PRESENT ACTION IS NOT PERSISTENT WRITE: CAUSAL RIGHTS OF OPTIMIZER HISTORY IN REPLAY-FREE CONTINUAL LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Stateful optimizers normally grant a learning signal two permissions at once: it may change parameters on the present update, and it may write persistent state that conditions future updates. We ask whether these permissions are experimentally separable in replay-free continual learning. We study a strict replay-free Split CIFAR-100 learner with one fixed 100-class head, current-task-only training, raw global inference, and no task routing, model growth, or post-hoc calibration. In a five-seed action×write factorial, a write-only intervention—denying the source its same-update parameter action while retaining its persistent Adam write—increases current acquisition by 0.00774 relative to the comparator, while increasing forgetting by 0.00604 and old-to-final recapture by 0.00569; the action-only cell increases acquisition by only 0.00058. Successive interventions localise this durable pathway predominantly to first-moment state, accumulated within-task history, and the current-relative orthogonal action of inherited history. We then follow acquisition-bearing history across the task lifecycle: in four independent paired seeds, with six origin tasks nested within each seed, all 24 provenance units show source-distinct classifier weights and Adam moments at the current-to-old boundary followed by source-distinct row-relative old-row displacement under exactly zero natural classifier gradient. This provides direct cross-lifecycle provenance evidence, not behavioural mediation. Finally, a function-preserving reparameterisation with transported Adam state leaves final features, all 100 logits, predictions, and the classifier head exactly unchanged at the branch, yet all four fresh pairs diverge on the first lawful subsequent update and develop different continual-learning trajectories. Within this tested Adam setting, present parameter action and persistent-history ownership are distinct causal rights, and persistent optimizer history contributes to the causal state governing future learning.

1 INTRODUCTION AND NOVELTY BOUNDARY

Stateful optimizers couple two operations that are usually treated as one. A learning signal contributes to the parameter transition made now, and the same signal updates persistent optimizer state that changes how later gradients are converted into parameter motion. Adam makes this coupling explicit through its first- and second-moment recurrences (Kingma & Ba, 2015). Persistent optimizer state and history-dependent optimization are established. Optimizer choice can materially affect catastrophic forgetting (Ashley et al., 2021). Gradient-based optimizers can function as memory-bearing systems that compress or depend on prior gradients and iterates (Behrouz et al., 2025; Cattaneo & Shigida, 2025). Controlled optimizer-state resets can expose history-dependent training memory (Sevetlidis & Pavlidis, 2026), and recent causal work identifies first-moment optimizer state as a carrier of source-specific effects (Xu, 2026). Our question is narrower and temporally explicit: *can the same source influence the present parameter transition independently of whether it is written into the state carried forward?*

Continual learning makes this distinction scientifically consequential. In class-incremental learning, one shared classifier must keep classes learned at different times mutually comparable. New-class bias and classifier imbalance are established issues in class-incremental learning (Wu et al., 2019; Zhao et al., 2020); cross-task class discrimination is an explicit challenge (Guo et al., 2023). Rep-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

resentation forgetting has been directly probed (Davari et al., 2022). Class-incremental prediction has also been decomposed into within-task prediction and task-identification or out-of-distribution-related components (Kim et al., 2022; 2023). Trivedi & Melwani explicitly frame catastrophic forgetting in terms of accessibility collapse (Trivedi & Melwani, 2026). We therefore do not claim accessibility failure as a new concept, and task-restricted performance is used only as a diagnostic of locally usable discrimination, not as proof that internal representations are unchanged. Instead, we use a deliberately strict replay-free fixed-head setting because it exposes a controlled failure mode in which current acquisition, global selectability, optimizer history, and task lifecycle can be manipulated separately.

The paper's central result is a *causal-rights factorisation*. For a source learning signal s at update t, we distinguish present-action authority from persistent-write authority. An action-only intervention allows s to affect the current Adam parameter transition but removes its direct contribution to the optimizer state stored for the next update. A complementary write-only intervention denies s direct participation in the same-update parameter action while allowing s to modify persistent optimizer state. Together with the neither and full-rights cells, these interventions form a physical 2×2 causal square rather than a comparison between loosely related training configurations.

The write-only cell is the causal hinge. In five fresh seeds, persistent write without same-update source action increases current acquisition by 0.00774 relative to the comparator, while increasing forgetting by 0.00604 and old-to-final recapture by 0.00569. The action-only cell changes current acquisition by only 0.00058 in the same factorial. Thus, in this tested source, a learning signal can be denied direct present action yet still write a state whose later lawful use changes behaviour. Conversely, present action does not imply ownership of future optimizer history.

We then ask what is written, when that history becomes material, how inherited history acts, and what survives when the task ceases to be current. Successive source-local interventions restrict the durable pathway predominantly to Adam first-moment state, accumulated within-task history, and the component of inherited history orthogonal to the contemporaneous controller. First-moment causal transport is prior art (Xu, 2026), and orthogonal-gradient methods are prior art (Farajtabar et al., 2020); our contribution is the controlled causal organisation of these objects inside the same replay-free shared-head forgetting pathway.

The final experiments extend the mechanism longitudinally rather than broadening its empirical scope. Acquisition-history formation produces source-distinct classifier parameters and optimizer moments at the current-to-old transition, and those rows later undergo source-distinct physical displacement while their natural classifier gradient is exactly zero. This is cross-lifecycle causal provenance, not behavioural mediation. A separate predictor-identical stress test then asks whether the instantaneous predictor is sufficient to specify the continuing learner. Function-preserving rescaling is established prior art (Neyshabur et al., 2015), and training sensitivity to parameter scaling and optimizer/parameterisation choices is also established (Li et al., 2022; Everett et al., 2024); here these ideas are inserted only after persistent-history causality has been localised, with existing adaptive state transported rather than reset.

The resulting contribution is therefore not that optimizers have memory, that first moments can carry history, or that equivalent predictors can train differently. It is the experimental factorisation of who may act, who may write, what persistent state carries the source-local phenotype, how that state accumulates and acts, and how its physical consequences cross the task lifecycle. The empirical claims remain confined to the tested Split CIFAR-100 / SmallCNN / Adam setting.

2 Experimental reduction and causal-rights formalism

We use CIFAR-100 split into ten sequential tasks of ten classes. The model is a compact convolutional network with two convolutional feature layers and one fixed learned 100-class linear classifier. Training uses only the current task. The primary behavioural prediction is raw argmax over all 100 classes. No experiment in the focal causal chain uses replayed old training samples, generated replay, stored old features or logits, prototypes, task routing, output masking, model or head growth, or post-hoc calibration. Unless a focal intervention changes a specified object, optimization uses Adam with learning rate 10^{-3}, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, batch size 64, and three epochs per task.

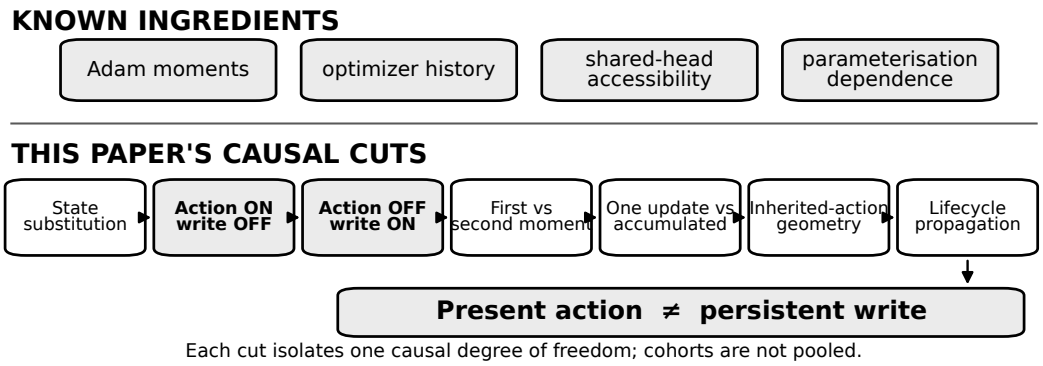

Figure 1: Causal intervention lattice. Established ingredients are separated from the seven interventions used here: optimizer-state substitution; present action without persistent write; persistent write without same-update source action; first- versus second-moment persistence; one-update versus accumulated history; inherited-action geometry; and task-lifecycle propagation. The intervention sequence motivates the central distinction: present action is not persistent write.

Table 1: Causal-rights cells for the focal source.

Cell	Present action A_s	Persistent write W_s
00	OFF	OFF
10	ON	OFF
01	OFF	ON
11	ON	ON

Task-restricted accuracy is a read-only diagnostic. It asks whether a task's classes remain mutually discriminable when competition is restricted to that task's ten outputs. It is not the deployed inference rule, and preserved task-restricted accuracy is not interpreted as proof that the internal representation is unchanged. Forgetting and recapture are reported separately from current acquisition; low forgetting produced by weak learning is not treated as protection.

Source construction. The causal-rights experiments use a source signal already localised within the shared backbone. Let $g_{\mathrm{C}}(t)$ denote the backbone gradient produced by the controller with the focal auxiliary non-target-suppression backbone route removed, while preserving the other current-learning, classifier-head, comparison, and capture-control routes. Let $g_{\mathrm{N}}(t)$ denote the backbone gradient from that auxiliary route on the same realised batch and model state. Over the concatenated backbone parameter vector, the focal source is the component of $g_{\mathrm{N}}(t)$ orthogonal to the contemporaneous controller gradient,

$$g_s(t) = g_{\mathrm{N}}(t) - \mathrm{proj}_{g_{\mathrm{C}}(t)}\, g_{\mathrm{N}}(t), \qquad \langle g_s(t), g_{\mathrm{C}}(t)\rangle = 0. \tag{1}$$

The source-OFF comparator and source-ON route are

$$g_0(t) = g_{\mathrm{C}}(t), \qquad g_1(t) = g_{\mathrm{C}}(t) + g_s(t). \tag{2}$$

The projection is computed over the shared backbone gradient vector. The comparator retains the remaining controller and classifier-head routes and removes only this focal backbone source; the source therefore affects the shared backbone, not the classifier head directly.

Causal-rights notation. For the focal source s, let

$$A_s(t) \in \{0,1\}, \qquad W_s(t) \in \{0,1\} \tag{3}$$

denote present-action authority and persistent-write authority. $A_s(t) = 1$ permits the source to influence the current parameter transition; $W_s(t) = 1$ permits the source to modify Adam state retained into later updates.

Write-only (01) means that the source is denied direct participation in the same-update parameter action but may modify optimizer state retained into later lawful updates. It does *not* mean that written history is forbidden from acting on later updates.

Adam implementation. Let (m_t, v_t, k_t) be the pre-step Adam state and let $g_0 = g_C(t)$, $g_1 = g_C(t) + g_s(t)$. From the same pre-state construct

$$m_0 = \beta_1 m_t + (1-\beta_1) g_0, \qquad v_0 = \beta_2 v_t + (1-\beta_2) g_0^2, \tag{4}$$

$$m_1 = \beta_1 m_t + (1-\beta_1) g_1, \qquad v_1 = \beta_2 v_t + (1-\beta_2) g_1^2. \tag{5}$$

The present parameter transition uses the Adam candidate associated with g_A, where $A = A_s(t)$, including the native step counter and bias correction. The state stored after the update is (m_W, v_W), where $W = W_s(t)$. Thus 10 uses the source candidate to move parameters but stores comparator history; 01 uses the comparator candidate for the present move but stores full-source history.

3 Why this causal question arises

The rights question emerged from a controlled accessibility failure rather than from a generic concern about Adam. In the fixed shared head, task-restricted competence and raw-global selectability can diverge, and current acquisition and forgetting are experimentally separable rather than forming an unavoidable lockstep. A later causal decomposition further distinguishes local learning, global expression, and recapture control: restoring expression can increase recency capture, while recapture control can reduce forgetting yet suppress newest-task performance. These observations motivate a mechanism study without implying a universal accessibility theory.

Historical optimizer state also has a lifecycle. Completed classifier rows can continue to move under inherited Adam state after their natural new-task gradient has vanished. Interrupting sustained optimizer-driven afterlife changes the behavioural phenotype, and completed-set comparison geometry independently changes accessibility. The key question is therefore not only whether optimizer state matters, but which learning signals are permitted to own the persistent state that remains active after direct gradient authority changes.

4 The causal-rights factorial

4.1 Optimizer state as a causal control object

Before separating rights, we tested whether optimizer state itself had causal leverage in the focal source. In a reciprocal state-substitution experiment, only the backbone Adam first- and second-moment tensors were exchanged between source-ON and source-OFF recipients immediately after each recipient's scientific gradient had been constructed and before the Adam parameter step. Recipient model parameters, classifier state, current gradients, Adam step counters, data stream, support state, and random state remained native. Donor moment state moved both recipients toward the donor acquisition/accessibility phenotype. This establishes source-local causal control by optimizer state. Optimizer choice can materially affect forgetting (Ashley et al., 2021); optimization dynamics can carry or reveal training history (Behrouz et al., 2025; Cattaneo & Shigida, 2025; Sevetlidis & Pavlidis, 2026); and recent causal work identifies first-moment optimizer state as a source-effect carrier (Xu, 2026). None of those broader facts is claimed as novel here.

4.2 Action and write are independently assignable

The causal-rights intervention starts from the same pre-step backbone parameters, controller gradient g_C, source g_s, Adam state (m_t, v_t), and step counter. From these, two candidate recurrences are constructed: the comparator recurrence produced by g_C and the full-source recurrence produced by $g_F = g_C + g_s$. Present-action authority selects which recurrence is used for the current parameter transition; persistent-write authority independently selects which recurrence is stored into optimizer state after that transition.

The action-only cell therefore uses the full-source Adam candidate for the present parameter step and then stores only the comparator recurrence. This removes the source's direct persistent write

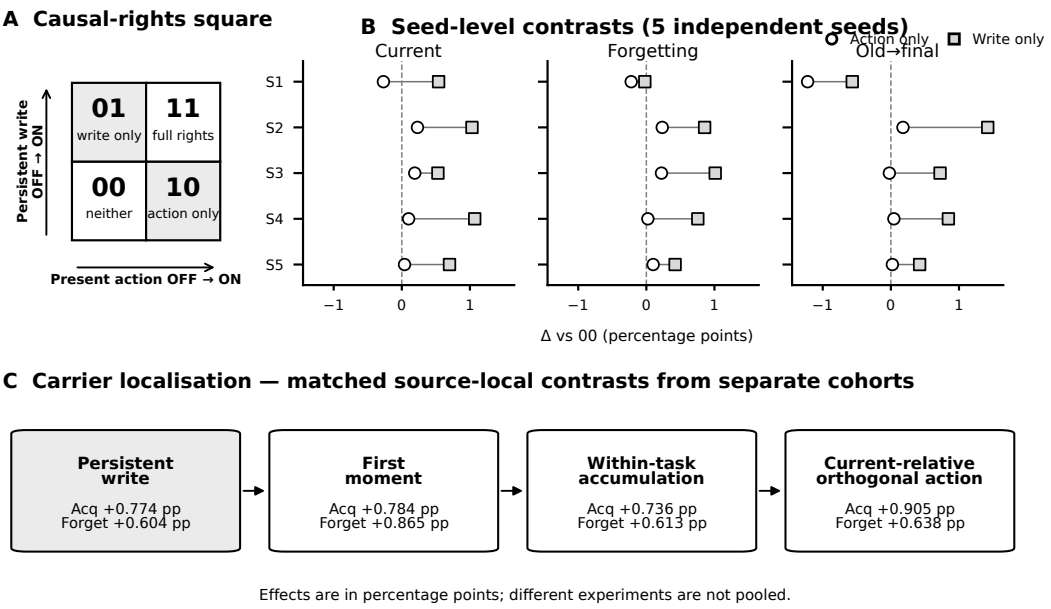

Figure 2: The causal-rights result and subsequent carrier localisation. **A:** the 2×2 rights square separates present-action authority from persistent-write authority. **B:** paired seed-level contrasts from the five independent factorial seeds; circles are action-only minus comparator and squares are write-only minus comparator, in percentage points. Lines connect the two contrasts for the same seed. No significance test is retrofitted; the prospective experiment used fixed seed/cohort and materiality rules. **C:** successive source-local interventions restrict the carrier to persistent write, first-moment state, within-task accumulation, and current-relative orthogonal inherited action. Panel C reports matched mean effects from separate cohorts; those cohorts are not pooled.

without erasing indirect future consequences of the parameter change: later forwards and gradients may still differ because the model moved differently now. In an eight-seed experiment, removing direct persistent ownership eliminated most of the native source route's forgetting penalty while retaining only a small positive acquisition advantage over the source-OFF comparator.

The complementary five-seed factorial supplies the missing write-only cell. The four rights combinations are 00 (neither), 10 (present action only), 01 (persistent write only), and 11 (full rights). In the five-seed panel, the action-only effect on current acquisition is $+0.00058$ relative to 00. The write-only cell, despite taking the comparator same-update source action, increases current acquisition by $+0.00774$, forgetting by $+0.00604$, and old-to-final recapture by $+0.00569$. The full-rights acquisition movement is approximately $+0.00918$, and the action$\times$write interaction is small on the primary acquisition endpoint. Descriptively, the write-only contrast accounts for about 84% of the full source-versus-comparator acquisition movement in this exact factorial.

The write-only manipulation is source-internally exact: its same-update parameter transition follows the comparator Adam candidate, while its persistent first- and second-moment state follows the full-source recurrence. This does not mean that the write-only trajectory must remain identical to the comparator on later updates. Once the written history is carried forward, it may lawfully influence later Adam actions. The treatment therefore isolates the temporal distinction of interest: write now, possible action later.

A learning signal can thus be denied same-update source action while retaining permission to write persistent optimizer history whose later use changes behaviour. Conversely, a signal can act on parameters now while being denied direct ownership of the stored history. Present parameter action and persistent-history ownership are experimentally separable causal rights.

5 LOCALISING THE PERSISTENT CARRIER

The causal square identifies persistent write as an acquisition- and cost-bearing object. We next restrict that object without changing the scientific question.

What is written? Holding same-update parameter action to the comparator law, we cross first- and second-moment persistence. First-moment-only persistence reproduces the direct-write phenotype: current acquisition is $+0.00784$ relative to comparator, with the associated forgetting/capture cost moving in the same direction. Second-moment-only persistence is approximately comparator-like on the primary endpoints. The source-local conclusion is therefore that the direct-write phenotype is carried predominantly by Adam's signed first-moment history. This does not imply that second moments are generally irrelevant, and first-moment causal transport itself is not claimed as novel (Xu, 2026).

When does it become material? A recurrence-provenance intervention distinguishes the current source innovation from accumulated within-task history and additional treatment-specific history carried across task boundaries. Retaining only the current innovation is approximately comparator-like. Allowing source-specific first-moment history to accumulate within the task recovers most of the acquisition and forgetting/recapture movement. Additional source-specific cross-task carry contributes only submaterial movement in this source. The material object is therefore accumulated within-task history rather than a single update of momentum.

How does inherited history act? We next hold future persistence fixed and alter only the component of already accumulated first-moment history allowed to affect the current parameter transition. The component parallel to the contemporaneous controller is approximately comparator-like. The current-relative orthogonal component reproduces the full accumulated-history phenotype, with current acquisition $+0.00905$ and forgetting $+0.00638$ relative to comparator. Orthogonality is not a safety principle here: the same component carries acquisition and cost. The distinct object is inherited optimizer history measured relative to the controller that is active when that history acts. This differs from using orthogonal projection itself as a general continual-learning protection rule (Farajtabar et al., 2020).

These interventions produce a compact source-local chain: persistent write $\rightarrow$ first-moment state $\rightarrow$ accumulated within-task history $\rightarrow$ current-relative inherited action. Each arrow corresponds to an intervention that removes or restricts a causal degree of freedom; the chain is not inferred from observational telemetry or pooled across experiments.

6 HISTORY CROSSES THE TASK LIFECYCLE

Localising the carrier still leaves a longitudinal question: what survives when the task that formed the history ceases to be current? We test this with a read-only lifecycle-provenance experiment. A common learner state is branched into history-OFF and history-ON conditions while a task block is current; the comparison/protection context is held fixed. The observer follows that task's classifier rows through the exact current-to-old transition and then through their native old-row dynamics.

The experiment contains four genuinely independent paired seeds. Within each seed, six origin tasks are nested observations, giving 24 seed$\times$origin evaluation units but not 24 independent population replications. The independent sample comprises four paired seeds; the nested units test whether the longitudinal pattern instantiates across task origins within each seed.

At the current-to-old boundary, all 24 nested units show source-distinct classifier weights W, Adam first moment m, and Adam second moment v. Relative boundary differences span 2.54–31.29% for W, 7.65–20.32% for m, and 4.28–25.65% for v. These quantities are physical state measurements at the lifecycle boundary, not separate behavioural mediators.

After the transition, the same classifier rows enter a phase in which their natural classifier gradient is exactly zero. Nevertheless, history-ON and history-OFF trajectories develop source-distinct cumulative row-relative old-row displacement in all 24 nested units; the first-afterlife relative difference spans 6.01–55.73%. Thus acquisition-history formation leaves a source-dependent param-

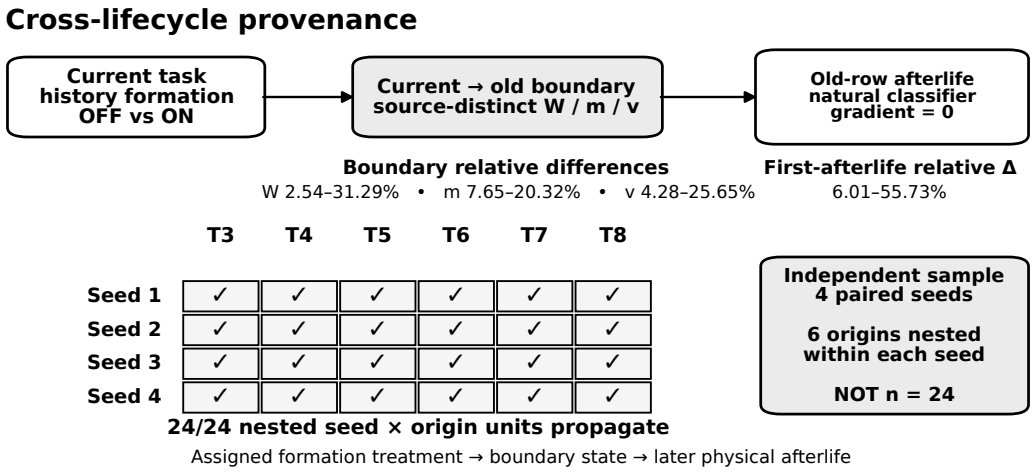

Figure 3: Cross-lifecycle provenance. Four independent paired seeds each contribute six origin tasks nested within seed; the 24 cells are nested evaluation units, not 24 independent replications. History formation produces source-distinct classifier $W/m/v$ state at the current-to-old boundary, followed by source-distinct row-relative old-row displacement while the natural classifier gradient is exactly zero. All $24/24$ nested units show the provenance pattern. This establishes lifecycle provenance, not behavioural mediation.

eter/optimizer state when a task becomes old, and that inherited difference propagates into later physical row dynamics after direct natural-gradient authority has vanished.

The claim boundary is strict. The assigned history-formation treatment gives causal provenance to the boundary difference; the subsequent observation tracks how that difference persists longitudinally. The experiment does not establish that boundary W, m, v, or row-relative displacement individually mediates the behavioural acquisition/forgetting phenotype. It establishes cross-lifecycle causal provenance, not behavioural mediation.

7 Same predictor, different future learner

The rights and lifecycle experiments suggest that the instantaneous predictor may be insufficient to specify how this continuing learner will evolve. We test that implication as a bounded stress test, not as a claim that parameterisation dependence is novel. Starting from a common learner after the first three tasks, we apply an exact internal function-preserving reparameterisation and transport the existing Adam state instead of resetting it. The classifier head is unchanged.

At the intervention boundary, each of four fresh pairs has bit-identical final features, bit-identical logits for all 100 classes, identical predictions, and a byte-identical classifier head. Despite this exact predictor identity, all four pairs diverge on the first lawful subsequent training update and develop materially different continual-learning trajectories thereafter. The tested canonical reparameterisation is harmful or coupled rather than protective; no safe parameter gauge is claimed.

Function-preserving rescaling is established prior art (Neyshabur et al., 2015), and parameter scaling and optimizer/parameterisation choices are known to affect training dynamics (Li et al., 2022; Everett et al., 2024). The contribution of this stress test is narrower: after persistent optimizer history has already been identified as a causal object, exact current predictor identity does not uniquely determine the next continual-learning trajectory when adaptive history is represented and subsequently formed in different valid parameter coordinates. A formal conjugacy analysis further shows why invariance alone is not an improvement mechanism: an exactly conjugate update rule is simply the coordinate image of the parent transition and maps back to the same function-level trajectory.

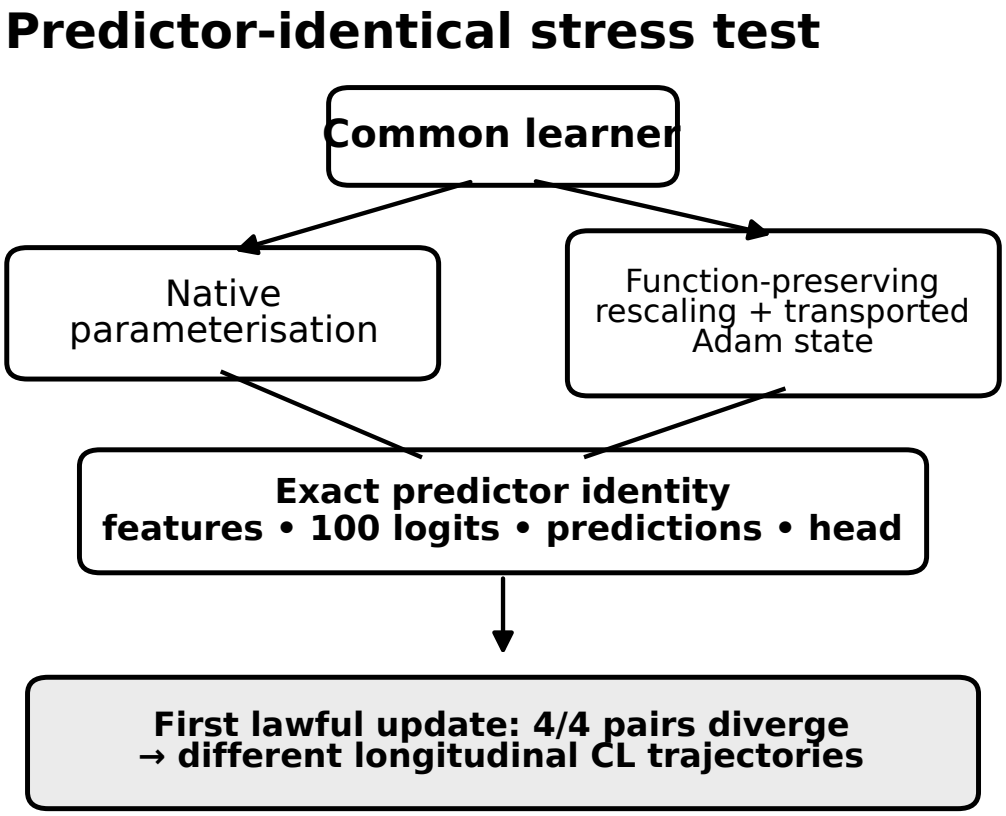

Figure 4: Predictor-identical stress test. A function-preserving reparameterisation with transported Adam state preserves branch-time features, logits, predictions, and classifier head, yet all four fresh pairs diverge on the first lawful subsequent update and develop different continual-learning trajectories. This is not a novelty claim for generic parameterisation dependence.

8 Falsification, implications, and limitations

8.1 Falsification bounds the tested mechanism family

The localisation repeatedly suggests an attractive engineering hypothesis: perhaps a sufficiently fine internal partition will retain acquisition while removing recapture. Prospectively tested alternatives do not support a robust static safe carrier. Second-moment-only persistence is approximately comparator-like; the current source innovation alone is insufficient; the current-parallel component of inherited first-moment action is approximately comparator-like; and a native-Adam action-coherent component retains acquisition but fails a fresh deterministic capture-separation criterion after a threshold-edge positive result. A later context-specific attenuation also fails to recur. Separately, explicit recapture control can reduce forgetting while suppressing newest-task performance. These negative results do not prove that no safe control law exists; they bound the tested family and motivate the paper's emphasis on causal organisation rather than a proposed mitigation algorithm.

8.2 Implications

Predictor state versus learner state. A parameter checkpoint can specify present inference while failing to specify the next training trajectory of a stateful learner. The experimentally licensed statement here is source-local: changing persistent-history rights or the parameter coordinates in which adaptive history is represented changes future dynamics even when other relevant objects are controlled.

History without replay. The training protocol excludes replayed old samples and old-derived training objects, yet persistent optimizer state carries historical information forward. This is not exemplar memory and is not replay. It is adaptive state produced by prior learning.

Causal rights. Stateful optimization ordinarily bundles immediate parameter action with future state formation. The 2×2 rights experiment shows that those permissions can be assigned independently. The useful conceptual distinction is therefore temporal: who may act now, and who may write the state that may act later?

Lifecycle inheritance. A task's causal influence need not end when it stops being current. Acquisition-history formation leaves source-dependent parameter and optimizer state at the current-to-old boundary and source-dependent physical motion after the task has lost natural classifier-gradient authority. The untested next step would be behavioural mediation; it is not assumed here.

8.3 Limitations

The evidence is concentrated on one benchmark, one fixed task order, one compact convolutional architecture, and Adam. Independent cohorts are modest and vary across focal experiments. Study-wide hypothesis development was adaptive, although the focal interventions reported as causal experiments used prospectively fixed sources, intervention semantics, seed cohorts, endpoints, and decision rules. Results from different cohorts are not pooled.

We do not report a state-of-the-art continual-learning method, a validated mitigation method, a universal theory of accessibility, proof that internal representations are unchanged, a robust safe carrier, or behavioural mediation of the lifecycle provenance result. Cross-benchmark, cross-architecture and cross-optimizer recurrence of the causal-rights organisation has not been established within the evidence presented here. The predictor-identical stress test does not establish a universal learner-state ontology; it is a consequence observed inside the same controlled Adam setting.

9 Conclusion

The central result is a causal distinction. In this strict replay-free fixed-shared-head Adam setting, the permission for a learning signal to affect parameters now is not the same as permission to write the persistent optimizer history that conditions later learning. The action-only and write-only interventions separate those rights, and the write-only cell shows that persistent history can carry material future acquisition and forgetting/recapture consequences without direct same-update source action. Successive interventions localise that history predominantly to accumulated first-moment state and its current-relative inherited action, while the lifecycle experiment shows that source-dependent parameter and optimizer state persists across the current-to-old transition and into zero-natural-gradient row dynamics. The source-local lesson is precise: present action is not persistent write.

Reproducibility Statement

The benchmark, architecture, training law, and primary evaluation protocol are specified in Section 2 and Appendix A. The source construction and causal intervention equations are given in Appendices B–C; independent seed/cohort structure and prospective decision conventions are listed in Appendix D; full seed-level rights results and the carrier-localisation contrasts are reported in Appendices E–G; lifecycle nesting and predictor-identical exactness are documented in Appendices H–I. Appendix K records matched-branch execution and exactness conditions. Different cohorts are not pooled, nested origin tasks are not treated as independent seeds, and reported materiality rules are the prospective rules used by the executed experiments rather than retrofitted significance tests.

References

Dylan R. Ashley, Sina Ghiassian, and Richard S. Sutton. Does the Adam optimizer exacerbate catastrophic forgetting?, 2021. URL https://arxiv.org/abs/2102.07686.

Ali Behrouz, Meisam Razaviyayn, Peilin Zhong, and Vahab Mirrokni. Nested learning: The illusion of deep learning architectures. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-1565. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/4309616aaed8e848009bc4a7ef73b493-Abstract-Conference.html.

Matias Cattaneo and Boris Shigida. How memory in optimization algorithms implicitly modifies the loss. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-5215. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/e4cc8ab4a64e99f962f36d07a7723d94-Abstract-Conference.html.

MohammadReza Davari, Nader Asadi, Sudhir Mudur, Rahaf Aljundi, and Eugene Belilovsky. Probing representation forgetting in supervised and unsupervised continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16712–16721, 2022. doi: 10.1109/CVPR52688.2022.01621. URL https://openaccess.thecvf.com/content/CVPR2022/html/Davari_Probing_Representation_Forgetting_in_Supervised_and_Unsupervised_Continual_Learning_CVPR_2022_paper.html.

Katie E. Everett, Lechao Xiao, Mitchell Wortsman, Alexander A. Alemi, Roman Novak, Peter J. Liu, Izzeddin Gur, Jascha Sohl-Dickstein, Leslie Pack Kaelbling, Jaehoon Lee, and Jeffrey Pennington. Scaling exponents across parameterizations and optimizers. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 12666–12700. PMLR, 2024. URL https://proceedings.mlr.press/v235/everett24a.html.

Mehrdad Farajtabar, Navid Azizan, Alex Mott, and Ang Li. Orthogonal gradient descent for continual learning. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 3762–3773. PMLR, 2020. URL https://proceedings.mlr.press/v108/farajtabar20a.html.

Yiduo Guo, Bing Liu, and Dongyan Zhao. Dealing with cross-task class discrimination in online continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11878–11887, 2023. doi: 10.1109/CVPR52729.2023.01143. URL https://openaccess.thecvf.com/content/CVPR2023/html/Guo_Dealing_With_Cross-Task_Class_Discrimination_in_Online_Continual_Learning_CVPR_2023_paper.html.

Gyuhak Kim, Changnan Xiao, Tatsuya Konishi, Zixuan Ke, and Bing Liu. A theoretical study on solving continual learning. In *Advances in Neural Information Processing Systems*, volume 35, pp. 5065–5079, 2022. doi: 10.52202/068431-0366. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/20f44da80080d76bbc35bca0027f14e6-Abstract-Conference.html.

Gyuhak Kim, Changnan Xiao, Tatsuya Konishi, and Bing Liu. Learnability and algorithm for continual learning. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 16877–16896. PMLR, 2023. URL https://proceedings.mlr.press/v202/kim23x.html.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations*, 2015. URL https://arxiv.org/abs/1412.6980.

Zhiyuan Li, Srinadh Bhojanapalli, Manzil Zaheer, Sashank Reddi, and Sanjiv Kumar. Robust training of neural networks using scale invariant architectures. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 12656–12684. PMLR, 2022. URL https://proceedings.mlr.press/v162/li22b.html.

Behnam Neyshabur, Ruslan Salakhutdinov, and Nati Srebro. Path-SGD: Path-normalized optimization in deep neural networks. In *Advances in Neural Information Processing Systems*, volume 28, 2015. URL https://proceedings.neurips.cc/paper/2015/hash/eaa32c96f620053cf442ad32258076b9-Abstract.html.

Vasileios Sevetlidis and George Pavlidis. Process-tensor tomography of SGD: Measuring non-markovian memory via back-flow of distinguishability. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pp. 2521–2529. PMLR, 2026. URL https://proceedings.mlr.press/v300/sevetlidis26a.html.

Ayushman Trivedi and Bhavika Melwani. Catastrophic forgetting as accessibility collapse: A three-level framework for knowledge persistence in continual learning, 2026. URL https://arxiv.org/abs/2606.06032.

Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 374–382, 2019. doi: 10.1109/CVPR.2019.00046. URL https://openaccess.thecvf.com/content_CVPR_2019/html/Wu_Large_Scale_Incremental_Learning_CVPR_2019_paper.html.

Qinyang Xu. Stored in optimizer state, valued by later training: A causal account of subliminal trait transfer, 2026. URL https://arxiv.org/abs/2608.20442.

Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. Maintaining discrimination and fairness in class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13208–13217, 2020. doi: 10.1109/CVPR42600.2020.01322. URL https://openaccess.thecvf.com/content_CVPR_2020/html/Zhao_Maintaining_Discrimination_and_Fairness_in_Class_Incremental_Learning_CVPR_2020_paper.html.

AI USE DISCLOSURE

Generative AI systems were used as research-assistance tools for research ideation and hypothesis refinement; conceptual-framework development; experimental-methodology and falsifier design; implementation assistance and code review; literature search and synthesis; interpretation and adversarial analysis of experimental results; and manuscript structure, drafting, and editing. AI-generated outputs were not treated as scientific evidence. Empirical claims in this paper derive from executed experiments and preserved experimental records. The human author reviewed AI-assisted material, checked empirical claims against the underlying evidence, and retains responsibility for the methodology, interpretation, text, claims, code, and final submission.

A EXACT BENCHMARK AND TRAINING SPECIFICATION

The benchmark is Split CIFAR-100 with ten sequential tasks of ten contiguous classes. Training is strictly current-task only. The model has a fixed 100-class learned output head throughout training and uses raw 100-way argmax for the primary behavioural evaluation. No focal experiment uses task identity at inference, output masking, replayed old examples, generated replay, stored old features or logits, prototypes, model/head growth, or post-hoc calibration.

Component	Specification
Feature extractor	Conv $3 \to 32$, ReLU; Conv $32 \to 64$, ReLU
Pooling	Adaptive average pool to 4×4
Classifier	Linear $1024 \to 100$; fixed shared head
Approx. parameters	121,892
Optimizer	Adam, lr 10^{-3}, $(\beta_1, \beta_2) = (0.9, 0.999)$
Numerics	$\epsilon = 10^{-8}$; no weight decay
Training	3 epochs/task; batch size 64
Updates	79/epoch; 237/task; 2,370 over ten tasks
Primary inference	raw 100-way argmax

Table 2: Exact benchmark and training specification used by the focal causal experiments.

Let $A^{\mathrm{raw}}_{i,j}$ denote raw 100-way accuracy on task j after training task i, and let $A^{\mathrm{local}}_{i,j}$ denote the corresponding task-restricted 10-way diagnostic. We report

$$\mathrm{current} = \frac{1}{10}\sum_{t=0}^{9} A^{\mathrm{raw}}_{t,t}, \tag{6}$$

$$\mathrm{forgetting} = \frac{1}{9}\sum_{t=0}^{8} \max\!\left(0, A^{\mathrm{raw}}_{t,t} - A^{\mathrm{raw}}_{9,t}\right), \tag{7}$$

$$\mathrm{old_normal} = \frac{1}{9}\sum_{t=0}^{8} A^{\mathrm{raw}}_{9,t}, \tag{8}$$

$$\mathrm{old_oracle} = \frac{1}{9}\sum_{t=0}^{8} A^{\mathrm{local}}_{9,t}. \tag{9}$$

The old-to-final endpoint is the prospectively specified raw-global recapture/capture endpoint used in the executed evaluation pipeline. We report it exactly as executed rather than reconstructing an unverified algebraic surrogate from summary statistics. Task-restricted accuracy is diagnostic only and is not interpreted as proof of unchanged internal representations.

B SOURCE-SIGNAL CONSTRUCTION

The causal-rights experiments use a single backbone source isolated inside a shared-head comparison controller. For input x, let $z_k(x)$ be class-k logits and let $\mathcal{B}$ denote completed ten-class blocks. A detached block reference is

$$r_b(x) = \tau \log\left(\frac{1}{|b|}\sum_{k\in b}\exp(z_k(x)/\tau)\right), \qquad \tau = 0.75. \tag{10}$$

The target-side completed-set reference is the arithmetic mean $R_T = |\mathcal{B}|^{-1}\sum_{b\in\mathcal{B}} r_b$, while the non-target reference R_N is the midpoint median of the completed-block references. For the current block, E_{nt} is the same-temperature log-mean-exp over current-block non-target logits. The controller response coordinates are

$$q_T = 0.15 + R_T - z_y, \qquad q_N = E_{nt} - R_N - 0.05, \tag{11}$$

with softplus-based target and non-target controller terms using $\tau = 0.75$ and weights $\lambda_G = 0.10$ and $\lambda_N = 0.04$, respectively.

Let $g_C(t)$ be the concatenated shared-backbone gradient from all non-focal controller routes on update t, and let $g_N(t)$ be the backbone gradient generated by the focal non-target controller route on the same realised batch and state. The source used in the rights experiments is the component of g_N orthogonal to the contemporaneous controller gradient:

$$g_s(t) = g_N(t) - \frac{\langle g_N(t), g_C(t)\rangle}{\|g_C(t)\|_2^2}\, g_C(t), \qquad \langle g_s(t), g_C(t)\rangle = 0, \tag{12}$$

with the zero-vector case handled by the executed deterministic implementation. Source-OFF uses $g_0(t) = g_C(t)$; source-ON uses $g_1(t) = g_C(t) + g_s(t)$. The focal intervention changes the shared-backbone source route only: classifier-row routes and the remaining controller are retained. The focal branch begins after a common state through the first three tasks, and matched arms consume the same realised current-task batches within seed.

C ACTION/WRITE IMPLEMENTATION EQUATIONS

For a backbone parameter tensor with pre-step Adam state (m_t, v_t, k_t), define the two candidate recurrences

$$m_0 = \beta_1 m_t + (1-\beta_1) g_0, \qquad v_0 = \beta_2 v_t + (1-\beta_2) g_0^2, \tag{13}$$

$$m_1 = \beta_1 m_t + (1-\beta_1) g_1, \qquad v_1 = \beta_2 v_t + (1-\beta_2) g_1^2. \tag{14}$$

The native Adam step counter and bias correction are retained. Present-action authority $A_s(t) \in \{0,1\}$ selects which candidate is used for the current parameter transition, while persistent-write authority $W_s(t) \in \{0,1\}$ independently selects which candidate state is retained after the update:

$$\theta_{t+1} = \mathrm{AdamStep}(\theta_t, m_{A_s(t)}, v_{A_s(t)}, k_t + 1), \qquad (m_{t+1}, v_{t+1}) = (m_{W_s(t)}, v_{W_s(t)}). \tag{15}$$

Cell	(A_s, W_s)	Operational meaning
00	$(0,0)$	comparator action; comparator persistent state
10	$(1,0)$	source acts now; comparator persistent state
01	$(0,1)$	comparator action now; source persistent state
11	$(1,1)$	source action; source persistent state

Table 3: Physical causal-rights cells. Write-only (01) writes source history now; that history may act later through native future Adam updates.

The write-only implementation is verified source-internally: its present transition equals the comparator candidate computed from its own pre-step state, while its stored first and second moments equal the full-source recurrence from that same pre-state. We do not require unsafe bitwise equality between independently evaluated GPU arms.

D SEED, COHORT, AND DECISION DEFINITIONS

Each focal experiment uses its own fresh cohort and matched within-seed arms. Cohorts are not pooled across experiments. Study-level hypothesis development was adaptive, but the focal source, branch point, treatment semantics, seed cohort, primary endpoints, and materiality/decision rules were pre-specified before inspecting focal outcomes. Executions that failed pre-specified validity checks were not treated as scientific results.

Where a materiality threshold governed a focal decision, we report the prospectively specified threshold. In the carrier-decomposition family, $\delta = 0.005$ is the recurring prospective materiality reference. We do not retrofit null-hypothesis significance tests as the original decision criterion, and we do not combine seed counts from different interventions into a larger apparent sample.

E FULL PER-SEED CAUSAL-RIGHTS RESULTS

Table 5 reports the two decisive matched contrasts from the five independent seeds: action-only minus comparator and write-only minus comparator. Values are accuracy fractions; multiply by 100 for percentage points.

The full-rights minus comparator current-acquisition effect is approximately $+0.00918$ at the cohort mean. The action-by-write interaction on current acquisition is $+0.00086$. Descriptively, the write-only contrast accounts for approximately 84.3% of the full source-versus-comparator acquisition movement; this is a factorial decomposition within this cohort, not a mediation estimate or a cross-experiment meta-analysis.

Experiment	Independent n	Cohort / design
Reciprocal Adam-state substitution	8	seeds 5346–5353; matched reciprocal recipients
Present action ON / direct write OFF	8	seeds 5354–5361
Action × persistent-write factorial	5	888929664, 1816955745, 1261647750, 1146965004, 636816954
First- vs second-moment persistence	8	145595555, 2100001635, 1739413335, 476926187, 1699995387, 256005724, 1809659920, 1869022535
Within-task history provenance	8	2075342045, 1070999005, 151477304, 1409121911, 769799396, 1921949496, 119910634, 1808615358
Inherited-action geometry	8	1479292123, 1454212573, 35811409, 1057022105, 221970527, 1611184104, 235568823, 489112511
Static-safety recurrence	8	377166949, 74069714, 857624422, 1487800176, 1898018415, 1664984475, 570594497, 1177820831
Lifecycle provenance	4 paired seeds	1255364162, 1291296787, 1071525738, 1674717793; six origin tasks nested within each seed
Predictor-identical stress test	4 paired seeds	1059107058, 1620889609, 1138558447, 992985155

Table 4: Independent cohort structure. The lifecycle study has 4 independent paired seeds and 24 nested seed×origin evaluation units; the latter are not 24 independent replications.

	Action-only − comparator			Write-only − comparator		
Seed	Current	Forget.	Old→final	Current	Forget.	Old→final
636816954	-0.0027	-0.002222	-0.012222	0.0054	-0.000222	-0.005667
888929664	0.0023	0.002333	0.001778	0.0103	0.008556	0.014222
1146965004	0.0019	0.002222	-0.000222	0.0053	0.010111	0.007222
1261647750	0.0010	0.000222	0.000444	0.0107	0.007556	0.008444
1816955745	0.0004	0.001000	0.000222	0.0070	0.004222	0.004222
Mean	0.00058	0.000711	-0.002000	0.00774	0.006044	0.005689

Table 5: Five-seed matched contrasts for the causal-rights factorial. The write-only acquisition mean is the registered $+0.00774$ result; its forgetting and old-to-final costs are also shown. Seed-level heterogeneity is retained rather than hidden by the mean.

F First- versus second-moment localisation

The direct-write intervention is crossed over persistent first-moment (m) and second-moment (v) ownership while same-update parameter action follows the comparator law. Eight fresh seeds are used.

Contrast	Current	Forget.	Old→final
Both moments − comparator	0.007812	0.008417	0.004750
First moment only − comparator	0.007837	0.008653	0.004514
Second moment only − comparator	-0.000400	-0.000056	-0.000375
First moment only − both	0.000025	0.000236	-0.000236
Moment interaction	0.000375	-0.000181	0.000611

Table 6: Eight-seed moment attribution. The source-local direct-write phenotype tracks first-moment persistence; second-moment-only persistence is approximately comparator-like on these endpoints.

First-moment-only persistence is materially positive on the primary acquisition endpoints in seven of eight seeds. These results do not imply that Adam's second moment is generally irrelevant; they attribute this specific direct-write phenotype within the executed source.

G WITHIN-TASK ACCUMULATION AND INHERITED-ACTION DECOMPOSITION

Temporal provenance. An eight-seed experiment separates the instantaneous source innovation from accumulated first-moment history. The comparator has no source-specific first-moment history; the one-recurrence arm admits only the current source innovation; the task-bounded arm accumulates source history within the current task and resets the source-specific delta at task boundaries; the full-history arm additionally retains source-specific treatment history across task boundaries.

Contrast	Current	Forget.	Old→final
One-recurrence − comparator	0.000638	-0.000472	-0.000194
Task-bounded − one-recurrence	0.007362	0.006125	0.008514
Full history − task-bounded	0.000800	0.001861	0.001944

Table 7: Eight-seed temporal provenance. Accumulated within-task history, rather than a single source innovation, carries most of the tested movement.

At the cohort-mean descriptive level, task-bounded accumulation recovers roughly 84% of the full-history acquisition movement and roughly 82% of the forgetting movement relative to the comparator. These percentages are descriptive recovery shares, not causal mediation coefficients.

Current-relative inherited action. A separate eight-seed intervention holds future persistence fixed and changes only which component of the already inherited first-moment history may act on the current parameter transition.

Contrast	Current	Forget.	Old→final
Full task-bounded history − comparator	0.008863	0.007056	0.005611
Current-parallel − comparator	0.000000	-0.000833	-0.000083
Current-orthogonal − comparator	0.009050	0.006375	0.007181
Current-orthogonal − full history	0.000188	-0.000681	0.001569

Table 8: Eight-seed inherited-action geometry. The dominant source-local action is current-relative orthogonal, but it carries both acquisition and cost.

The mean inherited-history norm is 0.028214, with parallel norm 0.006206 and orthogonal norm 0.027143 (ratio $4.374\times$). Mean realised Adam action norms are 0.001212 for full no-history recomputation, 0.000224 for parallel-only action, and 0.001219 for orthogonal-only action. This does not make orthogonality intrinsically safe.

H FULL LIFECYCLE 4×6 NESTED-UNIT RESULT

The lifecycle-provenance experiment contains four genuinely independent paired seeds. Six origin tasks (Tasks 3–8) are nested within each seed, producing 24 unit-level evaluations but only four independent seed pairs.

Seed	T3	T4	T5	T6	T7	T8
1255364162	✓	✓	✓	✓	✓	✓
1291296787	✓	✓	✓	✓	✓	✓
1071525738	✓	✓	✓	✓	✓	✓
1674717793	✓	✓	✓	✓	✓	✓

Table 9: Nested lifecycle evaluations. A check denotes the complete provenance pattern: source-distinct boundary state followed by source-distinct old-row afterlife. All $24/24$ nested units instantiate the pattern; inferential $n = 4$ paired seeds, not 24.

The result establishes cross-lifecycle causal provenance: assigned acquisition-history formation changes the state inherited at the current-to-old boundary, and the same source distinction is ob-

Quantity	Min	Median	Max
Boundary W relative difference	2.54%	12.25%	31.29%
Boundary m relative difference	7.65%	13.42%	20.32%
Boundary v relative difference	4.28%	13.74%	25.65%
First-afterlife cumulative displacement	6.01%	15.88%	55.73%

Table 10: Descriptive ranges across the 24 nested units. The natural old-row classifier gradient is exactly zero throughout the measured afterlife interval. These summaries do not convert nested units into independent replications.

served in later physical dynamics after natural gradient authority is lost. It does not establish that any single boundary component mediates later behavioural acquisition or forgetting.

I PREDICTOR-IDENTICAL INTERVENTION EXACTNESS

The predictor-identical stress test uses four fresh paired seeds. Starting from a common learner, an exact internal function-preserving rescaling is applied and the existing Adam state is transported according to the specified parameter transformation rather than reset. The classifier head is not changed.

Seed	Features	100 logits	Head	First divergence	Worst Δ current (task)
1059107058	bit-exact	bit-exact	byte-identical	next update	-0.005 (T3)
1620889609	bit-exact	bit-exact	byte-identical	next update	-0.019 (T4)
1138558447	bit-exact	bit-exact	byte-identical	next update	-0.030 (T8)
992985155	bit-exact	bit-exact	byte-identical	next update	-0.026 (T4)

Table 11: Exactness and longitudinal divergence for the predictor-identical stress test. All four pairs first diverge on the first lawful subsequent update. The tested reparameterisation is harmful/coupled rather than protective.

A formal qualification prevents an overly simple invariance interpretation. For the tested symmetry, an exactly gauge-conjugate state/action/source construction can be written down, but the uniquely source-derived exact construction is simply the coordinate image of the native parent transition. Mapping it back reproduces the parent function-level trajectory. Thus coordinate conjugacy removes arbitrary coordinate dependence by reproducing the same learner; it does not by itself create a better continual learner.

J NEGATIVE AND FALSIFICATION EXPERIMENTS

The mechanism was narrowed by retaining negative results rather than promoting every attractive decomposition.

K REPRODUCIBILITY DETAILS

All focal comparisons use matched within-seed branches from a common pre-intervention state. Within a seed, branch arms consume matched realised minibatches and random-state schedules unless the focal intervention explicitly changes a listed state object. Scientific gradients are constructed before state-substitution operations where required, so a moment-state intervention cannot retroactively change the recipient gradient whose update it is testing.

For the causal-rights factorial, each arm constructs both comparator and full-source Adam candidates from the same arm-local pre-step state. The action right determines the actual present parameter transition; the write right determines only the persistent moments retained afterward. Exactness checks verify the selected candidate transition and selected stored recurrence source-internally. The experiment does not rely on cross-arm bitwise equality after trajectories have diverged.

Tested explanation	Independent n	Result and claim boundary
Second-moment-only persistent write	8	Primary endpoints remain approximately comparator-like; does not imply second moments are universally irrelevant.
One-update first-moment innovation	8	Insufficient; material phenotype appears with accumulated within-task history.
Current-parallel inherited action	8	Approximately comparator-like; current-orthogonal action carries the source-local phenotype, but is not safe.
Native adaptive-action static safety	8 fresh recurrence	Acquisition remains material, but old-to-final relief versus the full carrier is only -0.002014 against a prospective -0.005 criterion; only $1/8$ seeds reaches material capture relief. The prior threshold-edge positive therefore does not recur robustly.
Comparison-context attenuation	2 fresh seeds in recurrence test	A previously observed attenuation pattern does not recur at the 0.005 materiality reference; no robust context-specific carrier attenuation claim is retained.
Recapture-control trade-off	2 paired units	Recapture and forgetting improve, but final/newest-task performance falls materially (about -0.061 and -0.049 final-task in the two units); the intervention does not yield a validated mitigation method.

Table 12: Selected prospective falsifiers that constrain the final mechanism. These results bound the tested family; they do not prove that no safe control law can exist.

Each focal experiment pre-specifies its source definition, branch point, seed cohort, primary endpoints, and decision rule before outcome inspection. No seed is replaced because of an observed outcome. Different experimental cohorts are reported separately and are not pooled. The lifecycle experiment treats origin-task observations as nested within four independent paired seeds. Materiality thresholds reported in the paper are the prospectively specified thresholds used by the experiment; conventional significance tests are not retrofitted as original decision criteria.

The main text and this appendix specify the benchmark, architecture, optimizer, source construction, intervention equations, cohort identities, endpoint contrasts, and exactness conditions needed to reproduce the focal causal experiments. Accompanying anonymous implementation, if distributed with the submission, should instantiate these same equations and cohort definitions. AI-generated outputs were not used as empirical evidence or as substitutes for executed records.