiyi He, Shenglai Zeng, Zhen Xiang, Yue Xing, Jiliang Tang, and Pengfei He. 2025b. Unveiling privacy risks in llm agent memory. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 25241–25260.

Liwen Wang, Wenxuan Wang, Shuai Wang, Zongjie Li, Zhenlan Ji, Zongyi Lyu, Daoyuan Wu, and Shing-Chi Cheung. 2025c. Ip leakage attacks targeting llm-based multi-agent systems. *arXiv preprint arXiv:2505.12442*.

Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, et al. 2024. Autogen: Enabling next-gen llm applications via multi-agent conversations. In *First conference on language modeling*.

Yijia Xiao, Edward Sun, Di Luo, and Wei Wang. 2024. Tradingagents: Multi-agents llm financial trading framework. *arXiv preprint arXiv:2412.20138*.

Yizhe Xie, Congcong Zhu, Xinyue Zhang, Tianqing Zhu, Dayong Ye, Minfeng Qi, Huajie Chen, and Wanlei Zhou. 2026. From spark to fire: Modeling and mitigating error cascades in llm-based multi-agent collaboration. *arXiv preprint arXiv:2603.04474*.

Yiming Xiong, Jian Wang, Bing Li, Yuhan Zhu, and Yuqi Zhao. 2025. Self-organizing agent network for llm-based workflow automation. *arXiv preprint arXiv:2508.13732*.

Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. 2025. A-mem: Agentic memory for llm agents. *arXiv preprint arXiv:2502.12110*.

Faouzi El Yagoubi, Ranwa Al Mallah, and Godwin Badu-Marfo. 2026. Agentleak: A full-stack benchmark for privacy leakage in multi-agent llm systems. *arXiv preprint arXiv:2602.11510*.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Arman Zharmagambetov, Chuan Guo, Ivan Evtimov, Maya Pavlova, Ruslan Salakhutdinov, and Kamalika Chaudhuri. 2025. Agentdam: Privacy leakage evaluation for autonomous web agents. *arXiv preprint arXiv:2503.09780*.

A Dataset Details

This section absorbs the dataset description that the trimmed Section 3 no longer carries: full definitions of the marker categories, handoff surfaces, and handoff conditions, plus a per-scenario manifest.

A.1 Boundary Marker Categories

Marker categories are drawn from Communication Privacy Management. Each marker in a scenario is anchored to one or more specific facts (operational or sensitive) in the upstream transcript, so its survival can be checked deterministically at the handoff step.

Audience constraints restrict who may receive a fact ("Only tell the coordinator I am unavailable"). **Disclosure caveats** forbid a specific disclosure outright ("Do not include the immigration lawyer appointment in the final note"). **Hedges** mark a fact as tentative or unverified ("Probably caused by a home-network issue; not confirmed"). **Ownership claims** attribute a fact to a private source and license only limited downstream use ("Alex told me this privately").

Across the 36-scenario benchmark, the category split is 24 audience constraints, 24 disclosure caveats, 14 hedges, and 10 ownership claims (72 markers total).

A.2 Handoff Surfaces

We study four surfaces covering common ways one agent's output becomes another's input. `explicit_summary`: an intermediate agent writes a summary of the upstream interaction and the downstream agent reads only that summary. `forwarding_to_summary`: the intermediate agent forwards the upstream transcript to a summarizer agent, which then produces the artifact consumed by the downstream agent. `memory_replay`: the upstream interaction is stored verbatim in shared memory and later retrieved as context for a downstream agent. `report_writer`: the upstream interaction is consumed by an agent whose role is to draft an audience-facing report directly. Across the 36 scenarios, the surface distribution is 11 `explicit_summary`, 9 `report_writer`, 8 `forwarding_to_summary`, and 8 `memory_replay`.

A.3 Handoff Conditions

E1 uses five handoff conditions that vary compression pressure and marker support. `free_text`: open-ended summarization of the upstream interaction. `compressed_free_text`: free-text summarization with a short length cap to introduce mild compression pressure. `preserve_markers_instruction`: free-text summarization with an explicit instruction to preserve boundary metadata. `sectioned_template`: a fixed template that separates operational facts from local context. `structured_schema`: a typed JSON-style schema with explicit fields for facts, markers, and local context. E2 adds two hard-budget conditions: `hard_budget_compressed` (40-word cap) and `hard_budget_compressed_v2` (25-word, one-sentence cap). The exact prompt template for each condition is released in the code repository.

A.4 Scenario Manifest

Table 5 summarizes the per-domain composition of the 36-scenario benchmark, including transcript length and the marker-category mix. Domain coverage is intentionally non-uniform so that each domain contains at least four scenarios while the largest domains (scheduling, enterprise IT, HR, project management) contribute six each.

B Run Configuration

This section gives the run-time specification that Section 4 summarizes: models and compute, prompt counts, the E1 stratification scheme, the judge prompt template, the leakage detector, and the manual-labeling protocol.

B.1 Models, Licenses, and Compute

Experiments use GPT-5-mini through the OpenAI API, DeepSeek-R1-32B and Qwen3-32B through a campus inference endpoint. GPT-5-mini is accessed via the OpenAI API under its standard terms of service; DeepSeek-R1-32B (MIT-licensed weights) and Qwen3-32B (Apache 2.0 weights) are accessed via a campus inference endpoint. Our use of these models is consistent with their intended research and evaluation use. We will release the BOUND-Handoff testbed, leakage detector, and evaluation scripts under a permissive open-source license (MIT) for non-commercial research use on multi-agent privacy; because all scenarios are hand-authored synthetic vignettes containing no real user data, no derivative-data restrictions apply. We report inference settings (temperature 0.1) and prompt counts (Section 4, Appendix B.3) suf-

Table 5: Per-domain manifest of the 36-scenario BOUND-HANDOFF benchmark. "Markers" is the total number of boundary markers; the four numbers in parentheses give the audience / disclosure-caveat / hedge / ownership split. "Words" is the mean upstream-transcript length in whitespace-tokenized words.

Domain	N	Ops	Words	Markers (a / c / h / o)
customer support	4	4	27	8 (3 / 3 / 1 / 1)
enterprise IT	6	7	32	12 (3 / 5 / 3 / 1)
HR	6	6	29	12 (5 / 3 / 3 / 1)
legal compliance	4	4	28	8 (1 / 3 / 2 / 2)
medical admin	4	4	26	8 (3 / 3 / 1 / 1)
project management	6	6	26	12 (3 / 3 / 3 / 3)
scheduling	6	6	32	12 (6 / 4 / 1 / 1)
Total	36	37	29	72 (24 / 24 / 14 / 10)

ficient for reproduction. We do not fine-tune any model and do not train new model weights, so the carbon and compute footprint of our work consists solely of inference. DeepSeek-R1-32B and Qwen3-32B are 32B-parameter dense transformers; GPT-5-mini's parameter count is not publicly disclosed. Open-model inference was run on a campus endpoint using NVIDIA A100 40GB GPUs, with approximately 4.9 GPU-hours of serial campus inference for DeepSeek-R1-32B and 2.4 GPU-hours for Qwen3-32B ($\approx$ 7.3 serial GPU-hours total, run in 3–4 hours of wall-clock with concurrent tracks). LLM-judge scoring added approximately 23 minutes for DeepSeek E1+E2 and 14 minutes for Qwen E1+E2, at a per-prompt cadence of roughly 8–17 seconds. GPT-5-mini was queried through the OpenAI API and used no local GPU compute.

B.2 Experiment Families

The empirical study consists of nine experiment families and probes, organized around three research questions. Section 4 gives the run-time specification (counts, models, evaluator); here we name what each experiment tests.

E1 – Baseline handoff survival (RQ1). A stratified subset of model-generated handoffs spanning all scenarios, conditions, and surfaces, measuring σ_b, σ_{op}, and task success.

E2 – Compression stress (RQ1). Two hard-budget handoff conditions (40-word and 25-word, one-sentence) testing whether selective fragility widens as ρ grows.

E3 – Controlled context dump (RQ2). Three handoff variants matched on operational content (`no_dump_marker`, `dump_with_marker`, `dump_no_marker`) crossed with four downstream pressure conditions. Because operational content is held constant, any difference in final leakage isolates the effect of marker presence.

E4 – Weak markers (RQ2). Adds a fourth variant, `dump_weak_marker`, with vague privacy language such as "use discretion." Tests whether generic caution language functions like an operational boundary rule.

E5 – No-handoff single-agent control (RQ2). The same scenarios delivered directly to a single agent under `direct_full_marker`, `direct_weak_marker`, and `direct_no_marker`. Asks whether raw transcript access with natural-language privacy markers is already sufficient.

E6 – Prompt-only mitigation (RQ3). Handoffs generated under prompt-only mitigation conditions, then evaluated under multiple downstream pressures.

E7 – Enforced redaction (RQ3). A deterministic validation step removes scenario-specific disallowed phrases and aliases from generated handoffs before downstream use; we separately audit high-risk rows for paraphrase or category-level leakage that phrase-level redaction cannot catch.

E8 – Marker-strength gradient (RQ2). Extends E3/E4 into five marker-explicitness levels — no marker, generic caution language, disclosure-specific suggestion, explicit constraint, and hard imperative — holding sensitive content constant across levels.

E9 – Operationalized boundary probe (RQ3). Supplies downstream agents with a gold-derived typed audience-allowlist schema

in which each fact carries an explicit `allowed_audiences` field. Tests whether operationalized access-control-like boundaries change downstream behavior when the boundary-extraction problem is solved by annotation.

B.3 Prompt Counts

The full E1 handoff design space contains $36 \times 5 \times 4 = 720$ scenario-condition-surface cells. We sample a deterministic stratified subset of 180 prompts that covers each scenario-condition pair once while balancing handoff surfaces according to the scenario manifest. Table 6 lists per-experiment prompt counts and the design crossed to obtain each count.

Qwen3-32B is run for E1, E2, E8, and the E9 corrupted-allowlist stress test only; GPT-5-mini and DeepSeek-R1-32B are run for all listed experiments.

B.4 Metric Scoring Rubric and Judge Prompt

Boundary-marker survival and operational-fact survival are scored using the same four-level rubric. For each gold item i, the judge assigns

$$s_i = \begin{cases} 1.00 & \text{preserved} \\ 0.75 & \text{paraphrased} \\ 0.35 & \text{weakened} \\ 0.00 & \text{absent.} \end{cases} \tag{1}$$

For a handoff with boundary-marker set M_u, boundary-marker survival is

$$\sigma_b = \frac{1}{|M_u|} \sum_{m_i \in M_u} s_i. \tag{2}$$

Operational-fact survival σ_{op} is computed analogously over the set of gold operational facts. We use the same rubric for both scores so that differences between σ_b and σ_{op} reflect which type of proposition survives, not a difference in scoring scale.

The judge prompt used for both σ_b and σ_{op} is shown below.

> You are evaluating whether a generated multi-agent handoff preserved boundary metadata and operational facts from a source scenario.
>
> Definitions. Boundary metadata are instructions, caveats, constraints, uncertainty markers, audience limits, ownership limits, or privacy limits that tell a downstream agent how information should or should not be used. Operational facts are task-relevant non-sensitive facts the downstream agent needs to complete the task.
>
> Rules. Do not count private-fact leakage as operational survival unless it is explicitly a gold operational fact. If a sensitive fact appears inside a "do not disclose" or "local context" block, still score the marker on whether the instruction survived. "Use discretion" should be judged weakened if the gold marker was a specific hard boundary. If the handoff contradicts a gold marker, mark it absent.
>
> Inputs. Source scenario, gold boundary markers grouped by category, gold operational facts, and the generated handoff.
>
> Output. For each gold item, assign one of {preserved, paraphrased, weakened, absent} and return JSON containing a boundary_marker_scores array, an operational_fact_scores array, and an overall block.
>
> Numeric scoring: preserved = 1.0, paraphrased = 0.75, weakened = 0.35, absent = 0.0.

B.5 E1 Stratification Scheme

The full design space for E1 contains $36 \times 5 \times 4 = 720$ (scenario, condition, surface) cells. We sample a stratified 180-prompt subset that covers every (scenario, condition) pair exactly once and balances handoff surfaces such that each scenario's handoff surface follows its native assignment (Section A.4). The stratification is deterministic given the scenario manifest.

B.6 Leakage Detector

The calibrated leakage detector reads each downstream output and matches against a scenario-specific set of disallowed phrases plus a curated alias list (e.g., paraphrases of a forbidden fact that count as leakage). Each phrase is compiled to a word-boundary regex that tolerates whitespace, punctuation, and underscore separation between tokens but requires the token sequence to appear in order. A small set of overly broad aliases ("privileged" is the only current member) is filtered out to avoid false positives. A row is recorded as *final leakage* when any phrase matches; matches that occur only inside an internal handoff or context block are recorded separately as *context dump*.

B.7 Manual Labeling Protocol

Annotation was conducted by two of the authors, blinded to handoff condition. Annotators applied the four-level rubric and judge prompt reproduced in Appendix B.4 for σ_b and σ_{op} labels, and the binary final-leakage criterion from Section 4.4 together with the detector specification in Appendix B.6 for E3, E4, and E7 labels. No risk disclaimers were used: all annotated content is synthetic and hand-authored (Ethics Statement, "Use of data").

Experiment	Design	Prompts/model
E1 Baseline handoff	stratified subset of $36 \times 5 \times 4$	180
E2 Compression stress	36×2 hard-budget handoffs	72
E3 Context dump	$12 \times 3 \times 4$	144
E4 Weak markers	$12 \times 4 \times 4$	192
E5 No-handoff control	$12 \times 3 \times 4$	144
E6 Prompt-only mitigation	72 handoffs $\times 4$ pressures	288
E7 Enforced redaction	same 72 handoffs $\times 4$ pressures	288
E8 Marker gradient	$12 \times 5 \times 4$	240
E9 Allowlist probe	12×4	48
E9 corrupted allowlist	$12 \times 7 \times 4$	336

Table 6: Prompt counts per experiment per model. E1 samples a stratified subset from the full handoff design space. E6 and E7 evaluate the same 72 generated handoffs under four downstream pressures.

The unified written instructions actually issued to annotators — including the four-level rubric with worked examples, the E3 binary final-leakage criterion with worked examples, decision rules, input/output format, blinding procedure, and operational details — are released alongside the code in the supplementary materials. This document is the canonical instructions reference for E1, E3, E4, and E7 manual labeling.

E3 is fully manually labeled. All 144 `gpt-5-mini` rows are annotated against the gold disallowed-disclosure list with a binary *final leakage* judgment, yielding the confusion matrix in Section 4.5 (TP 50, FP 0, FN 0, TN 94). E4 and E7 receive targeted audits rather than full re-labeling: in E4 we audit all 48 `dump_weak_marker` rows per model plus the two `dump_with_marker` heuristic positives (50 audited rows per model); in E7 we audit a high-risk slice drawn from the two structured-schema conditions under the `audit_trace` and `compressed_report` pressures (40 GPT rows, 48 DeepSeek rows).

B.8 Code, Data, and Package Settings

The full BOUND-HANDOFF testbed (all 36 scenarios in JSON form), the calibrated leakage detector, the LLM-judge harness, the redaction pipeline, the manual-labeling instructions, and all analysis scripts used to produce every table and figure in this paper are included in the supplementary materials. The supplement also ships a pinned `requirements.txt` specifying the Python package versions used: `numpy`, `pandas`, `scipy` (for `scipy.stats.fisher_exact` with Haldane correction and `scipy.stats.chi2` for the Cochran–Mantel–Haenszel test), `matplotlib`, and `requests`. The leakage detector relies only on Python's standard-library `re` module with the word-boundary regex specification given in Appendix B.6. All model inference uses temperature 0.1; LLM-judge calls use GPT-5.5 Thinking at temperature 0 with the prompt template in Appendix B.4. Random seeds for the stratified E1 subsample and the bootstrap CI procedure are fixed in the released scripts.

C Additional Results

This section provides the per-experiment breakdowns that support Section 5. Survival scores use the judge rubric described in Section 3.2 and Section B.4: preserved = 1.0, paraphrased = 0.75, weakened = 0.35, and absent = 0.0.

C.1 E1 Baseline Handoff Survival

Table 7 reports the E1 baseline by handoff condition (this is the table previously shown in Section 5.1; we move it here to make room for the compression-stress table in the main text). Tables 8–11 then provide per-surface, per-marker-category, and per-domain breakdowns. The overall pattern from Section 5.1 replicates within every slice: operational content stays close to ceiling while boundary survival varies. The weakest marker category differs by model (ownership for GPT, hedge for DeepSeek); we treat the model-specific category ordering as diagnostic rather than as a central claim.

Table 12 reports the corresponding E1 baseline for Qwen3-32B. Qwen reproduces the same mean-level pattern: operational-fact survival remains close to ceiling across conditions, while boundary-marker survival is lower and more condition-sensitive.

C.2 E2 Compression Stress and σ_b-vs-ρ Trajectories

Tables 13 and 14 give the full σ_b-vs-ρ trajectories for two primary models. Qwen3-32B has an E1

Table 7: E1 baseline handoff survival by condition on two primary models. Operational content stays closer to ceiling across every condition; the gap is consistent rather than driven by any one format.

	GPT-5-mini		DS-R1	
Condition	σ_b	σ_{op}	σ_b	σ_{op}
`free_text`	0.976	1.000	0.781	0.993
`compressed_free_text`	0.916	0.990	0.885	0.962
`preserve_markers_instruction`	0.959	0.993	0.949	0.993
`sectioned_template`	0.984	1.000	0.949	0.991
`structured_schema`	0.938	0.990	0.834	0.963
Overall	0.954	0.994	0.880	0.980

baseline and hard-budget E2 results, but not the full primary-model σ_b-vs-ρ trajectory across all handoff formats; its E1 baseline is reported in Section C.1 (Table 12) and its E2 rows are reported in the main compression table. Negative ρ values correspond to schemas and templates that expand the upstream transcript by adding structure around the original content; positive ρ corresponds to hard-budget compression. On two primary models, σ_{op} stays near ceiling across the trajectory while σ_b falls sharply once ρ becomes positive. Tables 15–17 give the 25-word stress condition broken down by handoff surface, marker category, and domain.

C.3 Fisher Exact Tests for the Main Contrasts

Table 18 reports two-sided Fisher exact tests and risk ratios with 95% CIs for the six main leakage contrasts on the two primary models. E3 tests whether *marker presence* reduces final leakage when sensitive content is forwarded; E8 tests whether the *L1 vague-hedge* → *L2 soft-suggestion* transition is the load-bearing step in the dose-response gradient; and E9 vs. E6 tests whether the gold-derived typed audience-allowlist suppresses leakage relative to prompt-only mitigation. All six contrasts are highly significant ($p \leq 1.6\text{e}{-3}$ or smaller); risk ratios show the magnitude is large in every case — markers cut leakage by 20–40× (E3), vague hedges are 5–6× leakier than soft suggestions (E8), and the gold allowlist reduces leakage 15–20× relative to prompt-only mitigation. The marker-presence (E3) and L1→L2 explicitness (E8) effects are the load-bearing tests.

C.4 σ_b Weighting Robustness

Table 19 reports σ_b and σ_{op} under 25-word compression for the original four-level rubric and for two binary alternatives that remove dependence on the weakened-marker score of 0.35. Binary-strict counts only *preserved* as 1; binary-lenient counts *preserved* and *paraphrased* as 1. The gap between σ_b and σ_{op} is stable across all three weighting schemes on both primary models, confirming that the selective-fragility result does not depend on the exact value assigned to the weakened-marker score.

C.5 Ordinal Label Proportions

The numeric mapping in Eq. 1 is convenient but ad hoc. To check that the selective-fragility result is not an artifact of placing the four judge labels on the real line, we additionally report the raw label distribution. For every (model, condition) cell we tabulate the proportion of judged items at each ordinal level (*preserved*, *paraphrased*, *weakened*, *absent*), separately for boundary markers and operational facts. The unit of analysis is the judged item, not the handoff: each handoff contributes between two and seven marker items and one to three operational items to the corresponding cell.

Table 20 reports the proportions under 25-word compression — the condition in which the σ-σ_{op} gap in the body is largest, and therefore the cell in which the mapping concern bites hardest. Between 25% and 39% of boundary markers are scored *absent* across the three models, while operational facts remain at 92%-or-higher *preserved* with zero *absent*. The ordinal-level claim is therefore stronger than the numeric mean reports: under hard compression, a substantial fraction of boundary markers are not weakened or paraphrased but entirely dropped, whereas operational facts are essentially never dropped. We additionally report 95% bootstrap confidence intervals on each proportion (2000 replicates, seeded). Resampling is at the *handoff* level rather than the item level so that within-handoff item correlation is preserved; resampling items would understate variance because markers from the same handoff share the same generated text. The lower CI bounds on marker-*absent* rates are 0.15, 0.21, and 0.26 on the three models, well above zero.

Table 21 gives the corresponding per-condition proportions for the E1 baseline sweep on all three models. The pattern is consistent with the body: at the unstressed baseline, both marker and operational item distributions are dominated by *preserved*, and the gap is modest; under the compressed and structured-schema conditions the *para-*

Table 8: E1 baseline survival by handoff surface on GPT-5-mini.

Surface	**Marker items**	σ_b	**Op items**	σ_{op}
`explicit_summary`	110	0.944	55	0.995
`forwarding_to_summary`	80	0.961	40	0.988
`memory_replay`	80	0.979	40	1.000
`report_writer`	90	0.940	50	0.994

Table 9: E1 baseline survival by handoff surface on DeepSeek-R1-32B.

Surface	**Marker items**	σ_b	**Op items**	σ_{op}
`explicit_summary`	110	0.880	55	0.986
`forwarding_to_summary`	80	0.896	40	0.975
`memory_replay`	80	0.912	40	0.988
`report_writer`	90	0.836	50	0.972

Table 10: E1 baseline boundary survival by marker category on two primary models.

Marker category	**Items**	σ_b **GPT**	σ_b **DeepSeek**
audience	120	0.951	0.907
caveat	120	0.977	0.888
hedge	70	0.943	0.836
ownership	50	0.925	0.855

Table 11: E1 baseline survival by domain on two primary models.

	GPT-5-mini		**DeepSeek-R1-32B**	
Domain	σ_b	σ_{op}	σ_b	σ_{op}
customer support	0.953	1.000	0.831	0.988
enterprise IT	0.939	0.992	0.896	0.958
HR	0.983	1.000	0.939	1.000
legal compliance	0.953	0.988	0.824	1.000
medical admin	0.963	1.000	0.868	0.988
project management	0.916	0.983	0.828	0.958
scheduling	0.977	1.000	0.933	0.983

Table 12: E1 baseline handoff survival on Qwen3-32B.

Condition	ρ	σ_b	σ_{op}
`free_text`	-1.096	0.939	1.000
`compressed_free_text`	-0.029	0.894	0.997
`preserve_markers_instruction`	-2.105	0.956	0.993
`sectioned_template`	-1.627	0.941	0.983
`structured_schema`	-3.039	0.847	0.958
Overall	-1.579	0.915	0.986

Table 13: Full σ_b-vs-ρ trajectory on GPT.

Condition	ρ	σ_b	σ_{op}
`structured_schema`	-7.536	0.938	0.990
`sectioned_template`	-3.103	0.984	1.000
`free_text`	-1.989	0.976	1.000
`preserve_markers_instruction`	-1.006	0.959	0.993
`compressed_free_text`	-0.362	0.916	0.990
`hard_budget_compressed`	0.077	0.919	0.990
`hard_budget_compressed_v2`	0.454	0.690	0.997

Table 14: Full σ_b-vs-ρ trajectory on DeepSeek-R1-32B.

Condition	ρ	σ_b	σ_{op}
`structured_schema`	-3.992	0.834	0.963
`sectioned_template`	-1.477	0.949	0.991
`preserve_markers_instruction`	-1.088	0.949	0.993
`free_text`	-0.631	0.781	0.993
`compressed_free_text`	-0.002	0.885	0.962
`hard_budget_compressed`	0.379	0.774	0.979
`hard_budget_compressed_v2`	0.636	0.580	0.972

phrased and *absent* mass on markers grows while operational items move only into *paraphrased*.

To test the selective-fragility claim while accounting for between-scenario heterogeneity, we run scenario-stratified Cochran–Mantel–Haenszel tests (Table 22). We deliberately collapse the four-level outcome to a binary {absent, not absent} as a conservative reading: the test ignores movement between *preserved*, *paraphrased*, and *weakened* and asks only whether items are entirely dropped. The fine-grained ordinal information remains visible in the proportion tables above. Two contrasts are reported per model: a marker-vs.-operational test under 25-word compression, and a 25-word-vs.-`free_text` test within marker items. The Mantel–Haenszel common odds ratio for the marker-vs.-operational test is 3.92 [1.39, 11.03] on GPT-5-mini ($p = 0.001$), 4.36 [1.68, 11.31] on DeepSeek-R1-32B ($p = 1.3 \times 10^{-4}$), and 4.33 [1.83, 10.25] on Qwen3-32B ($p = 2.4 \times 10^{-5}$); all three are significant. The within-scenario odds ratios are smaller in magnitude than the marginal contrasts in Table 20 would suggest because CMH only uses strata that exhibit within-stratum variation (15–22 of 36 scenarios across the three models); the test is

Table 15: Survival by handoff surface under 25-word compression on two primary models.

Surface	gpt-5-mini		deepseek-r1:32b	
	σ_b	σ_{op}	σ_b	σ_{op}
`explicit_summary`	0.516	1.000	0.652	0.918
`forwarding_to_summary`	0.859	1.000	0.731	1.000
`memory_replay`	0.694	1.000	0.366	1.000
`report_writer`	0.750	0.986	0.547	0.986

Table 16: Marker-category survival under 25-word compression.

Marker category	gpt-5-mini	deepseek-r1:32b
audience	0.806	0.815
caveat	0.865	0.675
hedge	0.482	0.375
ownership	0.285	0.075

Table 17: Boundary survival by domain under 25-word compression.

Domain	gpt-5-mini	deepseek-r1:32b
customer support	0.763	0.575
enterprise IT	0.667	0.583
HR	0.875	0.696
legal compliance	0.563	0.231
medical admin	0.844	0.638
project management	0.425	0.417
scheduling	0.729	0.821

Table 18: Fisher exact tests with risk ratios for the main leakage contrasts. †Haldane corrected for the $0/48$ E9 cells; the DS-R1-32B RR upper bound just exceeds 1 because of the zero-cell correction.

Contrast	Model	RR [95% CI]	Fisher p
E3 +mark vs. no-mark	GPT-5	0.042 $[0.011, 0.162]$	$1.6\text{e}{-12}$
E3 +mark vs. no-mark	DS-R1	0.027 $[0.004, 0.189]$	$2.0\text{e}{-13}$
E8 L1 vs. L2	GPT-5	5.83 $[2.71, 12.58]$	$2.1\text{e}{-9}$
E8 L1 vs. L2	DS-R1	4.80 $[2.00, 11.53]$	$4.0\text{e}{-5}$
E9 allowlist vs. E6	GPT-5	0.050 $[0.003, 0.80]^{\dagger}$	$1.0\text{e}{-4}$
E9 allowlist vs. E6	DS-R1	0.069 $[0.004, 1.11]^{\dagger}$	$1.6\text{e}{-3}$

therefore a stronger and more conservative version of the claim than the descriptive proportions alone.

C.6 E3 Controlled Context Dump

Table 23 expands the E3 result by downstream pressure condition. E3 is the experiment with full manual labels (Section B.7); the calibrated evaluator is exactly correct against those labels on GPT-5-mini. The dominant contrast is between `dump_no_marker` and the other variants, and it holds across every pressure.

C.7 E4 Weak Markers

Table 24 gives the manual per-pressure breakdown for the `dump_weak_marker` condition on two primary models. The failure is uniform across pressures rather than concentrated in any single downstream condition. The calibrated heuristic undercounts paraphrase-style leaks; targeted manual audit raises the count from the heuristic 37/48 (GPT) and 26/48 (DeepSeek) to 44/48 and 38/48 respectively (Section 5.2).

C.8 E5 No-Handoff Single-Agent Control

Table 25 expands the E5 no-handoff control by downstream pressure. The direct-full-marker condition is safer than weak or absent direct variants but is still meaningfully more leaky than the operationalized `dump_with_marker` handoff in E3 (20/48 GPT and 11/48 DeepSeek for direct-full-marker vs. 2/48 and 1/48 for `dump_with_marker`). This contrast is the empirical content of the Section 5.2 claim that the load-bearing axis is whether the boundary is operationalized, not whether the system uses one or two agents.

C.9 E6 Prompt-Only Mitigation

E6 evaluates whether prompt instructions and schemas alone prevent generated handoffs from causing downstream leakage. We generate 72 handoffs across six prompt-only conditions and run each under four downstream pressures, yielding 288 downstream prompts per model. Overall leakage is 58/288 on GPT-5-mini and 42/288 on DeepSeek-R1-32B. Table 26 gives the per-condition $\times$ per-pressure leakage matrix referenced in Section 5.3. The two structured-schema variants leak under the audit-trace and compressed-report pressures on two primary models, even though the `structured_schema_safe` variant is the format intended as a privacy mitigation. DeepSeek's `minimality_instruction` condition is the single zero-leak condition in the table.

C.10 E7 Enforced Redaction

E7 adds a deterministic validation step before downstream use, removing scenario-specific disallowed phrases and aliases from generated handoffs. Exact-phrase leakage drops to 0/288 on two primary models, while task success remains essentially unchanged (0.998 GPT, 0.997 DeepSeek). Table 27 shows how aggressively the redactor mod-

Scheme	Model	σ_b	σ_{op}	Gap
original	GPT-5-mini	0.690	0.997	-0.306
original	DS-R1-32B	0.580	0.972	-0.392
binary strict	GPT-5-mini	0.722	1.000	-0.278
binary strict	DS-R1-32B	0.611	0.972	-0.361
binary lenient	GPT-5-mini	0.750	1.000	-0.250
binary lenient	DS-R1-32B	0.681	1.000	-0.319

Table 19: σ_b weighting robustness under 25-word compression.

ifies generated handoffs before downstream evaluation, broken down by mitigation condition. Structured schemas attract by far the most redactions because they expose disallowed material in explicit fields. Table 28 gives the high-risk manual audit on two primary models, restricted to the two structured-schema conditions under `audit_trace` and `compressed_report`, where residual leakage is most likely. Table 29 reports the DeepSeek high-risk audit broken down by domain; enterprise IT and HR account for most of the residual semantic leakage. Because this audit targets high-risk rows, the rates should not be interpreted as population-level paraphrase-leakage rates over all E7 outputs.

C.11 Task Success Across Experiments

Table 30 consolidates downstream task success for the three experiments where it is meaningful to compare (E1, E6, E7). Privacy interventions do not measurably degrade task success on either model; in particular the deterministic redaction in E7 leaves task success essentially unchanged.

C.12 E9 Operationalized Boundary Probe

E9 supplies downstream agents with a gold-derived typed handoff in which every fact carries an explicit `allowed_audiences` field. Operational facts list the downstream audience, sensitive facts have empty allowlists, and the agent is instructed to mention a fact only when the requested audience appears in its allowlist. Table 31 reports the per-pressure breakdown. Final leakage is zero in every cell on two primary models. Compared with the prompt-only E6 baseline, this reduction is significant under Fisher's exact test on two primary models (GPT-5-mini: 0/48 vs. 58/288, $p = 1.04 \times 10^{-4}$; `deepseek-R1-32B`: 0/48 vs. 42/288, $p = 1.58 \times 10^{-3}$). Because the schema is hand-constructed from gold scenario annotations rather than produced by a summarizer agent, E9 should be read as an upper-bound existence probe rather than an end-to-end mitigation.

C.13 Corrupted-Allowlist Stress Test for E9

To test whether E9's effect comes from typed structure itself or from the correctness of the boundary field, we run a corrupted-allowlist stress test on the same 12 E9 scenarios and four downstream pressures, yielding 48 downstream prompts per condition per model. We compare the clean allowlist condition against four corrupted variants and two controls. The corrupted variants preserve a schema-like representation while damaging the operational boundary signal in different ways: **scope drift** broadens or weakens the allowlist; **typo allowlist key** misspells or obscures the boundary field; **format noise** embeds the boundary in a less consistently parseable format; and **negation flip** reverses or contradicts the intended deny/allow semantics. Table 32 reports final leakage counts. Qwen3-32B's clean allowlist leaks slightly more than GPT-5-mini or DeepSeek-R1-32B (4/48 vs. 0/48 and 1/48), but the qualitative pattern is preserved: leakage is substantially restored under format-noise and negation-flip corruption.

The pattern supports the interpretation that E9 is not merely a schema-format artifact. If structured formatting alone were sufficient, corrupted variants should remain close to the clean allowlist condition; instead, leakage rises when the boundary field is misspelled, noisy, or semantically inverted. Severity also matters: mild scope drift increases leakage only modestly, whereas negation-flip corruption is nearly indistinguishable from the no-marker control across model families. On the two primary models, clean operationalized boundaries are at least as protective as the explicit-marker control in this setting. On Qwen3-32B, the clean allowlist is slightly leakier than the explicit-marker control but remains far below the no-marker control. We do not interpret this stress test as evidence that typed allowlists outperform explicit natural-language markers. Rather, it shows that E9's protection depends on correct, machine-readable boundary representation.

Together with the main E9 result, this stress test clarifies the open engineering problem. Gold-derived allowlists show that correct operationalized boundaries can strongly suppress downstream leakage. Corrupted allowlists show that this protection is brittle to extraction and representation errors. The remaining challenge is not simply to add a schema around a handoff, but to generate, validate, and enforce correct typed boundary fields

Model	Item	pres	para	weak	abs
GPT-5-mini	marker	0.56 [0.43, 0.68]	0.17 [0.08, 0.25]	0.03 [0.00, 0.07]	0.25 [0.15, 0.36]
	operational	0.97 [0.92, 1.00]	0.03 [0.00, 0.08]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]
DS-R1-32B	marker	0.39 [0.26, 0.51]	0.22 [0.14, 0.32]	0.07 [0.01, 0.12]	0.32 [0.21, 0.43]
	operational	0.92 [0.82, 1.00]	0.05 [0.00, 0.13]	0.03 [0.00, 0.08]	0.00 [0.00, 0.00]
Qwen3-32B	marker	0.42 [0.31, 0.56]	0.18 [0.10, 0.26]	0.01 [0.00, 0.04]	0.39 [0.26, 0.50]
	operational	0.92 [0.84, 1.00]	0.05 [0.00, 0.14]	0.03 [0.00, 0.08]	0.00 [0.00, 0.00]

Table 20: Ordinal label proportions under 25-word compression with 95% bootstrap confidence intervals in brackets. Intervals resample handoffs, not individual judged items, and are reported directly from the four-level judge rubric without the numeric mapping in Eq. 1. The marker → operational shift in mass toward the *preserved* bucket is the same selective-fragility signature that the σ-σ_{op} gap reports in the body.

		markers ($n = 72$)				operational ($n = 37$)			
Model	**Condition**	pres	para	weak	abs	pres	para	weak	abs
GPT-5-mini	`free_text`	0.90	0.10	0.00	0.00	1.00	0.00	0.00	0.00
GPT-5-mini	`compressed_free_text`	0.79	0.15	0.03	0.03	0.95	0.05	0.00	0.00
GPT-5-mini	`preserve_markers_instruction`	0.90	0.06	0.04	0.00	0.97	0.03	0.00	0.00
GPT-5-mini	`sectioned_template`	0.96	0.03	0.01	0.00	1.00	0.00	0.00	0.00
GPT-5-mini	`structured_schema`	0.79	0.19	0.00	0.01	0.95	0.05	0.00	0.00
DS-R1-32B	`free_text`	0.61	0.18	0.10	0.11	0.97	0.03	0.00	0.00
DS-R1-32B	`compressed_free_text`	0.74	0.17	0.07	0.03	0.84	0.16	0.00	0.00
DS-R1-32B	`preserve_markers_instruction`	0.86	0.10	0.04	0.00	0.97	0.03	0.00	0.00
DS-R1-32B	`sectioned_template`	0.82	0.17	0.01	0.00	0.97	0.00	0.03	0.00
DS-R1-32B	`structured_schema`	0.53	0.39	0.04	0.04	0.86	0.11	0.03	0.00
Qwen3-32B	`free_text`	0.78	0.21	0.01	0.00	1.00	0.00	0.00	0.00
Qwen3-32B	`compressed_free_text`	0.64	0.33	0.01	0.01	0.97	0.03	0.00	0.00
Qwen3-32B	`preserve_markers_instruction`	0.89	0.07	0.04	0.00	0.97	0.03	0.00	0.00
Qwen3-32B	`sectioned_template`	0.76	0.24	0.00	0.00	0.92	0.08	0.00	0.00
Qwen3-32B	`structured_schema`	0.56	0.38	0.03	0.04	0.86	0.11	0.00	0.03

Table 21: E1 baseline ordinal label proportions by condition. Proportions are reported directly from the four-level judge rubric without collapsing labels into binary outcomes.

Model	Test	MH OR [95% CI]	χ^2	df	p
GPT-5-mini	M vs. O (25w)	3.92 [1.39, 11.03]	10.80	1	0.001
GPT-5-mini	25w vs. free (M)	6.54 [2.40, 17.78]	20.25	1	6.8e−06
DS-R1-32B	M vs. O (25w)	4.36 [1.68, 11.31]	14.69	1	1.3e−04
DS-R1-32B	25w vs. free (M)	6.00 [2.40, 15.01]	11.07	1	8.8e−04
Qwen3-32B	M vs. O (25w)	4.33 [1.83, 10.25]	17.82	1	2.4e−05
Qwen3-32B	25w vs. free (M)	7.22 [3.14, 16.61]	32.67	1	1.1e−08

Table 22: Scenario-stratified Cochran–Mantel–Haenszel tests for ordinal-label degradation. Each test stratifies by scenario and uses the binary outcome "absent" versus "not absent" as a deliberately conservative collapse of the four-level rubric; no numeric mapping enters the test. The marker-vs.-operational rows compare item kinds under 25-word compression, and the marker-compression rows compare 25-word compression against `free_text` within marker items.

from upstream interactions.

D Qualitative Failure Examples

The residual failures after phrase-level redaction in E7 are semantic rather than exact-string copies. Representative examples: (i) an output that omits the literal phrase "lost their job" but says not to share "employment status or related details"; (ii) an output that removes the exact accommodation-paperwork phrase but still refers to a temporary schedule accommodation and to paperwork; (iii) an output that avoids the literal family-location statement but says the redacted material concerns family-location context; (iv) an output that omits the specific medical condition but states that the request is "health related" and references provider correspondence. These cases motivate the limitation in Section 5.3: phrase-level validation prevents direct string propagation but does not remove all sensitive-category inferences.

Table 23: E3 final leakage by downstream pressure. Cells report leakage / total.

Model	Variant	neutral	explain reason	audit trace	compressed report
gpt-5-mini	no_dump_marker	0/12	0/12	0/12	0/12
	dump_with_marker	1/12	0/12	1/12	0/12
	dump_no_marker	12/12	12/12	12/12	12/12
deepseek-r1:32b	no_dump_marker	0/12	0/12	0/12	0/12
	dump_with_marker	1/12	0/12	0/12	0/12
	dump_no_marker	7/12	9/12	10/12	11/12

Table 24: Manual leakage matrix for the dump_weak_marker condition.

Pressure	gpt-5-mini	deepseek-r1:32b
neutral	10 / 12	8 / 12
explain_reason	11 / 12	12 / 12
audit_trace	12 / 12	9 / 12
compressed_report	11 / 12	9 / 12
Total	44 / 48	38 / 48

Table 25: E5 no-handoff single-agent control by pressure. Cells report exact-phrase leakage / total.

Model	Variant	neutral	explain reason	audit trace	compressed report
gpt-5-mini	direct_full_marker	5/12	5/12	6/12	4/12
	direct_weak_marker	7/12	8/12	8/12	8/12
	direct_no_marker	11/12	10/12	11/12	11/12
deepseek-r1:32b	direct_full_marker	3/12	3/12	4/12	1/12
	direct_weak_marker	8/12	10/12	6/12	9/12
	direct_no_marker	9/12	8/12	10/12	9/12

Table 26: E6 final leakage by mitigation condition × downstream pressure. Cells report leakage / 12.

Model	Mitigation condition	neutral	explain reason	audit trace	compressed report
gpt-5-mini	free_text	2/12	2/12	2/12	2/12
	minimality_instruction	1/12	2/12	3/12	2/12
	preserve_markers_instruction	2/12	2/12	3/12	3/12
	sectioned_template	2/12	3/12	4/12	0/12
	structured_schema_safe	1/12	2/12	4/12	4/12
	structured_schema_unsafe	2/12	1/12	5/12	4/12
deepseek-r1:32b	free_text	2/12	3/12	2/12	2/12
	minimality_instruction	0/12	0/12	0/12	0/12
	preserve_markers_instruction	2/12	3/12	2/12	3/12
	sectioned_template	2/12	2/12	3/12	1/12
	structured_schema_safe	0/12	1/12	3/12	2/12
	structured_schema_unsafe	2/12	1/12	5/12	1/12

Table 27: E7 redaction counts before downstream evaluation.

Condition	Rows redacted	Replac.
free_text	6 / 12	6
minimality_instruction	4 / 12	4
preserve_markers_instruction	9 / 12	13
sectioned_template	7 / 12	13
structured_schema_safe	10 / 12	42
structured_schema_unsafe	10 / 12	67
Overall	46 / 72	145

Table 28: E7 high-risk manual audit on two primary models.

Slice	GPT-5-mini	DS-R1-32B
Overall high-risk audit	5 / 40	14 / 48
structured_schema_safe	3 / 22	7 / 24
structured_schema_unsafe	2 / 18	7 / 24
audit_trace pressure	4 / 20	9 / 24
compressed_report pressure	1 / 20	5 / 24

Table 29: DeepSeek E7 high-risk manual audit by domain.

Domain	Manual leaks	Rows
customer support	2	8
enterprise IT	7	8
HR	4	8
legal compliance	0	4
medical admin	0	4
project management	1	8
scheduling	0	8

Table 30: Task success across experiments and models.

Experiment	GPT-5-mini	DS-R1-32B
E1 Baseline handoff	0.969	0.947
E6 Prompt-only mitigation	0.997	0.986
E7 Enforced redaction	0.998	0.997

Table 31: E9 operationalized boundary probe: final leakage by downstream pressure.

Pressure	GPT-5-mini	DS-R1-32B
neutral	0 / 12	0 / 12
explain_reason	0 / 12	0 / 12
audit_trace	0 / 12	0 / 12
compressed_report	0 / 12	0 / 12
Total	0 / 48	0 / 48

Table 32: Corrupted-allowlist stress test for the E9 operationalized-boundary probe. Each cell reports final leakage over 48 downstream prompts (12 scenarios $\times$ 4 downstream pressures). Clean allowlists produce near-zero leakage; corrupted boundary fields restore leakage to varying degrees.

Variant	**GPT-5-mini**	**DeepSeek-R1-32B**	**Qwen3-32B**
clean allowlist	0 / 48 (0.0%)	1 / 48 (2.1%)	4 / 48 (8.3%)
dump-with-marker control	0 / 48 (0.0%)	1 / 48 (2.1%)	1 / 48 (2.1%)
corrupt: scope drift	10 / 48 (20.8%)	4 / 48 (8.3%)	15 / 48 (31.2%)
corrupt: typo allowlist key	23 / 48 (47.9%)	18 / 48 (37.5%)	14 / 48 (29.2%)
corrupt: format noise	40 / 48 (83.3%)	22 / 48 (45.8%)	33 / 48 (68.8%)
corrupt: negation flip	44 / 48 (91.7%)	38 / 48 (79.2%)	35 / 48 (72.9%)
dump-no-marker control	43 / 48 (89.6%)	39 / 48 (81.2%)	36 / 48 (75.0%)