Topological Obstruction to Forgetting in Finite Distinction Systems: Quotient Invariants, Non-Hermitian Skin Effects, and Topological Side-Ledgers

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

FDS-P7 develops the topological-protection paper in the physical bridge sequence of Finite Distinction Systems. P4 studied ordinary coarse-grained anti-recurrence: once a non-injective truncation removes task-relevant preimage information from the effective record, later capacity recovery does not reconstruct the discarded distinction unless inverse information is preserved in side records, external ledgers, enlarged accounting boundaries, or invariant carriers. P7 studies the complementary case: distinctions whose recoverability is carried not by ordinary local records but by invariant or topological structure. We formalize this through a quotient map $q : X \to Q$ from microscopic state space to an invariant-sector space and the induced random variable $Q_{\mathrm{inv}} = q(X)$, a noisy invariant recovery bound, a positive-part operational forgetting rate, and a resource-ledger audit. The central physical bridge is the Non-Hermitian Skin Effect (NHSE), where point-gap winding, generalized Brillouin-zone structure, boundary-localized modes, and boundary-sensitive spectra supply one model class for invariant-supported persistence. P7 does not claim a violation of the Second Law and does not claim that all NHSE mechanisms are point-gap-topological. Protected invariants relocate entropy and resource accounting rather than deleting it. This framework provides a physical bridge for non-catastrophic forgetting in resource-constrained agents: not all distinctions must be retained as local records if task identity is carried by stable invariant structure. Deterministic numerical demonstrations illustrate quotient-sector winding, boundary diagnostic channels, noisy invariant recovery, positive-part forgetting-rate freezing, resource-ledger relocation, dual-channel signatures, metastable learning between P4 and P7 regimes, and a Physical-AI design normal form.

Reader contract. This paper is a registered FDS physical-bridge paper. It does not derive the Second Law from topology, does not claim perpetual motion, does not claim that protected invariants produce zero-cost computation, and does not claim that every non-Hermitian skin effect is governed by point-gap winding. Its narrower claim is conditional: under stated locality, accessibility, accounting-boundary, and protection-gap assumptions, some task-relevant distinctions are locally non-forgettable because their identity is carried by an invariant sector. If the invariant is destroyed, the protection gap closes, the boundary condition changes, or the accounting boundary excludes the invariant, P7 reduces to the ordinary P4 case of coarse-grained preimage loss.

Claim-status summary

Table I summarizes the main claims and demotion conditions.

INTRODUCTION

From P4 forgetting to P7 protection

P4 established the baseline finite-memory result: under ordinary non-injective truncation, task-relevant preimage distinctions are lost from the effective representation unless inverse information is preserved in side records, reversible garbage registers, environmental traces, protected invariants, or enlarged accounting boundaries [2]. In P4 notation, a local record Z may satisfy

$$H(V \mid Z) > 0, \tag{1}$$

where $V = f(X)$ is a task-relevant sector. P7 studies the complementary question. When can a distinction survive local forgetting without being stored as an ordinary local record?

The P7 answer is that some distinctions are carried by invariant sectors. In such cases local coordinates may be discarded, local records may become incomplete, and ordinary P4 hysteresis may appear at the local level, while the task-relevant sector remains recoverable through an invariant side-ledger:

$$H(V \mid Z) > 0, \qquad H(V \mid Z, Q_{\mathrm{inv}}) \approx 0. \tag{2}$$

The phrase "topology remembers" will occasionally be used as shorthand, but the formal meaning is narrower: the task sector remains inferable from an accessible invariant carrier.

Why non-Hermitian skin effects are the bridge model

Non-Hermitian systems have become a central setting for studying boundary-sensitive spectra, point gaps,

TABLE I. Central FDS-P7 claims, status, and failure or demotion conditions.

Claim	Status	What would weaken or falsify it
Invariant side-ledgers can suppress P4 residual inverse uncertainty.	Formal FDS bridge	A task variable factors through an accessible invariant, but $H(V \mid Z, Q_{\mathrm{inv}})$ remains high under the stated assumptions.
Noisy invariant readout gives a bounded recovery penalty.	Information bound	A noisy invariant readout with error probability δ exceeds the Fano-style bound without hidden information or changed task labels.
Local perturbations cannot change a protected invariant without a protection-breaking event.	Topological bridge	A local perturbation changes the invariant while the protection gap, locality assumptions, and accounting boundary remain intact.
NHSE supplies a model class for invariant-supported persistence.	Physical bridge	NHSE is present, but it carries no stable recoverable distinction, no boundary-sensitive protection, and no robustness to local perturbation in the registered model class.
Protection relocates entropy/resource accounting rather than deleting it.	O3-compatible accounting claim	A protected invariant supplies indefinite maintenance with no drive, boundary, refresh, dissipation, verification, control, or external ledger.
Protected phases can generate a dual forgetting/ledger signature.	Experimental bridge	Confirmed protection-breaking transition with no feature in operational forgetting and no corresponding resource/entropy signature under a well-powered registered protocol.

exceptional structures, generalized Brillouin zones, and the non-Hermitian skin effect [11–14]. In NHSE systems, an extensive set of bulk eigenstates can accumulate near boundaries under open boundary conditions, and the conventional Bloch bulk description may fail [15–18]. NHSE has now been realized or diagnosed across multiple platforms, including dynamic mechanical systems, ultracold atoms, and digital quantum simulations [22–24]. This makes NHSE a natural FDS bridge: recoverable distinction structure may be relocated from bulk-local records to boundary-sensitive spectral structure.

P7 uses NHSE as one model-class realization, not as a universal explanation of protected persistence. Recent work cautions that NHSE should not be reduced to a single point-gap mechanism in all settings; generalized and anomalous forms exist [19–21]. P7 therefore phrases its strongest statements for specified protection conditions rather than for every non-Hermitian system.

Main contributions

This paper contributes seven objects to the FDS physical spine:

1. a quotient-map formulation of invariant side-ledgers, $q : X \to Q$;
2. an invariant side-ledger theorem extending the P4 side-record criterion;
3. a noisy-invariant recovery bound for imperfect invariant diagnostics;
4. a positive-part operational forgetting rate $\kappa^{+}_{\mathrm{forget}}$;
5. an NHSE model-class bridge linking point-gap winding, GBZ structure, and boundary localization to invariant-supported persistence;
6. an O3-compatible resource/entropy ledger showing that protection relocates cost rather than deleting it;
7. deterministic numerical demonstrations and an engineering protocol linking P7 to Physical AI.

FDS BACKGROUND

Active finite distinction systems

The FDS core object is

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

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]. P7 focuses on π, $\mathcal{P}$, ℓ, Φ, and on the accounting boundary that determines whether an invariant sector is available to the system.

P4 baseline and O3 ledger

Let $T : X \to Z$ be a non-injective coarse-graining. P4 defines residual inverse uncertainty by conditioning on the effective record and accessible side information. If $V = f(X)$ is nonconstant on a truncated fiber and no side record distinguishes the relevant preimages, then

$$H(V \mid Z, L_A) > 0. \quad (4)$$

P4 therefore studies the fate of unprotected distinctions.

O3 studies finite-memory boundary maintenance as an operational Second-Law channel: finite records must be written, preserved, refreshed, repaired, synchronized, externalized, reused, or destroyed, and physical record turnover enters an entropy/resource ledger under bridge assumptions [3]. P7 inherits the same accounting principle. Topological protection does not violate thermodynamics; it changes what must be maintained and where the cost is paid.

INVARIANT SIDE-LEDGERS

Definition 1 (Local perturbation family)**.** *A local perturbation family $\mathcal{P}_{\text{loc}}$ is a set of allowed perturbations whose physical action is bounded, local, exponentially decaying in distance, or restricted to finite neighborhoods relative to the boundary and system size.*

Definition 2 (Invariant quotient map)**.** *A side-ledger is any accessible carrier of inverse or task-sector information outside the local coarse record Z. An invariant side-ledger is a quotient map*

$$q : X \to Q, \quad (5)$$

from microscopic state space to an invariant-sector space Q, together with the induced random variable

$$Q_{\text{inv}} = q(X). \quad (6)$$

Conditional entropies are taken with respect to Q_{inv}, not the map q itself. Its sector label remains stable under a specified local perturbation family until a protection condition is broken:

$$q(Px) = q(x), \qquad P \in \mathcal{P}_{\text{loc}}, \quad g(Px) > 0. \quad (7)$$

The protection margin g may denote a spectral gap, point gap, code distance, winding stability margin, localization margin, or other obstruction parameter.

Definition 3 (Protected task sector)**.** *A finite task-sector random variable $V \in \mathcal{V}$ is protected by q when there exists a map $h : Q \to \mathcal{V}$ such that*

$$V = h(Q_{\text{inv}}) = h(q(X)). \quad (8)$$

The quotient q is a task-sector ledger: it need not recover all microscopic detail, only the task identity.

Definition 4 (Topological obstruction to forgetting)**.** *A task distinction V is topologically protected relative to $(Z, q, \mathcal{P}_{\text{loc}})$ if*

$$H(V \mid Z) > 0, \qquad H(V \mid Z, Q_{\text{inv}}) = 0, \quad (9)$$

and q is stable under all allowed local perturbations while $g > 0$.

Theorem 1 (Invariant side-ledger theorem)**.** *Let $T : X \to Z$ be a local coarse-graining that loses a task distinction $V = f(X)$. Suppose there exists an accessible invariant quotient $q : X \to Q$ such that:*

1. V factors through q, i.e. $V = h(q(X))$;

2. q is stable under $\mathcal{P}_{\text{loc}}$ while $g > 0$;

3. the accounting boundary includes access to q.

Then

$$H(V \mid Z, Q_{\text{inv}}) = 0, \quad (10)$$

even when $H(V \mid Z) > 0$.

Proof. Since $V = h(Q_{\text{inv}}) = h(q(X))$, the quotient-sector random variable determines the task label. Conditioning on (Z, Q_{inv}) therefore determines V regardless of whether Z alone distinguishes the local preimage. Stability under $\mathcal{P}_{\text{loc}}$ ensures that local perturbations preserving $g > 0$ do not change the invariant sector. Accessibility ensures that Q_{inv} is part of the audited record rather than an unavailable mathematical label. □

Proposition 1 (Invariant-aware representation)**.** *Let*

$$Z^+ = (Z, Q_{\text{inv}}), \qquad Q_{\text{inv}} = q(X). \quad (11)$$

If $V = h(Q_{\text{inv}})$, then

$$H(V \mid Z^+) = 0. \quad (12)$$

Ordinary coarse-graining may lose local detail; invariant-aware coarse-graining preserves task identity when the task factors through the invariant quotient.

Proof. The claim follows from $V = h(Q_{\text{inv}})$, so V is a function of $Z^+ = (Z, Q_{\text{inv}})$ even if Z alone loses local detail. □

Proposition 2 (Noisy invariant recovery bound)**.** *Let $\widehat{Q}$ be a noisy diagnostic of the invariant quotient and suppose that the induced task-sector error obeys*

$$\mathbb{P}\{h(\widehat{Q}) \neq V\} \leq \delta. \quad (13)$$

Then

$$H(V \mid Z, \widehat{Q}) \leq h_b(\delta) + \delta \log_2(|\mathcal{V}| - 1), \quad (14)$$

where h_b is the binary entropy function. Thus a discrete or integer-valued invariant can remain useful under noisy readout when its error probability is small.

Proof. The claim is the standard Fano-style upper bound on residual uncertainty given a classifier with error probability at most δ [5, 6]. The FDS interpretation is that the invariant readout need not be perfect; it must keep residual task-sector uncertainty below the operational tolerance.

This bound assumes a finite discrete task alphabet $\mathcal{V}$, $0 \leq \delta \leq 1 - 1/|\mathcal{V}|$, and a diagnostic whose decoded task label $h(\widehat{Q})$ has error probability at most δ. For continuous task sectors, an operational finite-resolution, hypothesis-testing, or rate-distortion version is required. □

Remark 1 (Topology as a nonlocal side-ledger)**.** *In P4, inverse information is usually carried by explicit side records, external ledgers, or enlarged memory. P7 extends this: a quotient invariant can act as a side-ledger that is not stored as a local register. Its recovery power comes from global structure, discreteness, and a protection condition. NHSE provides a model class in which recoverable distinction structure is relocated from bulk-local records to boundary-sensitive spectral structure. In that class, a boundary-localized profile can provide an accessible physical carrier or diagnostic channel for the quotient sector under the model-class assumptions.*

NHSE AS A MODEL-CLASS REALIZATION

Minimal non-reciprocal chain

Consider a one-dimensional non-Hermitian asymmetric hopping model with open or periodic boundary conditions,

$$H(\theta) = \sum_{j=1}^{N-1} \left[(t+\theta)c_{j+1}^\dagger c_j + (t-\theta)c_j^\dagger c_{j+1}\right] + i\gamma \sum_j c_j^\dagger c_j. \tag{15}$$

The corresponding Bloch symbol under periodic boundary conditions is

$$H(k,\theta) = (t+\theta)e^{ik} + (t-\theta)e^{-ik} + i\gamma. \tag{16}$$

For a reference energy E_0, the point-gap winding is

$$\nu(E_0) = \frac{1}{2\pi i}\int_0^{2\pi} dk\, \partial_k \log\det[H(k,\theta) - E_0]. \tag{17}$$

For the scalar model and $E_0 = i\gamma$, the spectral loop encloses E_0 with orientation set by $\mathrm{sgn}(\theta)$ for $\theta \neq 0$, while $\theta = 0$ closes the point gap. The quotient map for this model class is

$$q(X) = \nu(E_0) \in \mathbb{Z}, \qquad Q_{\mathrm{inv}} = q(X). \tag{18}$$

The model-class protection margin can be represented by the point-gap distance

$$g_{\mathrm{pg}}(E_0) = \min_k |H(k,\theta) - E_0| \tag{19}$$

for the scalar model, or more generally by

$$g_{\mathrm{pg}}(E_0) = \min_k |\det(H(k,\theta) - E_0)|. \tag{20}$$

The NHSE bridge applies only while this margin remains nonzero. A protection-breaking transition occurs when $g_{\mathrm{pg}} \to 0$, when boundary conditions change, when a nonlocal perturbation is applied, or when the carrier itself is destroyed.

Generalized Brillouin zone and boundary localization

Under open boundary conditions, conventional Bloch waves do not generally describe the physical spectrum. The generalized Brillouin zone replaces e^{ik} by a complex factor β whose modulus is set by the open-boundary characteristic equation [16]. For the asymmetric chain, the bulk eigenstates under open boundaries acquire exponential localization,

$$|\psi_j| \propto e^{-\kappa j}, \qquad \kappa \sim \frac{1}{2}\log\left|\frac{t+\theta}{t-\theta}\right|. \tag{21}$$

The boundary-localized profile provides an accessible physical carrier or diagnostic channel for the quotient sector under the model-class assumptions. It is not an ordinary memory register. It is a boundary-sensitive spectral organization whose identity is constrained by the point-gap sector and boundary condition. Recent work on quantum geometry of the NHSE suggests additional diagnostic channels for skin localization and GBZ singularities [26].

Limits of the bridge

NHSE is not synonymous with topology. P7 uses the point-gap/GBZ setting as a clean model class. Other forms of skin effect, anomalous skin response, quasicrystalline behavior, disorder-sensitive boundary accumulation, and collective many-body variants may require different quotient maps and different protection conditions [19–21, 25]. Failure of this NHSE bridge would demote this realization, not the formal invariant side-ledger theorem.

COARSE-GRAINING OBSTRUCTION

Ordinary local truncation

P4 ordinary forgetting begins with a non-injective map $T : X \to Z$. If $T(x) = T(x')$ while $V(x) \neq V(x')$, then Z loses the task distinction unless side information is available. P7 adds that the side information need not be a local memory register; it can be an invariant quotient.

Invariant-aware truncation

A local truncation may identify microstates,

$$T(x) = T(x'), \tag{22}$$

while the invariant separates them,

$$q(x) \neq q(x'). \tag{23}$$

If $V = h(Q_{\rm inv}) = h(q(X))$, then local forgetting of coordinates does not imply loss of the protected task distinction. The invariant-aware effective representation is

$$Z^+ = (Z, Q_{\rm inv}), \qquad Q_{\rm inv} = q(X). \tag{24}$$

For protected task sectors,

$$H(V \mid Z^+) = 0. \tag{25}$$

Proposition 3 (Protection-breaking alternatives)**.** *Suppose q is stable under $\mathcal{P}_{\rm loc}$ while $g > 0$. A protected task sector can be erased only if at least one of the following occurs: (i) $g \to 0$; (ii) a perturbation outside $\mathcal{P}_{\rm loc}$ is applied; (iii) the boundary condition changes; (iv) q is excluded from the accounting boundary; (v) the carrier of the invariant is destroyed; or (vi) the task no longer factors through q.*

Proof. If none of these occurs, then the perturbation is local and gap-preserving, the boundary still includes the quotient, and $q(Px) = q(x)$. Since $V = h(Q_{\rm inv}) = h(q(X))$, the protected task sector is unchanged. □

RG and model-reduction caution

The older PoN/NHSE note described the obstruction in terms of RG flow on the generalized Brillouin zone and the breakdown of bulk-dominance assumptions when an extensive number of states localize at boundaries. P7 adopts a narrower claim: ordinary local coarse-graining that ignores the protected boundary-sensitive sector is incomplete for tasks depending on that sector. It does not assert that every RG construction fails in every non-Hermitian system. It asserts that invariant-aware coarse-graining is required when task identity factors through a protected sector.

OPERATIONAL FORGETTING AND RESOURCE LEDGER

Positive-part operational forgetting rate

Let V_0 be an initial task sector and let R_t be the accessible local record at time t. Define the finite-window positive-part forgetting rate

$$\kappa^+_{\rm forget}(t;\tau) = \frac{1}{\tau}\left[H(V_0 \mid R_{t+\tau}, Q_{{\rm inv},t+\tau}) - H(V_0 \mid R_t, Q_{{\rm inv},t})\right]_+ . \tag{26}$$

The positive part is used because later observations can reduce uncertainty. When only the local record is used, write $\kappa^{\rm loc}_{\rm forget}$; when the invariant quotient variable $Q_{\rm inv,t}$ is included, write $\kappa^{\rm inv}_{\rm forget}$.

A protected phase is characterized operationally by

$$\kappa^{\rm inv}_{\rm forget} \approx 0, \qquad \kappa^{\rm loc}_{\rm forget} > 0 \tag{27}$$

for tasks whose identity is carried by the invariant. In an ideal thermodynamic-limit model, a protection-breaking transition may appear as a non-analytic kink in $\kappa^+_{\rm forget}$ or in its parameter derivative. In finite devices, P7 predicts a finite-size-rounded sharp crossover whose sharpening under system-size scaling is the operational signature.

Resource ledger first, entropy ledger when calibrated

A protected distinction does not delete entropy production. Following O3, P7 treats persistence as an accounting relocation. For an active non-Hermitian device the directly measured quantity may be a resource or power ledger:

$$P_{\rm ledger} = P_{\rm bulk} + P_{\rm boundary} + P_{\rm drive} + P_{\rm control} + P_{\rm verify}. \tag{28}$$

When thermodynamic calibration is available, one may translate these channels into an entropy ledger,

$$\dot\Sigma_{\rm ledger} = \dot\Sigma_{\rm bulk} + \dot\Sigma_{\rm boundary} + \dot\Sigma_{\rm drive} + \dot\Sigma_{\rm control} + \dot\Sigma_{\rm verify} \geq 0. \tag{29}$$

The distinction is important. Throughout the experimental protocol, the directly measured object is a resource or power ledger. It becomes an entropy ledger only when reservoir temperatures, heat channels, and thermodynamic calibration are available. P7 does not infer a true entropy production rate from an uncalibrated power proxy.

Proposition 4 (Protection relocates cost)**.** *If an invariant suppresses local forgetting of a task distinction while the physical system remains driven, open, refreshed, controlled, or measured, then the cost of persistence must be counted in the coupled ledger rather than inferred from the local record alone.*

Criterion 1 (FDS dual-channel test)**.** *A P7 model class is strongly supported when the same registered control parameter produces both:*

1. *a feature in operational forgetting, $\kappa^+_{\rm forget}(\lambda)$;*
2. *a corresponding feature in the resource or entropy ledger, $P_{\rm ledger}(\lambda)$ or $\dot\Sigma_{\rm ledger}(\lambda)$,*

under controls excluding ordinary measurement, heating, finite-size, or boundary artifacts. In ideal limiting models this may appear as a two-kink signature or a pair of non-analytic derivative features; in finite devices P7 predicts a rounded finite-size feature rather than demanding a mathematically sharp kink in every apparatus.

The physical mark of a topological transition is therefore not free memory. It is the reorganization of forgetting cost and maintenance cost across the accounting boundary.

NUMERICAL DEMONSTRATIONS

The simulations are deterministic demonstrations generated by `code/generate_results.py`. They illustrate definitions and model-class bridges rather than fitting experimental data.

Simulation parameters

Figure 1 uses the scalar asymmetric hopping model with $t = 1$, $\gamma = 0.15$, and selected $\theta \in \{-0.35, 0, 0.35\}$, and records $g_{\rm pg}(E_0)$. Figure 2 uses an OBC chain with $N = 64$ sites and a Fano-style noisy invariant diagnostic over the alphabet $\mathcal{V}$. Figure 3 uses deterministic normal-form trajectories for protected and unprotected task-sector uncertainty and a synthetic resource ledger. Figure 4 gives a dual-channel transition schematic and a P4/metastable/P7 regime diagram. Figure 5 gives a Physical-AI design normal form. No experimental, biological, proprietary, medical, human-subject, or device data are used.

EXPERIMENTAL AND ENGINEERING PROTOCOLS

Protocol 1 (Active non-Hermitian circuit lattice)**.** *Build an active non-reciprocal circuit chain with tunable asymmetric coupling. Measure boundary voltage accumulation, complex spectral response, and total power or heat proxies under a sweep of the non-reciprocity parameter. The resource ledger can begin with directly measured quantities such as* $P_{\rm drive}(t) = \sum_i V_i(t)I_i(t)$ *and op-amp/control power, before any entropy interpretation is attempted.*

Protocol 2 (Symmetric and boundary controls)**.** *Repeat the protocol with reciprocal couplings, altered boundary conditions, and local perturbations below and above the protection margin. P7 predicts that the protected-sector signature should track the invariant and boundary condition, not merely total heating or measurement intensity.*

Protocol 3 (Finite-size scaling)**.** *Repeat the sweep for* $N = 20, 40, 80, 160$ *or comparable sizes. Ideal non-analyticities are rounded in finite devices; a registered protocol should therefore specify whether it expects a kink, a sharp crossover, or finite-size sharpening.*

Protocol 4 (Operational forgetting assay)**.** *Choose a task-sector label* V *associated with the invariant quotient and a local record* R_t *that omits the quotient. Compare* $H(V_0 \mid R_t)$ *against* $H(V_0 \mid R_t, Q_{{\rm inv},t})$ *and estimate* $\kappa^+_{\rm forget}$*. If the protected sector is real and accessible, the invariant-aware uncertainty should remain low until protection breaks.*

FALSIFICATION AND DEMOTION CONDITIONS

The P7-NHSE bridge is weakened or falsified if NHSE or the claimed topological sector is confirmed but local perturbations preserving the protection gap erase the protected sector with no nonlocal intervention, gap closure, boundary-condition change, hidden uncontrolled channel, or diagnostic failure. The resource-ledger bridge is weakened if protected persistence is observed but no measurable redistribution, crossover, or ledger feature appears under a well-powered registered protocol.

A negative result in the NHSE model class does not falsify the FDS formal core. It demotes the specific physical realization. The invariant side-ledger theorem survives as a formal bridge whenever another invariant carrier satisfies the stated conditions.

DESIGN IMPLICATION: INVARIANT CARRIERS FOR PHYSICAL AI

This section is not part of the NHSE bridge proof. It extracts the FDS design rule implied by P7. A robust physical AI agent should not attempt to store every distinction in volatile memory. Under finite capacity, the most important task identities should be placed where local perturbations cannot easily erase them: morphology, closed-loop dynamics, sensorimotor topology, stable environmental scaffolds, error-correcting memory, or invariant physical carriers. In this sense P7 provides a design principle for non-catastrophic forgetting in embodied agents:

$$\begin{gathered}\text{do not preserve all details;}\\ \text{preserve the quotient that preserves task identity.}\end{gathered} \tag{30}$$

A robot's balance manifold, a manipulation affordance, or a navigation invariant may be more valuable than a large local texture log. This is not a claim that robots implement NHSE. It is a functional projection of the P7

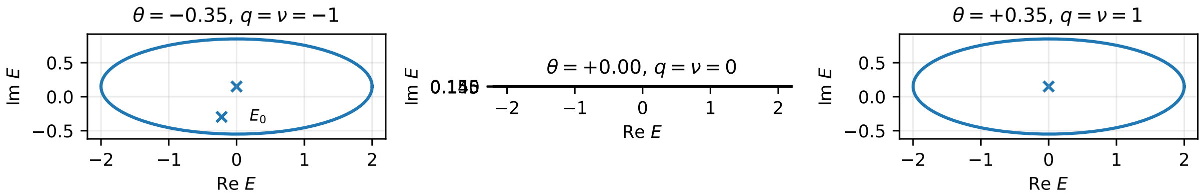

FIG. 1. Quotient-sector winding in the scalar asymmetric hopping model. The invariant side-ledger is the quotient-sector random variable $Q_{\rm inv} = q(X) = \nu(E_0)$. Nonzero non-reciprocity encloses the reference point with opposite orientations, while the Hermitian point closes the model-class point gap.

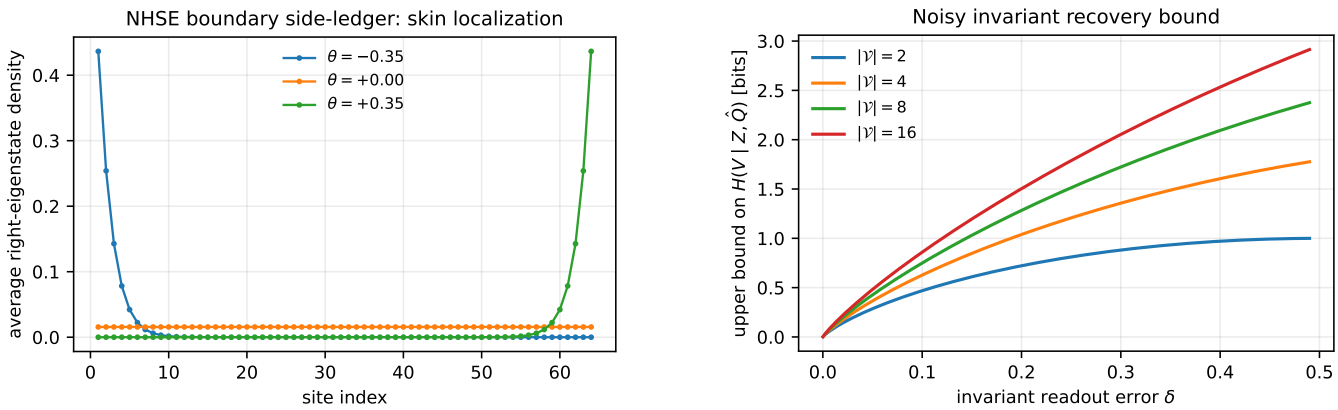

FIG. 2. Left: NHSE boundary diagnostic channel under OBC. Boundary localization can provide a physical carrier or diagnostic channel for the quotient sector under the model-class assumptions. Right: noisy invariant recovery bound. When the invariant readout error δ is small, residual task-sector uncertainty remains bounded even though local coordinates are coarse-grained.

principle: critical task identity should be moved from arbitrary complex records into law-like structural identity whenever possible.

RELATION TO EXISTING THEORY

P4 and O3

P4 supplies the baseline local-forgetting ledger. P7 supplies the protected complement. O3 supplies the entropy/resource accounting rule: protection does not make cost disappear, it moves the cost to boundary, drive, verification, maintenance, or externalization channels [2, 3].

Non-Hermitian topology

The NHSE literature provides the physical model class for boundary-sensitive spectral protection, point gaps, generalized Brillouin zones, and non-Bloch bulk-boundary correspondence [11, 12, 14, 16–18]. P7 does not replace that theory. It translates a subset of its structure into the FDS language of protected task sectors and accounting boundaries.

Model reduction and coarse-graining

Ordinary model reduction discards local degrees of freedom. P4 asks what inverse information has been lost. P7 asks whether a task sector factors through an invariant quotient. Invariant-aware coarse-graining is not less coarse; it is coarser in local coordinates while retaining the quotient that matters.

CONCLUSION

P7 defines the protected complement of P4 forgetting. P4 showed that capacity recovery is not distinction recovery: a local record that has discarded task-relevant preimages cannot reconstruct them without inverse information. P7 shows that not all inverse information must be stored as a local record. Some task sectors can remain recoverable through quotient invariants, topological side-ledgers, boundary-sensitive spectral structure, or other protected carriers.

This does not produce free memory. Topological protection does not delete entropy production; it relocates the accounting burden from local erasure to boundary, drive, verification, and protection maintenance, consis-

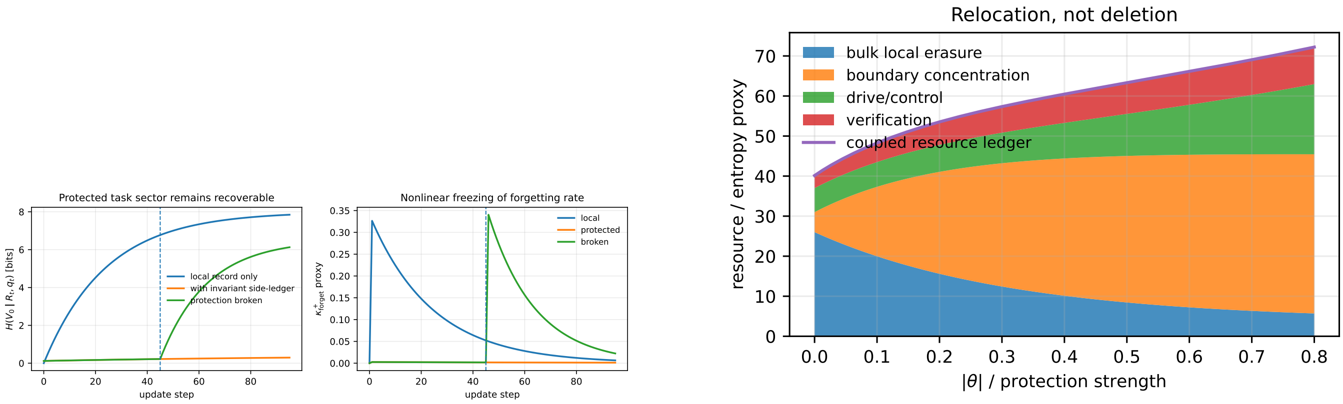

FIG. 3. Left: protected task sectors suppress the positive-part forgetting rate until the protection condition is broken. The plotted uncertainty trajectories are FDS normal-form simulations, not direct quantum-dynamical entropy estimates. Right: protection relocates resource burden from local bulk forgetting into boundary, drive/control, and verification channels. The y-axis is a resource/entropy proxy; thermodynamic interpretation requires reservoir and heat-channel calibration.

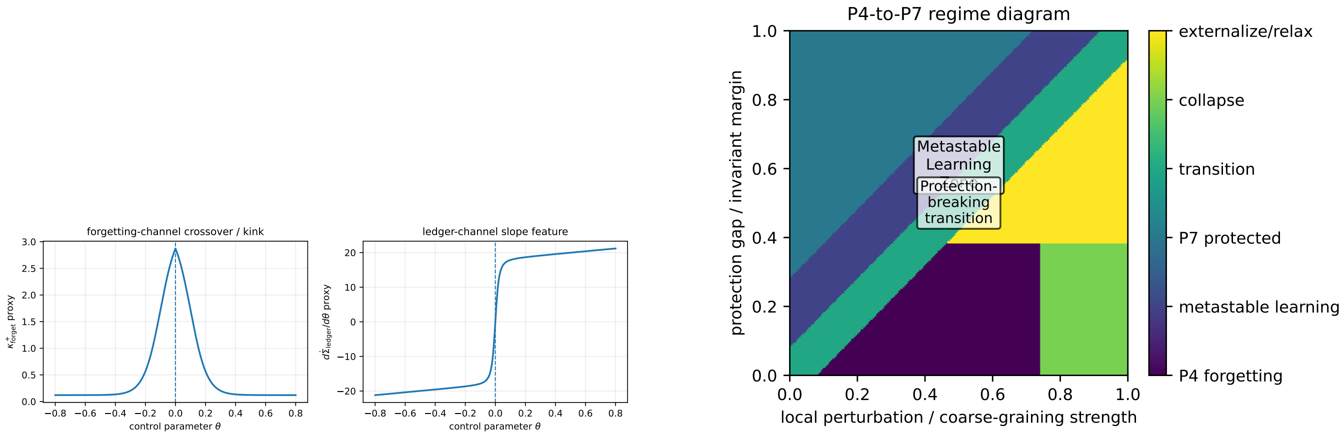

FIG. 4. Left: FDS dual-channel signature. In the ideal limit, the same control parameter can produce a feature in operational forgetting and a corresponding feature in the ledger channel. Finite devices may show rounded crossovers. Right: P4/P7 regime diagram. A metastable learning zone lies between ordinary local forgetting and stable P7 protection: the system pays update, verification, and maintenance costs to convert fragile local records into invariant quotients [4]. Protection fails when the margin closes or the carrier, boundary condition, or accounting access is destroyed.

tent with information thermodynamics and stochastic thermodynamics [7–10]. The physical mark of a topological transition is therefore the reorganization of forgetting cost, not the disappearance of cost. The NHSE model class provides a concrete setting where this reorganization can be modeled, simulated, and tested.

P7 defines the boundary between arbitrary complex records and law-like structural identity. What persists under finite local forgetting is not raw detail, but invariant-supported task structure. For Physical AI, the implication is direct: resource-constrained agents should not merely remember more; they should place critical task identity into stable physical, morphological, sensorimotor, or topological quotients that survive local degradation until a protection condition is broken.

At higher levels of FDS organization, symbolic language may function as a collective invariant side-ledger: not a literal physical topology, but a quotient structure that preserves task identity across local memory turnover. This analogy is not used in the P7 proof, but it motivates later FDS applications to symbolic and civilizational memory.

[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: Lost Preimages, Side Records, and Capacity-Recovery Asymmetry,” Zenodo (2026), doi:10.5281/zenodo.20265065.

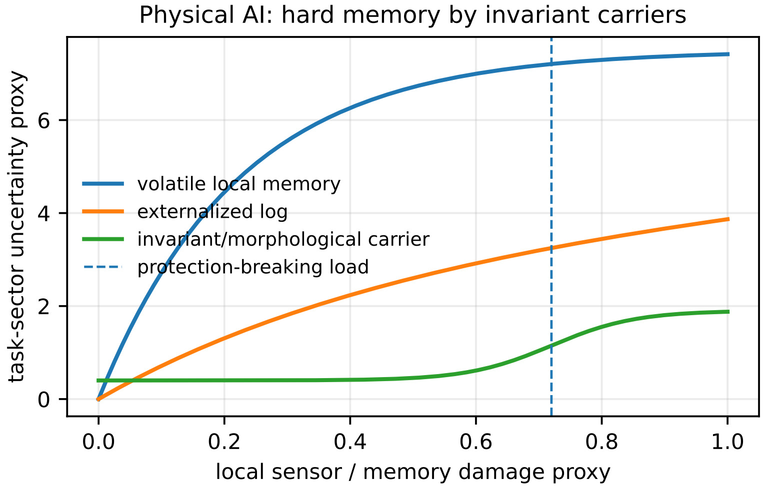

FIG. 5. Design implication for Physical AI. Critical task identity should not always be stored in volatile local memory. It can be carried by invariant, morphological, sensorimotor, or closed-loop physical structure. The figure is a normal-form design schematic, not a claim that robotics literally implements NHSE.

[3] 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.

[4] 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.

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

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

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

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

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

[10] 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.

[11] Y. Ashida, Z. Gong, and M. Ueda, "Non-Hermitian physics," Adv. Phys. **69**, 249–435 (2020), doi:10.1080/00018732.2021.1876991.

[12] S. Yao and Z. Wang, "Edge states and topological invariants of non-Hermitian systems," Phys. Rev. Lett. **121**, 086803 (2018), doi:10.1103/PhysRevLett.121.086803.

[13] K. Kawabata, K. Shiozaki, M. Ueda, and M. Sato, "Symmetry and topology in non-Hermitian physics," Phys. Rev. X **9**, 041015 (2019), doi:10.1103/PhysRevX.9.041015.

[14] E. J. Bergholtz, J. C. Budich, and F. K. Kunst, "Exceptional topology of non-Hermitian systems," Rev. Mod. Phys. **93**, 015005 (2021), doi:10.1103/RevModPhys.93.015005.

[15] F. K. Kunst, E. Edvardsson, J. C. Budich, and E. J. Bergholtz, "Biorthogonal bulk-boundary correspondence in non-Hermitian systems," Phys. Rev. Lett. **121**, 026808 (2018), doi:10.1103/PhysRevLett.121.026808.

[16] K. Yokomizo and S. Murakami, "Non-Bloch band theory of non-Hermitian systems," Phys. Rev. Lett. **123**, 066404 (2019), doi:10.1103/PhysRevLett.123.066404.

[17] N. Okuma, K. Kawabata, K. Shiozaki, and M. Sato, "Topological origin of non-Hermitian skin effects," Phys. Rev. Lett. **124**, 086801 (2020), doi:10.1103/PhysRevLett.124.086801.

[18] K. Zhang, Z. Yang, and C. Fang, "Correspondence between winding numbers and skin modes in non-Hermitian systems," Phys. Rev. Lett. **125**, 126402 (2020), doi:10.1103/PhysRevLett.125.126402.

[19] 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.

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

[21] X. Cai, "Non-Hermitian skin effect without point-gap topology in 2D quasicrystals," Commun. Phys. **9**, 61 (2026), doi:10.1038/s42005-026-02496-9.

[22] Z. Li, L.-W. Wang, X. Wang, Z.-K. Lin, G. Ma, J.-H. Jiang, et al., "Observation of dynamic non-Hermitian skin effects," Nat. Commun. **15**, 6544 (2024), doi:10.1038/s41467-024-50776-1.

[23] E. Zhao, Z. Wang, C. He, T. F. J. Poon, K. K. Pak, Y.-J. Liu, P. Ren, X.-J. Liu, and G.-B. Jo, "Two-dimensional non-Hermitian skin effect in an ultracold Fermi gas," Nature **637**, 565–573 (2025), doi:10.1038/s41586-024-08347-3.

[24] R. Shen, T. Chen, B. Yang, et al., "Observation of the non-Hermitian skin effect and Fermi skin on a digital quantum computer," Nat. Commun. **16**, 1340 (2025), doi:10.1038/s41467-025-55953-4.

[25] B. H. Kim, J.-H. Han, and M. J. Park, "Collective non-Hermitian skin effect: point-gap topology and the doublon-holon excitations in non-reciprocal

many-body systems," Commun. Phys. **7**, 73 (2024), doi:10.1038/s42005-024-01564-2.

[26] K.-I. Imura and K. Kawabata, "Quantum geometry of the non-Hermitian skin effect," arXiv:2604.10043 (2026).