Finite-Bath Memory, Markovianization, and Environmental Forgetting in Finite Distinction Systems:
Side Records, Memory Kernels, and the Loss of Recoverable Distinctions

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

Finite systems do not lose distinctions only through internal truncation. They also depend on environments, baths, instruments, logs, niches, and external carriers whose own memory is finite. FDS-P4 studied coarse-grained anti-recurrence inside an effective representation: once a non-injective truncation removes preimage information, later capacity recovery does not reconstruct the erased distinction unless inverse information is preserved elsewhere. FDS-P3 studies that "elsewhere." It develops a finite-bath account of environmental forgetting: environmental side records initially available for recovery can become operationally unrecoverable through mixing, dilution, finite-bath saturation, inaccessible correlations, coarse readout, finite update windows, or Markovianization. The relevant object is not the full bath state, but the accessible environmental readout $R_t^E = g_t(E_t)$ induced by the observation channel, latency, cost, and resolution. We define environmental side records, accessible and hidden bath memory, bath-retained inverse information, Markov closure error, memory-kernel burden, side-record recovery probability, finite-bath saturation, and an environmental resource ledger. The central claim is conditional: an effective Markovian description is valid for a finite active system only when accessible environmental history no longer improves prediction, recovery, or boundary maintenance at the registered tolerance. P3 does not deny microscopic reversibility, unitary evolution, Liouville dynamics, or Poincaré recurrence [24]. It identifies the conditions under which a finite system can no longer operationally recover a distinction from its environment. Deterministic normal-form simulations illustrate finite-bath memory decay, Markov closure error, memory-kernel burden, side-record recovery, bath saturation, Markovianization windows, environmental ledger costs, a robot side-record worked example, and P3/P4/P7 regimes.

Reader contract. This paper is a registered FDS physical-bridge paper. It is not a replacement for open-system physics, statistical mechanics, stochastic thermodynamics, hidden Markov models, projection-operator methods, or Markov-chain theory. It does not claim that environments literally annihilate microscopic information. It does not claim that all reservoirs are infinite, memoryless, or exactly Markovian. Environmental forgetting means loss of recoverable inverse information relative to an accounting boundary, an observation channel, an update window, and a task tolerance. The symbol R_t^E denotes an accessible environmental readout, not the full physical bath state. If a side record remains accessible and recoverable, P3 requires the accounting boundary to include it rather than reinterpreting the case as forgetting.

Claim-status summary

Table I summarizes the main claims, status, and failure conditions. The formal statements are about conditional information and projection; the physical bridge statements require specified bath access, finite update windows, and accounting boundaries.

Keywords: finite distinction systems; finite bath; environmental forgetting; Markovianization; side records; memory kernels; Mori-Zwanzig; Nakajima-Zwanzig; non-Markovianity; lumpability; open systems; finite memory; environmental ledger; Physical AI.

INTRODUCTION

From P4 internal loss to P3 environmental forgetting

P4 asks what happens when a finite effective representation identifies multiple microstates with one coarse state. If the internal record loses a task-relevant preimage, later capacity recovery does not reconstruct it unless inverse information was preserved in a side record, external ledger, reversible garbage register, environmental trace, or protected invariant [2]. P3 audits the most common rescue clause in that sentence: the environment.

The central question is:

> If the missing distinction was exported to the environment, for how long does the environment keep it recoverable for the finite system that needs it?

A microscopic universe may remain reversible while a finite system loses access to the environmental record it would need for exact recovery. Environmental forgetting is not annihilation of information. It is failure of operational recovery from the accessible bath.

TABLE I. Central FDS-P3 claims, status, and failure or demotion conditions.

Claim	Status	What would weaken or falsify it
Environmental side records have finite accessible recovery capacity.	Operational FDS bridge	A finite system recovers unbounded inverse information from the environment through a finite observation channel with no latency, cost, degradation, or boundary expansion.
Markovianization is an effective forgetting condition.	Model-class bridge	A projected process is treated as Markovian while accessible history measurably improves prediction or boundary maintenance under the same variables and accounting boundary.
Memory kernels measure unresolved environmental memory.	Projection-form bridge	Eliminated variables never reappear as memory, noise, or closure error in projected dynamics despite coupling and non-lumpable projection.
Finite baths can remember temporarily, forget operationally, and recur.	Physical/model-class caveat	A finite bath is always exactly Markovian and never returns correlations under any admissible finite-bath model.
Environmental forgetting complements P4 internal truncation.	FDS bridge	Internal preimages are lost, yet environmental side records remain fully accessible indefinitely with bounded cost and no accounting-boundary change.
Bath saturation forces collisions, compression, externalization, verification cost, or loss of recoverability.	Finite-record theorem	A finite accessible bath stores more distinguishable side records than its operational capacity without collision, compression, indexing, erasure, or hidden expansion.

Why finite baths matter

Ideal reservoirs are often modeled as infinite, memoryless sinks. Real side-record media are different. A finite bath can store correlations temporarily, hide them in inaccessible degrees of freedom, return them through recurrence or information backflow, dilute them by mixing, overwrite them through record collision, and make recovery too slow or costly for the task window. Thus,

$$\text{finite bath} \neq \text{perfect eraser} \tag{1}$$

and

$$\text{finite bath} \neq \text{infinite Markovian reservoir.} \tag{2}$$

P3 therefore treats the environment neither as a magical archive nor as an always-Markovian garbage dump. It treats environmental records as finite, partially accessible, and auditable.

P3 in the physical bridge ladder

P3 occupies the environmental side of the FDS physical bridge ladder. Table II summarizes the relation [3–5].

Main contributions

This paper contributes seven objects to the FDS physical spine:

1. an accessible/hidden decomposition of environmental side records;
2. an environmental side-record recovery criterion;
3. an operational Markovianization criterion based on history advantage and JS closure error;
4. a projection-operator interpretation of memory kernels as unresolved environmental memory;
5. a finite-bath recurrence caveat preventing naive monotone-forgetting claims;
6. a finite-bath saturation and record-collision theorem;
7. deterministic normal-form simulations and engineering protocols for testing side-record recoverability.

FDS BACKGROUND

Active finite distinction systems

An active finite distinction system is represented by

$$S = (X, E, B, M, Y, A, U, \pi, \ell, \Phi, \mathcal{P}, \tau), \tag{3}$$

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 [1]. P3 focuses on E, Y, M, U, π, ℓ, Φ, and τ.

TABLE II. P3 in the local FDS physical bridge ladder.

Paper	Core axis	P3 relation
P4	internal loss	what the effective record loses
P3	environmental forgetting	what side records fail to preserve
O3	record turnover	what environmental record reuse costs
P6	throughput	whether recovery happens in time
P7	protection	which records avoid ordinary forgetting

The active-boundary qualification matters. Environmental memory belongs in P3 only if access to environmental records can change prediction, recovery, action, or future boundary-maintenance loss. Passive correlations that cannot be read, used, or acted upon by the finite system are not operational side records.

Capacity deficit and externalization

For a task family Ψ, rate-distortion demand over update window τ is $R^{(\tau)}_{\min}(\epsilon;\Psi)$ and internal capacity is $C_S = \log_2|M|$. The FDS capacity deficit is

$$\Delta_\epsilon(\tau) = R^{(\tau)}_{\min}(\epsilon;\Psi) - C_S. \tag{4}$$

If $\Delta_\epsilon(\tau) > 0$, the system must approximate, externalize, relax the task, improve compression, or fail [1]. P3 adds a condition to the externalization branch: an external record helps only while it remains recoverable through the system's accounting boundary.

ENVIRONMENTAL SIDE RECORDS

Definition 1 (Environmental side record). *Let X_0 denote a task-relevant past state or preimage whose identity may later matter. An environmental side record at time t is an environmental carrier*

$$L^E_t = r(E_t) \tag{5}$$

that contains recoverable information about X_0 for the system under its observation channel and accounting boundary.

Definition 2 (Accessible environmental readout). *The environment may contain a full physical state E_t, but the finite system does not generally access E_t directly. Define the accessible environmental readout*

$$R^E_t = g_t(E_t), \tag{6}$$

where g_t includes the observation channel, finite resolution, latency, cost, readout protocol, and accounting boundary. Equivalently, R^E_t may be read as the accessible readout σ-algebra induced by the system. The hidden bath E^{hid}_t denotes environmental variables or correlations not available through R^E_t within the task window and resource budget.

Definition 3 (Accessible inverse information). *The accessible inverse information about X_0 at time t is*

$$I_{\text{acc}}(t) = \mathrm{I}(X_0; R^E_t \mid M_t, Y_{0:t}), \tag{7}$$

and the hidden inverse information is

$$I_{\text{hid}}(t) = \mathrm{I}(X_0; E^{\text{hid}}_t \mid M_t, Y_{0:t}). \tag{8}$$

P3 is concerned with $I_{\text{acc}}(t)$, not with metaphysical total information. The notation R^E_t, if used informally, means the accessible readout structure $R^E_t = g_t(E_t)$ rather than the full physical bath state.

Definition 4 (Environmental forgetting). *Environmental forgetting of X_0 in the residual-entropy sense occurs when*

$$\mathrm{H}(X_0 \mid M_t, R^E_t, Y_{0:t}) > h_{\text{tol}}, \tag{9}$$

where h_{tol} is a residual-entropy tolerance measured in bits. If the task tolerance is a distortion tolerance ϵ, the recovery condition should instead be expressed through a Bayes-risk, classification-error, or rate-distortion criterion.

Definition 5 (Bath record capacity). *For a finite accessible bath record space $\mathcal{E}_B$, define*

$$C_B = \log_2|\mathcal{E}_B|. \tag{10}$$

For continuous baths, C_B is replaced by a finite-resolution, noisy-channel, rate-distortion, or operational readout capacity.

Definition 6 (Side-record recovery probability). *The Bayes-optimal environmental recovery probability is*

$$P^E_{\text{rec}}(t) = \mathbb{E}_{M_t,R^E_t,Y_{0:t}}\left[\max_x p(x \mid M_t, R^E_t, Y_{0:t})\right]. \tag{11}$$

For a task alphabet $\mathcal{X}_0$, $P^E_{\text{rec}}(t) \ge 1-\delta$ is a classification-style recovery criterion. A distortion-style criterion is

$$\inf_{\widehat{X}_0 = f(M_t,R^E_t,Y_{0:t})} \mathbb{E}[d(X_0,\widehat{X}_0)] \le \epsilon. \tag{12}$$

These are environmental analogues of the P4 preimage recovery bound.

TABLE III. P3 environmental recovery decision table. Conditions are operational tests under a fixed accounting boundary, observation channel, and update window.

Condition	FDS interpretation
$\mathrm{H}(X_0 \mid M_t, R_t^E, Y_{0:t}) \leq h_{\mathrm{tol}}$	Environment recovers the distinction in residual-entropy form.
$\mathrm{H}(X_0 \mid M_t, R_t^E, Y_{0:t}) > h_{\mathrm{tol}}$ but $\mathrm{I}(X_0; E_t^{\mathrm{hid}} \mid M_t, Y_{0:t}) > 0$	Hidden record exists, but is not operationally accessible.
$P_{\mathrm{rec}}^E(t) \geq 1 - \delta$	Bayes recovery succeeds at error tolerance δ.
$\inf_{\hat{X}_0} \mathbb{E}[d(X_0, \hat{X}_0)] \leq \epsilon$	Distortion/Bayes-risk recovery succeeds.
$\mathrm{I}(Z_{t+1}; H_t \mid Z_t) \leq \eta$	Operational Markovianization by history-advantage test.
$\mathcal{K}_{\mathrm{burden}}(T) > \eta$	Unresolved bath memory still matters beyond cutoff T.
$r_{\mathrm{ext}} T > C_B$	Finite-bath saturation or record-collision pressure.
Invariant quotient exists	P7 protected bypass rather than ordinary side-record recovery.

MARKOVIANIZATION

Exact and operational Markovianity

Let $Z_t = \pi(X_t)$ be a projected system variable and let H_t denote accessible history, unresolved bath memory, or hidden variables relevant to prediction. Exact Markovianity requires

$$P(Z_{t+1} \mid Z_t, H_t) = P(Z_{t+1} \mid Z_t). \tag{13}$$

Operational Markovianization at tolerance η requires only that history adds negligible task-relevant information:

$$\mathrm{I}(Z_{t+1}; H_t \mid Z_t) \leq \eta. \tag{14}$$

Markov closure error

A distributional closure error can be defined by

$$\mathcal{M}_{\mathrm{env}} = \mathbb{E}_t \, \mathrm{D}_{\mathrm{JS}}\big(P(Z_{t+1} \mid Z_t, H_t) \,\|\, P(Z_{t+1} \mid Z_t)\big). \tag{15}$$

When $\mathcal{M}_{\mathrm{env}} \to 0$, accessible history has become operationally irrelevant for that projected task. This is environmental forgetting in the Markovianization sense.

Remark 1 (Relation to lumpability)**.** *For a Markov process on a microscopic state space, a projection remains Markovian exactly only under the usual lumpability conditions: projected transition probabilities must agree for all microstates in each coarse fiber [13, 14]. If lumpability fails, the projected process carries hidden-state memory unless environmental mixing or inaccessible bath dynamics makes the history operationally irrelevant. Recent work on memory loss in open quantum systems supports the view that Markovian behavior can arise as an operational loss of recoverable history, rather than as a primitive property of an ideal bath [26].*

Markovianization is not microscopic erasure

Markovianization is an effective boundary condition, not a claim that microscopic information has ceased to exist. A process can be effectively Markovian for a finite agent because relevant bath correlations are inaccessible, diluted, below tolerance, outside the update window, or too costly to recover.

FINITE-BATH MEMORY KERNELS

Projection-operator approaches write the resolved dynamics schematically as

$$\frac{d}{dt} P X_t = P L P X_t + \int_0^t K(t-s) P X_s \, ds + F_t, \tag{16}$$

where P is a projection, L is a generator, K is a memory kernel, and F_t is an orthogonal noise term [16–18]. Open quantum versions are often expressed through Nakajima-Zwanzig memory kernels [15, 17, 19–21].

In FDS terms, K is unresolved environmental memory. It is the price of pretending that the bath has forgotten while unresolved correlations still influence future boundary statistics. Define the kernel burden

$$\mathcal{K}_{\mathrm{burden}}(T) = \int_T^\infty \|K(s)\| \, ds. \tag{17}$$

The norm $\|K(s)\|$ is model-class dependent: it may be an operator norm, trace norm, total-variation influence norm, or task-loss-weighted kernel norm. If $\mathcal{K}_{\mathrm{burden}}(T) \leq \eta$, then for update windows longer than T the Markov approximation is operationally valid at tolerance η in that model class.

Finite baths complicate the picture. Kernels may decay, plateau, oscillate, or revive. Therefore finite-bath forgetting is a windowed operational claim, not a universal monotone theorem. Recent work connects Nakajima-Zwanzig memory kernels and influence-functional descriptions in a unified framework for open quantum dynamics with memory [25]. Structured long-memory kernels can also be represented through fractional open-system dynamics [29].

MAIN CLAIMS

Theorem 1 (Environmental side-record recovery criteria)**.** *Let X_0 be a task-relevant preimage absent from internal memory M_t, and let $R_t^E = g_t(E_t)$ be the accessible environmental readout. Environmental recovery at time t can be registered in any of the following task-dependent forms:*

$$\mathrm{H}(X_0 \mid M_t, R_t^E, Y_{0:t}) \leq h_{\mathrm{tol}}, \tag{18}$$

$$P_{\mathrm{rec}}^E(t) \geq 1 - \delta, \tag{19}$$

$$\inf_{\widehat{X}_0 = f(M_t, R_t^E, Y_{0:t})} \mathbb{E}[d(X_0, \widehat{X}_0)] \leq \epsilon. \tag{20}$$

If the registered recovery criterion fails under the fixed accounting boundary, then the distinction is environmentally forgotten relative to that boundary, even if microscopic information remains in E_t^{hid}.

Proof. The system's available recovery variables are M_t, R_t^E, and the observation history $Y_{0:t}$. The residual-entropy condition states that the remaining inverse uncertainty is below the allowed bit tolerance h_{tol}. The Bayes condition states that the maximum-posterior decoder succeeds with probability at least $1 - \delta$. The distortion condition states that some admissible decoder achieves the task loss tolerance ϵ. If the pre-registered criterion fails for all decoders using only these variables, successful recovery requires additional variables, a larger accounting boundary, a different task tolerance, an invariant quotient, or hidden side records. □

Theorem 2 (Operational Markovianization criterion)**.** *Let (Z_t, H_t) be a resolved variable and accessible history/bath-memory pair. If*

$$\mathrm{I}(Z_{t+1}; H_t \mid Z_t) \leq \eta \tag{21}$$

or, alternatively, if a registered distributional closure error such as $\mathcal{M}_{\mathrm{env}} \leq \eta_{\mathrm{D_{JS}}}$ is below tolerance, then Z_t is operationally Markovian for the task at tolerance η or $\eta_{\mathrm{D_{JS}}}$ respectively. The mutual-information criterion and the JS closure criterion are parallel operational tests, not strict mathematical equivalents.

Proof. The history variable H_t improves prediction of Z_{t+1} over Z_t by at most the registered information tolerance. A finite task whose loss is insensitive below that tolerance cannot operationally exploit the additional history. Thus the Markov closure is valid relative to the task, variables, and accounting boundary. □

Proposition 1 (Finite-bath non-monotonicity)**.** *For a finite closed or partially observed bath, accessible inverse information need not decrease monotonically. There can exist recurrence or backflow times t_r such that*

$$I_{\mathrm{acc}}(t_r) > I_{\mathrm{acc}}(t_r - \Delta t). \tag{22}$$

Therefore P3 cannot claim universal monotone environmental forgetting. It claims operational forgetting only over a specified window, observation channel, task tolerance, and accounting boundary. A recurrence that occurs after the task window or beyond the recovery budget does not operationally rescue the lost distinction.

Theorem 3 (Finite-bath saturation and record collision)**.** *Suppose an accessible bath has operational record capacity C_B and receives side records at rate r_{ext} bits per unit time. If over a time interval T,*

$$r_{\mathrm{ext}} T > C_B, \tag{23}$$

then at least one exit channel is required: overwrite, record collision, compression, indexing overhead, externalization into a larger bath, verification cost, loss of recoverability, or task relaxation.

Proof. The bath cannot host more than C_B distinguishable accessible record bits under the fixed resolution and accounting boundary. If the exported record load exceeds C_B and no exit channel occurs, the bath would distinguish more record states than its operational capacity permits. Thus some collision, compression, overwrite, boundary expansion, additional indexing, or recovery loss is forced. □

ENVIRONMENTAL FORGETTING LEDGER

Side records are not free merely because they are outside internal memory. P3 uses a resource-first ledger:

$$\begin{aligned}\dot{\mathcal{R}}_{\mathrm{P3}} =& \dot{\mathcal{R}}_{\mathrm{write}}^E + \dot{\mathcal{R}}_{\mathrm{mix/isolate}}^E + \dot{\mathcal{R}}_{\mathrm{monitor}}^E \\ &+ \dot{\mathcal{R}}_{\mathrm{retrieve}}^E + \dot{\mathcal{R}}_{\mathrm{verify}}^E + \dot{\mathcal{R}}_{\mathrm{refresh}}^E.\end{aligned} \tag{24}$$

When thermal reservoir, heat-channel, and calibration assumptions are available, this resource ledger may be connected to an entropy-production ledger. Without such calibration it remains a resource, latency, accessibility, or verification ledger [9–12]. Strong-coupling finite-bath models show that system and bath energetic quantities, entropy production, and work-like resources must be treated explicitly when the bath is finite and non-Markovian [27].

An environmental side record survives for the system only if it can be written, preserved, indexed, retrieved, verified, and used within the relevant update window and resource budget. Otherwise it may exist microscopically while failing operationally.

WORKED EXAMPLE: ROBOT EXTERNAL MEMORY IS NOT FREE

Consider a mobile robot that loses track of an object during occlusion. Let X_0 be the object identity and pose

TABLE IV. Core demonstration parameters. Full parameter values are exported to `data/simulation_parameters.json`.

Parameter	Value / role
Bath sizes	$16, 64, 256$ states (normalized)
Mixing rates	$0.018, 0.035, 0.060$
Task alphabet	$\lvert\mathcal{X}_{\text{task}}\rvert = 8$
Markov threshold	$\eta = 0.05$ JS proxy
Bath capacity	$C_B = 260$ normalized record bits
Figures	deterministic normal-form demonstrations

before occlusion, M_t the robot's internal memory at the current time, and R_t^E the accessible environmental readout available now: camera buffer, spatial map, external log, scene trace, environmental cue, or database response. P3 asks whether X_0 is recoverable from $(M_t, R_t^E, Y_{0:t})$ within the task window.

The robot may fail environmental recovery in several ways. The camera buffer may have been overwritten. The spatial map may be stale. The external log may be synchronized too slowly. The object or environment may have moved, destroying the cue. The database may contain the old record but be unreachable inside the control window. In all such cases, the information may exist somewhere in the world or in a larger data system, but it is not a usable side record for the embodied agent.

In P3 notation, successful recovery requires either

$$P_{\text{rec}}^E(t) \geq 1 - \delta \tag{25}$$

or a task-specific distortion condition such as

$$\mathbb{E}[d(X_0, \widehat{X}_0)] \leq \epsilon. \tag{26}$$

If these fail while $\mathrm{I}(X_0; E_t^{\text{hid}}) > 0$, the distinction has not been metaphysically destroyed; it is hidden outside the agent's accessible recovery boundary. This is the Physical-AI form of environmental forgetting.

NUMERICAL DEMONSTRATIONS

The simulations are deterministic normal-form demonstrations generated by `code/generate_results.py`. They illustrate definitions and failure modes rather than fitting empirical data. No experimental, biological, proprietary, medical, human-subject, or device data are used.

EXPERIMENTAL AND ENGINEERING PROTOCOLS

Protocol 1 (Hidden-memory Markov test)**.** *Compare a model using* $P(Z_{t+1} \mid Z_t)$ *with one using* $P(Z_{t+1} \mid Z_t, H_t)$. *If accessible history improves prediction or boundary maintenance beyond tolerance, the bath has not operationally forgotten.*

Protocol 2 (Side-record recovery assay)**.** *After internal truncation, attempt recovery of* X_0 *using environmental records. Estimate* $P_{\text{rec}}^E(t)$ *and* $\mathrm{H}(X_0 \mid M_t, R_t^E, Y_{0:t})$. *If recovery fails while hidden bath variables are excluded, P3 classifies the distinction as environmentally forgotten relative to the accounting boundary.*

Protocol 3 (Bath-size scaling)**.** *Vary bath size, mixing rate, and readout accessibility. Changing bath size and mixing rate changes the Markovianization window: faster mixing often reduces accessible history, while finite small baths may show recurrence, backflow, or non-monotone recovery.*

Protocol 4 (Environmental externalization audit)**.** *Write records into an external medium and vary refresh, verification, monitoring, and retrieval. Measure recovery probability, synchronization cost, readout latency, record collision, and ledger burden.*

Protocol 5 (Physical AI memory test)**.** *Compare local agent memory, external logs, environment-embedded memory, and invariant/morphological memory. Report recovery rate, Markov closure error, environmental ledger, latency, and boundary-maintenance loss.*

FALSIFICATION AND DEMOTION CONDITIONS

P3 is weakened if a finite accessible environment preserves arbitrarily many side records indefinitely with no cost, no degradation, no collision, no latency, and no boundary expansion. A Markovianization claim is weakened if accessible history continues to improve task prediction under the same variables and tolerance. A finite-bath forgetting model is weakened if measured recurrence or backflow is present but the model assumes monotone decay. A specific memory-kernel or saturation model may fail without falsifying the FDS formal core; it should then be demoted to a narrower model class. If side records survive, P3 does not fail. It says the accounting boundary must include them. Non-Markovianity can also act as a thermodynamic resource rather than merely as a burden; P3 therefore treats memory as an auditable channel, not as intrinsically good or bad [28].

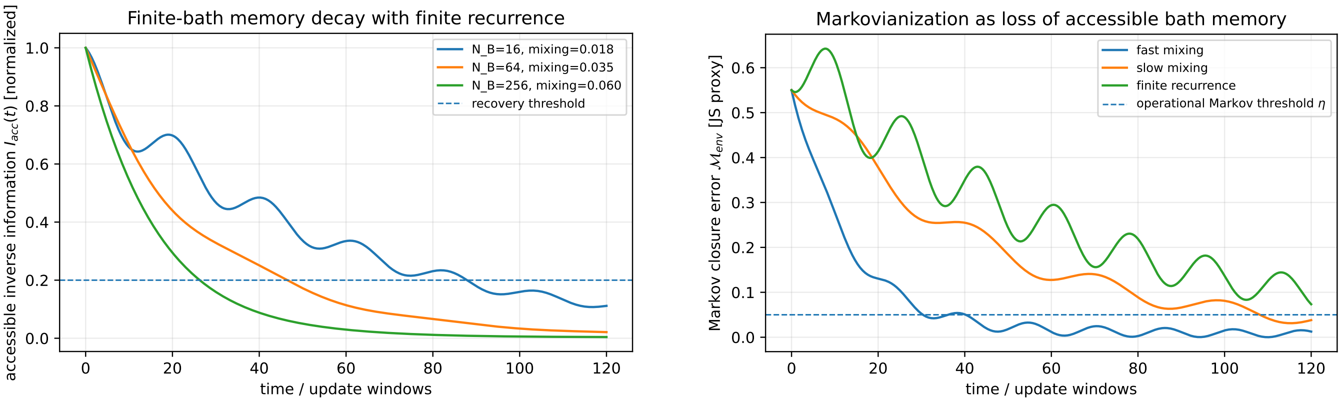

FIG. 1. Left: finite-bath memory decay. Accessible inverse information can decay while small baths exhibit recurrence-like revivals. Right: Markov closure error. Operational Markovianization occurs when accessible history no longer improves the projected transition beyond tolerance.

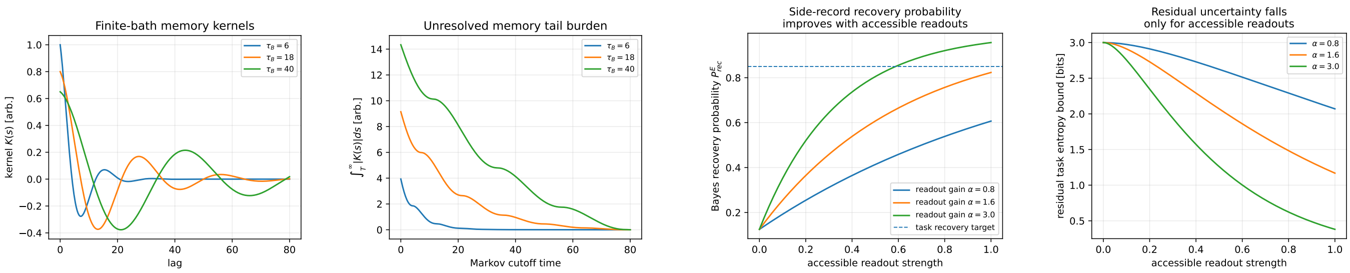

FIG. 2. Left: memory-kernel burden. Long kernel tails represent unresolved bath memory and increase the cost of Markov closure. Right: side-record recovery. Environmental records improve Bayes recovery only when they are accessible; hidden microscopic correlation does not by itself help the finite system.

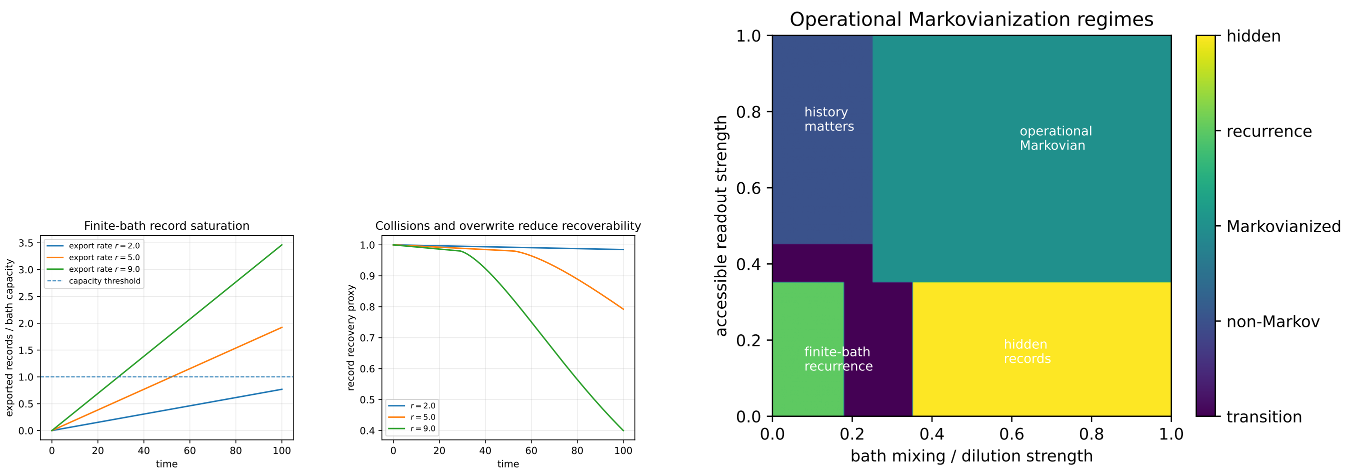

FIG. 3. Left: finite-bath saturation. When exported record load exceeds accessible bath capacity, collisions, overwrite, compression, or recovery loss begins. Right: operational Markovianization regimes. The map is a synthetic threshold diagram, not an empirical phase diagram.

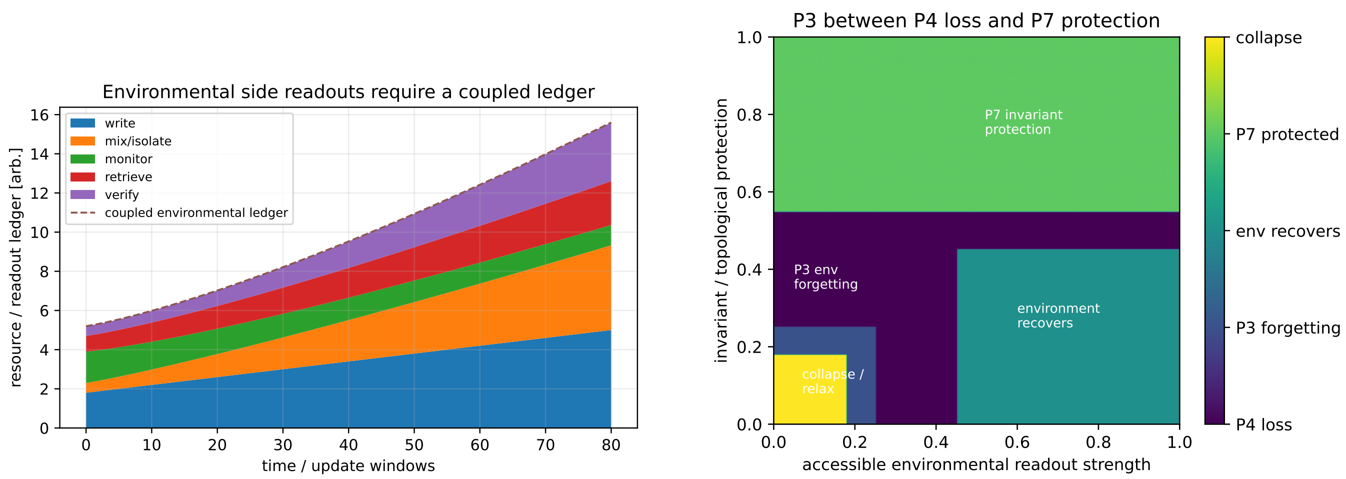

FIG. 4. Left: environmental side records require a coupled ledger for writing, isolation or mixing control, monitoring, retrieval, and verification. Right: P3 sits between P4 internal loss and P7 invariant protection: if local records lose preimages, the environment may recover them, forget them operationally, or be bypassed by invariant protection.

RELATION TO EXISTING THEORY

- **Open quantum systems** study system-bath coupling and non-Markovianity; P3 uses them as finite-bath memory and backflow model classes [19–21].
- **Projection-operator methods** study projected dynamics with memory kernels; P3 reads these kernels as unresolved environmental memory [15–18].
- **Markov-chain lumpability** gives exact closure conditions for projected dynamics [13, 14, 22, 23].
- **Information theory** supplies side-information and recovery bounds for accessible inverse information [6–8].
- **Stochastic thermodynamics** supplies resource and entropy-ledger bridges for environmental record maintenance [11, 12, 27].

P3 does not replace these theories. It packages their common operational implication inside FDS: a finite system can use the environment as memory only when the relevant environmental records remain accessible, recoverable, and affordable within the task window.

GLOSSARY

DESIGN IMPLICATION: EXTERNAL MEMORY IS NOT FREE MEMORY

This section is not part of the formal proof. It extracts the FDS design rule implied by P3. A physical AI agent cannot treat the environment, a database, a retrieval system, or a world state as an infinite perfectly indexed memory. External memory is useful only while it remains recoverable within the task window.

For embodied agents, P3 suggests four rules. First, external logs require refresh, synchronization, and verification. Second, environment-embedded memory can drift, mix, or become inaccessible. Third, high-speed physical tasks require local or invariant carriers rather than slow bath recovery. Fourth, robust agents need tests for when the environment has forgotten: if accessible history no longer improves prediction or recovery, the memory should be treated as Markovianized or lost for that task.

CONCLUSION

P4 showed that lost preimages do not return merely because internal capacity returns. P3 shows that exporting inverse information to the environment is not a free solution. A finite bath may temporarily preserve correlations, but those correlations can mix, saturate, recur, become inaccessible, or demand verification and retrieval costs beyond the system's budget. Markovianization is therefore not a metaphysical erasure of the past. It is an operational condition under which accessible history no longer improves boundary maintenance at the registered tolerance. A side record is memory only while it remains accessible, recoverable, affordable, and timely.

Together, P3 and P4 give the two halves of effective forgetting. The system may lose a distinction internally, and the environment may fail to return it externally. P7 identifies a protected exception through invariant quotients; P6 adds whether any recovery can occur within the required update window; O3 accounts for record reuse and environmental ledgers. P3 identifies the environmental cost of relying on side records.

TABLE V. Core P3 terms.

Term	P3 meaning
side record	environmental carrier of recoverable inverse information
accessible readout	$R_t^E = g_t(E_t)$ usable within boundary, cost, and latency
hidden bath	inaccessible, delayed, too costly, mixed, or unmeasured variables
environmental forgetting	loss of operational recovery, not microscopic erasure
Markovianization	accessible history no longer improves task prediction
memory kernel	unresolved environmental memory in projected dynamics
bath saturation	exported record load exceeds accessible bath capacity

[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, "Coarse-Grained Anti-Recurrence and Informational Hysteresis in Finite Memory Systems," Zenodo (2026), doi:10.5281/zenodo.20265065.

[3] Y. Wu, "Boundary Maintenance and the Second Law under Finite Memory," Zenodo (2026), doi:10.5281/zenodo.20255129.

[4] Y. Wu, "Speed, Precision, and Dissipation Bounds for Boundary Maintenance in Finite Distinction Systems: Finite Update Windows, Bottleneck Throughput, and Resource-Ledger Exit Theorems," Zenodo (2026), doi:10.5281/zenodo.20269733.

[5] Y. Wu, "Topological Obstruction to Forgetting in Finite Distinction Systems," Zenodo (2026), doi:10.5281/zenodo.20265386.

[6] C. E. Shannon, "A mathematical theory of communication," Bell Syst. Tech. J. **27**, 379–423 and 623–656 (1948).

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

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

[9] R. Landauer, "Irreversibility and heat generation in the computing process," IBM J. Res. Dev. **5**, 183–191 (1961), doi:10.1147/rd.53.0183.

[10] C. H. Bennett, "The thermodynamics of computation – a review," Int. J. Theor. Phys. **21**, 905–940 (1982).

[11] U. Seifert, "Stochastic thermodynamics, fluctuation theorems and molecular machines," Rep. Prog. Phys. **75**, 126001 (2012).

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

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

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

[15] S. Nakajima, "On quantum theory of transport phenomena," Prog. Theor. Phys. **20**, 948–959 (1958).

[16] H. Mori, "Transport, collective motion, and Brownian motion," Prog. Theor. Phys. **33**, 423–455 (1965).

[17] R. Zwanzig, "Ensemble method in the theory of irreversibility," J. Chem. Phys. **33**, 1338–1341 (1960).

[18] A. J. Chorin, O. H. Hald, and R. Kupferman, "Optimal prediction and the Mori-Zwanzig representation of irreversible processes," Proc. Natl. Acad. Sci. USA **97**, 2968–2973 (2000).

[19] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, Oxford, 2002).

[20] A. Rivas, S. F. Huelga, and M. B. Plenio, "Quantum non-Markovianity: characterization, quantification and detection," Rep. Prog. Phys. **77**, 094001 (2014).

[21] H.-P. Breuer, E.-M. Laine, J. Piilo, and B. Vacchini, "Colloquium: Non-Markovian dynamics in open quantum systems," Rev. Mod. Phys. **88**, 021002 (2016), doi:10.1103/RevModPhys.88.021002.

[22] J. L. Doob, *Stochastic Processes* (Wiley, New York, 1953).

[23] N. G. van Kampen, *Stochastic Processes in Physics and Chemistry*, 3rd ed. (North-Holland, Amsterdam, 2007).

[24] H. Poincaré, "Sur le problème des trois corps et les équations de la dynamique," Acta Math. **13**, 1–270 (1890).

[25] F. Ivander, L. P. Lindoy, and J. Lee, "Unified framework for open quantum dynamics with memory," Nat. Commun. **15**, 8087 (2024), doi:10.1038/s41467-024-52081-3.

[26] A. Ben-Asher, A. I. Fernández-Domínguez, and J. Feist, "Memory loss is contagious in open quantum systems," Commun. Phys. **8**, 400 (2025), doi:10.1038/s42005-025-02310-y.

[27] D. Tiwari, B. Bose, and S. Banerjee, "Strong coupling non-Markovian quantum thermodynamics of a finite-bath system," J. Chem. Phys. **162**, 114104 (2025), doi:10.1063/5.0254029.

[28] G. Zambon and G. Adesso, "Quantum processes as thermodynamic resources: The role of non-Markovianity," Phys. Rev. Lett. **134**, 200401 (2025), doi:10.1103/PhysRevLett.134.200401.

[29] B. Peng and Y. Zhang, "A fractional calculus framework for open quantum dynamics: From Liouville to Lindblad to memory kernels," J. Chem. Phys. **164**, 084103 (2026), doi:10.1063/5.0312309.