Coarse-Grained Anti-Recurrence and Informational Hysteresis in Finite Memory Systems:
Lost Preimages, Side Records, and Capacity-Recovery Asymmetry

Yining Wu
Independent Researcher
yining.wu@alumni.upenn.edu

Finite observers often operate with effective representations that identify many microscopic states with a single coarse state. This paper studies the operational consequence of such non-injective truncation: once preimage information is absent from the effective record, later capacity recovery does not reconstruct the discarded distinctions. Guaranteed exact preimage recovery is possible only when inverse information is preserved in side records, reversible garbage registers, environmental traces, or an enlarged accounting boundary. We formalize this effect through preimage entropy, a Bayes-optimal recovery bound, residual inverse uncertainty, a side-record recurrence criterion, and an informational hysteresis observable. We replace a purely metric closure error with an information-theoretic Jensen-Shannon/KL family of Markov-closure errors, and relate the resulting hidden-state burden to the Mori-Zwanzig memory-kernel form of projected dynamics. The resulting anti-recurrence claim is representation-relative and does not deny microscopic reversibility, unitary dynamics, Poincare recurrence, or fluctuation theorems [9, 11–13]. It states that finite effective systems face an asymmetry between losing a distinction and later regaining capacity to store distinctions. Deterministic synthetic simulations illustrate recovery bounds, capacity hysteresis, externalization break-even, memory-kernel burden, Markov closure error, benign versus malignant hysteresis, Phase-B low-cost survivor selection, and the transition from ordinary P4 truncation to P7 invariant/topological protection.

Reader contract. This paper is a registered FDS physical-bridge paper. It is not a replacement for statistical mechanics, stochastic thermodynamics, Markov-chain theory, reversible computation, model reduction, or quantum mechanics. It does not deny microscopic reversibility, Liouville evolution, Hamiltonian recurrence, reversible computation, unitary quantum dynamics, or fluctuation relations under their own hypotheses [9–13]. It gives a finite-representation account of exact preimage recovery after non-injective truncation. Whenever guaranteed exact preimage recovery is possible, P4 asks where the inverse information was stored and what boundary, memory, synchronization, verification, or thermodynamic ledger carries it.

Claim-status summary

Table I summarizes the main claims, status, and failure conditions. The status labels are internal to the FDS corpus: the formal information claims are low-risk, while the physical bridge claims depend on the stated representation and accounting boundary.

Keywords: finite distinction systems; coarse-graining; anti-recurrence; informational hysteresis; preimage entropy; side records; Mori-Zwanzig; lumpability; Markov closure; finite memory; externalization; invariant persistence.

INTRODUCTION

The problem

Microscopic laws may be reversible while finite effective descriptions are not. This is compatible with standard physics. Coarse-graining, measurement, finite memory, model reduction, and statistical inference all use records that identify many microstates with fewer effective states. The question studied here is not whether the full microscopic universe loses information. It is whether a finite effective system that has already applied a non-injective truncation can later recover the specific micro-distinction it discarded merely by recovering storage capacity.

P4 answers no. Capacity recovery restores the ability to store new distinctions. It does not reconstruct old inverse information that is absent from the record. Capacity-recovery asymmetry is the gap between restored representational capacity and unrecovered historical inverse information. Guaranteed exact preimage recovery requires that the inverse information was preserved somewhere: a side log, reversible garbage register, environmental trace, external archive, protected invariant, or enlarged accounting boundary. If no such inverse information exists inside the boundary being audited, the lost distinction cannot recur inside that effective account.

This gives a representation-relative form of anti-recurrence. It is not anti-recurrence in the Poincare sense. It is anti-recurrence of exact preimage recovery inside a truncated record.

TABLE I. Central FDS-P4 claims, status, and failure or demotion conditions.

Claim	Status	What would weaken or falsify it
Non-injective truncation creates preimage uncertainty relative to the effective record.	Formal information claim	A many-to-one map contains enough information, without side records or conventions, to distinguish all of its preimages.
Bayes-optimal exact recovery is bounded by the largest conditional preimage mass.	Decision-theoretic bound	A decoder using only Z exceeds the Bayes-optimal classifier bound for $X \mid Z$.
Capacity recovery does not recover distinctions erased during a bottleneck.	Informational hysteresis theorem	A finite system recovers exact task-relevant preimage distinctions after capacity restoration with no side record, no enlarged boundary, no external trace, and no hidden convention.
Non-lumpable coarse-graining creates hidden-state memory and effective stochasticity.	Markov projection bridge	A non-lumpable projection closes exactly on Z_t alone without hidden state, history, or extra variables.
Projection-induced memory burden has a Mori-Zwanzig analogue.	Relation to standard theory	Eliminated variables never reappear as memory, noise, or closure error in projected dynamics, even when lumpability fails.
Externalization restores inverse information only by moving it to a side ledger.	Accounting-boundary bridge	External logs restore guaranteed exact preimage recovery at no writing, retention, indexing, synchronization, retrieval, verification, or boundary-expansion cost.
Sustained truncation requires residual irrecoverability, side records, externalization, task relaxation, or failure.	Finite-memory exit theorem	A finite system repeatedly applies non-injective truncation to task-relevant distinctions while preserving guaranteed exact preimage recovery with no residual uncertainty and no extra ledger.

Main contributions

The paper contributes six objects to the FDS physical spine:

1. a preimage-entropy and Bayes-recovery bound for non-injective truncation;
2. a side-record recurrence criterion based on residual inverse uncertainty;
3. a strengthened hysteresis theorem for task variables nonconstant on a truncated fiber;
4. an information-theoretic Markov-closure error based on KL/JS divergence;
5. a bridge to Mori-Zwanzig memory kernels for projected dynamics;
6. an externalization break-even ledger and a benign/malignant hysteresis classification.

Terminological clarification: recurrence versus inverse recovery

The word “recurrence” has several meanings. P4 uses it only in the third sense below:

1. *Microscopic recurrence*: return of a full microstate or region in the complete state space under its full dynamics.
2. *Effective recurrence*: return of the same coarse state $z \in Z$.
3. *Preimage recovery*: reconstruction of which $x \in T^{-1}(z)$ occurred before the truncation.

P4 anti-recurrence means failure of exact preimage recovery inside a truncated representation. It does not mean failure of microscopic dynamical recurrence or failure of coarse-state recurrence. The same coarse state z may recur many times while the system remains unable to recover which prior preimage x produced it.

How P4 interfaces with T3, P5, and O3

P4 can be read independently as the representation-side truncation paper. In the broader FDS sequence, it supplies the inverse-information mechanism behind capacity overflow, entropy ledgers, and finite-memory record turnover. T3 studies overflow and effective stochasticity; P5 studies capacity deficit and entropy-production pressure; O3 studies finite-memory boundary maintenance and record reuse. P4 explains why record turnover is path-dependent: once preimages are absent from the accounting boundary, current capacity alone cannot reconstruct them.

FDS BACKGROUND

Finite distinction systems

An active finite distinction system is represented by

$$S = (X, E, B, M, Y, A, U, \pi, \ell, \Phi, \mathcal{P}, \tau), \tag{1}$$

where X is internal state, E environment, B boundary, M memory/model state, Y observation channel, A action space, U update rule, π finite projection, ℓ boundary-maintenance loss, Φ finite resource budget, $\mathcal{P}$ perturbation/pruning family, and τ update timescale. P4 focuses on π, M, U, ℓ, and the accounting boundary deciding whether inverse information is still available.

Finite capacity and truncation

For finite memory M, internal representational capacity is

$$C_S = \log_2 |M|. \tag{2}$$

For a task family Ψ, rate-distortion demand over window τ is $R_{\min}^{(\tau)}(\epsilon; \Psi)$. The capacity deficit is

$$\Delta_\epsilon(\tau) = R_{\min}^{(\tau)}(\epsilon; \Psi) - C_S. \tag{3}$$

When $\Delta_\epsilon > 0$, the system must approximate, externalize, prune, relax the task, or fail. P4 studies which distinctions disappear inside the effective record during such approximation.

EFFECTIVE TRUNCATION AND PREIMAGE RECOVERY

Definition 1 (Effective truncation map)**.** *An effective truncation map is a map*

$$T : X \to Z \tag{4}$$

from a high-resolution state space X to an effective state space Z. It is non-injective when there exist $x \neq x'$ with $T(x) = T(x')$. The preimage fiber of $z \in Z$ is

$$\mathrm{Pre}_T(z) = T^{-1}(z) = \{x \in X : T(x) = z\}. \tag{5}$$

Definition 2 (Preimage entropy)**.** *Given a distribution $p(x)$ and $Z = T(X)$, the preimage entropy of the effective record is*

$$H_T = \mathrm{H}(X \mid Z). \tag{6}$$

For a uniform finite fiber, $H_T(z) = \log_2 |T^{-1}(z)|$.

For continuous state spaces, preimage entropy should be interpreted through a finite-resolution, rate-distortion, or error-tolerant recovery measure. Differential entropy is not used here as a direct count of recoverable distinctions.

Lemma 1 (Coarse-grained anti-recurrence)**.** *In an effective representation containing a non-injective truncation map $T : X \to Z$, guaranteed exact recovery of a specific erased preimage is impossible from Z alone unless inverse information about the preimage is supplied from outside Z.*

Proof. Let $x \neq x'$ and $T(x) = T(x') = z$. The effective state z is compatible with at least two distinct preimages. Any map $R : Z \to X$ selecting x rather than x' uses information not contained in z unless a convention, side record, hidden state, environmental trace, or additional inverse state is supplied. Thus the distinction between x and x' cannot be recovered inside the effective representation alone. It can be recovered only in an enlarged representation containing inverse information. □

Proposition 1 (Bayes-optimal exact recovery bound)**.** *Let X be the preimage and $Z = T(X)$ the effective state. Any decoder $D : Z \to X$ has exact recovery probability bounded by*

$$P_{\mathrm{rec}}^* = E_Z \left[\max_{x \in T^{-1}(Z)} p(x \mid Z) \right]. \tag{7}$$

For a uniform fiber of size $N_z = |T^{-1}(z)|$,

$$P_{\mathrm{rec}}^*(z) = \frac{1}{N_z} = 2^{-\mathrm{H}(X|Z=z)}. \tag{8}$$

Proof. For each observed z, the optimal exact decoder chooses the maximum a posteriori preimage. Its conditional success probability is $\max_{x \in T^{-1}(z)} p(x \mid z)$. Averaging over Z gives Eq. (7). This bound is a direct consequence of Bayesian decision theory and standard information-theoretic coding limits [5, 6, 17]. If the fiber is uniform, $p(x \mid z) = 1/N_z$ and Eq. (8) follows. □

Corollary 1 (Microscopic reversibility compatibility)**.** *The anti-recurrence lemma does not contradict microscopic reversibility. If the full microstate x is preserved in an enlarged system, or if reversible computation retains garbage registers sufficient to reconstruct the preimage, guaranteed exact preimage recovery may occur in that enlarged account. P4 applies to the reduced representation that lacks the preimage information.*

Example 1 (Four-state toy model)**.** *Let*

$$X = \{00, 01, 10, 11\}, \qquad T(x) = \textit{first bit of } x. \tag{9}$$

Then $T(00) = T(01) = 0$ and $T(10) = T(11) = 1$. If a later task asks for the second bit, the effective record

TABLE II. P4 in the physical-operational FDS sequence.

Paper	Core object	P4 relation
T3	capacity overflow and effective stochasticity	P4 formalizes erased preimages behind projected randomness
P5	capacity deficit and entropy production	P4 supplies residual inverse-information loss for entropy/resource ledgers
O3	finite memory and operational Second-Law channel	P4 explains why record turnover remains path-dependent after capacity re
P7	invariant-supported persistence / topology	P4 defines ordinary local forgetting; P7 asks when local forgetting is obstru

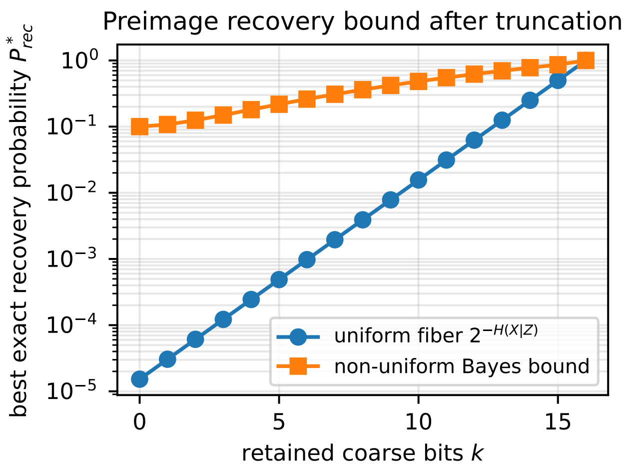

FIG. 1. Preimage recovery bound after non-injective truncation. For uniform fibers, exact micro-recovery decays as $2^{-H(X|Z)}$. Non-uniform fibers improve the Bayes-optimal bound only when one preimage dominates the posterior.

Z cannot recover it. Even if the system later expands from one-bit to two-bit capacity, the historical second bit is absent unless it was stored in a side record.

ACCOUNTING BOUNDARIES AND SIDE RECORDS

Definition 3 (Accounting boundary)**.** *An accounting boundary $\mathcal{A}$ specifies which records, reservoirs, side logs, reversible garbage registers, environmental traces, external memories, and protected invariants are counted as accessible to the system. Information outside $\mathcal{A}$ is not operationally available for recovery inside $\mathcal{A}$.*

Definition 4 (Side-record inverse information)**.** *Let $L_\mathcal{A}$ denote side records accessible within accounting boundary $\mathcal{A}$. The inverse information retained by that boundary after truncation is*

$$I^{\mathcal{A}}_{\text{inv}} = \mathrm{I}(X; L_\mathcal{A} \mid Z). \tag{10}$$

The residual inverse uncertainty is

$$H^{\mathcal{A}}_{\text{res}} = \mathrm{H}(X \mid Z, L_\mathcal{A}). \tag{11}$$

Proposition 2 (Side-record recovery bound)**.** *For side records $L_\mathcal{A}$ accessible within boundary $\mathcal{A}$, the Bayes-optimal exact-recovery probability satisfies*

$$P^*_{\text{rec}}(Z, L_\mathcal{A}) = \mathbb{E}_{Z,L_\mathcal{A}}\left[\max_x p(x \mid Z, L_\mathcal{A})\right], \tag{12}$$

and $P^*_{\text{rec}}(Z, L_\mathcal{A}) \geq P^*_{\text{rec}}(Z)$, with strict improvement whenever $L_\mathcal{A}$ carries inverse information with $\mathrm{I}(X; L_\mathcal{A} \mid Z) > 0$.

Criterion 1 (Side-record recovery criterion)**.** *Guaranteed exact preimage recovery within boundary $\mathcal{A}$ is possible only if accessible side records contain enough inverse information to identify the relevant preimage up to task tolerance:*

$$\mathrm{H}(X \mid Z, L_\mathcal{A}) \leq \epsilon_{\text{inv}}. \tag{13}$$

For exact finite recovery, $\epsilon_{\text{inv}} = 0$.

Proof. An exact recovery operation is a decoder $D : Z \times L_\mathcal{A} \to X$. If $\mathrm{H}(X \mid Z, L_\mathcal{A}) > 0$, multiple preimages remain compatible with the available record, so guaranteed exact recovery is impossible without additional information or arbitrary guessing. If the conditional entropy is zero, then X is a function of $(Z, L_\mathcal{A})$ almost surely, so an exact decoder exists in principle. □

Remark 1 (Externalization is not free)**.** *Side records can reduce residual uncertainty only by expanding the accounting boundary. External records must be written, preserved, indexed, synchronized, verified, retrieved, protected, and eventually pruned. Guaranteed exact preimage recovery may be restored, but the burden moves from lost preimage uncertainty to side-record maintenance.*

CAPACITY-RECOVERY ASYMMETRY AND INFORMATIONAL HYSTERESIS

Capacity recovery is not distinction recovery

Let C_t be available effective capacity. A system can lose capacity and later regain it. P4's claim is that this does not imply recovery of distinctions erased during the low-capacity interval. The system expands from the post-truncation record, not from the original preimage, unless inverse information was preserved.

Definition 5 (Informational hysteresis)**.** *Let* X_{past} *be a past high-resolution state or task-relevant variable,* Z_t *the current effective record,* C_t *current capacity, and* L_{ext} *accessible inverse records. Informational hysteresis is*

$$H_{\text{irr}}(t) = \text{H}(X_{\text{past}} \mid Z_t, C_t, L_{\text{ext}}). \tag{14}$$

The system exhibits informational hysteresis if, after capacity recovers to a previous value, H_{irr} *remains larger than it would have been along a path that never crossed the bottleneck.*

Theorem 1 (Task-relative informational hysteresis)**.** *Let* $V = f(X_{t_0})$ *be a task variable. Suppose a non-injective truncation* T *maps* $x, x' \in X_{t_0}$ *to the same effective state but* $f(x) \neq f(x')$*, and that both preimages remain admissible with nonzero posterior probability under the final record:*

$$p(x \mid Z_f, C_f, L_f) > 0, \qquad p(x' \mid Z_f, C_f, L_f) > 0. \tag{15}$$

If no accessible side record distinguishes x *from* x'*, then after any later capacity recovery that does not add inverse information,*

$$\text{H}(V \mid Z_f, C_f, L_f) > 0. \tag{16}$$

Thus any future task requiring V *suffers residual uncertainty despite restored capacity.*

Proof. Since $T(x) = T(x')$ but $f(x) \neq f(x')$, the effective record identifies two task-distinct preimages. By assumption, L_f does not distinguish them. Therefore V has at least two possible values conditioned on (Z_f, C_f, L_f), giving positive conditional entropy. Restoring storage capacity changes the size of possible future records but does not supply the missing inverse information about which preimage occurred at t_0. □

Corollary 2 (Path dependence)**.** *Two systems may have the same final capacity and effective update rule but different recoverable past distinctions if one crossed a truncation bottleneck and the other did not.*

Definition 6 (Hysteresis gap)**.** *For a bottleneck path* γ_b *and a non-bottleneck path* γ_0 *with the same final capacity, define*

$$\Delta H_{\text{hyst}} = \text{H}_{\gamma_b}(X_{\text{past}} \mid Z_f, L_f) - \text{H}_{\gamma_0}(X_{\text{past}} \mid Z_f, L_f). \tag{17}$$

A positive value indicates path-dependent inverse-information loss.

Benign versus malignant hysteresis

Not every erased distinction is harmful. A finite system must discard many details to act at all. P4 therefore distinguishes two regimes.

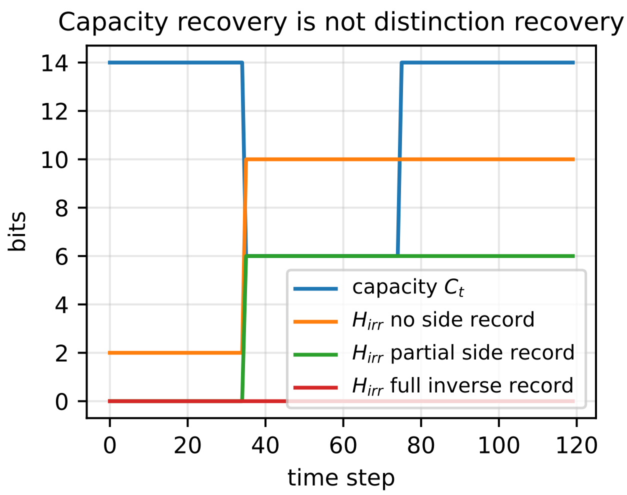

FIG. 2. Capacity recovery is not distinction recovery. Capacity drops and later returns, but residual irrecoverable information remains unless inverse records are retained. Full logs remove residual uncertainty only by expanding the accounting boundary.

Definition 7 (Benign informational hysteresis)**.** *Hysteresis is benign relative to a task family* Ψ *when erased distinctions are task-irrelevant or replaced by cheaper invariant variables without increasing expected task loss:*

$$\mathbb{E}[\ell_\Psi \mid H_{\text{irr}} > 0] \approx \mathbb{E}[\ell_\Psi \mid H_{\text{irr}} = 0]. \tag{18}$$

Definition 8 (Malignant informational hysteresis)**.** *Hysteresis is malignant relative to* Ψ *when erased distinctions later become task-relevant and residual inverse uncertainty increases expected loss:*

$$\mathbb{E}[\ell_\Psi \mid H_{\text{irr}} > 0] > \mathbb{E}[\ell_\Psi \mid H_{\text{irr}} = 0]. \tag{19}$$

This distinction is crucial for later FDS applications. Abstraction, learning, model reduction, and scientific law formation often require benign hysteresis: details are discarded to reveal stable patterns. Catastrophic forgetting, institutional memory loss, fragile AI agents, and degenerative memory failure are malignant cases: distinctions that later matter cannot be recovered from the effective record.

PROJECTED DYNAMICS, MARKOV CLOSURE, AND MEMORY KERNELS

Lumpability and hidden-state memory

Let (X_t) be a Markov chain on X with transition kernel $P(x' \mid x)$ and let $Z_t = T(X_t)$. The projected process is Markovian in Z_t alone only under lumpability: for any

x, x' in the same fiber $T^{-1}(z)$,

$$\sum_{y:T(y)=z'} P(y \mid x) = \sum_{y:T(y)=z'} P(y \mid x') \tag{20}$$

for every projected state z'.

Proposition 3 (Non-lumpability creates hidden-state memory)**.** *If Eq. (20) fails, then (Z_t) is not Markovian with respect to Z_t alone. Any Markovian representation requires additional hidden state, history, or an enlarged coarse variable.*

Proof. If Eq. (20) fails, two microstates in the same projected class have different projected transition probabilities. Since Z_t cannot identify which microstate is present, $\mathbb{P}(Z_{t+1} \mid Z_t)$ depends on the conditional distribution over hidden preimages. That conditional distribution generally depends on prior observations. Hence the projected dynamics carries hidden-state memory. Failure of a projected dynamics to remain Markovian is a standard feature of lumped Markov processes and coarse-grained stochastic thermodynamics [7, 8, 18, 19]. □

Definition 9 (Information-theoretic Markov closure error)**.** *For a projection T, define*

$$\mathcal{M}_{\mathrm{D_{JS}}}(T) = \mathbb{E}_z \mathbb{E}_{x,x' \sim p(\cdot|z)} \left[\mathrm{D_{JS}}\big(P_T(\cdot \mid x) \| P_T(\cdot \mid x')\big) \right], \tag{21}$$

where

$$P_T(z' \mid x) = \sum_{y:T(y)=z'} P(y \mid x). \tag{22}$$

A KL-family version is

$$\mathcal{M}_{\mathrm{D_{KL}}}(T) = \mathbb{E}_z \mathbb{E}_{x \sim p(\cdot|z)} \mathrm{D_{KL}}\big(P_T(\cdot \mid x) \| \bar{P}_T(\cdot \mid z)\big), \tag{23}$$

where $\bar{P}_T(\cdot \mid z) = \mathbb{E}_{x \sim p(\cdot|z)} P_T(\cdot \mid x)$. Both errors vanish for strongly lumpable projections.

Relation to Mori-Zwanzig memory kernels

Projection-operator methods make the same structural fact visible in continuous statistical mechanics. Let L be a microscopic generator and P a projection onto resolved variables. Mori-Zwanzig theory yields a generalized equation of the schematic form

$$\frac{d}{dt} PA_t = PLPA_t + \int_0^t K(t-s) PA_s \, ds + F_t, \tag{24}$$

where K is a memory kernel and F_t is an orthogonal noise term [14–16, 21]. P4 does not rederive Mori-Zwanzig. It uses it as a standard expression of the same burden: eliminated variables do not become irrelevant merely because they are absent from the resolved representation. They re-enter as memory, noise, or closure error unless the projection is dynamically closed.

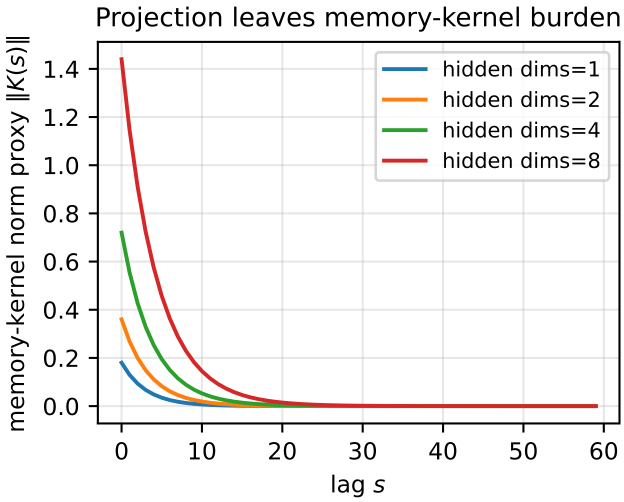

FIG. 3. Stylized Mori-Zwanzig memory-kernel burden. Eliminated hidden dimensions leave longer or larger memory-kernel tails in the resolved dynamics. The figure is synthetic and illustrates the bridge rather than fitting data.

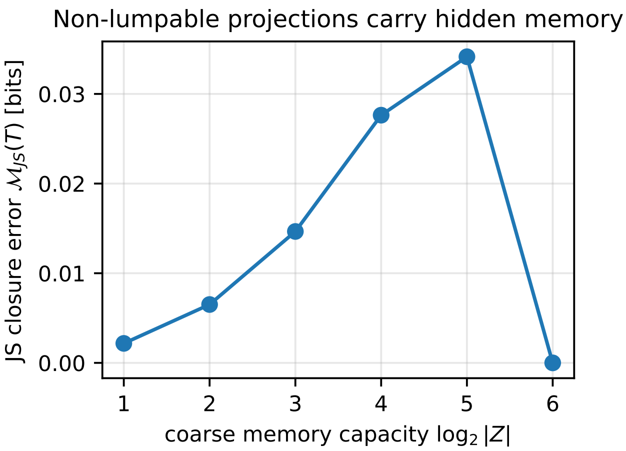

FIG. 4. Jensen-Shannon Markov closure error for a synthetic 64-state chain under increasingly fine projections. Coarser non-lumpable projections carry larger hidden-state memory burden.

In FDS terms, the memory kernel is an operational shadow of inverse-information deficit. A finite observer that deletes variables but still faces tasks depending on their influence must either carry memory of them, accept stochastic residuals, enlarge the representation, externalize records, or relax the task. The long tail of the memory kernel measures the additional computational and dissipative burden required to compensate for preimage information removed from the resolved representation.

EXTERNALIZATION AND BREAK-EVEN AUDITS

Externalization as inverse-information relocation

External memory can prevent anti-recurrence by storing inverse information [9–11]. It does not make recovery free. It moves the ledger from internal preimage uncertainty to external writing, synchronization, verification, retrieval, indexing, and protection.

Let D denote erased task-relevant inverse information. Let L denote accessible side-record inverse information. A simple residual model is

$$H_{\rm res} = \max\{D - L, 0\}. \tag{25}$$

External cost may be decomposed as

$$C_{\rm ext}(L) = C_{\rm write}(L)+C_{\rm sync}(L)+C_{\rm verify}(L)+C_{\rm retrieve}(L)+C_{\rm protect}(L). \tag{26}$$

Definition 10 (Externalization surplus)**.** *Let $C_{\rm int}^{\rm recover}(D)$ be the expected internal cost of reacquiring or reconstructing lost inverse information without a side record. Define*

$$\Pi_{\rm ext}(D, L) = C_{\rm int}^{\rm recover}(D) - C_{\rm ext}(L). \tag{27}$$

Externalization is favorable when

$$\Pi_{\rm ext}(D, L) > 0 \quad \textit{and} \quad H_{\rm res}(D, L) \le \epsilon_{\rm task}. \tag{28}$$

Remark 2. *The break-even point is not merely economic. It appears in biology, cognition, AI memory, organizations, science, law, and civilization. Writing a log, preserving an archive, delegating to a tool, or encoding a standard is favorable only when the reduced internal burden exceeds the maintenance and verification cost of the external record.*

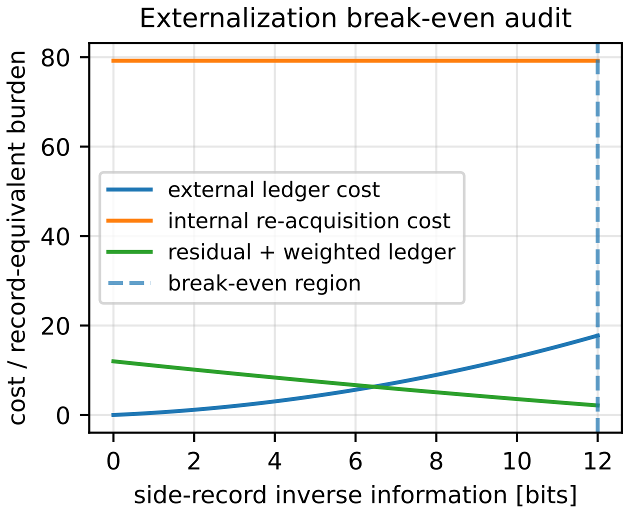

FIG. 5. Externalization break-even audit. Side records can reduce residual inverse uncertainty, but they impose writing, synchronization, retrieval, and verification costs. Externalization is favorable only beyond a break-even region.

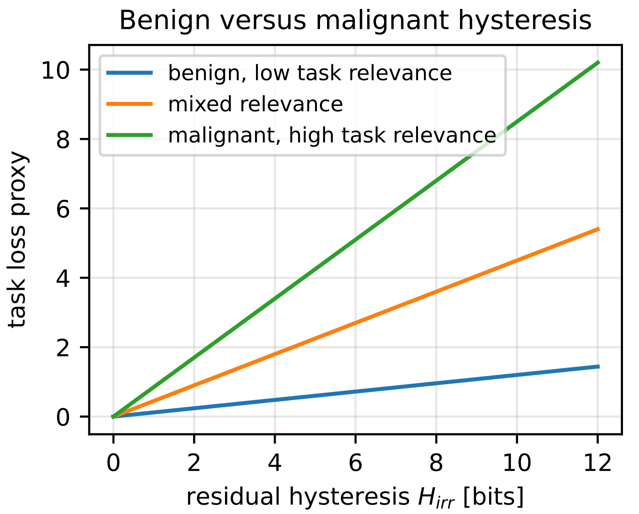

FIG. 6. Benign versus malignant informational hysteresis. The same residual inverse uncertainty has small effect when erased distinctions are task-irrelevant, but large effect when future tasks depend on them.

PHASE-B SURVIVORS AND THE P7 BRIDGE

Low-cost predictive survivors

P4 is not only about loss. Finite systems also gain abstraction by discarding details that do not matter for future maintenance. A coarse variable is useful when it remains predictive while being cheap to update, maintain, and dynamically close.

Definition 11 (Coarse-variable survival score)**.** *Let ϕ be a candidate coarse variable. Define*

$$S(\phi) = \frac{\mathrm{I}(\phi_t; \phi_{t+\tau})}{C_{\rm upd}(\phi) + C_{\rm maint}(\phi) + \lambda \mathcal{M}_{\rm DJS}(\phi) + \epsilon_0}. \tag{29}$$

High-score variables are Phase-B candidates: low-cost predictive residues that remain useful under finite capacity.

Remark 3. *This is a heuristic selection criterion, not a universal law. Its role in P4 is to separate destructive loss from abstraction. Benign hysteresis often occurs when erased detail is replaced by a lower-cost, more stable predictive variable.*

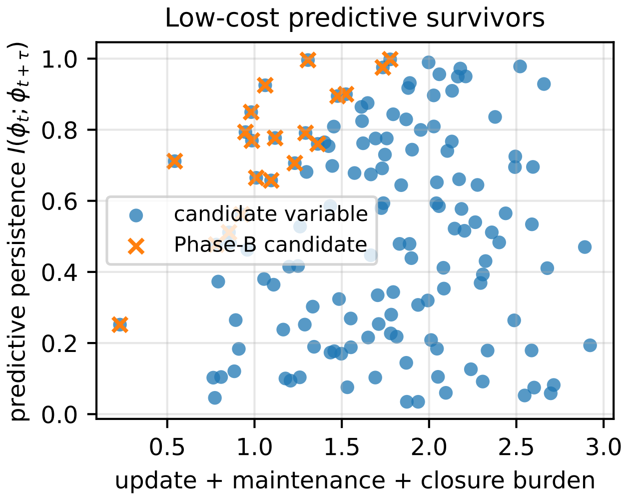

FIG. 7. Synthetic Phase-B candidate selection. Variables with high predictive persistence per update, maintenance, and closure burden are low-cost survivors under finite capacity.

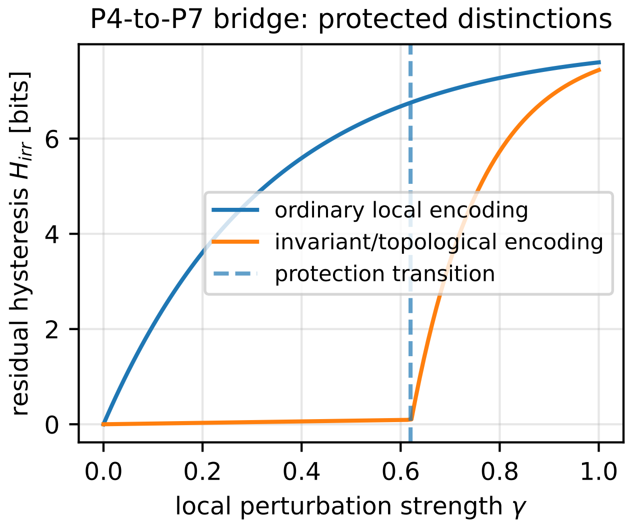

FIG. 8. P4-to-P7 bridge. Ordinary local encoding accumulates residual hysteresis under perturbation. Invariant or topological encoding can keep inverse uncertainty near zero until protection is broken. The figure is a schematic synthetic bridge, not an experimental claim. It is not a derived topological model.

Bridge to P7: invariant and topological protection

P4 studies distinctions lost under ordinary local truncation. P7 studies distinctions protected against local truncation by invariant or topological encoding. Phase-B survivors are not arbitrary compressed residues; the strongest survivors are variables whose identity is carried by invariants, often topological or global, so that local truncation leaves little residual preimage uncertainty. If $\mathcal{I}_{\rm top}$ denotes an accessible invariant sufficient to identify the relevant equivalence class, then the protected inverse uncertainty may vanish:

$$\mathrm{H}(X_{\rm fiber} \mid Z, \mathcal{I}_{\rm top}) = 0, \qquad H(X_{\rm task} \mid Z, \mathcal{I}_{\rm inv}) \approx 0. \quad (30)$$

The transition from P4 to P7 is therefore natural: P4 defines the baseline fate of unprotected distinctions; P7 asks when the boundary, topology, or invariant structure prevents local forgetting without crossing a protection-breaking transition.

FINITE-MEMORY EXIT THEOREM

Theorem 2 (Finite-memory anti-recurrence exit theorem)**.** *Let an active finite distinction system maintain a task at tolerance ϵ under finite effective memory C_t. Suppose it undergoes non-injective truncations erasing task-relevant inverse information with cumulative load $D(t) > 0$. If the system must later recover those micro-distinctions to keep the same tolerance, then at least one of the following must occur:*

1. *residual irrecoverability is accepted and task loss increases;*
2. *inverse information is maintained in side records inside an enlarged boundary;*
3. *inverse information is externalized to an environment, archive, tool, institution, or record system;*
4. *the task is relaxed so erased distinctions are no longer required;*
5. *the system fails to maintain the boundary or task.*

A finite system cannot repeatedly apply task-relevant non-injective truncation while preserving guaranteed exact preimage recovery with no side records, no externalization, no task relaxation, no enlarged boundary, and no failure.

Proof. By the anti-recurrence lemma, erased preimage distinctions are not recoverable from the truncated state alone. If guaranteed exact preimage recovery is required, the side-record criterion demands sufficient inverse information within the accounting boundary. If this information is not maintained internally, it must be externalized, or the task must be relaxed. If none of these occurs, the system lacks information required to maintain the original tolerance; loss increases or the boundary/task fails. □

Corollary 3 (O3 connection)**.** *If maintaining side records requires physical record turnover, refresh, verification, synchronization, overwrite, or cleanup, then the recovery-preserving channel enters the O3 entropy/resource ledger. P4 supplies the inverse-information term behind finite-record irreversibility.*

OPERATIONAL TESTS AND FAILURE CONDITIONS

P4 test template

A P4 experiment, simulation, or case study should specify:

1. high-resolution state X;
2. effective projection $T : X \to Z$;
3. task variable $V = f(X)$;
4. whether V is constant or nonconstant on relevant fibers;
5. accounting boundary $\mathcal{A}$;
6. side records $L_{\mathcal{A}}$;
7. capacity bottleneck path C_t;
8. recovery condition after capacity restoration;
9. residual uncertainty $\mathrm{H}(V \mid Z, L_{\mathcal{A}})$;
10. Markov closure error $\mathcal{M}_{\mathrm{D_{JS}}}$ or $\mathcal{M}_{\mathrm{D_{KL}}}$;
11. externalization cost and break-even condition;
12. failure, demotion, or task-relaxation outcome.

Falsification and demotion

P4 would be weakened or falsified by a verified case satisfying all of the following:

1. a non-injective truncation identifies task-distinct preimages;
2. no side record, hidden state, external log, environmental trace, protected invariant, or enlarged boundary contains inverse information;
3. after capacity recovery, the system exactly recovers the erased task-relevant preimage;
4. no hidden convention, stochastic lucky guess, additional physical degree of freedom, or task redefinition supplies the missing information.

Demotion occurs if the projection was misidentified, the task variable was constant on the fiber, or the erased distinction was not task-relevant.

Boundary of applicability

P4 applies to finite effective representations containing non-injective truncation. It does not apply to complete microscopic descriptions that retain all preimage information. It does not apply to tasks for which erased distinctions are irrelevant. It does not apply when external logs are counted as part of the system and contain sufficient inverse information.

NUMERICAL DEMONSTRATIONS

The simulations are deterministic synthetic demonstrations generated by `code/generate_results.py`. They illustrate definitions and theorems rather than fitting empirical data.

1. Figure 1 shows the Bayes-optimal exact recovery bound after non-injective truncation.
2. Figure 2 shows capacity-recovery asymmetry after a bottleneck.
3. Figure 3 shows the Mori-Zwanzig memory-kernel burden from eliminated hidden dimensions.
4. Figure 4 computes Jensen-Shannon Markov closure error for a synthetic 64-state chain.
5. Figure 5 compares external ledger cost with internal re-acquisition cost for $D = 12$ erased inverse bits.
6. Figure 6 classifies benign versus malignant hysteresis using task-loss sensitivity.
7. Figure 7 samples Phase-B survivor candidates by predictive utility, update cost, and closure error.
8. Figure 8 illustrates the P4-to-P7 transition from ordinary side records to invariant-supported protection.

Simulation parameters

The script uses fixed deterministic models. Figure 1 uses $n = 16$ binary distinctions and varies retained bits. Figure 2 uses a capacity schedule with a bottleneck from $t = 35$ to $t = 75$. Figure 5 compares external ledger cost with internal re-acquisition cost for $D = 12$ erased inverse bits. Figure 3 uses exponentially decaying memory-kernel proxies for eliminated hidden dimensions. Figure 4 constructs a synthetic 64-state Markov chain and computes Jensen-Shannon closure error under modular partitions. Figure 6 varies task relevance for residual hysteresis. Figure 7 samples candidate coarse variables with predictive information, update cost, maintenance cost, and closure error. Figure 8 is a schematic protected-versus-unprotected distinction model. No proprietary, biological, medical, organizational, quantum-device, or human-subject data are used.

Reproducibility checklist

1. Code availability: all simulation code is included in the replication package.
2. Deterministic execution: fixed seeds are used where synthetic randomness appears.
3. Figure reproduction: run `python code/generate_results.py`.
4. Data status: all numerical outputs are synthetic demonstrations generated from the stated model.
5. Platform independence: the code uses standard Python scientific libraries.

AI assistance disclosure

AI-assisted tools were used for language polishing, structural feedback, LaTeX drafting support, and code-debugging assistance. The author reviewed and edited all content and remains responsible for all claims, references, simulations, and conclusions. No AI system is listed as an author.

RELATION TO EXISTING THEORY

Statistical mechanics

P4 is compatible with statistical mechanics. Coarse-graining and entropy increase are familiar in Gibbsian and Boltzmannian frameworks. P4 adds an operational finite-system statement: capacity recovery after truncation does not reconstruct preimage distinctions unless inverse information is preserved inside the accounting boundary.

Reversible computation

Bennett-style reversible computation avoids erasure costs by retaining enough history to reverse computations. P4 agrees. Recovery is possible when inverse information is retained. It becomes costly when finite memory forces cleanup, overwrite, compression, externalization, or task abandonment.

Stochastic thermodynamics

Stochastic thermodynamics tracks entropy production along trajectories [12, 13, 18–20]. P4 does not replace trajectory-level accounting. It identifies a representation-level source of hidden variables, effective stochasticity, and residual inverse uncertainty that must be included in the ledger when guaranteed exact preimage recovery is demanded.

Markov state models and model reduction

The lumpability condition is standard in Markov-chain theory [7, 8]. Mori-Zwanzig projection is standard in nonequilibrium statistical mechanics and model reduction [14–16]. P4 uses both as bridges: projected variables are dynamically clean only when closure holds; otherwise, discarded distinctions return as memory kernels, noise, or closure error.

Topological protection

P4 defines ordinary local forgetting. P7 will study the complementary case: distinctions protected by global invariants or topology, including recent developments in non-Hermitian skin effects and generalized Brillouin-zone physics [22–24]. The P4 claim is therefore not that all distinctions are equally forgettable. It is that unprotected distinctions erased by non-injective truncation cannot be recovered without inverse information.

CONCLUSION

P4 establishes a finite-system account of effective anti-recurrence. A non-injective truncation does not merely hide detail; it removes inverse information from the effective record. The system may later regain capacity, but capacity recovery is not distinction recovery: capacity-recovery asymmetry is the fundamental gap between restored storage capacity and unrecovered preimage information. Guaranteed exact preimage recovery requires side records, reversible garbage registers, environmental traces, external archives, protected invariants, or an enlarged accounting boundary.

The paper's central object is therefore not the triviality that many-to-one maps lack inverses. It is the ledger that follows from that fact in finite systems: which task distinctions were erased, whether they later matter, what side records preserve them, what costs those records impose, whether the projected dynamics closes, and when hysteresis is benign abstraction rather than malignant loss. In this sense, forgetting is not merely failure. It is the price finite systems pay to produce stable macroscopic variables. Under repeated capacity pressure, finite systems are driven toward pruning, externalization, task relaxation, or invariant-supported persistence. P4 supplies the representation-side bridge for that transition.

[1] Y. Wu, "Active Finite Distinction Systems: A Formal Core for Boundary Maintenance under Finite Capacity," Zenodo (2026), doi:10.5281/zenodo.20158923.

[2] Y. Wu, "Capacity Overflow, Effective Stochasticity, and Phase-B Invariants: Critical Deficit, Markov Closure, and Invariant Selection under Finite Projection," Zenodo (2026), doi:10.5281/zenodo.20250367.

[3] Y. Wu, "Capacity Deficit and Entropy Production: A Finite Distinction Systems Bridge from Representation Pressure to Physical Dissipation," Zenodo (2026), doi:10.5281/zenodo.20254259.

[4] Y. Wu, "Boundary Maintenance and the Second Law under Finite Memory: Irreversible Record Reuse, Entropy Ledgers, and Operational Time Arrows," Zenodo (2026), doi:10.5281/zenodo.20255129.

[5] C. E. Shannon, "A mathematical theory of communication," Bell System Technical Journal 27, 379–423 and 623–656 (1948).

[6] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. (Wiley, Hoboken, 2006).

[7] J. G. Kemeny and J. L. Snell, *Finite Markov Chains* (Springer, New York, 1976).

[8] J. R. Norris, *Markov Chains* (Cambridge University Press, Cambridge, 1997).

[9] R. Landauer, "Irreversibility and heat generation in the computing process," IBM Journal of Research and Development 5, 183–191 (1961), doi:10.1147/rd.53.0183.

[10] C. H. Bennett, "Logical reversibility of computation," IBM Journal of Research and Development 17, 525–532 (1973).

[11] C. H. Bennett, "The thermodynamics of computation - a review," International Journal of Theoretical Physics 21, 905–940 (1982).

[12] J. M. R. Parrondo, J. M. Horowitz, and T. Sagawa, "Thermodynamics of information," Reviews of Modern Physics 87, 45–77 (2015), doi:10.1103/RevModPhys.87.45.

[13] U. Seifert, "Stochastic thermodynamics, fluctuation theorems and molecular machines," Reports on Progress in Physics 75, 126001 (2012).

[14] R. Zwanzig, "Ensemble method in the theory of irreversibility," Journal of Chemical Physics 33, 1338–1341 (1960).

[15] H. Mori, "Transport, collective motion, and Brownian motion," Progress of Theoretical Physics 33, 423–455 (1965).

[16] A. J. Chorin, O. H. Hald, and R. Kupferman, "Optimal prediction and the Mori-Zwanzig representation of irreversible processes," Proceedings of the National Academy of Sciences 97, 2968–2973 (2000).

[17] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cambridge, 2003).

[18] J. Degünther, J. van der Meer, and U. Seifert, "Fluctuating entropy production on the coarse-grained level: Inference and localization of irreversibility," Phys. Rev. Research **6**, 023175 (2024), doi:10.1103/PhysRevResearch.6.023175.

[19] J. van der Meer and K. Saito, "Thermodynamic bounds and error correction for faulty coarse graining," Phys. Rev. Research **7**, 043260 (2025), doi:10.1103/p3zx-7txy.

[20] U. Seifert, "Universal bounds on entropy production from fluctuating coarse-grained trajectories," arXiv:2512.07772 (2025).

[21] Y. Yu, J. Zhang, Q. Wang, and J. Shen, "Learning coarse-grained dynamics on graph," Chaos, Solitons & Fractals **194**, 116133 (2025).

[22] J. T. Gohsrich, A. Banerjee, and F. K. Kunst, "The non-Hermitian skin effect: A perspective," Europhys. Lett. **150**, 60001 (2025), doi:10.1209/0295-5075/addf77.

[23] Z. Wei, J.-Y. Fan, K. Cao, X.-R. Ma, and S.-P. Kou, "Generalized non-Hermitian skin effect," arXiv:2505.10252 (2025).

[24] X. Cai, "Non-Hermitian skin effect without point-gap topology in 2D quasicrystals," Commun. Phys. **9**, 61 (2026).