

Digital Persistence, Finite Memory, and Personal Identity: A Higher-Categorical Semantics

Anonymous author

19 October 2026 | Version 1.0

Prepared for *Higher Categorical Semantics*

Abstract

We formalize the claim that digital immortality requires indefinitely increasing storage density. The claim is valid only under additional retention and resource hypotheses. Exact archives require enough distinguishable states to encode all admissible histories, whereas task-relative behavior can persist with finite memory. We define observed dynamics in an infinity-topos and prove preservation under coherent equivalence migrations and effective descent. We distinguish arbitrarily long finite realizations from a compatible infinite continuation, and characterize positive survival probability by summability of conditional fatal-loss hazards. A bounded-capacity assumption yields an obstruction to indefinitely growing archives, including a separate information-theoretic hypothesis for bounded-error recovery. A noncontractible realization fiber shows why equivalent specifications need not determine unique physical realizers. These mathematical results do not supply a criterion of consciousness or personal survival. The contribution is a synthesis of established tools into an explicit account of the distinct obligations of digital persistence.

Keywords: digital persistence; information capacity; automata; infinity-topoi; descent; personal identity.

1. Scope and logical status

Digital immortality can mean preservation of records, continued execution of a cognitive process, or survival of a particular person. These are different specifications. We study the implication from indefinite persistence to increasing storage density, and identify the hypotheses under which it is valid. We make no assumption that a human mind is digitally realizable, and no prediction of a technological singularity.

The contribution is a synthesis of information bounds, automata semantics, higher-categorical migration, descent, and reliability. The counting and minimization arguments are classical [1,2]; the categorical framework uses standard constructions [3,4]. Their application here separates mathematical guarantees from physical and personal-identity assumptions. No new law of physics or criterion of consciousness is claimed.

We work externally in classical set theory with choice and appropriate universes. The higher-categorical ambient category is an arbitrary infinity-topos $\mathcal{X}$: an accessible left-exact localization of

a presheaf infinity-category of spaces. All categorical limits, colimits, and pullbacks are homotopy coherent. A point of $S \in \mathcal{X}$ means a global section $1 \rightarrow S$; local inhabitation need not supply such a point. Numerical storage budgets are external resource assignments, not intrinsic cardinalities of arbitrary objects of $\mathcal{X}$.

2. Exact retention and its resource requirements

Definition 2.1 (archive). At stage n , let H_n be a finite nonempty set of admissible records. An exact archive is an encoder $e_n : H_n \rightarrow M_n$ and decoder $d_n : M_n \rightarrow H_n$ satisfying $d_n e_n = \text{id}_{H_n}$. A b_n -bit memory has $|M_n| \leq 2^{b_n}$. Every record-dependent resource available to the decoder, including external storage or a varying codebook, is included in M_n . A fixed decoder may be shared among all records.

Theorem 2.2 (exact archive bound). An exact b_n -bit archive exists abstractly if and only if

$$b_n \geq K_n := \lceil \log_2 |H_n| \rceil.$$

Proof. The decoding identity makes e_n injective, so $|H_n| \leq |M_n| \leq 2^{b_n}$. Conversely, enumerate H_n and encode its elements by distinct binary strings of length K_n . Define a decoder arbitrarily on unused strings. This establishes representability, not an efficient or physically realizable implementation. $\square$

Corollary 2.3 (density requirement). If usable capacity is at most $\rho_n V_n$ bits, exact retention requires $\rho_n V_n \geq K_n$. If $\liminf_{n \rightarrow \infty} K_n/n \geq h > 0$ and $V_n \leq V_0 < \infty$, then, for every $0 < \epsilon < h$, eventually

$$\rho_n \geq \frac{(h - \epsilon)n}{V_0}.$$

Proof. Apply Theorem 2.2 and the definition of the lower limit. $\square$

For $H_n = H_0 \times A^n$, with $|A| = q \geq 2$, one has $K_n = \lceil \log_2 |H_0| + n \log_2 q \rceil$. This gives linear memory growth for freely combinable distinguishable inputs. A single predetermined history need not have this storage requirement. The archive must distinguish the admissible alternatives. Positive information growth is a sufficient condition for unbounded memory, not a consequence of elapsed time alone.

Proposition 2.4 (bounded-error retention). Let W be uniform on $N \geq 2$ records, let the stored variable Z take at most 2^b values, and let a deterministic decoder have error probability p . Then

$$b \geq \log_2 N - h_2(p) - p \log_2(N - 1),$$

where $h_2(p) = -p \log_2 p - (1 - p) \log_2(1 - p)$ with $0 \log 0 = 0$. In particular, $p \leq \epsilon < 1$ implies $b \geq (1 - \epsilon) \log_2 N - 1$. **Proof.** Let E indicate decoding error. Given Z and $E = 0$, W is known; given Z and $E = 1$, at most $N - 1$ alternatives remain. The entropy chain rule gives $H(W | Z) \leq h_2(p) + p \log_2(N - 1)$. Hence $\log_2 N = I(W; Z) + H(W | Z) \leq b + h_2(p) + p \log_2(N - 1)$. Use $h_2(p) \leq 1$ for the final inequality. $\square$

This is the elementary Fano-type bound. Neither lossless compression nor a fixed nonunit tolerated error removes the divergence when $\log_2 |H_n| \rightarrow \infty$. These are worst-case representational or specified-distribution claims, not estimates of the information in human experience.

3. Continued behavior without complete history

Definition 3.1 (task-relative memory). For a finite input alphabet A and output set O , fix a task $\tau : A^* \rightarrow O$. Define

$$u \sim_\tau v \iff \forall w \in A^*, \tau(uw) = \tau(vw).$$

A deterministic realization consists of a state set S , initial state s_0 , updates $\delta : S \times A \rightarrow S$, and observation $o : S \rightarrow O$ with $o(\delta^*(s_0, u)) = \tau(u)$. The empty word is denoted ε .

Theorem 3.2 (minimal task memory). The quotient $Q = A^*/\sim_\tau$ has a well-defined realization with initial state $[\varepsilon]$, update $[u] \mapsto [ua]$, and observation $[u] \mapsto \tau(u)$. Every realization has at least $|Q|$ reachable states. Consequently a finite-state realization exists exactly when Q is finite. **Proof.** If $u \sim_\tau v$, taking $w = \varepsilon$ gives equal observations, and taking continuations aw shows $ua \sim_\tau va$. This proves well-definedness and correctness. If two histories reach the same state in any realization, every continuation produces equal task outputs, so those histories are equivalent. The map from reachable states to their history classes is therefore a well-defined surjection. $\square$

This is the Moore-output version of the Myhill-Nerode argument [2], not a new minimization theorem. For binary inputs and $\tau(u)$ equal to parity, there are exactly two classes. One bit supports arbitrarily many logical updates while discarding most of the history. By contrast, the task $\tau(u) = u$ distinguishes all histories. Abstract indefinite execution therefore does not imply an ever-growing archive. A finite autonomous deterministic system eventually cycles; an open finite-state system can have nonperiodic outputs supplied by its environment. Neither fact proves physical eternal operation.

4. Higher-categorical dynamics and migration

Let BA^* be the free category with one object and one generating loop for each $a \in A$, regarded as an infinity-category. Composition is successive input processing, so the action maps satisfy $f_{uv} \simeq f_v f_u$. Fix an observation object $O \in \mathcal{X}$.

Definition 4.1 (observed dynamics). Put

$$\mathcal{D}_{A,O}(\mathcal{X}) = \text{Fun}(BA^*, \mathcal{X}) \times_{\mathcal{X}} \mathcal{X}_{/O},$$

using evaluation at the unique object and the domain functor. An object is a coherent action on a state object S together with $o : S \rightarrow O$. A migration $m : D \rightarrow D'$ is a coherent natural transformation whose underlying $m : S \rightarrow S'$ also satisfies $o'm \simeq o$. Thus $mf_a \simeq f'_a m$, with the coherences required by the functor category. For ordinary sets the same data are a coalgebra $S \rightarrow O \times S^A$ [3]. A chosen initial state is a point of S , not necessarily a fixed point of the action.

Lemma 4.2 (preservation under migration). A migration preserves observations after every finite input word. If its underlying state map is an equivalence, it is an equivalence in $\mathcal{D}_{A,O}(\mathcal{X})$.

Proof. Naturality for each word gives $mf_w \simeq f'_w m$; composing with o' gives $o'f'_w m \simeq of_w$. Equivalences in a functor infinity-category are objectwise, and equivalences in a slice are detected on their underlying arrows. Equivalences are also detected componentwise in the displayed pullback. $\square$

Theorem 4.3 (coherent migration). Suppose $D : \mathbb{N} \rightarrow \mathcal{D}_{A,O}(\mathcal{X})$ is a coherent sequence all of whose transition maps are equivalences. Its colimit D_∞ exists, and every canonical map $D_n \rightarrow D_\infty$ is an equivalence. A point of S_0 is transported to a point of the colimit state object with the same

finite-word observations. **Proof.** Colimits of actions are formed objectwise. Colimits in $\mathcal{X}/O$ are created on domains, so these colimits give compatible colimits in the defining pullback. The arrows $D_0 \rightarrow D_n$ form an objectwise equivalence from the constant diagram at D_0 to D . The nerve of the indexing poset $\mathbb{N}$ is contractible; hence its constant-diagram colimit is D_0 . The same argument applies to every cofinal tail. Point and observation preservation follow from the colimit cocone and Lemma 4.2. $\square$

This theorem concerns a diagram already supplied with coherent migrations. It does not construct future hardware, guarantee a physical transfer, or identify a colimit with an event at an infinite time. Pairwise equivalences chosen independently are not the full specified diagram. Equivalence is a sufficient preservation condition; behavioral quotients such as Theorem 3.2 may preserve selected observations while losing state information.

5. Descent and compatible continuation

Theorem 5.1 (descent for observed dynamics). Let $U \rightarrow 1$ be an effective epimorphism in $\mathcal{X}$, and let $U_\bullet$ be its Čech nerve, so $U_n = U^{\times(n+1)}$. Write $O_U = O \times U \rightarrow U$. Then

$$\mathcal{D}_{A,O}(\mathcal{X}) \simeq \lim_{[n] \in \Delta} \mathcal{D}_{A,O_{U_n}}(\mathcal{X}_{/U_n}).$$

Consequently local dynamics with a full coherent descent datum determine global dynamics, uniquely up to a contractible space of choices. **Proof.** Effective descent in an infinity-topos gives $\mathcal{X} \simeq \lim_{\Delta} \mathcal{X}_{/U_n}$ [4, Theorem 6.1.3.9 and Section 6.2.3]. Forming functor categories with fixed source BA^* preserves limits. Forming slices over the compatible family O_{U_n} also preserves this limit: objects, arrows to that family, and their coherences are computed componentwise. Taking the pullback in Definition 4.1 proves the equivalence. The statement about choices is the homotopy-fiber characterization of an equivalence. $\square$

Example 5.2 (local sections do not supply a global point). In the infinity-topos of space-valued sheaves on S^1 , take the discrete-valued sheaf P of local sections of $z \mapsto z^2$. It is locally inhabited, so $P \rightarrow 1$ is an effective epimorphism. A global point would give a continuous section s of this double covering. On fundamental groups it would satisfy $2s_* = \text{id}_{\mathbb{Z}}$, which is impossible. Thus global objects obtained by descent must not be confused with global sections of those objects. Theorem 5.1 requires coherent compatibility; it does not assert that arbitrary local choices glue. Only Čech descent is used, with no hypercompleteness assumption.

Definition 5.3 (compatible horizons). Let R_n be the set of admissible realizations through stage n , with restriction maps $r_n : R_{n+1} \rightarrow R_n$. An indefinitely extendible realization in this model is an element of $\lim_n R_n$.

Proposition 5.4 (finite horizons and compactness). Nonemptiness of every R_n does not generally imply nonemptiness of the inverse limit. It does imply it when every R_n is finite. **Proof.** For the negative result take $R_n = \{m \in \mathbb{N} : m \geq n\}$ and inclusion maps. A compatible sequence would have a common value at least every natural number. For the positive result form the tree of compatible finite tuples $(x_0, \dots, x_n)$. Every level is nonempty, since an element of R_n restricts to such a tuple. Each level is finite. Starting at a root with arbitrarily long extensions, at least one of its finitely many children has arbitrarily long extensions; otherwise their finite height bounds would bound the parent. Repeating produces an infinite branch, hence a compatible sequence. This is the classical finite-branching compactness argument. $\square$

The positive statement is external and classical. It supplies neither a computable selection algorithm nor the physical resources to run its branch. Thus $\forall n \exists$ a finite continuation and $\exists$ a compatible continuation for all n are distinct claims.

6. Physical capacity and reliability

Assumption 6.1 (bounded accessible capacity). A specified resource envelope admits at most 2^{B_*} perfectly recoverable classical records, for some finite B_* . For bounded-error claims we additionally require $I(W; Z) \leq B_*$ for every allowed record ensemble and decoder-accessible observation Z . The envelope includes all record-dependent decoder resources. A classical memory with at most 2^{B_*} states satisfies both clauses; the exact clause alone need not imply the information clause.

Theorem 6.2 (bounded-resource archive obstruction). Under the exact clause of Assumption 6.1, exact archives for all n are impossible if $\sup_n K_n = \infty$. Under its information clause, uniform records cannot be retained with error bounded by a fixed $\epsilon < 1$ when $\log_2 |H_n| \rightarrow \infty$. **Proof.** Exact recovery would give $\log_2 |H_n| \leq B_*$ for every n , contradicting unbounded K_n . For bounded error, the proof of Proposition 2.4 uses $I(W; Z) \leq B_*$ in place of $H(Z) \leq b$, giving $(1 - \epsilon) \log_2 |H_n| - 1 \leq B_*$. $\square$

The theorem is independent of a particular physical bound. A conditional motivation is the Bekenstein entropy bound for a complete weakly self-gravitating system of circumscribing radius R and total energy E [6]. If accessible classical information is bounded by maximal entropy in bits, one may take

$$B_* = \frac{2\pi ER}{\hbar c \ln 2}.$$

This is an additional physical applicability assumption, not a theorem about all cosmological spacetimes or all information tasks. The operational distinctions studied in [7] caution against equating every information quantity with storage capacity. At bounded E and R in the stated regime, endlessly denser media cannot accommodate an unbounded exact archive. Distributed expansion may enlarge the resource envelope; it does not refute the theorem.

Capacity and dissipation are different budgets; reversible computation forbids assigning an erasure cost to every logical operation [8,9].

For ideal isothermal reset of one unbiased classical bit in the usual symmetric-memory setting, Landauer's bound is $k_B T \ln 2$ of dissipated heat. If at least $e > 0$ such bits are erased per stage at $T \geq T_0 > 0$, cumulative heat after n stages is at least $nek_B T_0 \ln 2$. A finite cumulative dissipation budget would then bound the number of stages. This conditional argument requires the stated reset operations; it is not a general per-operation energy bound.

Theorem 6.3 (cumulative failure). Let L_n be the event of no fatal loss through stage n , with L_0 certain and $L_n \subseteq L_{n-1}$. Set $p_n = 1 - \Pr(L_n \mid L_{n-1})$, assuming $0 \leq p_n < 1$. Then

$$\Pr \left(\bigcap_{n \geq 1} L_n \right) = \prod_{n \geq 1} (1 - p_n).$$

This probability is positive exactly when $\sum_n p_n < \infty$, and equals one exactly when every $p_n = 0$. **Proof.** Iterated conditioning gives $\Pr(L_N) = \prod_{n=1}^N (1 - p_n)$, without independence. Continuity of probability for decreasing events gives the infinite product. Since $-\log(1 - p) \geq p$, a divergent

hazard sum forces product zero. If the sum converges, eventually $p_n \leq 1/2$ and $-\log(1 - p_n) \leq 2p_n$; the finite initial factors are positive, so the product is positive. A single positive hazard makes it strictly less than one. $\square$

For $p_n = 1/(n+1)$ survival through N stages is $1/(N+1)$; for $p_n = 1/(n+1)^2$ it is $(N+2)/(2(N+1))$, tending to $1/2$. Improving reliability can therefore give either zero or positive indefinite survival probability. Positive probability is not certainty. These are illustrative hazards, not measurements of a proposed technology.

7. Semantic equivalence and personal identity

Let $\mathcal{P}$ be a category of physical realizations and allowed transformations, and let $F : \mathcal{P} \rightarrow \mathcal{D}_{A,O}(\mathcal{X})$ be a chosen semantic functor. This functor is a modeling assumption; the paper does not construct it for brains. For a specification D , define its space of framed realizations by

$$\text{Real}(D) = \mathcal{P}^{\simeq} \times_{\mathcal{D}_{A,O}(\mathcal{X})^{\simeq}} \{D\},$$

where the superscript denotes the maximal infinity-groupoid. Its points are realizations with specified equivalences to D .

Proposition 7.1 (nonuniqueness of realizers). A unique chosen semantic specification need not have a contractible realization space, even when every realization has exactly equivalent dynamics.

Proof. Take $\mathcal{P} = \mathcal{D}_{A,O}(\mathcal{X}) \times K$, where K is the discrete two-point category, and let F be projection. Then $\text{Real}(D) \simeq K$. The two realizations have identical semantics and distinct invariant tags, with no equivalence between them in $\mathcal{P}$. $\square$

The example is a logical countermodel to an implication, not evidence for a particular metaphysics. Univalence identifies equivalent types within its specified universe [5]; it does not make the fiber of an arbitrary realization functor contractible. In particular it does not identify two distinct elements merely because a map sends them to the same value.

Proposition 7.2 (scope of survival claims). If a satisfiable formal theory of realizations and semantics has no axioms involving an added binary predicate *Survives*, then its structural axioms alone cannot determine whether that predicate holds of a specified pair of realizations. **Proof.** For any model and specified pair, expand the model once with the predicate interpreted as the full relation and once as the empty relation. Both expansions satisfy the original axioms and disagree at that pair. $\square$

This modest model-theoretic observation does not prove that personal survival is unknowable or nonphysical. It identifies a missing bridge assumption. The distinct questions of consciousness and identity under uploading are analyzed in [10, Sections 9-10]. Neither a compressed archive nor a coherent semantic migration supplies that bridge by definition.

8. Consequences and limitations

Ever-increasing storage density is necessary for an ever-growing retained information requirement within a fixed spatial budget; shrinking that budget strengthens the requirement. Indefinite execution alone has no such implication. Complete retention, bounded accessible capacity, and unbounded information growth cannot all hold. Coherent migration and descent establish structural preservation and compatibility, while the inverse-limit and hazard results isolate additional temporal obligations.

The model specifies exact dynamics and selected observations. It gives no neural encoding, consciousness test, working upload, engineering forecast, or cosmological survival proof. Established counting, minimization, descent, and compactness results are used with attribution; the contribution is their explicit separation and application to digital persistence. The accompanying appendices discuss approximate migrations, observational adequacy, and reproducibility.

Authorship and assistance. Anonymous author. AI assistance was used to draft the exposition, develop and check arguments, locate references, and generate the accompanying files. The manuscript supplies mathematical proofs, but no proof-assistant certification or empirical validation of digital immortality is claimed. Editorial acceptance and independent review are not presumed.

Appendix A. Approximate migration and accumulated error

The exact equivalences of Theorem 4.3 are strong requirements. Approximate implementations need quantitative hypotheses, which are not supplied merely by placing objects in an infinity-category.

Definition A.1 (metric implementation). Let (S, d_S) and (T, d_T) be metric spaces, with deterministic input updates $f_a : S \rightarrow S$ and $g_a : T \rightarrow T$. Let $m : S \rightarrow T$ be a proposed migration. Suppose

$$d_T(mf_a(x), g_a m(x)) \leq \eta \quad (x \in S, a \in A),$$

and suppose every g_a is L -Lipschitz for a common $L \geq 0$. A pair of trajectories driven by the same inputs is denoted $x_{k+1} = f_{a_k}(x_k)$ and $y_{k+1} = g_{a_k}(y_k)$.

Proposition A.2 (finite-horizon error propagation). If $e_k = d_T(m(x_k), y_k)$, then

$$e_n \leq L^n e_0 + \eta \sum_{j=0}^{n-1} L^j.$$

For $L = 1$ this is $e_n \leq e_0 + n\eta$. For $0 \leq L < 1$ it is $e_n \leq L^n e_0 + \eta/(1 - L)$.

Proof. The triangle inequality gives

$$\begin{aligned} e_{k+1} &= d_T(mf_{a_k}(x_k), g_{a_k}(y_k)) \\ &\leq d_T(mf_{a_k}(x_k), g_{a_k} m(x_k)) + d_T(g_{a_k} m(x_k), g_{a_k}(y_k)) \\ &\leq \eta + L e_k. \end{aligned}$$

Induction yields the geometric sum. When $L = 0$, take the usual convention $L^0 = 1$ in the finite sum; the formula for positive n remains valid. $\square$

If observations take values in another metric space, o_T is C -Lipschitz, and $d(o_T m(x), o_S(x)) \leq \zeta$, the observation error is at most $C e_n + \zeta$. This follows by another triangle inequality. Uniformly small one-step error therefore implies uniformly small long-run error only under additional stability conditions. The estimate provides an upper bound; growth of that bound is not a proof that every trajectory's error grows. The system $f(x) = x$, $g(y) = y + \eta$, $m = \text{id}$ on the real line, starting at zero, realizes the linear bound exactly.

For a sequence of migration maps with an additive error recurrence $e_{k+1} \leq e_k + \eta_k$, summability of η_k gives a bounded accumulated error. It does not by itself establish exact equivalence or personal survival. This deterministic error budget is also conceptually distinct from the conditional fatal-loss probabilities of Theorem 6.3.

Stages count logical transitions. An interpretation as an unbounded physical lifespan must also supply execution times $t_n \rightarrow \infty$ and a progress guarantee excluding permanent stalling at a finite stage. The present estimates alone do not provide these witnesses.

Appendix B. Observation-relative archives and the system boundary

Definition B.1 (query family). Fix a finite record set H and a collection of queries $\mathcal{Q}$, with $q : H \rightarrow O_q$ for each $q \in \mathcal{Q}$. Put $h \equiv_{\mathcal{Q}} h'$ if every query gives the same answer on h and h' .

Proposition B.2 (query-sufficient storage). The smallest abstract fixed-length storage capable of answering every query exactly has

$$\lceil \log_2 |H/\equiv_{\mathcal{Q}}| \rceil$$

bits.

Proof. If two nonequivalent records had the same encoding, a query distinguishing them could not be answered correctly from that encoding. Thus the number of memory values is at least the number of equivalence classes. Conversely, encode the equivalence class and answer each query by its well-defined value on that class. $\square$

Theorem 3.2 strengthens this static construction by closing the query family under all future input continuations. This closure is necessary if a summary is to support future updates without the discarded history. A quotient preserving only present observations can fail after one further input.

For example, let H consist of binary words and let the present observation be the final bit, with a special output for the empty word. Two nonempty words with the same final bit agree on this observation. If a future input asks for the penultimate bit of the stored word, they need not agree. A migration certified only against the original query is not certified against that enriched interface. Adequacy of the observation family is a substantive part of a specification.

The examples make no claim that parity, verbal reports, or a finite behavioral battery captures personal identity. Increasing the observation family can refine the equivalence relation and increase required memory. Choosing a coarse quotient can make an implementation small by weakening what must be preserved; this is a change in the specification.

Side information. If a decoder is given a record-dependent database Y , the actual archive state is (Z, Y) . The counting argument applies to jointly accessible states. With uniform records and exact reconstruction, Shannon's identities give $H(W | Y) \leq H(Z | Y) \leq H(Z)$. Thus side information can reduce local storage, but its contribution must appear in the resource boundary. A remote archive does not become resource-free by being omitted from a device's dimensions. A clock used to recover record length is also side information when its state distinguishes admissible records.

Compression and one particular history. Theorems 2.2 and 2.4 quantify over a declared family or distribution of alternatives. They do not say that each individual history requires its uncompressed length. A computable repetitive sequence can have a short description, although specifying an arbitrary finite prefix length may itself require growing information. A decoder containing the entire selected record cannot be treated as a shared fixed decoder while comparing alternatives. These distinctions avoid turning a uniform coding lower bound into an unjustified assertion about every particular life.

Information versus the number of physical microstates. Assumption 6.1 concerns distinguishable, accessible classical records. It does not assert that quantum pure states form a finite set. Nor does an abstract cardinality bound quantify latency, addressability, error-correction overhead,

manufacturing feasibility, or robustness. The converse in Theorem 2.2 is strictly combinatorial. Physical encodings require additional witnesses.

Appendix C. Descent, types, and realizations

The descent theorem is an assertion about a category of structured objects. In more detail, write $C_n = \mathcal{X}_{/U_n}$ and $O_n = O_{U_n}$. Effective descent gives $\mathcal{X} \simeq \lim_n C_n$ in the infinity-category of categories. An object on the right consists of objects over every Čech level with all the specified pullback identifications and higher compatibilities. It is not just a list of locally available candidates.

The functor $\text{Fun}(BA^*, -)$ preserves the limit because maps into its value are equivalent to maps out of a product with the fixed source. Slicing the limit over the compatible object O also gives the limit of slices: an arrow into O is precisely a compatible family of arrows into the O_n . These equivalences respect evaluation and domains. Limits commute with limits, so the pullback defining observed dynamics commutes with the Čech limit. This proves Theorem 5.1 at the level of objects, maps, and all higher homotopies.

For ordinary discrete data this specializes to familiar gluing: local state spaces, their transition maps, and observations agree after changes of coordinates on overlaps, and the changes of coordinates satisfy a cocycle condition. For higher data the cocycle law is accompanied by higher homotopies. A chosen point must itself satisfy descent to produce a global point; Example 5.2 shows why local existence cannot replace that requirement.

In dependent-type language, a semantic realization problem can be represented by a family $\text{Real} : \mathcal{D}^\simeq \rightarrow \mathcal{S}$ of spaces. Transport along a semantic identity or equivalence relates the corresponding fibers. It does not prove that a fiber is inhabited, that it is contractible, or that two of its points are equal. A physical witness and a proof of uniqueness are different data. No assertion here requires adding univalence as an axiom to the external set-theoretic metatheory; the type-theoretic reading explains the scope of the interpretation in [5].

The separate predicate in Proposition 7.2 is a bookkeeping device for a proposition that has not been given a mathematical or empirical interpretation. If a theory defines personal survival in terms of specified physical or psychological relations, that definition and its supporting arguments must be stated. The proposition does not refute such definitions. Conversely, redefining survival to mean observational equivalence makes a preservation statement true by the chosen convention, without showing that the convention answers a person’s original question about continued experience.

Appendix D. Worked examples and verification

All numbers in the accompanying scripts are synthetic. The binary archive illustration uses $H_n = \{0, 1\}^n$, so exact history retention needs n bits. The parity task uses one bit. Its two reachable classes are distinguished already by their present output, and appending a bit acts by addition modulo two.

An exact migration relabels the states 0, 1 by the strings `even` and `odd`. Updates toggle both representations on input 1 and leave them unchanged on input 0. Observation maps return the original parity. The migration is bijective, commutes with each update, and preserves observations. A proposed map that swaps the labels while retaining their original observation meanings fails the observation condition. This finite example tests the full specification, not only bijectivity.

For the hazard examples, exact rational arithmetic gives

$$\prod_{n=1}^N \left(1 - \frac{1}{n+1}\right) = \frac{1}{N+1},$$

$$\prod_{n=1}^N \left(1 - \frac{1}{(n+1)^2}\right) = \frac{N+2}{2(N+1)}.$$

The first telescopes directly. In the second, factor each term as $n(n+2)/(n+1)^2$ and cancel. Their infinite limits follow from these closed forms, not from a finite simulation. A constant hazard $p > 0$ gives $(1-p)^N \rightarrow 0$.




Figure 1: Synthetic examples. Left: exact binary history retention requires n bits; parity uses one bit. Right: three conditional hazard sequences yield different survival curves. These examples are not forecasts or estimates of human memory.

The file `verify_examples.py` checks finite encodings, parity updates, migration conditions, exact finite hazard products, and the metric recurrence. It writes a structured report; `make_figure.py` regenerates Figure 1. `build.py` generates the PDF, HTML, LaTeX source, and submission JSON from the Markdown sources. The checks corroborate finite examples and document preparation. They do not certify the general theorems in Lean, Coq, or another proof assistant, and do not test consciousness, physical uploading, or infinite-time behavior.

The main article is supplied separately as `journal-body.md` to satisfy the journal’s 20,000-character body limit. The full manuscript adds these appendices. Both use the anonymous byline. The package’s submission JSON is a prepared payload, not evidence of submission, acceptance, or publication.

References

- [1] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal* 27 (1948), 379-423, 623-656. [Part I DOI](#); [full text](#).
- [2] A. Nerode. Linear automaton transformations. *Proceedings of the American Mathematical Society* 9 (1958), 541-544. [DOI](#).

- [3] J. J. M. M. Rutten. Universal coalgebra: a theory of systems. *Theoretical Computer Science* 249 (2000), 3-80. [DOI](#); [author's institutional repository](#).
- [4] J. Lurie. *Higher Topos Theory*. Annals of Mathematics Studies 170, Princeton University Press, 2009. [Author-maintained text](#); [preprint](#).
- [5] The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study, 2013. [Book and corrections](#).
- [6] J. D. Bekenstein. How does the entropy/information bound work? *Foundations of Physics* 35 (2005), 1805-1823. [DOI](#); [preprint](#).
- [7] P. Hayden and J. Wang. What exactly does Bekenstein bound? *Quantum* 9 (2025), 1664. [DOI](#); [preprint](#).
- [8] R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development* 5 (1961), 183-191. [DOI](#).
- [9] C. H. Bennett. Notes on Landauer's principle, reversible computation, and Maxwell's Demon. *Studies in History and Philosophy of Modern Physics* 34 (2003), 501-510. [Full text and publication metadata](#).
- [10] D. J. Chalmers. The singularity: a philosophical analysis. *Journal of Consciousness Studies* 17, nos. 9-10 (2010), 7-65. [Author's text](#).