

Quantum Mechanics as Information Preservation

Channels, Recoverability, and Higher Descent

Anonymous author

20 October 2026

Abstract

Quantum mechanics admits an operational interpretation as a theory of information-bearing systems and their permitted transformations. Its storage interpretation becomes precise only after distinguishing physical states, descriptions of states, and recoverable information. We define a memory by an encoding, a physical channel, and a recovery channel. Exact storage is a retract equation; approximate storage is controlled by diamond distance. Observable algebras specify which classical, quantum, or hybrid information must survive. We construct a stack of memory implementations in an infinity-topos and prove descent under changes of frame. Its unitary exact-memory sector is a classifying stack of projective unitary frames, yielding perfectly preserving families without global frames. We distinguish this obstruction from measurement contextuality and prove that arbitrary continuous deformation cannot serve as operational equality. The contribution is a synthesis and an explicit semantics of memory families, not a reconstruction of quantum mechanics from storage or a claim of unlimited readable classical information.

1. Scope and the informational thesis

The thesis studied here is: **quantum mechanics constrains the preparation, preservation, transformation, and retrieval of information carried by physical systems.** Storage is the preservation of a specified logical system or observable algebra through an intervening process. A storage claim must state what information, which process, which recovery operations, what duration, and what error tolerance are intended.

Spatial size and Hilbert-space dimension are different parameters. No spatial-size variable occurs in the definitions below. Deriving a storage-density bound requires a physical model connecting dimension, energy, locality, and matter. Mathematical abstraction makes these distinctions explicit; small physical objects are not an axiom of the categorical semantics.

Categorical quantum mechanics supplies structural semantics of quantum protocols [1]. Completely positive maps [2] and categories incorporating classical and quantum channels [3] provide related foundations. We use their operational setting, not an identification of the channel category with a topos. The normalized trace-preserving subcategory need not inherit the dagger compact structure of a larger category of completely positive maps.

Classical probability also supports storage. Therefore storage alone does not select complex quantum mechanics. Informational reconstructions require further principles [13]. Here finite-dimensional

channel theory and the Born rule are assumptions. Known quantum-information results are distinguished from the elementary proofs and memory-stack construction supplied below.

2. Registers, channels, and exact storage

Write $M_d = \text{End}_{\mathbb{C}}(\mathbb{C}^d)$. A state is $\rho \in M_d$ with $\rho \geq 0$ and $\text{Tr } \rho = 1$. A channel $\Phi : M_d \rightarrow M_e$ is a completely positive, trace-preserving complex-linear map. Complete positivity requires $\Phi \otimes \text{id}_{M_r}$ to preserve positivity for every r . Channels compose and tensor to form a symmetric monoidal category QChan .

A measurement has positive operators F_y with $\sum_y F_y = I$, giving $p(y \mid \rho) = \text{Tr}(F_y \rho)$. Classical outputs can be included through finite-dimensional commutative algebras [3,4].

Definition 2.1 (memory). A memory consists of channels

$$M_\ell \xrightarrow{\mathcal{E}} M_p \xrightarrow{\mathcal{N}} M_q \xrightarrow{\mathcal{R}} M_\ell.$$

Its effective logical channel is $\mathcal{T} = \mathcal{R}\mathcal{N}\mathcal{E}$. The memory is exact when

$$\boxed{\mathcal{R}\mathcal{N}\mathcal{E} = \text{id}_{M_\ell}.} \quad (1)$$

The physical process may be $\mathcal{N}_t$ for a specified duration t . A claim uniform in time must specify whether recovery may depend on t .

Categorically, (1) makes M_ℓ a retract of M_q , with section $\mathcal{N}\mathcal{E}$. The encoding is also split monic, with left inverse $\mathcal{R}\mathcal{N}$. More explicitly, let $\mathbb{M}$ be the category presented by three objects L, P, Q , arrows $e : L \rightarrow P$, $n : P \rightarrow Q$, $r : Q \rightarrow L$, and relation $rne = 1_L$. Exact memories are functors $\mathbb{M} \rightarrow \text{QChan}$.

Proposition 2.2 (finite exactness test). For $E_{ab} = |a\rangle\langle b|$, the following are equivalent: (i) $\mathcal{T} = \text{id}$; (ii) $\mathcal{T}(E_{ab}) = E_{ab}$ for all a, b ; (iii) $\mathcal{T} \otimes \text{id}$ preserves every joint state with every finite reference register.

Proof. Matrix units are a complex-linear basis, giving (i) iff (ii). Tensoring the identity gives (iii). Conversely, (iii) includes preservation of all logical states; density matrices span M_ℓ , giving (i). $\square$

Preserving computational-basis states alone does not suffice: off-diagonal matrix units carry coherence. Reference preservation also prevents a criterion that ignores correlations with other registers.

3. Approximate storage and composition

For a linear map Λ , set

$$\|\Lambda\|_\diamond = \sup_{r \geq 1} \sup_{X \neq 0} \frac{\|(\Lambda \otimes \text{id}_{M_r})(X)\|_1}{\|X\|_1}.$$

Let $d_\diamond(\Phi, \Psi) = \frac{1}{2}\|\Phi - \Psi\|_\diamond$. An ε -memory satisfies $\delta(\mathcal{T}) := d_\diamond(\mathcal{T}, \text{id}) \leq \varepsilon$, with $0 \leq \varepsilon \leq 1$. The reference system in this norm makes the criterion stable in larger experiments [4].

Proposition 3.1 (composition and perturbation). Compatible channels satisfy

$$\delta(\mathcal{T}_2 \mathcal{T}_1) \leq \delta(\mathcal{T}_1) + \delta(\mathcal{T}_2), \quad \delta(\mathcal{T}_1 \otimes \mathcal{T}_2) \leq \delta(\mathcal{T}_1) + \delta(\mathcal{T}_2).$$

For two memory triples with the same dimensions,

$$d_\diamond(\mathcal{RN}\mathcal{E}, \mathcal{R}'\mathcal{N}'\mathcal{E}') \leq d_\diamond(\mathcal{R}, \mathcal{R}') + d_\diamond(\mathcal{N}, \mathcal{N}') + d_\diamond(\mathcal{E}, \mathcal{E}').$$

Proof. Expand the differences by telescoping. For example,

$$\mathcal{T}_2\mathcal{T}_1 - \text{id} = \mathcal{T}_2(\mathcal{T}_1 - \text{id}) + (\mathcal{T}_2 - \text{id}),$$

and

$$\mathcal{T}_1 \otimes \mathcal{T}_2 - \text{id} \otimes \text{id} = (\mathcal{T}_1 - \text{id}) \otimes \mathcal{T}_2 + \text{id} \otimes (\mathcal{T}_2 - \text{id}).$$

The triple difference expands into three terms by replacing one factor at a time. Apply the triangle inequality and submultiplicativity; every channel has diamond norm one. $\square$

4. Specifying the information to preserve

Define the Heisenberg adjoint by $\text{Tr}(\mathcal{T}(\rho)a) = \text{Tr}(\rho\mathcal{T}^\dagger(a))$.

Definition 4.1 (retained algebra). For a unital C^* -subalgebra $\mathcal{A} \subseteq M_\ell$, a memory preserves $\mathcal{A}$ if $\mathcal{T}^\dagger(a) = a$ for every $a \in \mathcal{A}$. Equivalently, it preserves all expectation values specified by $\mathcal{A}$ for every input.

The decomposition $\mathcal{A} \cong \bigoplus_j M_{d_j}$ describes classical sector information and quantum information inside each sector. A commutative algebra specifies classical data; a full matrix algebra specifies a complete quantum register.

For a channel $\mathcal{F}(\rho) = \sum_i K_i \rho K_i^\dagger$, the established full-support correctability criterion is

$$\mathcal{A} \text{ is correctable for } \mathcal{F} \iff [a, K_i^\dagger K_j] = 0 \text{ for all } a \in \mathcal{A}, i, j. \quad (2)$$

This is the full-support specialization of operator algebra quantum error correction [5, Theorem 9]; it is an imported result. To analyze an encoded memory, take $\mathcal{F} = \mathcal{N}\mathcal{E}$ and its recovery to have output M_ℓ .

Example 4.2 (classical information survives dephasing). Let $\Delta_Z(\rho) = (\rho + Z\rho Z)/2$. Its adjoint fixes $\text{span}\{I, Z\}$, so computational-basis probabilities survive. Yet

$$\Delta_Z(|+\rangle\langle+|) = \Delta_Z(|-\rangle\langle-|) = I/2.$$

No recovery can send this one output to both required input states. Hence the full qubit is not recoverable.

Example 4.3 (a recoverable quantum register). Set $c_0 = 000$, $c_1 = 111$, and encode $V|a\rangle = |c_a\rangle$. For $S = \{000, 001, 010, 100\}$, define $X_e|x\rangle = |x \oplus e\rangle$. Consider

$$\mathcal{N}(\rho) = \sum_{e \in S} p_e X_e \rho X_e^\dagger, \quad p_e \geq 0, \quad \sum_e p_e = 1.$$

Define $D_e = \sum_{a=0}^1 |a\rangle\langle c_a \oplus e|$ and $\mathcal{R}(\sigma) = \sum_e D_e \sigma D_e^\dagger$. The eight strings $c_a \oplus e$ are distinct, so $\sum_e D_e^\dagger D_e = I_8$. Also $D_f X_e V = \delta_{fe} I_2$. Therefore

$$\mathcal{RN}(V\rho V^\dagger) = \sum_e p_e \rho = \rho. \quad (3)$$

This preserves all matrix units, including coherences. It illustrates the established theory of quantum error-correcting codes [6]. The code corrects the specified single-bit-flip error set; it does not correct arbitrary single-qubit errors.

5. Constraints on copying and classical readability

Proposition 5.1 (no universal pure-state cloning). No channel copies every pure state when $d \geq 2$.

Proof. If the two-copy output is pure, a Stinespring isometry must have $W|\psi\rangle = |\psi\rangle|\psi\rangle|e_\psi\rangle$. Choose inputs with $0 < s = |\langle\psi|\phi\rangle| < 1$. Preservation of inner products implies $s = s^2|\langle e_\psi|e_\phi\rangle| \leq s^2$, a contradiction. $\square$

No-broadcasting extends the obstruction to noncommuting mixed states [7]. Error correction distributes an encoding over a larger register; it does not clone an unknown logical state.

Proposition 5.2 (distinguishability bound). A d -dimensional register encodes at most d perfectly distinguishable classical messages without additional correlated resources.

Proof. Suppose a POVM distinguishes states ρ_x perfectly. Then $\text{Tr}(F_x\rho_x) = 1$ and $\text{Tr}(F_x\rho_y) = 0$ for $x \neq y$. Positivity implies $\text{supp } \rho_x \subseteq \ker(I - F_x)$ and $\text{supp } \rho_y \subseteq \ker F_x$. These supports are pairwise orthogonal. At most d nonzero mutually orthogonal supports fit in dimension d . $\square$

Corollary 5.3 (storage bottleneck). Exact storage implies $\ell \leq \min\{p, q\}$.

Proof. The ℓ logical basis inputs remain perfectly distinguishable at either intermediate stage: apply the remaining channels and measure the recovered basis. Proposition 5.2 applies to each intermediate register. $\square$

The established Holevo bound further gives, for an ensemble followed by a measurement,

$$I(X : Y) \leq S\left(\sum_x p_x \rho_x\right) - \sum_x p_x S(\rho_x) \leq \log_2 d, \quad (4)$$

without preshared entanglement or correlated side information [8]. Here $S(\rho) = -\text{Tr}(\rho \log_2 \rho)$. Continuous amplitudes are not an arbitrarily long classical record readable from one specimen. Tomography on repeated preparations consumes additional specimens and is a different resource model.

6. Operational equality and deformation

Proposition 6.1 (path collapse). The topological space of channels of fixed input and output dimensions is contractible.

Proof. It is convex: $(1 - t)\Phi + t\Psi$ remains completely positive and trace preserving. Interpolating to a fixed channel gives a contraction. $\square$

Thus treating arbitrary continuous deformation as operational equality identifies all channels of fixed dimensions. In particular it identifies the identity channel with the depolarizing channel $\rho \mapsto \text{Tr}(\rho)I/d$, despite their different storage behavior. Deformation can be meaningful family data, but cannot automatically certify equality of physical operations.

We retain channels as coefficient data and use specified changes of frame as representation equivalences. A continuous family of coefficient values is not an internal identity between its endpoint operations.

7. Higher semantics of memory families

Fix ℓ, p, q, ε . In the finite-dimensional topology of Choi matrices let

$$S_\varepsilon = \{(\mathcal{E}, \mathcal{N}, \mathcal{R}) : d_\diamond(\mathcal{R}\mathcal{N}\mathcal{E}, \text{id}) \leq \varepsilon\}.$$

Let $G = PU(\ell) \times PU(p) \times PU(q)$. Projective unitaries suffice because scalar phases do not change conjugation channels. Write $\text{Ad}_u(a) = uau^\dagger$. Define

$$\begin{aligned}\mathcal{E}' &= \text{Ad}_v \mathcal{E} \text{Ad}_{v^{-1}}, \\ \mathcal{N}' &= \text{Ad}_w \mathcal{N} \text{Ad}_{w^{-1}}, \\ \mathcal{R}' &= \text{Ad}_u \mathcal{R} \text{Ad}_{u^{-1}}.\end{aligned}\tag{5}$$

Then $\mathcal{T}' = \text{Ad}_u \mathcal{T} \text{Ad}_{u^{-1}}$. Unitary invariance of diamond norm makes S_ε invariant.

Choose a small skeleton $\mathcal{C}$ of second-countable manifolds with continuous maps, equipped with the open-cover topology. Work in the infinity-topos $\mathcal{X} = \text{Shv}(\mathcal{C}; \mathcal{S})$ of space-valued sheaves [9]. Define the internally 0-truncated coefficient sheaves

$$\mathbf{S}_\varepsilon(T) = C(T, S_\varepsilon), \quad \mathbf{G}(T) = C(T, G).$$

Definition 7.1 (memory stack). The stack of ε -memories is the quotient

$$\boxed{\mathbf{Mem}_\varepsilon = [\mathbf{S}_\varepsilon / \mathbf{G}]}.\tag{6}$$

An object over T is a locally trivial principal G -bundle $P \rightarrow T$ with an equivariant map $f : P \rightarrow S_\varepsilon$, using $f(pg) = g^{-1}f(p)$. Arrows are equivariant bundle isomorphisms preserving f . This is a 1-stack inside an infinity-topos. It does not posit additional higher physical state variables. Its families have higher topology through their topological classifying spaces.

The operational category supplies processes; the infinity-topos supplies a universe for varying specifications, equivalences, and descent. These roles must remain distinct. In the exact case, unitary natural isomorphisms of the functors in Section 2 give the change-of-frame arrows.

Theorem 7.2 (effective descent). For an open cover $\{T_i \rightarrow T\}$, let $m_i : T_i \rightarrow S_\varepsilon$ and $g_{ij} : T_i \cap T_j \rightarrow G$ satisfy $m_i = g_{ij} \cdot m_j$, $g_{ii} = 1$, and $g_{ij}g_{jk} = g_{ik}$. These data determine an object of $\mathbf{Mem}_\varepsilon(T)$, unique up to equivalence compatible with the specified descent data.

Proof. The cocycle glues trivial principal bundles to a principal G -bundle. The transformation equations glue the local coefficient maps to an equivariant map into S_ε . Conversely, local trivializations recover these data, and changes of trivialization give their equivalences. Gauge invariance of the storage inequality makes correctness well defined throughout the glued family. $\square$

Pairwise equivalences alone are insufficient. On a nonempty triple overlap, $g_{12} = g_{23} = 1$ and $g_{13} = h \neq 1$ can preserve the same local object when h is its automorphism, but fail the cocycle equation. This invalidates those identifications, not necessarily the existence of a different gluing.

Tensor products give coherent operations between the stacks: two tolerances ε, η give tolerance at most $\min\{1, \varepsilon + \eta\}$, by Proposition 3.1. Tensoring unitary frames defines the required group homomorphisms.

8. Perfect storage without a global frame

Theorem 8.1 (unitary exact-memory sector). For $\ell = p = q = d$, the substack where all three channels are unitary and their composite is the identity is equivalent to $BPU(\mathbf{d})$.

Proof. The $PU(d)^3$ action is transitive: choose frames successively along the unitary encoding and evolution; exactness makes recovery the identity too. The stabilizer of $(\text{id}, \text{id}, \text{id})$ is the diagonal $PU(d)$. The quotient stack of a transitive action is the classifying stack of its stabilizer. $\square$

Corollary 8.2 (nontrivial qubit families). Over S^4 there are perfectly preserving memory families, locally identity triples, with no global simultaneous frame.

Proof. Theorem 8.1 identifies them with principal $PU(2)$ -bundles. The double cover $SU(2) \cong S^3 \rightarrow PU(2) \cong SO(3)$ gives $\pi_3(PU(2)) \cong \mathbb{Z}$. Clutching along the equator of S^4 yields nontrivial bundles. A global frame would trivialize the bundle. $\square$

The associated topological classifying space has $\pi_4(BPU(2)) \cong \mathbb{Z}$. This higher homotopy classifies representation families. The obstruction concerns framing, not information loss: a global bundle-valued memory exists. The example permits a varying logical register bundle. If a global logical frame is fixed and encoding is unitary, that encoding trivializes the other register bundles, so this particular obstruction disappears. The topology supplies no extra readable classical capacity for an individual qubit.

9. Measurement contextuality is another gluing problem

For a set U of binary observables, let $\mathcal{O}(U) = \{+1, -1\}^U$, with restriction by forgetting values. An empirical model supplies distributions and support constraints on jointly measurable contexts. For two qubits consider

$$\begin{array}{ccc} X \otimes I & I \otimes X & X \otimes X \\ I \otimes Z & Z \otimes I & Z \otimes Z \\ X \otimes Z & Z \otimes X & Y \otimes Y. \end{array} \tag{7}$$

Here X, Y, Z are Pauli matrices. Every row and column is commuting. All row products and the first two column products equal I ; the final column product equals $-I$.

Proposition 9.1 (parity obstruction). There is no context-independent assignment of values ± 1 to these nine observables satisfying all six product constraints.

Proof. Multiplying the constraints counts each assigned value twice, giving $+1$ on the left and -1 on the right. $\square$

This is the Mermin parity argument [10], within the sheaf account of contextuality [11]. The unrestricted assignment presheaf has global assignments; the quantum support constraints exclude every one. This does not imply absence of a quantum state: density matrices supply compatible Born distributions on the commuting contexts.

Topos approaches to algebraic quantum theory organize commutative observational contexts, with internal spectra and probability valuations [12]. Their context-based constructions differ from the parameter-based memory stack above. Contextuality obstructs a global noncontextual outcome assignment; a nontrivial frame bundle obstructs a global coordinate choice. They are not interchangeable obstructions.

10. Interpretation, evidence, and limitations

A state is an information-bearing physical state relative to specified preparations and observations. Its coordinates are not automatically a readable classical record. Storage is relational: a logical specification, encoding, intervening dynamics, and recovery determine what survives. Higher semantics records representation equivalences and their coherence across families.

The assumptions here do not derive the Born rule, single out quantum theory among all information theories, or establish an ontology of information. Infinite-dimensional systems, quantum fields, energy budgets, locality, thermal stability, and fault-tolerant engineering require further hypotheses. No storage density or lifetime is inferred from Hilbert-space dimension alone.

Exact integer complex arithmetic checked 16 repetition-code compression identities, 16 matrix-unit recovery identities, 18 commuting pairs, and the six context-product signs $(+, +, +, +, +, -)$. Exhaustion of all 512 assignments found zero satisfying all contextuality constraints. The attached standard-library Python script reproduces these checks. They verify the finite examples, not the general proofs or the higher stack construction. This is not a machine-checked formalization.

Authorship and automation. Anonymous author is the public byline of the responsible human contributor requesting publication. An AI assistant assisted with drafting, proof development, source checking, and finite calculations. The contribution is a synthesis and explicit memory-family semantics; no priority claim is made for the established results cited below.

References

1. S. Abramsky and B. Coecke. *A categorical semantics of quantum protocols*. LICS, 2004. <https://arxiv.org/abs/quant-ph/0402130>.
2. P. Selinger. *Dagger compact closed categories and completely positive maps*. QPL 2005 preliminary version. <https://www.mscs.dal.ca/~selinger/papers/dagger.pdf>.
3. B. Coecke, C. Heunen, and A. Kissinger. *Categories of Quantum and Classical Channels*. <https://arxiv.org/abs/1305.3821>.
4. J. Watrous. *The Theory of Quantum Information*. Cambridge University Press, 2018. <https://cs.uwaterloo.ca/~watrous/TQI/>.
5. C. Beny, A. Kempf, and D. W. Kribs. *Quantum Error Correction of Observables*. Physical Review A **76**, 042303, 2007. <https://arxiv.org/abs/0705.1574>.
6. E. Knill and R. Laflamme. *A Theory of Quantum Error-Correcting Codes*. 1996 preprint. <https://arxiv.org/abs/quant-ph/9604034>.
7. H. Barnum, C. M. Caves, C. A. Fuchs, R. Jozsa, and B. Schumacher. *Noncommuting mixed states cannot be broadcast*. Physical Review Letters **76**, 2818-2821, 1996. <https://arxiv.org/abs/quant-ph/9511010>.
8. A. S. Holevo. *The Capacity of Quantum Channel with General Signal States*. IEEE Transactions on Information Theory **44**, 269-273, 1998. <https://arxiv.org/abs/quant-ph/9611023>.
9. J. Lurie. *Higher Topos Theory*. <https://arxiv.org/abs/math/0608040>.
10. N. D. Mermin. *Simple unified form for the major no-hidden-variables theorems*. Physical Review Letters **65**, 3373-3376, 1990. <https://doi.org/10.1103/PhysRevLett.65.3373>.
11. S. Abramsky and A. Brandenburger. *The Sheaf-Theoretic Structure of Non-Locality and Contextuality*. New Journal of Physics **13**, 113036, 2011. <https://arxiv.org/abs/1102.0264>.
12. C. Heunen, N. P. Landsman, and B. Spitters. *A topos for algebraic quantum theory*. Communications in Mathematical Physics **291**, 63-110, 2009. <https://arxiv.org/abs/0709.4364>.
13. G. Chiribella, G. M. D'Ariano, and P. Perinotti. *Informational derivation of Quantum Theory*. Physical Review A **84**, 012311, 2011. <https://arxiv.org/abs/1011.6451>.