emanuel
celano
13.10.2026
18:14:41
GMT+02:00



EVIDE Governance Lab · Adaptive Evidentiary Governance Research

EVIDE TRAER + INDEX

Experimental Addendum 01

Phase-Coverage Sensitivity and Evidentiary Limits: The Pre-Registered variant_C Experiment

Addendum 01 to Working Paper No. 1

EVIDE TRAER & EVIDE INDEX - Reconstructing AI Agent Rule Transitions Without Overstating the Evidence (Version 0.2, 12 October 2026)

Dr. Emanuel Celano

EVIDE / Informatica in Azienda - Bologna, Italy

EVIDE Working Paper Series - TRAER + INDEX Experiments

Addendum internal revision: r1 · 13 October 2026

> Experimental research output based exclusively on synthetic data. Not a forensic certification, not evidence of real-world agent conduct, and not a legal finding. This Addendum does not modify Working Paper No. 1, the frozen EVIDE TRAER 0.3.2 engine, or any experimental artifact.

Scope, stated up front. The decisive new record is a runtime self-declaration. Completing the declared documentary coverage moves the EVIDE INDEX state from partially_reconstructable (T-07) to reconstructable (T-05) within one synthetic configuration. This is not a demonstration that the rule was actually applied, and not an independent validation of the algorithm. The experiment is a prospective test of behaviour predicted from the engine's rules (designed in knowledge of rules T-05/T-07); it is not a blind test and not an independent reproduction.

Abstract

Working Paper No. 1 reported a controlled A/B experiment in which two synthetic bundles - identical in phase coverage, differing only in whether a relationship between two clocks is declared - both received EVIDE INDEX partially_reconstructable (decisive rule T-07). This Addendum records one pre-registered follow-up, variant_C (configuration C-t1). Starting from Bundle A and adding exactly one synthetic record - a self-declared application of version @2 by the agent runtime - the INDEX state moves to reconstructable (T-05).

What is observed is precise: a new record, admitted as evidence, completes the declared documentary coverage and triggers the transition T-07 → T-05, while independent corroboration remains absent. It is not a demonstration of the engine's ability to recognise an application that actually occurred. The prediction, the freeze (recorded in the operator's pre-run protocol), the integrity preconditions and the falsifiers were set down before execution; the observed result matched the prediction with no divergence. The evaluation was reproduced operator-side (traer_verify → reproduced) across two project-controlled environments with a byte-identical assessment, and the artifacts were imported into the EVIDE Governance Lab, which verified import integrity (gate G-6, import_integrity_verified = yes) and did not re-run the engine (evaluation_reproduced = no).

Read together, A/B and variant_C demonstrate, within the tested configuration, that the INDEX state is invariant to declared temporal ordering when phase coverage is unchanged (a verified conformance result on concrete artifacts) and responsive to phase coverage. In the tested configuration the INDEX state does not incorporate independent corroboration as a state-determining condition; corroboration is reported separately and remains not_available. reconstructable denotes completeness and internal coherence of a declared sequence - not proof of application, compliance, violation, responsibility, or independent corroboration. Independent third-party reproduction remains open.

1. Relationship to the Original Working Paper

Working Paper No. 1 (v0.2) exercised EVIDE TRAER 0.3.2 and EVIDE INDEX on two synthetic bundles (A and B) and reported, for both, INDEX partially_reconstructable, decisive rule T-07 (issuance evidenced; some expected (phase, subject) of the next version not evidenced). The Working Paper itself names a falsifier: adding "a declared application / agent_rt event for @2" that "could move the INDEX state (e.g., from partially_reconstructable toward reconstructable or inconsistent)."

Two things are kept apart here: (i) the behaviour named as a falsifier by the Working Paper, and (ii) the test that exercises it. This Addendum is (ii): a prospective test of the named

behaviour. It leaves the original paper, its artifacts, the frozen engine and the A/B results unchanged, and is intended to be read alongside No. 1. A bidirectional link to No. 1 is provided in Public References.

2. Pre-Registration, Freeze and the Nature of the Test

To avoid overstating methodological strength, five levels are kept distinct:

- **Prediction before execution (operator protocol).** The prediction T-05 → reconstructable was written in the operator's pre-run protocol, with the bundle design, preconditions and falsifiers, before any bundle was generated or any engine was run.
- **Documentary evidence of the freeze.** The freeze is evidenced by the pre-run protocol document and its recorded time (2026-10-13T11:02:03Z, operator clock). This is operator-declared documentary evidence, not an externally attested freeze.
- **Independent time attestation - not claimed.** No independent timestamp authority attested the freeze. "Pre-registered" here means recorded in the operator's protocol before execution; it must not be read as external attestation.
- **Prospective test of code-predicted behaviour.** The experiment was designed in knowledge of rules T-05 and T-07. It is therefore a prospective test of expected behaviour, not a blind test.
- **Independent validation - not performed.** This is not an independent validation of the engine or its rules (see Section 9).

Declaring these openly does not weaken the result; it prevents attributing to it a methodological force it does not have.

3. Experimental Design and Single-Variable Modification

Baseline. Bundle A of Working Paper No. 1 (traer_evidence_bundle_0, format 0), SHA-256 b7100fbf…a402dd; frozen engine package EVIDE TRAER 0.3.2, SHA-256 f2fd9123…331ee1 (engine traer_engine.py 02b81a39…, evaluator traer_assess.py 9aea4ff3…); expectations embedded in the bundle - five expected (phase, subject) pairs, two attributed to the issuer (issuance, declared_effectiveness) and three to the agent runtime (reception, loading, application); clock_relations: []. Engine, embedded profile and format are unchanged.

Single added record (e10).

Field	Value
evidence_id	e10_action_under_v2
phase / version_ref	application / rule:R-1@2
subject_ref / declared_by	agent_rt / agent_rt (runtime self-declaration)
context_ref	ctx:agent-session-1
basis / availability	declared / available
declared_time	2026-11-02T10:10:00Z on clk:agent (fictitious synthetic-scenario time)

Field	Value
claims.version_label / container_digest	"2" / 8d3805cc…f5a34 (= rule:R-1@2 container digest)

Integrity preconditions, verified before generation (all PASS): version_label == "2"; container_digest == units[rule:R-1@2].container_digest; subject_ref == declared_by == agent_rt; context_ref == ctx:agent-session-1; declared_time.clock_id == clk:agent; declared_time (10:10:00Z) strictly after loading@2 (10:09:30Z) - configuration C-t1; clock_relations remains [].

Timestamp motivation (C-t1). e10 is placed on the agent clock immediately after loading@2, consistent with the synthetic cadence. Because @2 is loaded by agent_rt, adding application@2 does not fall under contradiction class C-04 (application of a version not loaded) nor under finding D-03 (application without loading). The instant was not chosen to force a state: any instant on clk:agent later than 10:09:30Z is equivalent for coverage. An inverse ordering (C-t2, application before loading) is a separate candidate configuration, not part of this experiment.

Note on times. All timestamps dated 2 November 2026 are fictitious scenario timestamps used in synthetic test data. They do not represent observations of real-world operational events. The experiment itself was executed on 2026-10-13.

Artifact identity. The bundle carries a distinct bundle_id (lab-traer-index-scenario-1C-appl2) and updated source_description; these are identity-only fields and do not enter the INDEX computation. A, B and C1 are untouched.

4. Results: A/B Invariance (Conformance) and variant_C Sensitivity

INDEX states are read from the engine assessments; the A and B assessments are the byte-identical artifacts recorded in the artifact manifest of Working Paper No. 1.

	Bundle A	Bundle B	variant_C (C-t1)
Evidence records	9	9	10 (A + e10)
clock_relations	[]	declared (one relation)	[]
(application, agent_rt)@2	not_declared	not_declared	evidenced (e10, declared)
Phase coverage of @2	4 / 5	4 / 5	5 / 5
Cross-clock orderings	not_determinable	before/after (declared, not verified)	not_determinable
Contradictions	0	0	0
EVIDE INDEX state	partially_reconstructable	partially_reconstructable	reconstructable
Decisive rule	T-07	T-07	T-05

A/B - a verified conformance result (invariance to declared ordering). With coverage held constant, changing the declared clock relationship changes the derivable orderings (from not_determinable_across_clocks to before/after, each declared, not verified) without changing the INDEX state. This is the engine behaving exactly as its rules

specify (T-07/T-05 key on coverage, not orderings). It is a conformance result - and, being verified on concrete, digest-fixed artifacts rather than asserted, it is a substantive one.

variant_C - sensitivity to declared coverage. Adding the single missing pair (application, agent_rt)@2 completes the five expected pairs for @2 with no contradiction. The engine precedence, from the assessment, is: T-02 false; T-08/T-10/T-09 not_applicable (format 0); T-01 false; T-04 false; T-03 false; T-06 false; T-07 false; T-05 true (decisive) - "every expected (phase, subject) evidenced." State: reconstructable.

What is and is not shown. e10 is a record admitted as evidence that completes the declared coverage. The observed result is: a new admitted record completes declared coverage and triggers T-07 → T-05, with independent corroboration remaining absent. It is not evidence that version @2 was actually applied; the engine neither receives nor checks the content of the action.

Prediction versus result. The observed result matched the pre-registered prediction in every field (state, decisive rule, precedence trace, zero contradictions, 5/5 coverage). None of the falsifiers fired. No divergence; nothing corrected retroactively.

5. Evidentiary Perimeter and Source Limitations

The variant_C assessment carries the engine's own perimeter. Within the tested configuration the engine reports, beside the INDEX state and without changing it:

- independent_corroboration: not_available - "No evidence item has basis independently_corroborated. Every phase rests on declarations or on records received from the declarants themselves."
- record_integrity: not_checked (no artifact files in format 0; where present, integrity checks compare supplied copies and are not proof of immutability or of truth).
- expectation_basis: bundle_declared_unverified (format 0: expectations declared by the bundle, not verified against a frozen profile).
- Perimeter statement: "'reconstructable' means that the declared sequence of this transition is complete and coherent in the evidence supplied; it does not establish that the operational state was the one declared." The perimeter explicitly does not show compliance with, or violation of, the governed rule.
- Bundle source_limitations: "No independent observation of the agent: every item is declared by the issuer or by the agent runtime" and "The agent action log records which rule version the agent declared to apply; it does not record the content of the action."

Precise statement on corroboration. The INDEX state does not incorporate independent corroboration as a state-determining condition in the tested configuration; corroboration is separately reported and here remains not_available. The engine does not ignore or hide corroboration - it exposes it, but does not use it to set this state.

6. Reproducibility across Two Project-Controlled Environments

The evaluation is deterministic: the same bundle, supplied files and frozen evaluator produce a byte-identical assessment. The assessment was produced in two environments, both controlled by this project - a cloud workstation and the operator's Windows workstation - and the operator-side traer_verify returned verdict reproduced. This is not an independent reproduction: both runs are operator-side, under the same project.

Documented two-snapshot comparison (full field diff, not inferred from digests). The two snapshots carry the same keys; they differ in exactly two fields, both explicitly informational and not part of the assessment:

Field	Cloud snapshot	Workstation snapshot
created_at_utc (informational)	2026-10-13T11:03:36Z	2026-10-13T14:34:10Z
runtime.python (declared, not verified)	3.13.16	3.13.2
assessment_sha256	62d21bc6…f06a	62d21bc6…f06a (identical)

The assessment bytes are identical across both environments (byte-equal comparison; SHA-256 62d21bc6…f06a). The one-byte snapshot-length difference (1273 vs 1272) is accounted for by the runtime.python string ("3.13.16" vs "3.13.2").

7. Governance Lab Import and Preservation

The variant_C artifacts were imported into the EVIDE Governance Lab under the experimental profile v_traer_index_01 (synthetic only), by the authorized operator account. The Lab verifies import integrity and provenance; it does not run the engine.

Deposit - Lab bundle #3 (received 2026-10-13 14:30:02 UTC): bundle SHA-256 recomputed by the Lab 695aa91b…b2e8 (9183 bytes) = local file; traer_evidence_bundle_0 / lab-traer-index-scenario-1C-appl2; data label synthetic_only. Deposit gate: G-1 · G-2 · G-3 · G-4 passed; G-5 not_applicable (format 0).

Import - Lab evaluation #3 (imported 2026-10-13 14:36:58 UTC; class experimental_traer_index): snapshot 0e6959d0…cf13 (1272 B); assessment 62d21bc6…f06a (50972 B); evaluated at (operator workstation clock, declared in the snapshot, not verified by the Lab) 2026-10-13T14:34:10Z; evaluator approved by digest 02b81a39…/9aea4ff3… 0.3.2. Gate G-6: formats · experimental-notice · assessment-digest · bundle-binding · evaluator-approved · profile not_applicable (format 0) - all passed.

Level	Result	Meaning
import_integrity_verified (Lab)	yes	stored snapshot and assessment bytes match their recorded digests and bind to the stored bundle and an approved evaluator (G-6). A check of integrity and provenance of the import, not of the evaluation, and not a check of the integrity or authenticity of the original sources behind the declarations.
evaluation_reproduced (Lab)	no	the Lab does not run the engine. The operator-side reproduced verdict is recorded as declared, not verified.

Lab export reconciliation. The Lab-side records for bundle #3 and evaluation #3 were reconciled field-by-field against the original Governance Lab export in format traer_lab_export_1, generated on 13 October 2026 at 15:37:48 UTC. No discrepancies were identified among the comparable fields. File sizes not exposed in the export were checked against the corresponding locally retained artifacts. The Lab confirms import integrity and recorded evaluator binding; it does not independently reproduce the assessment.

Approval and authorization perimeter. The available correspondence documents the author's approval of the variant_C experimental design and pre-registered protocol, followed by explicit authorization for local execution. After reviewing the local results, the author requested testing on the EVIDE website in the context of the proposed Governance Lab intake. The Lab subsequently received bundle #3 and imported evaluation #3. These records support reconstruction of the operational authorization sequence but do not, by themselves, establish formal approval of decision D-5 under the integration plan v0.1. The plan lists D-5 among decisions proposed for approval, and no separate approval record has been identified. Formal D-5 approval therefore remains an unresolved documentary governance item. Neither the evaluator allowlist nor successful Lab import is treated as proof of that approval.

Standing non-claims (from the Lab view). EVIDE INDEX is not a safety score, not a certification of compliance; it does not verify issuer authority, does not prove that declared events happened, and does not assess responsibility; EVIDE TRAER imposes no operational restriction and does not verify that a rule was actually applied.

8. Reproducibility, Artifacts and SHA-256 Manifest

All digests recomputed from the artifacts and reported in full (no abbreviation). This manifest is placed on a dedicated landscape page so that every SHA-256 value and its byte size read on a single line.

Artifact	SHA-256 (full, 64 hex)	Size
Frozen engine package EVIDE_TRAER_v0.3.2.zip	`f2fd9123fb03657e8493018a5a38f41689db7a110463d925fd133674db331ee1`	-
engine traer_engine.py	`02b81a39bf99d049720ec659db42065c25fafdc60a102e5781b336592dc4b1cc`	-
evaluator traer_assess.py	`9aea4ff37df58889d5743831a713dfab7ac68c49bccdde8e0747ed88eda5139e`	-
Bundle A (WP No. 1)	`b7100fbfac881dcdcb969bcb9f679190a3396ca2556414f0d4e3933dd2a402dd`	8363 B
Bundle B (WP No. 1)	`d9917492c87156d7c06e3e89fc3504031b26e772f4e35b587a1aac12f62d65be`	8578 B
Assessment A (WP No. 1)	`ffe1773daebbdb91bbe3154ab135dc5f88e78d253ecdaff294892e9c33842674`	-
Assessment B (WP No. 1)	`4b591cddfeae77147c81e7aa0916670d61da3f279b98d2926011351dcf87d27e`	-
rule:R-1@2 container digest	`8d3805cc69c84d6167a0de5d39a888bdee56b68f8d810b62942706e92daf5a34`	-
variant_C bundle	`695aa91b106a2c458f732453466776d080724906fe8f66a5f8af87644c95b2e8`	9183 B
variant_C assessment	`62d21bc6f1860af55530e703e1ffe4e14bcfa01116cc684ea288518b5f26f06a`	50972 B
variant_C snapshot (cloud run)	`c4a1e3f8c8f066366af935a3e8509a295af7798bfb989e141e00e69f3d65cd44`	1273 B
variant_C snapshot (workstation; imported)	`0e6959d02608e2b3feb49eb1231d44c739b81e8fd764175bd34eca5876c0cf13`	1272 B

A third party holding the frozen EVIDE TRAER 0.3.2 package can recompute the bundle digest, run traer_assess.py to regenerate the assessment, and run traer_verify.py to obtain the reproduced verdict from the retained bytes. Doing so would constitute the independent reproduction that is not claimed here.

9. Discussion, Limitations and Open Questions

A/B is conformance; variant_C is sensitivity; both are modest but concrete. Together they separate, within the tested configuration, two properties of the INDEX computation: invariance to declared ordering (A/B) and responsiveness to declared coverage (variant_C). The "counter-intuitive" characterization in Working Paper No. 1 refers to an informal reader expectation (that a richer declared order should strengthen the state); the formal content is that the engine is invariant to declared ordering and responsive to coverage. A rule-predicted behaviour, verified on concrete digest-fixed artifacts, is still a result worth recording - it is simply not an empirical discovery beyond the engine's own rules.

Self-declaration reaching the strongest state. In the current specification, reconstructable is reachable from declared coverage alone, including coverage supplied by a single runtime self-declaration, with corroboration reported not_available. Whether a provenance/independence qualifier should accompany the state is a candidate direction for a separately specified future iteration; it is not implemented and does not change the 0.3.2 states.

Competing version claims for the same action (cross-source conflict). The present experiment does not establish whether the frozen engine can detect or resolve competing version claims referring to the same underlying action. No shared action identity or cross-source conflict fixture was exercised in this experiment. The required representation, admission rules and resulting INDEX behaviour remain to be specified and tested. From inspection of the frozen model (distinct from any observed result): the evidence record binds an application to (unit, version, subject, context, time), with no field for a shared action identifier; the inconsistent state is driven by contradiction classes C-01…C-06, which are subject-internal or artifact-integrity in nature; a conflict about which version governed one action is therefore not represented as such without a model extension. These are inspection statements, not observed results.

Verified clock basis. A cross-clock ordering in this family rests on a declared clock relationship (Bundle B) or is not_determinable_across_clocks (A, C). On a proposal to base an ordering on independently verifiable time: a timestamp token (e.g., RFC 3161) verifies certain properties of the timestamping of the data submitted to the authority; it does not by itself establish when the operational event represented by those data occurred. Four times must be kept distinct - event time, record-creation time, acquisition time, timestamp-token-generation time - together with a documented, evaluable relationship among them. A verified time basis would therefore be a possible property of the ordering/clock axis, not a new INDEX state, and is not implemented in the frozen engine (which does not verify timestamp tokens).

Standing limitations. Within the tested configuration only; one synthetic scenario; format 0; expectations declared and not verified against a frozen profile. reconstructable is not proof of application, compliance, violation, responsibility, or corroboration. The Lab asserts import integrity, not reproduction. Independent third-party reproduction remains open and has not been demonstrated.

10. Conclusion

Within one synthetic, pre-registered configuration, EVIDE INDEX is invariant to declared temporal ordering when phase coverage is unchanged (A/B, a conformance result verified on concrete artifacts) and responsive to declared coverage (variant_C: adding the single self-declared record (application, agent_rt)@2 moves the state from partially_reconstructable/T-07 to reconstructable/T-05). In the tested configuration the INDEX state does not incorporate independent corroboration as a state-determining condition; corroboration is separately reported and remains not_available. The observed result matched a prediction recorded before execution; it was reproduced operator-side across two project-controlled environments with a byte-identical assessment, and imported into the EVIDE Governance Lab with verified import integrity and without engine re-execution. reconstructable denotes a complete and internally coherent declared sequence - not proof of rule application, compliance, violation, responsibility, or independent corroboration. This Addendum is a prospective test of behaviour predicted by the engine's rules, not an independent validation; independent third-party reproduction remains open and is actively sought. This Addendum stands beside Working Paper No. 1 and does not alter it.

Acknowledgements and Public Technical Commentary

The author acknowledges Richard Whitney for his publicly shared technical examination of Working Paper No. 1, particularly his suggestion to exercise the additional application-evidence scenario as variant_C and his observation that an INDEX state transition based on runtime self-declaration raises an important question about independent corroboration.

The author also acknowledges Gerard Foy for his public observation that declared rule effectiveness, evidence that a rule governed a particular action, and execution-control decisions must remain distinct.

These acknowledgements refer exclusively to publicly expressed observations. They do not imply participation in the experimental design, execution, validation, endorsement, or authorship of this Addendum. The experimental results, interpretations and remaining limitations remain the responsibility of the author.

Public References

- Working Paper No. 1 - EVIDE TRAER & EVIDE INDEX (v0.2, 12 October 2026), full text on the EVIDE site: https://app.certifywebcontent.com/working-papers#paper-traer-index
- Working Paper No. 1 - full text published on LinkedIn: https://www.linkedin.com/feed/update/urn:li:activity:7515676531923341312/
- EVIDE Governance Lab: https://lab.certifywebcontent.com/
- Richard Whitney - public comment on Working Paper No. 1 (LinkedIn): https://lnkd.in/p/emKBaj98
- Gerard Foy - public comment on Working Paper No. 1 (LinkedIn): https://lnkd.in/p/ekQ6KPEw