

Readability and Recoverability of Decohered Information

Sina Iman *Independent Researcher*

Abstract

Under unitary quantum mechanics, decoherence redistributes information into system-environment correlations but does not destroy it. The statement that information remains “in the environment” is operationally incomplete: it may be globally present without being locally readable or actively recoverable under realistic constraints. We distinguish three notions—**presence**, **readability**, and **recoverability**—and show that they are operationally distinct.

In controlled simulations of four decoherence models (random circuits, Trotterized Ising, isotropic Heisenberg, CNOT fan-out), we find that recovery is **inverse-specific**: the correct inverse of the decoherence process succeeds while generic circuits of comparable depth fail. Coupling geometry further shapes accessibility: anisotropic Ising and isotropic Heisenberg dynamics exhibit sharply different local capture structure despite comparable scalar decoherence indicators.

On superconducting (Rigetti Ankaa-3) and trapped-ion (IonQ Forte) hardware, exact inverse recovery remains observable on both platforms. On IonQ, the geometry-dependent capture separation between Ising and Heisenberg dynamics reappears ($\sim 10 \times$ in capture efficiency), while on Rigetti it is suppressed to the noise floor by routing overhead—indicating that local accessibility is more hardware-sensitive than global recoverability.

The central empirical result is not that exact inversion can restore decohered correlations—a consequence expected under unitary control—but that decoherence models with comparable scalar indicators and identical global recoverability can produce sharply different local accessibility structures, and that this difference is selectively visible depending on hardware compilation overhead. These results do not establish asymptotic hardness barriers or distinguish between interpretations of quantum mechanics. They support a narrower conclusion: **the operational accessibility of decohered information is determined by coupling geometry, not merely by the degree of decoherence**, and whether it remains readable or recoverable depends jointly on interaction structure, observer access, and available control operations.

Keywords decoherence, recoverability, computational complexity, quantum Darwinism, inverse-specificity, operational accessibility, quantum information

Published Apr 02, 2026

Code, Data & Updates

dpid.org/1077

1. INTRODUCTION

Decoherence redistributes locally accessible coherence into system-environment correlations Schlosshauer (2019); Zurek (1981); Zurek (2003). Under unitary quantum mechanics, this process preserves information globally. But global preservation does not settle the operational question: under what conditions can decohered information still be accessed or recovered?

This paper distinguishes three notions that are often conflated:

1. **Presence:** information exists in the global quantum state.
2. **Readability:** information is passively extractable from accessible fragments.
3. **Recoverability:** information can be actively restored by an allowed class of operations.

These need not coincide. Information may be globally preserved yet unreadable from local fragments, or readable only as a classical pointer observable, or recoverable only through highly structured control. This paper studies how recoverability depends on both the **control class** available to an observer and the **interaction geometry** of the decoherence process.

1.a. Scope and Claim Status

To maintain evidentiary discipline, we separate three levels of claim:

- **Claim A (standard):** Under unitary evolution, information is globally preserved during decoherence.
- **Claim B (operational):** Recovery of decohered correlations is **inverse-specific**: the correct inverse succeeds while generic circuits of comparable depth fail.
- **Claim C (complexity-theoretic):** The required inverse may have large circuit complexity, making recovery computationally inaccessible.

Claim A is standard quantum mechanics. Claim B is the central claim of this paper, supported analytically and experimentally across four decoherence models. Claim C is established only in a restricted setting (random circuit decoherence, Theorem 3, Appendix A) and remains open for physical Hamiltonians, where the Haah-Liu-Tan no-go theorem Haah et al. (2025) poses a nontrivial obstruction. “Hard to recover” can mean **structurally constrained** (only the right operation works) or **computationally intractable** (the right operation exists but is prohibitively costly). This paper argues for the former and only conditionally for the latter. Accordingly, the scientific question addressed here is not whether a known unitary can in principle be inverted, but how operational accessibility changes under restricted access and structured decoherence, and which of those accessibility signatures survive under realistic control.

1.b. Relation to existing work

The framework draws on three bodies of work. **Decoherence and classicality:** standard decoherence theory explains the emergence of pointer states via

environmental monitoring Schlosshauer (2019); Zurek (1981); Zurek (2003); quantum Darwinism Brandão et al. (2015); Zurek (2009) shows how classical observables become redundantly encoded in environmental fragments. **Quantum recovery:** decoupling theorems Brandão et al. (2016), Petz recovery maps Junge et al. (2018); Petz (1986); Wilde (2015), and entanglement structure formalize when and how information can be recovered from subsystems. **Complexity of recovery:** Harlow and Hayden Harlow & Hayden (2013) showed that information encoded in Hawking radiation requires a computation of enormous complexity to extract. The operational pattern—information present but inaccessible without a structured inverse—is shared with the present setting, but the analogy is motivational only: the scrambling regime, the environment structure, and the available hardness tools differ substantially, and this paper does not claim to extend Harlow-Hayden results to ordinary decoherence.

Statement	Status
Information preserved globally under unitary decoherence (A)	Proven (standard QM)
Recovery requires the specific inverse of the decoherence process (B)	Demonstrated for four models; theoretically expected
Random-circuit inverse complexity grows linearly with depth (C, restricted)	Proven (Theorem 3, Appendix A)
Physical decoherence induces generic recovery hardness (C, general)	Open conjecture; obstructed by Haah-Liu-Tan
Recoverability structure depends on coupling geometry	Demonstrated (Ising vs. Heisenberg, simulator + trapped-ion hardware)

The framework is formulated in operational terms and does not depend on any particular interpretation of quantum mechanics. All interpretations preserving unitarity make identical predictions for the observables measured here.

1.c. Terminology

Four core terms recur throughout:

- **Presence:** information exists in the global state (guaranteed by unitarity).
- **Readability:** information is passively extractable from local or environmental fragments.
- **Recoverability:** information is actively restorable under a specified operation class \mathcal{O} , quantified by $\mathcal{R}(\mathcal{O}, A)$.
- **Inverse-specificity:** the empirical finding that only the correct inverse operation succeeds; generic circuits of comparable depth fail.

These terms map onto established formalisms (Holevo quantity, Petz recovery, accessible information; see Section 3.5). The individual concepts are not new. The contribution is showing that they **take qualitatively different values simultaneously** for a single post-decoherence state, with the pattern determined by coupling geometry, and that this separation survives onto hardware.

1.d. Contributions

1. **Operational synthesis.** We show that globally preserved information, passively accessible information, and actively recoverable information can diverge qualitatively for a single post-decoherence state. These concepts are individually established; the contribution is the joint demonstration that they separate simultaneously in controlled models, with the pattern of separation controlled by coupling geometry.
2. **Geometry-dependent accessibility.** Ising and Heisenberg decoherence produce comparable scalar indicators (capture efficiency η , entanglement growth) and identical global recoverability, yet sharply different local capture and recovery curves. The distinguishing variable is coupling geometry—anisotropic vs. isotropic—not the degree of decoherence. This is the paper’s clearest empirical distinction.
3. **Hardware evidence.** On trapped-ion hardware (IonQ Forte), the geometry-dependent capture separation reappears ($\sim 10 \times$ in η), while on superconducting hardware (Rigetti Ankaa-3), it is suppressed by routing overhead. Inverse-specificity is robust on both platforms. The fact that local accessibility is more hardware-sensitive than global recoverability is itself an informative finding.
4. **Restricted complexity result.** In Appendix A, we establish that for random-circuit decoherence, exact recovery inherits the circuit complexity of the decohering dynamics.

1.e. Structure of the Paper

Section 2 reviews the physical background. Section 3 introduces the recoverability framework and formal definitions. Section 4 presents simulation and hardware results across four decoherence models on two hardware platforms. Section 5 discusses the results, limitations, and outlook. Section 6 concludes. Appendix A establishes a restricted complexity result for random-circuit decoherence. Appendix B discusses connections to objective collapse models.

2. PHYSICAL FOUNDATIONS

2.a. Unitarity and Information Conservation

Quantum mechanical time evolution is unitary: $|\psi(t_0)\rangle \mapsto |\psi(t)\rangle = U(t, t_0) |\psi(t_0)\rangle$ preserves the inner product and hence the total information content of the state. For our purposes, unitarity implies that the total quantum state $|\Psi\rangle$ contains the complete information of all its components at all times. The question is whether and how that information can be operationally recovered.

2.b. Decoherence and Einselection

Decoherence is the dynamical process by which a quantum system interacting with an environment loses phase coherence between certain states Schlosshauer (2019); Zurek (1981); Zurek (2003). The states that survive environmental monitoring—the *pointer states*—form the preferred basis via environment-induced superselection (einselection). Off-diagonal elements of the system’s reduced density matrix in the

pointer basis are suppressed on a timescale τ_D that depends on the coupling strength and environment size: $\tau_D \sim 10^{-20}$ s for macroscopic objects Schlosshauer (2019), microseconds to milliseconds for engineered quantum systems.

Decoherence does not destroy information. It redistributes it: coherence information is transferred to system-environment correlations. The total state $|\Psi_{SE}\rangle$ remains pure even as $\rho_S = \text{Tr}_E(|\Psi_{SE}\rangle\langle\Psi_{SE}|)$ becomes diagonal in the pointer basis Wallace (2012).

We adopt the einselection resolution for the preferred basis: the relevant decomposition is the pointer basis selected by the interaction Hamiltonian Zurek (2003). Our operational results—inverse-specificity, geometry-dependent recoverability—depend on the structure of the decoherence map, not on the ontological status of the components.

3. OPERATIONAL RECOVERABILITY UNDER DECOHERENCE

3.a. From Information Preservation to Operational Accessibility

That decoherence preserves global information does not imply that it remains operationally accessible. We separate three notions:

- **Presence:** information is encoded in the global quantum state.
- **Readability:** information can be passively extracted from an accessible subsystem or environmental fragment.
- **Recoverability:** information can be actively restored by an allowed class of operations.

A decohered state may preserve information globally while rendering it unreadable from any small subsystem. Conversely, pointer-state information may be redundantly exposed to many observers while coherent correlations remain inaccessible.

3.b. Coherence-Sensitive Recovery

Definition 1 (Coherence-sensitive recovery). A quantum operation \mathcal{E} performs *coherence-sensitive recovery* if its output depends on interference between components of a superposition in a way that cannot be reproduced by acting independently on decohered components and classically mixing the results. Let $\{|B_i\rangle\}$ denote the pointer decomposition. Then \mathcal{E} is coherence-sensitive if there exist output statistics from $\mathcal{E}(|\Psi\rangle)$ that cannot be reproduced by any classical mixture $\sum_i p_i \mathcal{E}_i(|B_i\rangle\langle B_i|)$.

Not every coherent operation is a recovery operation. Recovery requires that coherence be used to reconstruct information that no single decohered component carries alone—as in quantum computation, where algorithms produce outputs depending on interference between information distributed across decoherent components Deutsch (1985); Deutsch (1997).

3.c. The Recoverability Functional

Consider a system S initially correlated with a reference system R , and let decoherence couple S to an environment E . After decoherence, the joint state is ρ_{SER} . Suppose an experimenter has access only to a subsystem $A \subseteq SE$ and is restricted to an allowed operation class \mathcal{O} .

Definition 2 (Operational recoverability). The *recoverability* of the original system-reference correlation under operation class \mathcal{O} acting on accessible subsystem A is:

$$\mathcal{R}(\mathcal{O}, A) = \max_{\mathcal{E} \in \mathcal{O}} I(R : S)_{\mathcal{E}_A(\rho_{SER})} \quad (1)$$

where $I(R : S)$ is the mutual information between the reference and system after the recovery operation is applied. This quantity measures how much of the original correlation can be restored, given both *what degrees of freedom are accessible* and *what kinds of operations are physically allowed*.

Decoherence may leave the total information content of the global state unchanged while dramatically reducing $\mathcal{R}(\mathcal{O}, A)$ for realistic choices of A and \mathcal{O} . The central object of this paper is therefore not global information retention, but the constrained recoverability landscape induced by (\mathcal{O}, A) and the decoherence map.

3.d. Three Operational Accessibility Regimes

The behavior of $\mathcal{R}(\mathcal{O}, A)$ organizes naturally into three regimes.

3.d.i. Regime I: Coherent Access:

Coherence is preserved sufficiently well that interference can be directly exploited. \mathcal{R} remains high under modest coherent control. This is the regime exploited by quantum computation, bounded by the Holevo bound Wootters & Zurek (1982), the no-cloning theorem, and BQP Aaronson (2013).

3.d.ii. Regime II: Inverse-Specific Recoverability:

Information remains globally present but is no longer directly accessible from the system alone. The defining feature is **inverse-specificity**: recovery succeeds only when the applied operation is matched to the decoherence process. A generic circuit of comparable depth fails, even when substantial coherent control is available.

This separates two senses of inaccessibility: (1) **structural**—only the correct inverse works, and (2) **computational**—the correct inverse exists but is too costly to implement. For random circuit decoherence, these align: the required inverse acquires increasing circuit complexity with depth (Theorem 3, Appendix A). For structured physical Hamiltonians, inverse-specificity persists empirically, but asymptotic complexity remains open Haah et al. (2025).

Inverse-specificity (Claim B) is the robust result of this paper, demonstrated across four decoherence models. Computational hardness (Claim C) is proven for random circuits but open for physical Hamiltonians.

3.d.iii. Regime III: Classical Objectivity:

Environmental fragments redundantly encode pointer-state information Brandão et al. (2015); Wagner et al. (2021); Zurek (2009). Multiple independent observers access the same classical observable by probing different fragments. In terms of $\mathcal{R}(\mathcal{O}, A)$: high passive readability of classical pointer observables, negligible coherence-sensitive recoverability of quantum correlations. What remains available is a stable classical shadow of the full quantum state.

3.e. Geometry and Structure Matter

Recoverability is not determined solely by how much information has spread into the environment, but by **how it has spread**. Two decoherence processes can produce similar scalar indicators (entanglement growth, mutual-information decay, capture efficiency) while exhibiting sharply different recoverability profiles, because $\mathcal{R}(\mathcal{O}, A)$ depends on the structure of information flow, not only its magnitude. In the experimental section, anisotropic Ising-type couplings and isotropic Heisenberg couplings exhibit qualitatively different recoverability despite comparable scalar capture signatures.

3.f. Framework Summary

Regime	$\mathcal{R}(\mathcal{O}, A)$	Recovery structure	Example
I. Coherent Access	High under BQP operations	Polynomial-cost retrieval via interference	Quantum computer
II. Inverse-Specific	High only with correct inverse; ≈ 0 otherwise	Inverse-specific; complexity model-dependent	Mesoscopic system
III. Classical Objectivity	$O(1)$ for pointer data; ≈ 0 for quantum correlations	Only classical shadow readable	Macroscopic object

3.g. Formal Correspondence to Existing Frameworks

The three-notion distinction organizes existing quantum information concepts under a shared operational lens. We make the mapping to established formalism explicit.

This paper	Existing formalism	Key reference
Presence	Global purity of ρ_{SER} ; unitarity guarantees $S(\rho_{SER}) = 0$	Standard QM
Readability	Accessible information / Holevo quantity χ for measurements on subsystem A ; POVM-extractable correlations	Holevo Wootters & Zurek (1982); accessible info Zurek (2009)

This paper	Existing formalism	Key reference
Recoverability	Channel recovery fidelity; Petz recovery map \mathcal{R}_σ ; fidelity of recovery $F(\rho, \mathcal{R} \circ \mathcal{N}(\rho))$	Petz Petz (1986); universal recovery Junge et al. (2018); Wilde (2015)
Inverse-specificity	Structure-dependence of optimal recovery channel; failure of universal (channel-independent) recovery	Relates to conditions for exact Petz recovery
Capture efficiency η	Normalized area under partial-trace mutual information curve; measures redundancy of classical information across fragments	Quantum Darwinism redundancy Brandão et al. (2015); Zurek (2009)

The contribution is not the individual concepts—which are well-established—but the observation that **these quantities can take qualitatively different values simultaneously** for a single post-decoherence state, with the pattern determined by coupling geometry. The recoverability functional $\mathcal{R}(\mathcal{O}, A)$ is a constrained version of the channel recovery fidelity: when \mathcal{O} is the set of all quantum channels on A , it reduces to the standard Petz recovery fidelity. The regimes correspond to different qualitative behaviors as \mathcal{O} and A vary.

4. SIMULATION AND HARDWARE RESULTS

The experiments below are proof-of-principle probes of the recoverability hierarchy, not demonstrations of computational hardness. At 6–12 qubits, all systems are classically simulable. The results demonstrate *inverse-specific recovery structure* (Claim B) and geometry-dependent accessibility, not asymptotic complexity separation. We therefore interpret the experiments as probes of finite-size operational accessibility structure—what is locally readable, what is actively recoverable, and how those differ across models—rather than as evidence for asymptotic computational separation.

All experiments use a 1-qubit system S , a 1-qubit reference R (initialized in the Bell state $|\Phi^+\rangle_{SR}$), and k environment qubits E (initialized to $|0\rangle^{\otimes k}$). Mutual information $I(R : S)$ is estimated from correlators measured in all nine Pauli basis combinations ($\{X, Y, Z\}^{\otimes 2}$), with 4096 shots per circuit and 3–5 independent random circuit instances per depth.

4.a. Information Decay Under Decoherence




Figure 1: Mutual information $I(R : S)$ between reference and system qubits as a function of random circuit decoherence depth d , with $k = 4$ environment qubits. At $d = 0$, the Bell state yields $I(R : S) = 2$ bits. As depth increases, $I(R : S)$ decays toward zero, consistent with the decoupling theorem prediction (Theorem 1): for $d \geq \Omega(n + k)$, the reduced state of S is exponentially close to maximally mixed.

Figure 1 shows $I(R : S)$ decaying toward zero with circuit depth for $k = 4$ environment qubits, consistent with the decoupling theorem: information about the reference becomes inaccessible from the system alone as environmental entanglement accumulates.

4.b. Geometry-Dependent Partial Recoverability

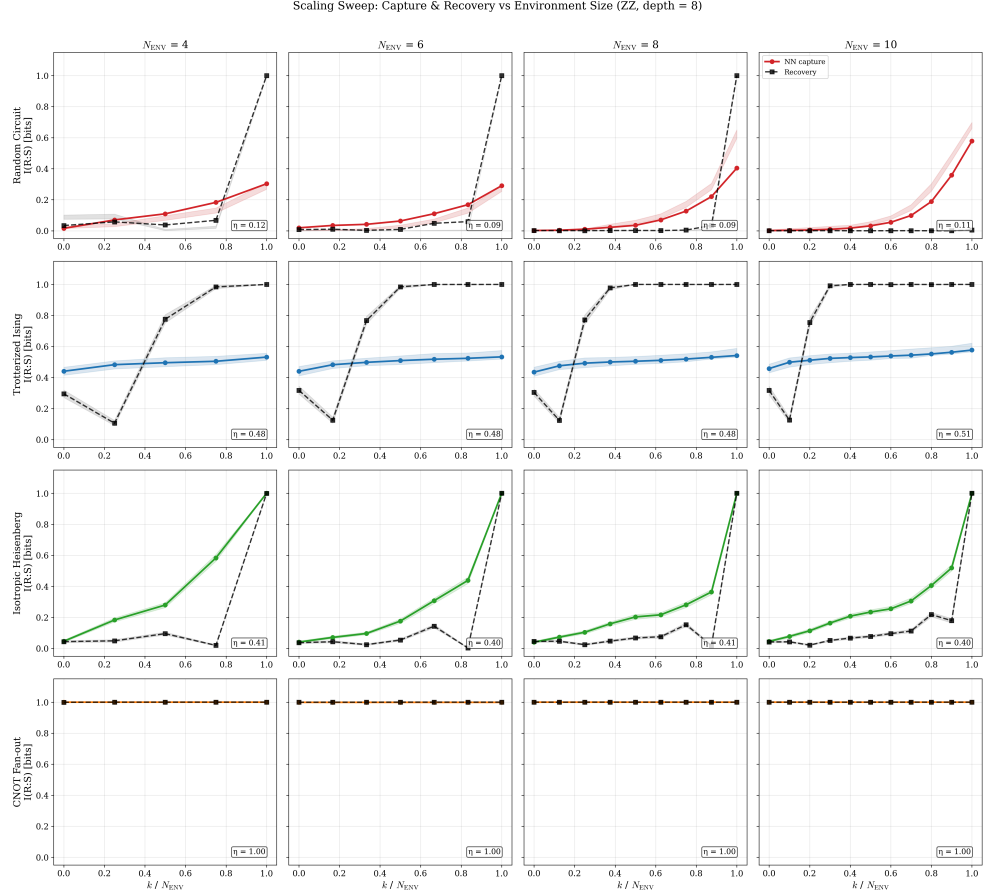


Figure 2: **Scaling sweep.** Nearest-neighbor capture curve (colored solid) and recovery curve (black dashed) for all four decoherence models at environment sizes $N_{\text{ENV}} = 4, 6, 8, 10$ (total qubits 6–12). Depth 8, ZZ basis, 4096 shots. Shaded bands show 95% bootstrap confidence intervals (200 resamples). The x -axis is normalized (k/N_{ENV}) to enable cross-size comparison. Capture efficiency η (inset) is computed from the all-subset curve. The Ising model shows recovery saturation at $k \approx 3$ regardless of environment size—consistent with nearest-neighbor locality. Heisenberg recovery remains flat as the environment grows. Random and Darwinism maintain their characteristic barrier signatures across all sizes.

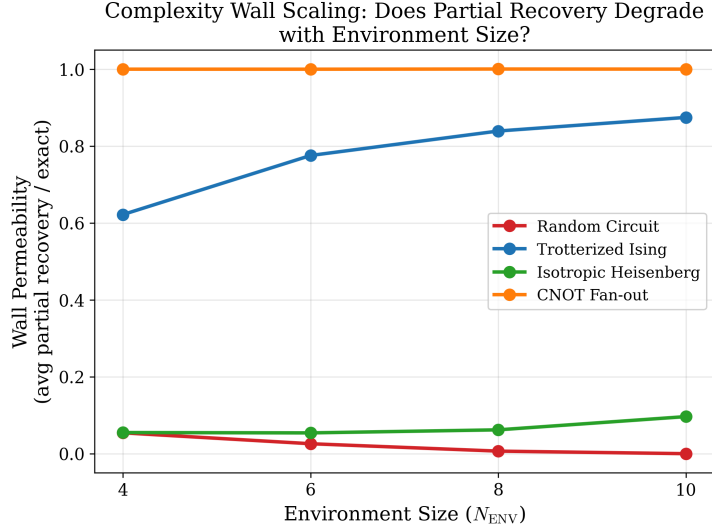


Figure 3: Partial recoverability vs. environment size. Average partial- k recovery ($k = 1, \dots, N_{ENV} - 1$) divided by exact recovery, for each model across $N_{ENV} = 4, 6, 8, 10$. Darwinism partial recoverability ≈ 1 (no access restriction). Ising partial recoverability *increases* with N_{ENV} because recovery saturates at $k \approx 3$, and additional environment qubits beyond the locality radius are irrelevant. Heisenberg and random partial recovery remains near zero across all tested sizes.

[Figure 2](#) and [Figure 3](#) test whether the barrier structure persists as the environment grows, across $N_{ENV} = 4, 6, 8, 10$ for all four models at depth 8. All values are robust to shot-noise fluctuations (95% bootstrap CIs, 200 resamples).

Three distinct scaling behaviors emerge. **Random circuits:** partial recovery at all $k < N_{ENV}$ approaches zero with increasing environment size, consistent with the $2^{(n-k)/2}$ suppression factor (Theorem 1). **Ising:** partial recoverability *increases* from 0.65 to 0.85 because recovery saturates at $k \approx 3$ nearest neighbors regardless of total environment size—the anisotropic $ZZ + X$ coupling creates a finite locality radius. η remains stable at ≈ 0.48 across all sizes. **Heisenberg:** partial recoverability remains near zero at all sizes (≤ 0.09) despite stable $\eta \approx 0.41$, confirming that isotropic coupling prevents local recovery even as the environment grows.

The most structurally informative contrast is the divergence between Ising and Heisenberg despite comparable scalar capture signatures. The barrier has model-dependent permeability determined by coupling geometry, not merely by how much decoherence has occurred. The Ising locality saturation is consistent with Lieb-Robinson bounds Chen et al. (2023): for short-range interactions, partial time-reversal on a region of radius $R \approx v_{LR}\tau$ recovers all locally accessible information regardless of total system size.

4.c. Inverse-Specific Recovery




Figure 4: Recovery of reference-system correlations under four strategies, comparing ideal simulator (*left*) and IBM Quantum hardware (*ibm_marrakesh*, 156-qubit Heron r2, *right*). *Exact inverse* (C^\dagger , green): perfect recovery on the simulator; degrading with depth on real hardware due to accumulated gate errors. *No recovery* (gray): information decays as in Experiment 1. *Random circuit* ($V \neq C^\dagger$, red): fails at all depths. *Partial inverse* (50% of layers, orange): intermediate recovery. The separation between exact and random recovery shows that recovery requires the *specific* inverse circuit (Claim B, inverse-specificity).

Figure 4 serves as a control establishing inverse-specificity (Claim B); the main scientific content of this paper comes from how partial accessibility differs across decoherence models (Sections 4.5–4.6). On the ideal simulator, exact inverse (C^\dagger) restores $I(R : S) = 2.0$ bits at all depths. On real hardware (IBM Quantum *ibm_marrakesh*, Heron r2), exact recovery degrades from ≈ 1.85 to ≈ 0.45 bits at depth 12 due to accumulated gate errors, but remains an order of magnitude above baseline. A random circuit of the same depth produces $I(R : S) \approx 0$ at all depths beyond $d = 0$. Partial inverse (50% of layers) yields intermediate recovery.

Recovery is strongly concentrated at the correct inverse. Circuits of similar depth and gate count but incorrect structure fail completely, producing $I(R : S) \approx 0$ rather than partial recovery. This is the operational content of inverse-specificity: recovery requires the *specific* inverse, not merely any operation of comparable complexity. The more informative question—addressed in the following subsections—is what remains recoverable when access is only partial or spatially local, since exact inversion is expected whenever the full decohering unitary is known and implementable.

4.d. Classical Objectivity and Quantum Darwinism




Figure 5: Quantum Darwinism signatures with $k = 6$ environment qubits. *Left*: Mean mutual information between system and individual environment fragments, measured in the Z-basis (classical pointer observable, blue) and X-basis (quantum coherence, red). Classical correlations increase and saturate; quantum correlations remain suppressed. *Center*: Pointer-state information per environment qubit as a function of depth, showing progressive broadcasting. *Right*: Redundancy—classical information accessible from each environment fragment at depth 16 saturates at ~ 1 bit, confirming redundant pointer-state broadcasting.

Figure 5 demonstrates Regime III emergence. Z-basis mutual information (classical pointer correlations) increases from 0 to ~ 1 bit and saturates; all six environment qubits independently carry the full classical information. X-basis mutual information (quantum coherence) remains near zero for all depths beyond $d = 1$. Redundancy saturates: each fragment independently encodes the same classical data, consistent with quantum Darwinism Brandão et al. (2015); Wang et al. (2025); Zurek (2009).

These results confirm the three regimes for random circuit decoherence. Do they persist for structurally different dynamics? Passive Darwinism signatures have been confirmed on hardware Wang et al. (2025), but passive accessibility (readability) is distinct from active recoverability. The following subsection addresses this gap across four qualitatively distinct models.

4.e. Capture Geometry and Recovery Structure Across Models

To test whether the hierarchy extends beyond random circuits, we compare four decoherence models with qualitatively different entanglement structures on a common 6-qubit architecture (1 reference, 1 system, 4 environment qubits, Bell state initialization):

- **Random circuit**: Haar-random single-qubit rotations and random CNOT pairings on SE (uniform scrambling, matching the Appendix A model).
- **Trotterized Ising**: Nearest-neighbor ZZ coupling ($J = 1.0$) and transverse field ($h = 0.5$) on a linear chain SE , with Trotter step $\Delta t = 0.3$ (anisotropic, near-integrable dynamics).
- **Isotropic Heisenberg**: $XX + YY + ZZ$ coupling ($J = 0.8$) on a nearest-neighbor chain with Trotter step $\Delta t = 0.3$ (isotropic scrambling, qualitatively distinct from the anisotropic Ising model).
- **CNOT fan-out**: System broadcasts its computational-basis state to environment qubits via sequential CNOTs (minimal model of quantum Darwinism).

4.e.i. Partial Capture Curves and Capture Efficiency:

For each model and depth, we measure all 6 qubits in either the ZZ basis (classical pointer-state correlations) or XX basis (quantum coherence correlations) and compute, via post-processing of marginals, the *partial capture curve*: $I(R : S + k \text{ env qubits})$ as a function of $k = 0, 1, 2, 3, 4$. We compute two variants: the *all-subset* capture curve (averaged over all binom $4k$ subsets) and the *nearest-neighbor* (NN) capture curve (using only the first k environment qubits in chain order). The

NN capture curve is directly comparable to the partial- k recovery curve, which also uses the first k nearest-neighbor qubits.

The capture curve reveals how decohered information is distributed across the environment. We define the *capture efficiency*:

$$\eta = \frac{\text{area under capture curve}}{\text{maximum possible area}} \quad (2)$$

A concave capture curve ($\eta \rightarrow 1$) indicates that information is accessible from small environmental fragments—characteristic of quantum Darwinism. A convex curve ($\eta \rightarrow 0$) indicates that information is locked in global correlations—characteristic of scrambling. Capture efficiency η is not intended as a universal monotone, but as a coarse operational summary of how quickly reference correlations become accessible from progressively larger environmental fragments. The main conclusions do not rest on η alone; they are supported by the full capture and recovery curve shapes, which carry strictly more information than the scalar summary.

Partial Capture Curves: Information Geometry of Decoherence

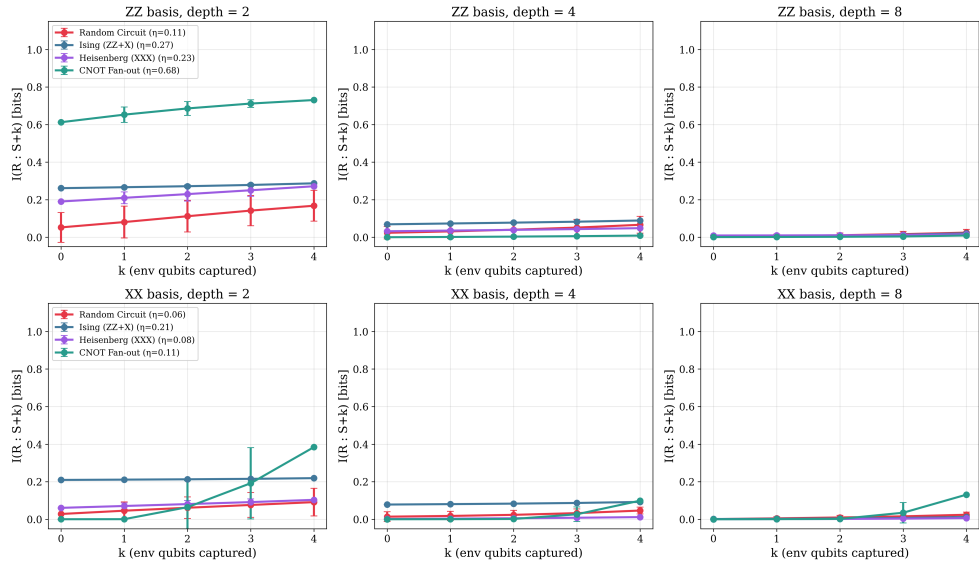


Figure 6: Partial capture curves for four decoherence models in the ZZ basis (top, classical information) and XX basis (bottom, quantum coherence). CNOT fan-out (green) shows $\eta = 1.0$ in ZZ : classical information is immediately accessible from any fragment. Random circuits (red) show $\eta \approx 0.11$: information requires nearly all environment qubits. Ising (blue, $\eta \approx 0.3$ – 0.6) and Heisenberg (purple, $\eta \approx 0.4$) are intermediate, with Ising showing stronger depth-dependence due to its anisotropic coupling. In the XX basis, all four models show low η : quantum coherence is redistributed into global correlations across all four tested models.

Figure 6 reveals sharply distinct information geometries. In the ZZ (classical) basis, the four models span from $\eta = 0.11$ (random circuits) to $\eta = 1.0$ (CNOT fan-out), with Ising ($\eta \approx 0.31$) and Heisenberg ($\eta \approx 0.44$) at intermediate values. Despite similar η , Ising and Heisenberg produce distinctly shaped capture curves. In the XX

(quantum) basis, all models show $\eta < 0.3$: quantum coherence information is redistributed into global correlations regardless of structure.

4.e.ii. Model-Dependent Recovery: Partial- k Sweep:

We apply model-appropriate recovery strategies, sweeping the spatial extent: a partial- k strategy applies time-reversal on the system plus the first k environment qubits in chain order. Gates coupling accessible to inaccessible qubits are omitted. This produces a *recovery curve*— $I(R : S)$ as a function of k —directly comparable to the NN capture curve from Phase 1.

Model	Exact recovery ($k = 4$)	Wrong recovery	Partial- k ($k = 0, 1, 2, 3$)
Random circuit	Full inverse C^\dagger	Different random circuit V	Partial C^\dagger on $S + k$ env qubits
Ising	Time reversal ($-H$)	Wrong- H ($J' = 0.3, h' = 1.5$)	Time reversal on $S + k$ env chain
Heisenberg	Time reversal ($-H$)	Wrong- J ($J' = 0.2$)	Time reversal on $S + k$ env chain
CNOT fan-out	CNOT uncopy	Random circuit V	Uncopy first k env qubits




Figure 7: Recovery of reference-system correlations for four decoherence models. Exact model-specific recovery (green) succeeds for all models. Wrong recovery (red) fails for all models. Partial inverse with $k = 2$ environment qubits (orange) reveals the key distinction: Ising retains $I(R : S) \approx 0.99$ at depth 4 (nearest-neighbor info recoverable from fragments); Heisenberg retains only ≈ 0.14 (isotropic coupling delocalizes info); CNOT fan-out retains 0.5 (must uncopy all coupled qubits); random circuits retain ≈ 0.10 .

Figure 7 shows the standard recovery comparison with a representative partial- k strategy ($k = 2$). The four models produce sharply different partial-recovery baselines, directly demonstrating that the recoverability barrier depends on the decoherence structure. But the full picture emerges when we compare the *complete* recovery curve (varying k from 0 to 4) against the capture curve.

4.e.iii. Central Result: Capture and Recovery Curves Jointly Characterize the Recoverability Barrier:




Figure 8: **Central result.** Nearest-neighbor (NN) capture curve (colored solid, from Phase 1) and all-subset capture curve (colored dotted) vs. recovery curve (black dashed, from Phase 2) for each model in the ZZ basis at depth 8. The NN capture curve measures $I(R : S + \text{first } k \text{ env})$; the recovery curve measures $I(R : S)$ after time-reversal on $S +$ the first k environment qubits. These use the same nearest-neighbor subset selection and are measured from *different circuits*. Inset: Pearson r between NN capture and recovery curves, with p -values from a permutation null model (1000 shuffles). Random and Heisenberg show strong shape agreement ($r > 0.84$); Ising shows partial agreement ($r = 0.70$); Darwinism is undefined (flat NN capture curve).




Figure 9: NN capture vs. recovery curves (ZZ basis) for all four models across three decoherence depths ($d = 2, 4, 8$). Shape similarity r is annotated per panel. Random and Heisenberg maintain high r across depths. Ising r varies with depth (0.58 at $d = 2$, 0.97 at $d = 4$, 0.79 at $d = 8$), reflecting the depth-dependent interplay between nearest-neighbor propagation speed and partial-inverse effectiveness. CNOT fan-out shows a characteristic mismatch: flat capture at 1.0 vs. step-function recovery.

Figure 8 and Figure 9 constitute the central result. For each model, the nearest-neighbor (NN) capture curve (Phase 1, no recovery) and recovery curve (Phase 2, model-specific partial inverse) use the same subset selection but are measured from different circuits, so any correspondence is empirical, not tautological. Shape

similarity is quantified by Pearson r against a permutation null model (1000 shuffles). The nearest-neighbor restriction is imposed for comparability between passive accessibility and partial time-reversal under local control, not because it maximizes agreement. Indeed, the models exhibit qualitatively different agreement patterns under this same protocol—including systematic mismatch in the Darwinism limit (r undefined, flat capture vs. step-function recovery)—which indicates that the protocol is not biased toward correspondence.

That the correct inverse outperforms the wrong one is unsurprising. What is nontrivial is that models with indistinguishable scalar signatures produce qualitatively different capture–recovery geometries—the scalar indicators alone do not predict the accessibility structure. The four models produce four qualitatively distinct capture–recovery signatures:

1. **Random circuit** ($r = 0.87$ -- 0.93): Both curves flat near zero for $k < 4$. Information is uniformly distributed; recovery jumps to ~ 1.0 only at $k = 4$ (full exact inverse). No partial subset suffices.
2. **Ising** ($r = 0.58$ -- 0.97): The recovery curve rises steeply, reaching $I(R : S) \approx 0.92$ at $k = 3$ (depth 8). The anisotropic ZZ coupling keeps information localized, making nearest-neighbor time-reversal more effective than passive capture suggests—a low, permeable barrier.
3. **Heisenberg** ($r = 0.76$ -- 0.97): Despite comparable all-subset η to Ising (0.41 vs. 0.46), both NN capture and recovery remain flat near zero for $k < 4$. Isotropic coupling distributes information into non-local correlations that neither passive capture nor partial time-reversal can access. **The Ising–Heisenberg contrast is the sharpest result: similar scalar capture efficiency, but radically different NN capture and recovery curves**, determined by coupling geometry.
4. **CNOT fan-out** (r undefined): NN capture is flat at $I \approx 1.0$ for $k \geq 1$ (Darwinism), but recovery remains at ≈ 0.5 until all coupled qubits are uncopied. This demonstrates that **readability \neq recoverability**: redundant copies are individually readable, but quantum recovery requires the full inverse.

Three implications follow. First, local recoverability is **geometrically determined**: coupling geometry, not just scalar η , determines the capture–recovery relationship. Second, quantum coherence inaccessibility is **consistent across all tested models**: in the XX basis, both capture and recovery are flat near zero for all four models—an empirical finding, not proven by Appendix A. Third, the hierarchy **persists with scaling**: the scaling sweep (Figure 3) shows these signatures persisting from $N_{\text{ENV}} = 4$ to 10, with Ising partial recoverability increasing (locality saturation at $k \approx 3$) and Heisenberg partial recovery remaining near zero at all tested sizes.

All simulator results use the Braket local state-vector simulator (4096 shots, 5 circuit instances, seed 42). Phase 1 and Phase 2 use identical per-circuit RNG seeding. Bootstrap confidence intervals (95%, 200 resamples) confirm robustness.

4.e.iv. *Hardware Results: Superconducting (Rigetti Ankaa-3) and Trapped-Ion (IonQ Forte):*

The role of hardware here is not to establish large-scale complexity separation, but to test which operational accessibility signatures survive realistic native-gate compilation and noise. We executed both Phase 1 (capture curves) and Phase 2 (partial- k recovery) on two hardware platforms with complementary error profiles: Rigetti Ankaa-3 (84-qubit superconducting, limited connectivity, accessed via Amazon Braket) and IonQ Forte 1 (36-qubit trapped-ion, all-to-all connectivity, accessed via Amazon Braket). The cross-platform comparison isolates which features are physically robust from those that are hardware-fragile.

Rigetti Ankaa-3 ($N_{\text{ENV}} = 4$, all four models). We executed 336 circuits at 4096 shots each on Ankaa-3. The native gate set does not include ZZ , XX , or YY gates; all two-qubit Hamiltonian evolution gates were decomposed into CNOT + single-qubit rotation sequences. At depth 0, all models and strategies yield $I(R : S) \approx 0.61$ – 0.64 , reflecting the hardware noise floor.

Inverse-specificity is confirmed on hardware for the three structured decoherence models:

Model	Strategy	$d = 2$	$d = 4$	$d = 8$
Ising	exact	0.640	0.620	0.606
	wrong	0.034	0.021	0.001
	none	0.256	0.064	0.000
Heisenberg	exact	0.673	0.651	0.642
	wrong	0.027	0.004	0.000
	none	0.144	0.019	0.005
Darwinism	exact	0.636	0.652	0.618
	wrong	0.066	0.000	0.001
	none	0.611	0.000	0.000

Exact recovery maintains $I(R : S) \approx 0.62$ – 0.67 at all depths for all three models, while wrong recovery decays to < 0.03 by $d = 4$. The separation between exact and wrong recovery exceeds $20 \times$ at all depths, confirming that recovery requires the *specific* inverse operation.

The partial- k recovery sweep reveals qualitative differences between coupling geometries:

Model ($d = 4$)	partial-0	partial-1	partial-2	partial-3	exact
Ising	0.156	0.181	0.105	0.031	0.620
Heisenberg	0.038	0.007	0.000	0.002	0.651
Darwinism	0.000	0.487	0.617	0.569	0.652

Darwinism partial recovery increases monotonically with k : each environment qubit holds an independent copy. Heisenberg partial recovery remains near zero at all $k < 4$, confirming that isotropic coupling distributes information non-locally.

Rigetti Ankaa-3 ($N_{\text{ENV}} = 6$, Ising scaling). At the larger environment size, Rigetti’s routing overhead substantially degraded the signal. The nearest-neighbor capture curves for both Ising and Heisenberg collapsed to near the noise floor ($\eta \approx 0.009$ for Ising, 0.006 for Heisenberg, compared to simulator baselines of

0.461 and 0.405 respectively). However, the Ising *exact recovery* remained strong: $I(R : S) = 0.653$, with wrong (0.000) and partial- k recovery (< 0.007 for all $k < 6$) near zero. This confirms that inverse-specificity persists at larger sizes on superconducting hardware even when the local capture signal is lost to compilation noise.

IonQ Forte 1 ($N_{\text{ENV}} = 6$, Ising vs. Heisenberg diagnostic). To determine whether Rigetti’s null capture result reflects a genuine absence of the geometry signal or platform-specific routing artifacts, we executed a targeted diagnostic on IonQ Forte 1 trapped-ion hardware (512 shots, $N_{\text{ENV}} = 6$, depth 8). The all-to-all connectivity of the trapped-ion architecture eliminates SWAP routing entirely, providing a cleaner test of whether the coupling geometry contrast survives hardware.

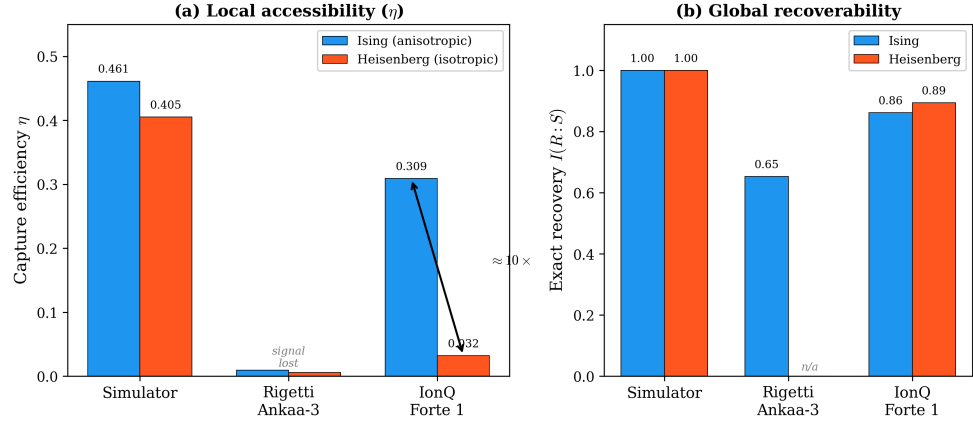


Figure 10: **Cross-platform comparison** at $N_{\text{ENV}} = 6$. **(a)** Local accessibility (η): Ising and Heisenberg show comparable η in simulation, but on IonQ the geometry-dependent separation emerges ($\approx 10 \times$). On Rigetti, both collapse to the noise floor. **(b)** Global recoverability: exact inverse recovery remains strong on all platforms for both models, confirming that global information is preserved regardless of local accessibility.

Platform	Model	η	Exact $I(R : S)$	Wrong $I(R : S)$
Simulator	Ising	0.461	1.000	0.001
	Heisenberg	0.405	1.000	0.000
Rigetti	Ising	0.009	0.653	0.000
	Heisenberg	0.006	—	—
IonQ	Ising	0.309	0.861	0.032
	Heisenberg	0.032	0.894	—

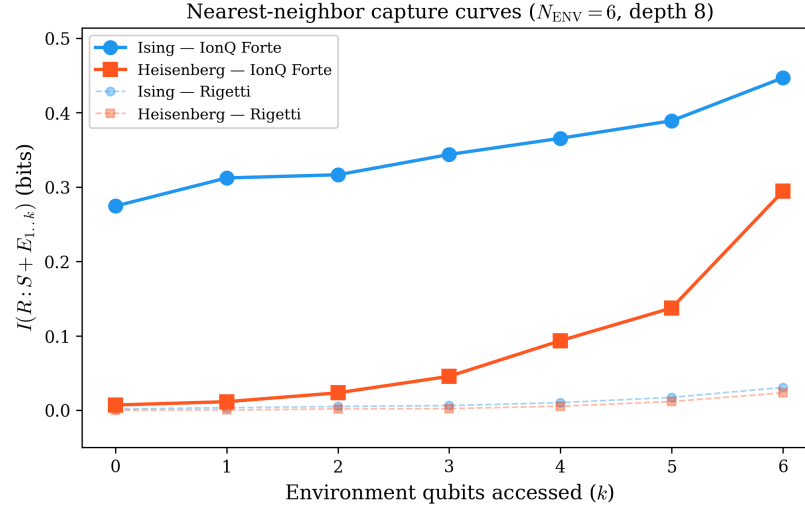


Figure 11: **Nearest-neighbor capture curves** on IonQ Forte and Rigetti Ankaa-3 at $N_{\text{ENV}} = 6$, depth 8. On IonQ (solid lines), Ising information rises steeply with nearest-neighbor access while Heisenberg remains low—a clear geometry-dependent separation. On Rigetti (dashed), both models are suppressed to the noise floor by routing overhead.

The IonQ results recover the geometry-dependent capture signal that Rigetti suppressed:

1. **Capture efficiency separation.** On IonQ, Ising capture efficiency ($\eta = 0.309$) exceeds Heisenberg ($\eta = 0.032$) by approximately $10 \times$. On Rigetti, both collapsed to < 0.01 . The simulator baseline shows $\eta_{\text{Ising}} = 0.461$ and $\eta_{\text{Heisenberg}} = 0.405$ —similar magnitudes but qualitatively different NN capture curve shapes. IonQ preserves the shape distinction: Ising information is locally capturable from nearest-neighbor fragments, while Heisenberg information is not.
2. **Exact recovery is model-independent.** Both Ising ($I(R : S) = 0.861$) and Heisenberg (0.894) achieve near-perfect exact recovery on IonQ, confirming that global information is present in both cases. The geometry contrast is not in *whether* information survives, but in *how it is distributed*.
3. **Presence \neq readability \neq recoverability.** Heisenberg decoherence produces a state where information is globally present (exact recovery succeeds strongly), but locally inaccessible ($\eta = 0.032$ —little information is capturable from nearest-neighbor fragments). Ising decoherence produces a state where information is both globally recoverable *and* locally capturable ($\eta = 0.309$). This is direct hardware evidence that the three notions in the paper’s conceptual hierarchy—presence, readability, and recoverability—are operationally distinct.

Cross-platform interpretation. The comparison across platforms yields a coherent picture. Inverse-specificity (Claim B) is robust: exact recovery succeeds and wrong recovery fails on both superconducting and trapped-ion hardware. The geometry-dependent capture structure is physically real but hardware-fragile: it is suppressed to the noise floor on Rigetti Ankaa-3, where SWAP routing introduces $O(n)$ additional two-qubit gates, but recovered on IonQ Forte, where all-to-all connectivity preserves the circuit structure as designed. The fact that the capture signal is *platform-dependent* while the inverse-specificity signal is *platform-independent* is itself an informative result: it suggests that local accessibility is more sensitive to compilation overhead than global recoverability, consistent with the expectation that local observables are lower-order correlations more easily disrupted by noise.

5. DISCUSSION

The central finding is not that the correct inverse outperforms the wrong one—that is expected. The nontrivial result is the **joint pattern**: decoherence processes with comparable scalar signatures (capture efficiency, entanglement growth, mutual-information decay) and identical global recoverability produce sharply different local accessibility structures, determined by coupling geometry rather than by the degree of decoherence. This pattern is not a relabeling of known facts; it is an empirical demonstration that scalar decoherence indicators systematically underdetermine the operational accessibility of the resulting state. On hardware, inverse-specificity is robust across platforms, while the geometry-dependent capture signal is platform-sensitive—

suppressed on Rigetti but recovered on IonQ—indicating that local accessibility is more fragile than global recoverability under compilation overhead.

5.a. *What This Work Establishes*

(i) Scalar decoherence indicators underdetermine accessibility. Ising and Heisenberg decoherence produce comparable capture efficiencies and identical global recoverability, yet sharply different local capture and recovery curves. The distinguishing variable is coupling geometry, not the degree of decoherence.

(ii) Presence, readability, and recoverability separate simultaneously. On IonQ hardware, Heisenberg decoherence produces a state where information is globally present (exact recovery $I(R : S) = 0.894$), locally inaccessible ($\eta = 0.032$), and recoverable only with the full inverse. Ising decoherence on the same hardware is both globally recoverable and locally capturable ($\eta = 0.309$). The three notions take different values for different coupling geometries.

(iii) Local accessibility is more hardware-fragile than global recoverability. The geometry-dependent capture signal survives on trapped-ion hardware but is suppressed to the noise floor on superconducting hardware by routing overhead, while inverse-specificity is robust on both platforms.

5.b. *Limitations*

This work does **not** establish that realistic macroscopic decoherence involves an asymptotic computational hardness barrier. Appendix A does not establish the general physical claim, but provides a rigorous anchor showing how inverse-specificity can coincide with growing inverse complexity in a canonical scrambling model. That restricted result does not extend directly to physical Hamiltonians; the Haah-Liu-Tan no-go theorem Haah et al. (2025) constrains direct extension.

The hardware results are in the **small-scale NISQ regime** ($N_{\text{ENV}} \leq 6$, classically simulable). They demonstrate qualitative features of the hierarchy but cannot probe the computationally non-trivial regime. Scaling to ≥ 20 qubits on lower-error hardware would be needed.

The reported signatures are robust across the tested depth and environment-size sweeps (bootstrap 95% CIs, 200 resamples), but should be interpreted as finite-size operational phenomena rather than fully universal asymptotic diagnostics.

The approximate recovery literature Junge et al. (2018); Petz (1986); Wilde (2015) shows that partial recovery is sometimes possible without the full inverse. Characterizing recoverable information at polynomial cost as a function of decoherence depth is an important open direction.

5.c. *Relation to Existing Work*

Relative to **Quantum Darwinism** Brandão et al. (2015); Zurek (2009): the CNOT fan-out model serves as an operational endpoint where readability is high but coherence-sensitive recoverability is absent. It demonstrates that redundant classical readability can coexist with restricted quantum recoverability—“objectivity” and “recoverability” are distinct operational properties. This model also serves a methodological role: its systematic mismatch between capture and recovery curves (r undefined) shows that the NN comparison protocol does not inherently favor correspondence.

Relative to **Hayden-Preskill** and **Harlow-Hayden** Harlow & Hayden (2013): the operational pattern (information present but inaccessible without a structured inverse) is shared, but extending those results from scrambling/black-hole settings to ordinary physical decoherence requires additional work. The structural parallel is suggestive, not proven.

5.d. *Interpretational Scope*

The operational results depend on the structure of the decoherence map, not on the ontological status of decoherent components. All interpretations preserving unitary evolution make identical predictions for the observables measured here.

5.e. *Outlook*

(i) Larger-scale hardware. Probe these observables at ≥ 20 qubits on architectures with lower routing overhead.

(ii) **Physical Hamiltonians.** Whether the complexity barrier extends beyond random circuits remains open; state-level pseudorandomness Cotler et al. (2023) and complexity-constrained thermodynamics Faist et al. (2025) may provide a route.

(iii) **Recovery basin analysis.** Systematically perturbing the inverse circuit (gate corruption, coupling mismatch, gate reordering) would convert the current qualitative inverse-specificity into a measurable geometric property of the recovery landscape.

(iv) **Resource-theoretic formulation.** Observer constraints, subsystem access, and recovery classes naturally define partial orders over accessible information.

6. CONCLUSION

Different decoherence mechanisms with comparable scalar signatures can produce sharply different operational accessibility structures. Ising and Heisenberg dynamics yield similar capture efficiencies and identical global recoverability, yet their local capture and recovery curves diverge qualitatively—a distinction determined by coupling geometry, confirmed in simulation across environment sizes $N_{\text{ENV}} = 4\text{--}10$, and recovered on trapped-ion hardware with $\sim 10 \times$ separation in capture efficiency.

This geometry dependence is the paper’s central empirical contribution. It demonstrates that the operational accessibility of decohered information is underdetermined by scalar decoherence indicators and requires the finer-grained characterization provided by capture–recovery curve comparison. The computational cost of recovery is a separate question: provably increasing for random circuits (Appendix A), but open for physical Hamiltonians Haah et al. (2025). The main conclusion is accordingly operational rather than asymptotic: decohered information may remain globally present while differing sharply in local readability and constrained recoverability, and that separation is shaped not only by how much information has spread, but by *how* it has spread.

7. APPENDIX A: RESTRICTED COMPLEXITY RESULT FOR RANDOM CIRCUIT DECOHERENCE

This appendix establishes a restricted complexity result for a specific model of decoherence in which the system-environment interaction is represented by a sufficiently deep random local quantum circuit. The goal is not to prove that all physically realistic decoherence processes generate computationally inaccessible information. Rather, the goal is to show that in a well-defined random-circuit setting: (1) information initially localized in the system becomes inaccessible from the system alone, (2) recovery requires the specific inverse of the decohering circuit, and (3) that inverse acquires increasing circuit complexity with circuit depth. This provides a controlled model in which the operational phenomenon described in the main text—inverse-specific recoverability—admits a formal complexity-theoretic realization. The result should be interpreted as a **restricted existence proof**, not as a theorem about generic decoherence in nature.

7.a. A.1 Setup

Consider a system S of n qubits and an environment E of k qubits. The system is initially entangled with a reference R (also n qubits) in the maximally entangled state:

$$|\Phi\rangle_{SR} = \frac{1}{\sqrt{2^n}} \sum_{i=1}^{2^n} |i\rangle_S |i\rangle_R \quad (3)$$

The environment begins in a product state $|0\rangle_E^{\otimes k}$. Decoherence is modeled by a random quantum circuit C of depth d acting on the joint system SE , composed of 2-local gates drawn independently from the Haar measure on $U(4)$. The reference R is untouched. After decoherence, the global state is:

$$|\Psi\rangle_{SER} = (C_{SE} \otimes I_R) |\Phi\rangle_{SR} |0\rangle_E \quad (4)$$

The entanglement between S and R represents the reference-system correlations that decoherence redistributes. Recovery means restoring $I(R : S)$ from the post-decoherence state.

7.b. A.2 Inaccessibility Without Environment (Decoupling)

Theorem 1 (Information loss from reduced system). After a random circuit of depth $d \geq \Omega(n + k)$ on SE , the reduced state of S is exponentially close to the maximally mixed state:

$$\left| \rho_S - \frac{I}{2^n} \right|_1 \leq 2^{-(k-n)/2+\delta} \quad (5)$$

for small $\delta > 0$, with high probability over the choice of circuit C .

Proof. This follows from the decoupling theorem Brandão et al. (2016). After $d = \Omega(n+k)$ layers of random 2-local gates, the circuit C_{SE} forms an approximate unitary 2-design on the 2^{n+k} -dimensional Hilbert space of SE . For a Haar-random unitary U on SE , the expected trace distance between the reduced state $\rho_S = \text{Tr}_E(U |\psi\rangle\langle\psi| U^\dagger)$ and the maximally mixed state satisfies:

$$\mathbb{E}_U \left| \rho_S - \frac{I}{2^n} \right|_1 \leq \sqrt{\frac{2^{2n}}{2^{n+k}}} = 2^{(n-k)/2} \quad (6)$$

For $k > n$ (environment larger than system, which is the physically relevant regime), this is exponentially small. The 2-design property of random circuits at depth $\Omega(n+k)$ ensures the same bound holds for random circuits, not just Haar-random unitaries. \square

Corollary 1. The mutual information between S and R after decoherence satisfies:

$$I(S : R) \leq 2n \cdot 2^{-(k-n)/2+\delta} \quad (7)$$

By the Holevo bound, no measurement on S alone can extract more than $I(S : R)$ bits of information about R . For $k \gg n$, this is negligible. **Recovery from the system alone is information-theoretically impossible.**

7.c. A.3 Complexity of Recovery With Environment Access

An agent with joint access to S and E can in principle recover the reference-system correlations by applying C_{SE}^\dagger , undoing the decoherence. We now show this recovery has high circuit complexity.

Theorem 2 (Complexity of recovery). Let C be a random quantum circuit of depth d on $n+k$ qubits. With probability $\geq 1 - 2^{-\Omega(n+k)}$ over the choice of C , any quantum circuit implementing C^\dagger (up to error $\varepsilon < 1/4$ in diamond norm) requires circuit size at least $\Omega(d)$.

Proof. Haferkamp et al. Haferkamp et al. (2022) proved that for random quantum circuits on m qubits, the circuit complexity $\mathcal{C}(C)$ satisfies:

$$\Pr[\mathcal{C}(C) \geq \Omega(d)] \geq 1 - 2^{-\Omega(m)} \quad (8)$$

for all $d \leq 2^{\Omega(m)}$, where $m = n+k$ is the total number of qubits. Since $\mathcal{C}(C^\dagger) = \mathcal{C}(C)$ (reversing the gate sequence), the recovery operation C^\dagger has circuit complexity $\Omega(d)$ with overwhelming probability. \square

Remark. Theorem 2 establishes a lower bound on the complexity of implementing C^\dagger specifically. It does not rule out the possibility that some other unitary $V \neq C^\dagger$ could achieve the same recovery task (restoring the SR correlations) with lower complexity. Proving a *task-complexity* lower bound—that no recovery strategy, regardless of method, can succeed with sub- $\Omega(d)$ complexity—remains open. However, for random circuits, the decoupling result (Theorem 1) shows that without environment access, no strategy works at all; and with environment access, the circuit C^\dagger is the natural recovery channel. Whether alternative recovery strategies can circumvent the complexity of C^\dagger for random circuits is a question closely related to the general hardness of unitary synthesis.

7.d. A.4 Complexity Scaling

Combining Theorems 1 and 2 yields the complexity barrier for random circuit decoherence:

Theorem 3 (Complexity barrier for random circuit decoherence). For random circuit decoherence with depth d on $n+k$ qubits ($k > n$):

1. **No access to E :** Recovery is information-theoretically impossible for $d \geq \Omega(n+k)$.
2. **With access to E :** Recovery requires circuit complexity $\Omega(d)$.
3. **Saturation:** For $d \geq 2^{\Omega(n+k)}$, the complexity saturates at $\Omega(2^{n+k}/(n+k))$, which is exponential in the total number of qubits.

In the physically relevant regime where $k \gg n$, this gives:

- For **shallow decoherence** ($d = \text{poly}(k)$): recovery complexity is $\Omega(\text{poly}(k))$ —superpolynomial but not exponential.

- For **deep decoherence** ($d \geq 2^{\Omega(k)}$): recovery complexity is $\Omega(2^k/k)$ —exponential in the number of environmental degrees of freedom.
- For **macroscopic decoherence** ($k \sim 10^{23}$): even at polynomial depth, the complexity $\Omega(\text{poly}(10^{23}))$ is far beyond any conceivable computational resource.

7.e. A.5 Limitations of the Random-Circuit Result

The preceding theorems do **not** establish that generic physical decoherence processes produce computationally inaccessible information in full generality. In particular, they do not show that:

- arbitrary local Hamiltonian dynamics generate pseudorandom unitaries,
- physically realistic environments generically realize the same complexity growth as random circuit ensembles,
- or exact recovery of naturally decohered macroscopic systems is generically exponential-time hard under physically realistic locality constraints.

The random-circuit model should be interpreted as a **controlled model of structured recoverability loss**, not as a universal theorem about decoherence in nature. Two specific gaps remain:

1. **Locality:** Physical Hamiltonians are constant-local (2-body interactions), and Haah, Liu, and Tan Haah et al. (2025) showed that constant-local Hamiltonian evolution cannot produce pseudorandom unitaries. Our random circuit model uses Haar-random 2-local gates, which is a stronger form of randomness than physical dynamics provides.
2. **Conservation laws:** Physical interactions obey conservation laws (energy, particle number) that restrict the set of achievable unitaries Marvian (2022). Random circuits with symmetry constraints produce designs of lower order and converge more slowly.

7.f. A.6 Obstruction to Direct Extension

A major obstacle to extending the above result to realistic physical Hamiltonians is the Haah-Liu-Tan no-go theorem Haah et al. (2025), which shows that constant-local Hamiltonian evolution does not generate pseudorandom unitaries or even approximate 2-designs in the sense required by random-circuit hardness arguments. This sharply limits any direct inference from random-circuit complexity growth to generic physical decoherence.

Two possible extension routes remain:

1. **Unitary-level extension:** physically realistic dynamics generate sufficiently complex effective unitaries through many-body interactions at polylog locality, or
2. **State-level extension:** even if full unitary pseudorandomness fails, the resulting state structure may suffice to produce strong recoverability barriers. Cotler et al. Cotler et al. (2023) showed that eigenstates of strongly interacting Hamiltonians produce approximate k -designs for reduced density matrices, providing partial evidence for this route.

This paper does not resolve this issue. It identifies it as a central open problem.

7.g. A.7 Corollary (Restricted Operational Interpretation)

Within the random-circuit model, decoherence can produce a regime in which information remains globally present but operationally inaccessible except under the correct high-structure recovery operation. This provides a controlled formal example of the main thesis of the paper: **presence does not imply recoverability**.

8. APPENDIX B: CONNECTIONS TO OBJECTIVE COLLAPSE MODELS

The recoverability framework is compatible with standard unitary quantum mechanics and could, in principle, inform experimental programs that distinguish unitary decoherence from objective collapse models. Under unitary QM, decoherence suppresses local coherence but preserves information globally. In objective collapse models (GRW Bassi et al. (2013), CSL), the wave function undergoes genuine, irreversible localization events that destroy superposition components. This distinction is testable at the boundary between Regimes I and II.

Residual coherence prediction. For a mechanical oscillator of mass m , unitary QM predicts that off-diagonal density matrix elements decay as $|\rho_{ij}(t)| = |\rho_{ij}(0)| \exp(-\Gamma_{\text{dec}} \cdot t)$, where Γ_{dec} depends only on known environmental sources. CSL predicts an additional contribution $\Gamma_{\text{CSL}} = \lambda(m/m_0)^2 [1 -$

$\exp(-\Delta x^2/4r_C^2)]$ that saturates for large spatial separations Bassi et al. (2013)—a distinctive signature. Current experiments have excluded CSL with $\lambda \geq 10^{-11}$ Hz Donadi et al. (2021); reaching the original GRW parameter space ($\lambda \approx 10^{-16}$ Hz) requires extending coherence measurements to the 10^6 – 10^9 amu range Fein et al. (2019).

Quantum error correction prediction. Under unitary QM, the overhead for fault-tolerant QEC scales with noise complexity but encounters no fundamental ceiling. Objective collapse models predict that QEC will fail at some mass/energy scale because collapse cannot be corrected. Each successful demonstration of fault-tolerant QEC at larger scale constitutes incremental evidence for unitary QM.

These predictions are shared with all unitary interpretations; the recoverability framework provides conceptual context, not empirically distinctive consequences beyond standard QM.

REFERENCES

- Aaronson, S. (2013). *Quantum Computing Since Democritus*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511979309>
- Bassi, A., Lochan, K., Satin, S., Singh, T. P., & Ulbricht, H. (2013). Models of wave-function collapse, underlying theories, and experimental tests. *Reviews of Modern Physics*, 85(2), 471–527. <https://doi.org/10.1103/RevModPhys.85.471>
- Brandão, F. G. S. L., Harrow, A. W., & Horodecki, M. (2016). Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346(2), 397–434. <https://doi.org/10.1007/s00220-016-2706-8>
- Brandão, F. G. S. L., Piani, M., & Horodecki, P. (2015). Generic emergence of classical features in quantum Darwinism. *Nature Communications*, 6, 7908. <https://doi.org/10.1038/ncomms8908>
- Chen, C.-F., Lucas, A., & Yin, C. (2023). Speed limits and locality in many-body quantum dynamics. *Reports on Progress in Physics*, 86, 116001. <https://doi.org/10.1088/1361-6633/acfaae>
- Cotler, J., Mark, D. K., Huang, H.-Y., Choi, S., & others. (2023). Emergent quantum state designs from individual many-body wave functions. *PRX Quantum*, 4(1), 10311. <https://doi.org/10.1103/PRXQuantum.4.010311>
- Deutsch, D. (1985). Quantum theory, the Church–Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London a*, 400(1818), 97–117. <https://doi.org/10.1098/rspa.1985.0070>
- Deutsch, D. (1997). *The Fabric of Reality*. Allen Lane.
- Donadi, S., Piscicchia, K., Curceanu, C., Diósi, L., Laubenstein, M., & Bassi, A. (2021). Underground test of gravity-related wave function collapse. *Nature Physics*, 17(1), 74–78. <https://doi.org/10.1038/s41567-020-01089-5>
- Faist, P., Younger Halpern, N., Eisert, J., Kothakonda, N. B. T., Haferkamp, J., & Munson, A. (2025). Complexity-constrained quantum thermodynamics. *PRX Quantum*, 6, 10346. <https://doi.org/10.1103/PRXQuantum.6.010346>
- Fein, Y. Y., Geyer, P., Zwick, P., Kialka, F., Pedalino, S., Mayor, M., Gerlich, S., & Arndt, M. (2019). Quantum superposition of molecules beyond 25 kDa. *Nature Physics*, 15(12), 1242–1245. <https://doi.org/10.1038/s41567-019-0663-9>
- Haah, J., Liu, Y., & Tan, X. (2025). Random unitaries from Hamiltonian dynamics. *Arxiv Preprint Arxiv:2510.08434*.
- Haferkamp, J., Faist, P., Kothakonda, N. B. T., Eisert, J., & Younger Halpern, N. (2022). Linear growth of quantum circuit complexity. *Nature Physics*, 18, 528–532. <https://doi.org/10.1038/s41567-022-01539-6>
- Harlow, D., & Hayden, P. (2013). Quantum computation vs. firewalls. *Journal of High Energy Physics*, 2013(6), 85. [https://doi.org/10.1007/JHEP06\(2013\)085](https://doi.org/10.1007/JHEP06(2013)085)
- Junge, M., Renner, R., Sutter, D., Wilde, M. M., & Winter, A. (2018). Universal Recovery Maps and Approximate Sufficiency of Quantum Relative Entropy. *Annales Henri Poincaré*, 19(10), 2955–2978. <https://doi.org/10.1007/s00023-018-0716-0>
- Marvian, I. (2022). Restrictions on realizable unitary operations imposed by symmetry and locality. *Nature Physics*, 18, 283–289. <https://doi.org/10.1038/s41567-021-01464-0>
- Petz, D. (1986). Sufficient subalgebras and the relative entropy of states of a von Neumann algebra. *Communications in Mathematical Physics*, 105(1), 123–131. <https://doi.org/10.1007/BF01212345>
- Schlosshauer, M. (2019). Quantum decoherence. *Physics Reports*, 831, 1–57. <https://doi.org/10.1016/j.physrep.2019.10.001>

- Wagner, R., Baldijão, R. D., Duarte, C., Amaral, B., & Cunha, M. T. (2021). Emergence of noncontextuality under quantum Darwinism. *PRX Quantum*, 2(3), 30351. <https://doi.org/10.1103/PRXQuantum.2.030351>
- Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199546961.001.0001>
- Wang, Z., Touil, A., Song, C., Li, H., Guo, Q., Bao, Z., Zhang, P., Wang, H., Zhu, Z., Mondaini, R., Song, Z., & Salice, K. (2025). Observation of quantum Darwinism and the origin of classicality with superconducting circuits. *Science Advances*, 11. <https://doi.org/10.1126/sciadv.adx6857>
- Wilde, M. M. (2015). Recoverability in quantum information theory. *Proceedings of the Royal Society a*, 471(2182), 20150338. <https://doi.org/10.1098/rspa.2015.0338>
- Wootters, W. K., & Zurek, W. H. (1982). A single quantum cannot be cloned. *Nature*, 299(5886), 802–803. <https://doi.org/10.1038/299802a0>
- Zurek, W. H. (1981). Pointer basis of quantum apparatus: Into what mixture does the wave packet collapse?. *Physical Review D*, 24(6), 1516–1525. <https://doi.org/10.1103/PhysRevD.24.1516>
- Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75(3), 715–775. <https://doi.org/10.1103/RevModPhys.75.715>
- Zurek, W. H. (2009). Quantum Darwinism. *Nature Physics*, 5(3), 181–188. <https://doi.org/10.1038/nphys1202>