

Axioms of Research Continuation: A State-Restoration Theory for Transcript-Sufficient Research Systems

Peter Bell

2026

Axioms of Research Continuation

A State-Restoration Theory for Transcript-Sufficient Research Systems

Peter Bell

ORCID: 0009-0004-0415-7774

New Reflexive Laboratory / Reflexive Laboratory Research Program

Draft status: v03 DOI-bearing local build / noncanonical candidate

VERSION_DOI: 10.5281/zenodo.20501949

CONCEPT_DOI: 10.5281/zenodo.20500050

Abstract

Research records may preserve enough information to replay a process without preserving enough state to resume it. This distinction is familiar in systems contexts, where logs and checkpoints solve different problems, but it remains underformalized in discussions of AI-assisted research infrastructure. This paper develops a state-restoration account of research continuation for transcript-sufficient research systems. It defines an operative research state, an externalized record, a record projection, a replay operator, a restore operator, and a continuation-equivalence relation based on admissible next operations. The central claim is modest: a record can support replay while failing to support resumption when projection collapses states that are not continuation-equivalent. The paper then specifies which state distinctions matter for research rather than generic computation: evidence-mediated working state, re-entry and non-erasure, artifact identity, proposal/admissibility/update separation, enacted authority, and validation-gated advancement. These are organized as a tiered axiom system: a mathematical spine, a transcript-state spine, and a governed-artifact spine. Graph representations, provenance standards, research-object packages, process mining, checkpoint/restore mechanisms, and software configuration management are treated as important analogues or implementation substrates, not as rivals. The contribution is not broad mathematical novelty. It is a formal semantics for research-state restoration: a way to state what a record must preserve if research is to be continued rather than merely replayed.

Keywords: research continuation; state restoration; record projection; transcript sufficiency; artifact authority; validation; canonicity; Reflexive Laboratory.

1. Introduction

Research continuation is a state-restoration problem. A project may leave behind transcripts, logs, files, packages, drafts, generated artifacts, and publication records. Those records can be useful for reconstructing what happened. They do not necessarily preserve enough state to determine what can responsibly happen next.

This distinction is not new in general systems terms. Logs, traces, checkpoints, workflow provenance, version-control histories, and release packages already distinguish different kinds of recoverability. A log may support replay. A checkpoint may support restoration. A package may preserve files and metadata. A workflow record may support rerun. Checkpoint/restore systems such as CRIU explicitly aim to freeze a running process or container, save its state, and later restore execution from that saved state; recent scientific-computing work similarly describes checkpointing as stopping a long-running job and continuing it later on another resource (CRIU Project 2025; Andrijauskas et al. 2024).

The problem addressed here is narrower than the general checkpointing problem. It asks what state restoration requires when the object to be restored is not only a computation, but a research project with evidence, artifacts, authority, validation, and publication status. Research continuation depends not only on process state, but also on scholarly state: which claims are evidence-supported, which artifacts are current, which proposals have been admitted, which objects are authoritative, which outputs have been checked, and which publication objects function as interfaces rather than complete research objects.

The published article *Replay Is Not Resumption: AutoResearch and the Architecture of Research Continuation* identified this problem conceptually. It argued that replayable traces, managed memory, provenance, packages, generated artifacts, and process logs are valuable but insufficient for accountable continuation unless composed into governed research state (Bell 2026, *Replay Is Not Resumption*). The present paper supplies the formal kernel for that claim. It does not attempt a full unified theory of the Reflexive Laboratory. It asks what record projection must preserve if research is to be resumed rather than merely replayed.

The formal spine comes from the record-sufficiency model developed in P10. That paper distinguishes execution sufficiency from state sufficiency by modeling operative state, externalized record, replay, and restoration. It shows that a record may support replay while failing to preserve enough operative state for continuation (Bell 2026, *Execution Sufficiency and State Sufficiency*). The present paper keeps that spine, but narrows its contribution. It does not claim that the underlying projection result is mathematically surprising. It uses that elementary result to define a research-specific restoration target.

That restoration target is **continuation equivalence**. A record does not need to restore every hidden detail of a research process. It must restore enough state to preserve the same admissible next operations. If two states license different next actions, then a record that collapses them is not sufficient for continuation. If two states differ only in irrelevant details but allow the same responsible next moves, they may be equivalent for continuation.

This shift matters because research state contains components not captured by ordinary action replay. In transcript-sufficient research, current state depends on the retained corpus, evidence selected from that corpus, working-state derivations, artifact identity, governance rules, validation records, authority states, open tasks, and re-entry links. P7 supplies the transcript-to-state chain from corpus to indexed units, evidence, and working state (Bell 2026, *From Chat Corpus to Working State*). P8 supplies proposal/admissibility/update separation for bounded automation (Bell 2026, *Bounded Autoresearch*). P11 and P20 supply graph and artifact-registry structure (Bell 2026, *Research Lab as a Graph*; Bell 2026, *Corpus as Graph*). P15 supplies enacted authority and canonicity distinctions (Bell 2026, *Canonicity Is Not Presence*). P22 and P34 supply artifact-integrity and sanity-check layers (Bell 2026, *Process Transparency Is Not Artifact Integrity*; Bell 2026, *Sanity-Check Operators*). P16 and P17 support publication-as-projection and process-legible object formation (Bell 2026, *Research as a Media Pipeline*; Bell 2026, *Constructive Closure and Reflexive Realization*).

The paper’s contribution is therefore compositional and semantic. It identifies the research-state distinctions that must survive record projection if a project is to be restored for continuation. It does not replace workflow provenance, W3C PROV, RO-Crate, process mining, checkpoint/restore systems, or software configuration management. Those are close analogues and possible implementation substrates. The claim is that research continuation adds state semantics concerning evidence, artifact identity, authority, validation, publication projection, and governed update.

The axiom system is tiered. Tier 1 contains the mathematical spine: record projection and continuation equivalence. Tier 2 contains the transcript-state spine: evidence-mediated working state, re-entry, and non-erasure. Tier 3 contains the governed-artifact spine: artifact identity, proposal/admissibility/update separation, enacted authority, and validation-gated advancement. Tiering reduces overclaiming. It makes clear which axioms support the elementary projection result and which axioms specify the research-domain restoration target.

The paper proceeds in ten sections. Section 2 defines research continuation as state restoration. Section 3 introduces state, record projection, and continuation equivalence. Section 4 states the tiered axiom system. Section 5 presents the projection-collapse lemma and the replay/resumption corollary. Section 6 identifies the research-state components required for governed continuation. Section 7 relates the model to existing infrastructure. Section 8 interprets the Reflexive Laboratory as a worked transcript-sufficient case. Section 9 states limits and nonclaims. Section 10 concludes.

Plain-language summary: this paper does not say that logs are useless. It says that a log can be enough to show what happened and still not be enough to continue the work. To continue research, a system must preserve what is current, supported, proposed, accepted, checked, authoritative, and still open.

2. Research Continuation as State Restoration

Research continuation is the task of restoring enough state to continue a project under its own rules. It is not the same as reconstructing a past sequence. It is not the same as preserving every document. It is not the same as publishing a final artifact. Continuation requires a current operational target: what may be done next?

Let X_t denote the operative research state at time t . Let $\text{AdmNext}(X_t)$ denote the set of admissible next operations from that state. These operations may include drafting, revising, checking, integrating, validating, registering, publishing, deprecating, or reopening a branch. Which operations are admissible depends on more than the action history. It depends on evidence, artifact state, governance rules, validation status, and authority relations.

The target relation is continuation equivalence.

Equation E1. Continuation equivalence

$$X_t \equiv_{\text{cont}} X'_t \iff \text{AdmNext}(X_t) = \text{AdmNext}(X'_t).$$

Two states are continuation-equivalent when they support the same admissible next operations. This relation is weaker than equality. It is too strong to require a restored record to reproduce every hidden feature of the original research situation. Some details may be irrelevant to continuation. But it is too weak to say that any replayable trace is sufficient. If two states differ in which artifact is authoritative, which proposal is enacted, which figure is validated, or which evidence supports the active claim, then they may license different next actions. They are not continuation-equivalent.

This definition also clarifies the role of governance. Governance does not appear here as bureaucracy. It appears as part of the state that determines $\text{AdmNext}(X_t)$. In P8, an automated operator may generate a proposal, but that proposal does not update working state unless admitted through a review path (Bell 2026, *Bounded Autoresearch*). In P15, a visible artifact may be present or under review without being canonical; canonicity requires governance-visible transition (Bell 2026, *Canonicity Is Not Presence*). Both cases show that admissible continuation depends on state distinctions beyond what an action log alone records.

State restoration therefore has a bounded target. The target is not perfect reconstruction. The target is restoration up to continuation equivalence. A record is state-sufficient if it supports reconstruction of a state that is equivalent for admissible continuation.

Plain-language example: if a project pauses after three candidate figures were produced, it is not enough to know that all three figures exist. To continue, one must know whether any of them was validated, which one the manuscript currently uses, whether one was rejected, and whether another remains under review. Those distinctions change the next valid action.

3. State, Record Projection, and Continuation Equivalence

The operative research state is modeled as follows.

Equation E2. Operative research state

$$X_t = (C_t, E_t, W_t, A_t, G_t, \\ V_t, K_t, \Omega_t, L_t).$$

The tuple is intentionally compressed. An earlier planning draft included U_t and P_t in the main tuple. In this polished draft, U_t is kept as part of the transcript-state derivation path rather than a top-level state component, and publication objects are treated through a publication-projection proposition rather than a core state axiom. This keeps the model focused on restoration.

The components are as follows. C_t is the retained transcript corpus. E_t is the evidence register selected from the corpus. W_t is working state derived from evidence. A_t is the artifact set. G_t is the governance layer. V_t is validation state. K_t is authority state. Ω_t is open continuation context, including unresolved tasks and active branches. L_t is the set of re-entry links from state or artifact entries back to source.

The externalized record is produced by projection.

Equation E3. Record projection

$$R_t = \rho(X_t).$$

The projection ρ determines what becomes durable record and what is lost, compressed, or left implicit. The central failure mode arises when ρ maps two continuation-distinct states to the same record.

Replay and restoration are different operators. Replay is represented by the replay operator.

Equation E4. Replay operator

$$P_{\text{replay}}(R_t).$$

Restoration is represented by the restore operator.

Equation E5. Restore operator

$$Q_{\text{restore}}(R_t).$$

The replay operator returns a replayable path or execution class. The restore operator attempts to recover a continuation-equivalent state. The two targets differ. A record can support the first without supporting the second.

A record is resumable when the restore operator recovers a continuation-equivalent state.

Equation E6. Resumability target

$$Q_{\text{restore}}(R_t) \equiv_{\text{cont}} X_t.$$

Execution sufficiency and state sufficiency can now be stated as follows.

Equation E7. Execution sufficiency

$$\text{ExecSuff}(R_t) \iff R_t \rightarrow P_{\text{replay}}(R_t).$$

Equation E8. State sufficiency

$$\text{StateSuff}(R_t) \iff Q_{\text{restore}}(R_t) \equiv_{\text{cont}} X_t.$$

The important shift is that state sufficiency is not defined as perfect memory. It is defined as restoration sufficient for governed continuation.

Table 1. Notation Table

| Symbol | Meaning | Role |
|------------------------|-----------------------------------------------|--------------------------------------|
| X_t | operative research state at time \mathbf{t} | restoration target |
| R_t | externalized record at time \mathbf{t} | projected record |
| ρ | record projection | possible site of state collapse |
| P_{replay} | replay operator | reconstructs or reruns a path |
| Q_{restore} | restore operator | attempts state restoration |
| AdmNext | admissible-next-operation function | defines continuation target |
| \equiv_{cont} | continuation-equivalence relation | restoration equivalence |
| C_t | retained transcript corpus | source archive |
| U_t | indexed transcript units | derivation layer |
| E_t | evidence register | admissible support |
| W_t | working state | operational state |
| A_t | artifact set | governable objects |
| G_t | governance state | rules and transitions |
| V_t | validation state | checking status |
| K_t | authority state | role-relative authority |
| Ω_t | open continuation context | active branches and unresolved tasks |
| L_t | re-entry links | paths back to source |
| Δ | governed update operator | accepted state change |
| α | authority function | assigns artifact authority |
| ν | validation function | assigns validation state |
| η | enactment relation | records authority transition |

Plain-language summary: X_t is the current condition of the research project. R_t is what the record preserves. Replay asks what can be rerun from the record. Resumption asks whether the project state can be restored well enough to continue.

4. A Tiered Axiom System

The axiom system is organized into three tiers. The tiers are not ornamental. They prevent the manuscript from pretending that all axioms have the same status. Tier 1 contains the elementary mathematical spine. Tier 2 specifies transcript-state requirements. Tier 3 specifies governed-artifact requirements.

Table 2. Tiered Axiom Set

| Tier | Axiom | Formal core | Plain-language meaning | Scope |
|------|------------------------------------|-------------------------------------------------------|--------------------------------------------|-----------------------|
| 1 | A1 Record Projection | $R_t = \rho(X_t)$ | record is a projection of state | general |
| 1 | A2 Continuation Equivalence | AdmNext equality | resumption preserves admissible next moves | general |
| 2 | A3 Evidence-Mediated Working State | $C_t \rightarrow U_t \rightarrow E_t \rightarrow W_t$ | working state is evidence-derived | transcript-sufficient |

| Tier | Axiom | Formal core | Plain-language meaning | Scope |
|------|----------------------------------|---------------------|----------------------------------------|--------------------------|
| 2 | A4 Re-entry and Non-erasure | $x \rightarrow C_t$ | entries preserve source path | transcript-sufficient |
| 3 | A5 Artifact Identity | identity fields | artifacts must be identifiable | artifact-governed |
| 3 | A6 Proposal-Admissibility-Update | proposal not update | suggestions do not change state alone | automation-bearing |
| 3 | A7 Enacted Authority | enacted transition | authority changes visibly | governed-artifact |
| 3 | A8 Validation-Gated Advancement | validation gate | high-reliance artifacts require checks | trustworthy continuation |

Tier 1: mathematical spine

A1. Record Projection Axiom

There exists an operative research state X_t and an externalized record R_t such that Equation E3 holds. The record is a projection of the research state, not necessarily the full state. This axiom creates the formal site where state distinctions can be preserved or lost. Without it, projection collapse cannot be stated.

A2. Continuation Equivalence Axiom

Two states are equivalent for continuation when they license the same admissible next operations, as stated in Equation E1. The restored state must support the same responsible next moves. This defines the restoration target without requiring exact equality of all hidden details.

Tier 2: transcript-state spine

A3. Evidence-Mediated Working State Axiom

Working state is derived from admissible evidence selected from retained corpus.

Equation E9. Transcript-to-state derivation

$$C_t \rightarrow U_t \rightarrow E_t \rightarrow W_t.$$

For each load-bearing state entry $w \in W_t$, there exists supporting evidence.

Equation E10. Evidence support

$$\forall w \in W_t, \quad \exists e \in E_t \quad \text{with} \quad \text{supports}(e, w).$$

The current working state must be evidence-mediated, not free summary. This axiom separates memory from operational working state. P7 supplies the source model for transcript-to-state derivation (Bell 2026, *From Chat Corpus to Working State*).

A4. Re-entry and Non-erasure Axiom

For each load-bearing state entry or artifact, there exists a re-entry path back to source.

Equation E11. Re-entry

$$\forall x \in W_t \cup A_t, \quad \exists \ell \in L_t \quad \text{such that} \quad \ell : x \rightarrow C_t.$$

Governed updates preserve prior trace or explicitly mark deprecation, supersession, or correction. Important state and artifact entries must preserve a path back to source. This makes transcript sufficiency reconstructable rather than decorative.

Tier 3: governed-artifact spine

A5. Artifact Identity Axiom

Every state-bearing artifact has identity, type, lineage, and status.

Equation E12. Artifact identity

$$\forall a \in A_t, \quad \exists (\text{id}(a), \text{type}(a), \text{lineage}(a), \text{status}(a)).$$

A research system cannot govern or validate artifacts it cannot identify. P11 and P20 provide the artifact graph and directed artifact-representation basis for this axiom.

A6. Proposal-Admissibility-Update Separation Axiom

Let o be a candidate-generating operator. Generated proposals do not automatically change state.

Equation E13. Proposal is not update

$$o(W_t, E_t) = p \quad \not\Rightarrow \quad \Delta(X_t, p) = X_{t+1}.$$

For load-bearing updates, admission is required.

Equation E14. Admission condition

$$\Delta(X_t, p) = X_{t+1} \quad \Rightarrow \quad \text{Admit}(p, X_t, G_t).$$

A proposal is not an accepted update. This prevents generated material from becoming state merely by being produced. P8 supplies the bounded-operator model behind this axiom.

A7. Enacted Authority Axiom

Artifact authority changes only through governance-visible enactment.

Equation E15. Enacted authority

$$\begin{aligned} \alpha(a, r, t+1) &\neq \alpha(a, r, t) \\ \Rightarrow \exists g \in G_t \text{ such that } &\eta(g, a, \alpha(a, r, t), \alpha(a, r, t+1)). \end{aligned}$$

Authority is enacted, not inferred from visibility, recency, polish, or inclusion. This axiom preserves the distinction between presence and canonicity. P15 is the source basis.

A8. Validation-Gated Advancement Axiom

Let D_t be the diagnostic-axis set at time t , and let s_{high} denote a high-reliance target state such as released, canonical, or publication-facing. High-reliance advancement requires validation state.

Equation E16. Validation-gated advancement

$$\begin{aligned} &\text{Advance}(a, s_{\text{high}}) \\ \Rightarrow \nu(a, D_t) &\in \text{Allowed}(s_{\text{high}}). \end{aligned}$$

Artifacts should not advance into high-reliance state without recorded checking. This axiom separates generation from validation. P22 gives the artifact-integrity problem; P34 supplies sanity-check operators as a candidate operational layer.

Plain-language summary: A1 and A2 say what restoration means. A3 and A4 say how transcript-derived state remains evidence-based and re-enterable. A5 through A8 say how artifacts become governable, reviewable, authoritative, and checkable.




Figure 1: Record Projection and State Collapse. A record may preserve enough information to replay a path while collapsing states that differ for continuation. State-sufficient records must discriminate continuation-relevant distinctions.

5. Projection Collapse and the Replay/Resumption Corollary

The central result is elementary but useful. It follows from non-injective projection over continuation-relevant distinctions.

Lemma 1. Record Projection Collapse

Assume that two research states project to the same record.

Equation E17. Projection collapse

$$\rho(X_t) = \rho(X'_t) = R_t.$$

Assume also that they are not continuation-equivalent.

Equation E18. Continuation distinction

$$X_t \not\equiv_{\text{cont}} X'_t.$$

Then the shared record is not state-sufficient.

Equation E19. Record projection collapse lemma

$$\rho(X_t) = \rho(X'_t) \wedge X_t \not\equiv_{\text{cont}} X'_t \Rightarrow \neg \text{StateSuff}(R_t).$$

Proof sketch: if the same record is compatible with two states that license different admissible next operations, then the record alone cannot determine which continuation state should be restored. State sufficiency requires restoration up to continuation equivalence. Therefore the record is not state-sufficient.

This lemma is not deep mathematics. It is the projection-collapse mechanism that makes the research-continuation problem precise.

Corollary 1. Replayability Does Not Imply Resumability

Assume the conditions of Lemma 1 and assume the replay operator is defined. Then replayability does not imply resumability.

Equation E20. Replayability does not imply resumability

$$\text{Replayable}(R_t) \not\Rightarrow \text{Resumable}(R_t).$$

Proof sketch: replayability requires that the record determine a replayable path or execution class. Resumability requires that the restore operator recover a continuation-equivalent state. Since the same record may collapse continuation-distinct states, replayability may hold while resumability fails.

This is best presented as a corollary or design theorem, not as a surprising mathematical discovery. Its value is methodological. It identifies why action logs, transcripts, memory entries, and provenance traces may still fail as continuation records.

Design Criterion 1. Continuation-Discriminating Projection

Let \mathcal{X} be a target class of research states. If a record projection is state-sufficient for \mathcal{X} , then it must preserve continuation-relevant distinctions.

Equation E21. Continuation-discriminating projection criterion

$$\begin{aligned} \text{StateSuff}_{\mathcal{X}}(R_t) \Rightarrow \forall X, Y \in \mathcal{X}, \quad \rho(X) = \rho(Y) \\ \Rightarrow X \equiv_{\text{cont}} Y. \end{aligned}$$

This is the main design criterion. A state-sufficient record projection need not preserve all distinctions. It must preserve distinctions that change admissible next action.

This criterion is more useful than the headline corollary. It tells system designers what to ask: does the record distinguish states that require different next moves? If not, the record is not sufficient for continuation.

Table 3. Theorem / Proposition Register

| ID | Type | Statement | Status | Limit |
|----|----------------------------|-------------------------------------------------------------------|--------------------------|---------------------------------|
| T1 | Lemma | projection collapse prevents state sufficiency | core formal lemma | elementary |
| T2 | Corollary / design theorem | replayability does not imply resumability | headline result | replay remains useful |
| T3 | Design criterion | state sufficiency requires continuation-discriminating projection | main practical criterion | target-relative |
| P1 | Proposition | memory does not imply working state | derived | memory can be extended |
| P2 | Proposition | provenance does not imply authority | derived | provenance may inform authority |
| P3 | Proposition | presence does not imply canonicity | derived | artifact-dense systems |
| P4 | Proposition | generation does not imply validation | derived | proof-carrying exceptions |

| ID | Type | Statement | Status | Limit |
|----|--------------------|----------------------------------------------|------------------|--------------------------------------|
| P5 | Proposition | proposal does not imply enactment | derived | governed auto-admission possible |
| P6 | Proposition | publication does not exhaust research object | projection claim | simple papers may approximate object |
| P7 | Design proposition | governed state supports bounded continuation | synthesis | no truth/optimalty guarantee |

Plain-language summary: the core lemma is simple. If two different project states require different next actions but leave the same record, the record cannot tell you how to resume. That is the formal meaning of “replay is not resumption.”

6. Research-State Components: Evidence, Artifacts, Authority, Validation, and Re-entry

The projection criterion says that records must preserve continuation-relevant distinctions. The next question is which distinctions matter for research systems.

6.1 Evidence-mediated working state

In transcript-sufficient research, the current working state should be derived from evidence rather than from ungrounded memory or summary. P7 supplies the formal chain in Equation E9. This chain matters because transcript corpora are too large to function directly as working state and too valuable to erase through uncontrolled summarization. Evidence mediation creates a layer between archive abundance and operational state.

Let M_t denote a managed memory layer at time t . A memory layer may preserve or organize information, but working state requires operational selection and evidence support. Therefore memory does not imply working state.

Proposition equation P1. Memory does not imply working state

$$M_t = M'_t \not\Rightarrow W_t = W'_t.$$

6.2 Artifact identity

Research continuation depends on artifacts. A later user must know which draft, figure, table, registry, package, transcript, or publication object is being referenced. P11 supplies registry-state logic, and P20 supplies directed artifact graph representation.

Artifact presence is not enough. The system needs identity, type, lineage, and status. Without those, artifacts cannot be governed, validated, or restored into a continuation-equivalent state.

6.3 Authority and canonicity

Authority determines whether an artifact may be relied on for a role. P15’s core claim is that canonicity is not presence. An artifact can exist, be registered, be reviewed, or be included in a package without becoming authoritative. Authority changes through enactment.

Authority can be represented as follows, where \mathbf{a} is an artifact, \mathbf{r} is a role, and \mathbf{k} is an authority state.

Equation E22. Authority function

$$\alpha(a, r, t) = k.$$

Canonicity can be treated as a special authority state.

Equation E23. Canonical authority state

$$\alpha(a, r, t) = \text{canonical}.$$

The separate symbol κ is unnecessary in the main text. It can be defined in an appendix if needed.

Let $\text{prov}(a, b)$ denote a generic provenance relation from artifact **a** to artifact **b**. Provenance does not imply authority.

Proposition equation P2. Provenance does not imply authority

$$\text{prov}(a, b) \not\Rightarrow \alpha(a, r, t) = \text{canonical}.$$

Presence also does not imply canonicity.

Proposition equation P3. Presence does not imply canonicity

$$a \in A_t \not\Rightarrow \alpha(a, r, t) = \text{canonical}.$$

6.4 Validation state

Validation state records whether an artifact has passed specified checks. It does not guarantee truth. P22 establishes the non-equivalence between process transparency and artifact integrity. P34 turns this into an operator layer: sanity-check procedures for claims, methods, equations, citations, figures, tables, reproducibility conditions, and corpus relations.

Validation may be represented as follows.

Equation E24. Validation function

$$\nu(a, D_t) \in V_t.$$

Generation does not imply validation.

Proposition equation P4. Generation does not imply validation

$$a = o(X_t) \not\Rightarrow \nu(a, D_t) = \text{pass}.$$

A system may produce a fluent manuscript, plausible table, or attractive figure without checking whether the artifact is structurally reliable.

6.5 Proposal, admissibility, and update

P8 provides the key distinction. Operators may propose, but proposals do not update state unless admitted.

Proposition equation P5. Proposal does not imply enactment

$$p \in \text{Proposal}_t \not\Rightarrow \Delta(X_t, p) = X_{t+1}.$$

Proposal does not imply enactment. This is essential in automation-bearing systems because generated outputs can otherwise mutate state without review.

6.6 Publication as interface

Publication projection is defined as follows.

Proposition equation P6. Publication projection

$$P_t = \pi_{\text{pub}}(X_t).$$

In general, the publication object does not exhaust the research object. P5 treats the paper as a reader-facing projection of a deeper stack of transcripts, registers, rendered objects, containers, and continuity records. P16 models research as a media pipeline in which publication render, formal object, and trace are distinct layers.

6.7 Re-entry

Re-entry links preserve the path back from state or artifact to corpus. Without re-entry, a restored state may be usable but not auditable. P7 treats re-entry as a structural requirement, and P17’s constructive-closure criterion also depends on governed, reconstructable transformation paths.

Plain-language summary: research-state restoration is not generic checkpointing. It requires preserving the distinctions that determine scholarly continuation: evidence, artifacts, authority, validation, proposal status, publication interface, and source re-entry.

Figure 2. Tiered Axiom System

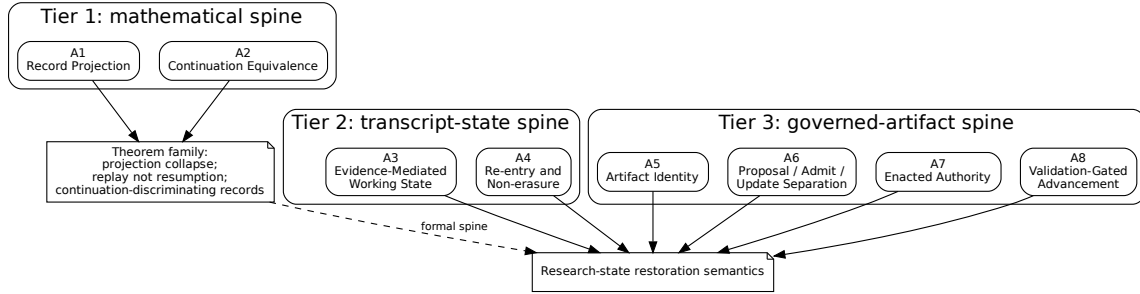


Figure 2: Tiered Axiom System. The axiom set separates the mathematical spine from transcript-state and governed-artifact requirements. Tiering prevents the elementary projection theorem from being overclaimed as a complete theory of research infrastructure.

7. Relation to Existing Infrastructure

The axiom system should not be presented as a replacement for existing infrastructure. Its best use is as a semantics layer that clarifies what existing infrastructures must preserve when the target is research continuation.

Workflow provenance already records important execution histories, inputs, outputs, and dependencies. Workflow Run RO-Crate, for example, is described by its working group as an extension of RO-Crate and Schema.org for capturing workflow execution provenance at different granularities and bundling associated products such as inputs, outputs, and code; the associated 2024 PLOS ONE article is *Recording provenance of workflow runs with RO-Crate* (Leo et al. 2024). The present paper treats workflow-run provenance as a strong substrate for replay and partial restoration, while adding that scholarly continuation also requires authority state, validation state, proposal status, evidence support, and publication-interface status.

W3C PROV is also a close substrate. PROV defines provenance as information about entities, activities, and people involved in producing a data item or thing, and its family of documents supports interoperable provenance exchange through models and serializations (Groth and Moreau 2013). The present paper does not replace PROV. It says that provenance alone does not decide which artifact is authoritative, which recommendation was enacted, which object passed validation, or which next action is admissible.

Figure 3. Layer Ladder as State Refinement

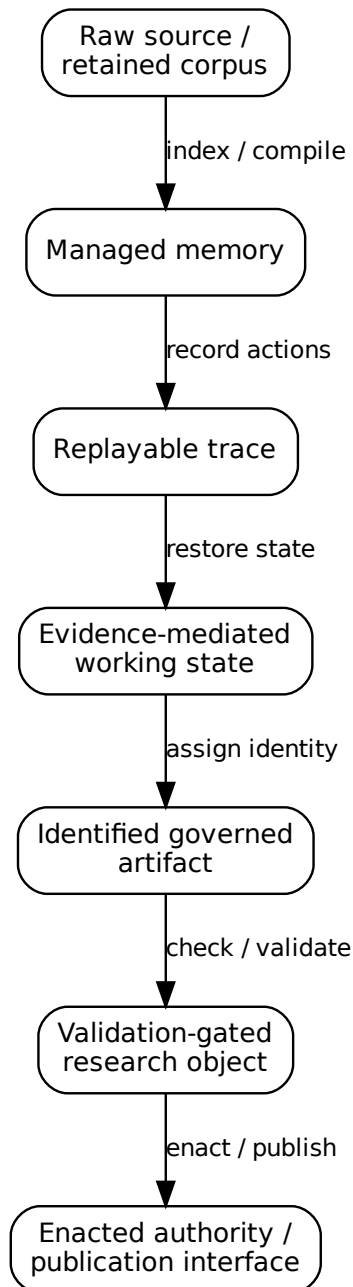


Figure 3: Layer Ladder as State Refinement. Each layer adds structure not guaranteed by the previous layer. The ladder expresses non-equivalence claims rather than a mandatory universal workflow.

RO-Crate is a research-object packaging substrate. RO-Crate Metadata Specification 1.2 is the newest release, published on 2025-06-04, and its specification includes sections for data entities, contextual entities, provenance of entities, profiles, workflows and scripts, and RO-Crate JSON-LD (Sefton et al. 2025). A governed research-state package could plausibly be encoded using RO-Crate-like structures. But the package format does not itself define the restoration semantics proposed here.

Process mining provides tools for event logs, conformance checking, process discovery, and operational support. The IEEE Task Force on Process Mining describes the Process Mining Manifesto as intended to promote process-mining maturity and improve process design, control, and support (IEEE Task Force on Process Mining 2012). A transcript-sufficient research system could emit process-minable events such as capture, segment, select evidence, propose, review, validate, enact, publish, deprecate, and restore. Process mining can help analyze behavior. It does not by itself define artifact-authority semantics.

Checkpoint/restore systems are the closest mathematical analogue. The paper’s central result is not novel relative to checkpointing. A log is not a checkpoint; replay is not restoration. CRIU’s documentation states that it can freeze a running application or container, checkpoint its state to disk, and later use the saved data to restore the application and run it as it was at the time of the freeze (CRIU Project 2025). The contribution here is to specify what checkpoint-like restoration means for research state rather than process state alone.

Software configuration management is the strongest practical analogue. Git’s official documentation describes commits as objects pointing to snapshots of staged content and parent commits, and it explains that branch switching can revert files in a working directory to the snapshot a branch points to (Chacon and Straub 2014). Pull requests, review, tests, merge, tags, releases, and issue histories already model many distinctions similar to proposal, validation, enactment, and release. The present paper should therefore treat software configuration management as a close analogue, while noting that research systems add evidence, scholarly authority, publication interface, transcript sufficiency, and epistemic validation constraints.

Table 4. Comparator Infrastructure Matrix

| Infrastructure | Preserves well | Overlap | State-restoration gap | Use in paper |
|-----------------------------------|----------------------------------------------------------------|-----------------------------------|-----------------------------------------------------------------|----------------------|
| Workflow provenance | executions, inputs, outputs, dependencies | replay/provenance | scholarly authority and validation state | substrate |
| W3C PROV | entities, activities, agents, derivations | lineage | admissible next action and canonicity | provenance substrate |
| RO-Crate / Workflow Run | package, metadata, workflow-run | research object record | semantics of restoration not intrinsic | package substrate |
| RO-Crate | provenance | | | |
| Process mining | event logs, conformance, process discovery | event-layer analysis | artifact authority not intrinsic | instrumentation |
| Checkpoint/restore | saved state and process restoration | closest formal analogue | computational state, not scholarly state | analogue |
| Software configuration management | snapshots, branches, history, review/release analogues | proposal/review/release analogues | evidence and scholarly authority semantics need domain modeling | practical analogue |
| Reflexive Laboratory | corpus, evidence, artifacts, governance, validation, authority | worked interpretation | not universal | case |

Plain-language summary: existing infrastructure is not the enemy of this paper. It is the implementation landscape. The paper’s contribution is to say what state semantics those systems need when the object is research

continuation.

8. Worked Interpretation: Reflexive Laboratory

The Reflexive Laboratory is a worked transcript-sufficient case. It should not be treated as the only possible implementation of the axiom system.

The mathematical spine is already present in P10. That paper defines operative state, record projection, replay, restoration, execution sufficiency, and state sufficiency. It also states that a record may be replayable without being resumable when the projection collapses state.

The transcript-state spine is present in P7. It defines a retained corpus, indexed transcript units, admissible evidence, working state, re-entry links, and an update operator.

The proposal/update spine is present in P8. Bounded autoresearch operators can surface evidence, detect discrepancies, flag low-confidence gaps, generate state-update proposals, and check bounded consistency. But they do not replace human admissibility judgment or silently update working state.

The artifact-identity spine is present in P11 and P20. P11 treats the research lab as a governed artifact graph and uses the Master Global Artifact Index to record artifact identities, statuses, dependencies, and authority relations. P20 treats the corpus as a directed artifact graph with papers, transcripts, and support artifacts linked by typed relations.

The authority spine is present in P15. That paper distinguishes source presence, registration, proposal/review, enactment, and canonicity. It argues that artifact authority depends on governance-visible transition rather than mere inclusion in an archive or transcript.

The validation spine is present in P22 and P34. P22 distinguishes process transparency from artifact integrity; P34 proposes sanity-check operators as an operational layer before stabilization or canonicalization.

The publication-projection spine is present in P5 and P16. P5 treats the paper as one interface to a deeper research object, while P16 models research as a media pipeline in which source, intermediate knowledge, formal object, publication render, and trace are distinct layers.

In this interpretation, the Reflexive Laboratory satisfies the axiom system as a candidate architecture. It preserves record projection and state-sufficiency distinctions; derives working state from transcript evidence; preserves re-entry; identifies artifacts; separates proposal from update; enacts authority; records validation; and treats publication as interface.

This does not prove that the system is optimal. It does not prove that its outputs are true. It does not prove that all research should be organized this way. It only shows that the Reflexive Laboratory supplies a concrete case where the axiom system has an interpretation.

9. Limits and Nonclaims

The paper's main theorem family is elementary. This must be stated directly. A non-injective projection cannot be inverted into a unique state. A replayable log is not a checkpoint. A record that collapses continuation-distinct states is not state-sufficient. The paper does not claim otherwise.

The contribution is not mathematical novelty in the abstract. It is a domain-specific formalization for research systems. The paper defines continuation equivalence as the restoration target and identifies the research-state components that should be preserved when continuation is the goal.

The axiom system is target-relative. A small project may not require all tiers. A simple computational workflow may only need workflow provenance and checkpointing. A manuscript-only project may not need a full artifact graph. A transcript-sufficient, AI-assisted, artifact-dense research program requires more.

Authority is not truth. A canonical artifact may later be corrected, deprecated, or superseded. Canonicity records enacted authority for a role.

Figure 4. Research-State Restoration Semantics

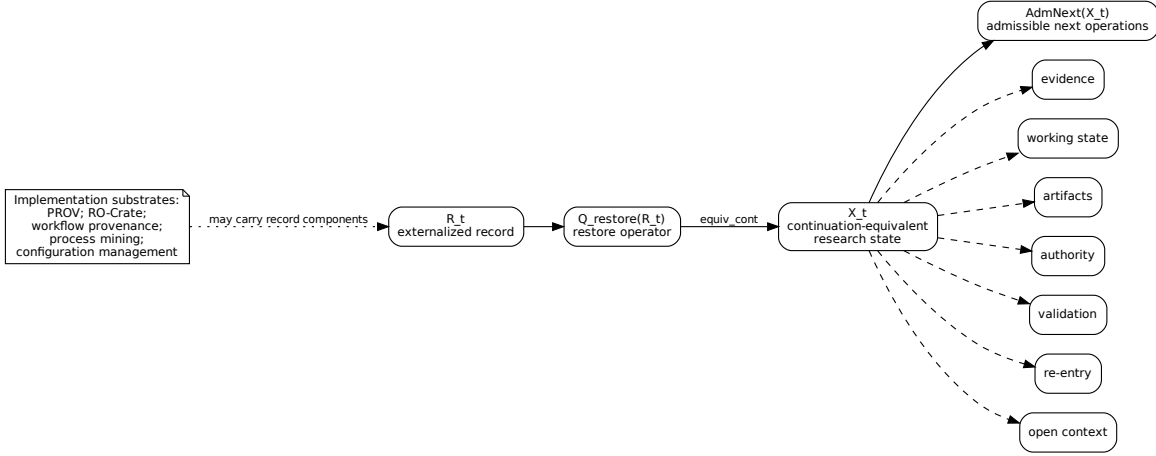


Figure 4: Research-State Restoration Semantics. The restoration target is not the full past process but a continuation-equivalent research state. Existing infrastructure can carry parts of that state, but the record must preserve the distinctions that determine admissible next operations.

Validation is not correctness. A validated artifact has passed specified checks. It may still be wrong. Validation state records review status, not infallibility.

Publication is not the whole research object. It is a reader-facing projection. The deeper research object may include transcripts, evidence registers, artifact graphs, validation records, source specifications, packages, and governance traces.

The Reflexive Laboratory is not universal. It is a worked case. Other systems may implement the same semantics through workflow engines, semantic graphs, research-object packages, version control, process logs, or specialized databases.

The axiom system also remains incomplete. It does not yet include a full uncertainty model, a quantitative validation model, a full deontic logic of governance, or a category-theoretic operator algebra. Those may be later papers. They are not necessary for the present contribution.

Plain-language summary: the paper should not oversell itself. Its durable claim is small but useful: a record for continuing research must preserve the distinctions that change what the next valid research move can be.

10. Conclusion

This paper has repositioned the axioms of research continuation as a state-restoration theory for research systems. The core problem is simple: a record can preserve enough information to replay a process without preserving enough state to resume the research.

The formal spine contains two basic ideas. First, a research record is a projection of operative state:

Equation E3. Record projection

$$R_t = \rho(X_t).$$

Second, restoration should be judged by continuation equivalence:

Equation E1. Continuation equivalence

$$X_t \equiv_{\text{cont}} X'_t \iff \text{AdmNext}(X_t) = \text{AdmNext}(X'_t).$$

From these, the projection-collapse result follows. If two states are not continuation-equivalent but project to the same record, the record is not state-sufficient. If the same record still supports replay, then replayability does not imply resumability. The mathematically elementary result becomes useful when applied to research systems.

The domain-specific claim is that research states contain more than action history. They include evidence-mediated working state, artifact identity, proposal/admissibility/update status, authority, validation, re-entry, publication interface, and open continuation context. These distinctions determine what can responsibly happen next.

The revised axiom system is tiered. The first tier supplies the mathematical spine. The second supplies transcript-state requirements. The third supplies governed-artifact requirements. This tiering makes the contribution more precise and reduces overclaiming.

Existing infrastructures remain central. Workflow provenance, W3C PROV, RO-Crate, process mining, checkpoint/restore, and software configuration management are not displaced by this model. They are close analogues and implementation substrates. The axiom system specifies the research-state semantics that such infrastructures must carry when the goal is accountable continuation.

The Reflexive Laboratory provides a worked case because its prior papers have already developed the required pieces: transcript-to-state derivation, bounded automation, record sufficiency, artifact graph governance, canonicity, media projection, constructive closure, corpus graphing, and validation operators. The present paper composes those pieces into a compact formal kernel.

The next step is not more formal expansion. The next step is review of the built artifact and any final release decision. The paper should remain narrow: continuation-discriminating record projection and research-state restoration semantics.

Appendix Notes

Appendix A should contain the full v01/v02 axiom crosswalk. Its purpose is to preserve the reduction path from the broader fifteen-axiom exploratory system to the tiered eight-axiom v03 system.

Appendix B should contain formal module details for the graph module, governance module, validation module, and publication-projection module.

Appendix C should contain the expanded comparator matrix, including workflow provenance, W3C PROV, RO-Crate / Workflow Run RO-Crate, process mining, checkpoint/restore, software configuration management, managed memory systems, AutoResearch loops, and the Reflexive Laboratory.

Appendix D should contain the math-rendering protocol, including equation IDs, LaTeX-safe source syntax, rendered-page checks, symbol-definition checks, and visual PDF audit requirements.

References

Internal Reflexive Laboratory sources

- Bell, Peter. 2026. *The Reflexive Laboratory as a Scientific Computing Stack for Research: From Transcript-Governed Inquiry to Containerized Research Objects*. Zenodo. Version DOI: 10.5281/zenodo.18991940. Concept DOI: 10.5281/zenodo.18902418.
- Bell, Peter. 2026. *From Chat Corpus to Working State: Machine-Readable Methods for Transcript-First Research*. Zenodo. Version DOI: 10.5281/zenodo.19020479. Concept DOI: 10.5281/zenodo.18902418.
- Bell, Peter. 2026. *Bounded Autoresearch: Operator-Constrained Automation in Transcript-Derived Research Systems*. Zenodo. Version DOI: 10.5281/zenodo.19026495. Concept DOI: 10.5281/zenodo.18902418.
- Bell, Peter. 2026. *Execution Sufficiency and State Sufficiency: Record Sufficiency in Reflexive Research Systems*. Zenodo. Version DOI: 10.5281/zenodo.19042930. Concept DOI: 10.5281/zenodo.18902418.
- Bell, Peter. 2026. *The Research Lab as a Graph: Snapshot Governance in a Transcript-Sufficient Research System*. Zenodo. Version DOI: 10.5281/zenodo.19045093. Concept DOI: 10.5281/zenodo.18902418.
- Bell, Peter. 2026. *Canonicity Is Not Presence: Artifact Authority in Transcript-Sufficient Research Systems*. Zenodo. Version DOI: 10.5281/zenodo.19104311. Concept DOI: 10.5281/zenodo.18902418.
- Bell, Peter. 2026. *Research as a Media Pipeline: Compression, Reconstruction, and the Role of Trace in Mathematical Objects*. Zenodo. Version DOI: 10.5281/zenodo.19119710. Concept DOI: 10.5281/zenodo.18902418.
- Bell, Peter. 2026. *Constructive Closure and Reflexive Realization: A Minimal Criterion for Process-Legible Research Objects*. Zenodo. Version DOI: 10.5281/zenodo.19142333. Concept DOI: 10.5281/zenodo.18902418.
- Bell, Peter. 2026. *The Corpus as Graph: Directed Artifact Representation in the Reflexive Laboratory*. Zenodo. Version DOI: 10.5281/zenodo.19153448. Concept DOI: 10.5281/zenodo.18902418.
- Bell, Peter. 2026. *Process Transparency Is Not Artifact Integrity: Why Automated Research Systems Need Structural Validation*. Zenodo. Version DOI: 10.5281/zenodo.19157570. Concept DOI: 10.5281/zenodo.18902418.
- Bell, Peter. 2026. *Sanity-Check Operators for Buggy Research in a Reflexive Laboratory*. Zenodo. Version DOI: 10.5281/zenodo.19226488. Concept DOI: 10.5281/zenodo.18902418.
- Bell, Peter. 2026. *Replay Is Not Resumption: AutoResearch and the Architecture of Research Continuation: Managed Memory, Governed State, and the Reflexive Laboratory*. Zenodo. Version DOI: 10.5281/zenodo.20500050. Concept DOI: 10.5281/zenodo.18902418.

External infrastructure sources

- Andrijauskas, Fabio, Igor Sfiligoi, Diego Davila, Aashay Arora, Jonathan Guiang, Brian Bockelman, Greg Thain, and Frank Wuerthwein. 2024. *CRIU – Checkpoint Restore in Userspace for Computational Simulations and Scientific Applications*. arXiv:2402.05244. DOI: 10.48550/arXiv.2402.05244.
- Chacon, Scott, and Ben Straub. 2014. *Pro Git*. 2nd ed. Apress / Git official online book. <https://git-scm.com/book/en/v2>.
- CRIU Project. 2025. *Checkpoint/Restore In Userspace Documentation*. https://criu.org/Main_Page.
- Groth, Paul, and Luc Moreau, eds. 2013. *PROV-Overview: An Overview of the PROV Family of Documents*. W3C Working Group Note, 30 April 2013. <https://www.w3.org/TR/prov-overview/>.
- IEEE Task Force on Process Mining. 2012. *Process Mining Manifesto*. In *Business Process Management Workshops*, Lecture Notes in Business Information Processing 99, 169-194.
- Leo, Simone, Michael R. Crusoe, Laura Rodriguez-Navas, Raul Sirvent, Alexander Kanitz, Paul De Geest, Rudolf Wittner, Luca Pireddu, Daniel Garijo, Jose M. Fernandez, et al. 2024. *Recording Provenance of Workflow Runs with RO-Crate*. *PLOS ONE* 19(9): e0309210. DOI: 10.1371/journal.pone.0309210.

Sefton, Peter, Eoghan O Carragain, Stian Soiland-Reyes, Oscar Corcho, Daniel Garijo, Raul Palma, Frederik Coppens, Carole Goble, Jose Maria Fernandez, Kyle Chard, et al. 2025. *RO-Crate Metadata Specification 1.2*. researchobject.org community. Version DOI: 10.5281/zenodo.13751027.