

Fiber Observability: Coherent Reconstruction, Truncation Barriers, and Statistical Identification

Anonymous author

20 October 2026 · Version 1.0

Abstract

We formulate fiber observability in an infinity-topos as the resolution of coherent realizations of an observation. Full observability is a monomorphism criterion, and truncation bounds the information retained by admissible probes. We establish refinement, test-adequacy, locality, and task-descent results and exhibit hidden-symmetry and local-choice counterexamples. A finite controlled specialization stabilizes by a bounded observation horizon. For finite stochastic experiments, equal observation laws obstruct identification, while a separation margin yields an explicit confidence bound. The paper combines standard categorical foundations with elementary dynamical and statistical proofs. Exact finite computations illustrate the scope of the statements.

1. Scope and contribution

An observation usually leaves a family of compatible hidden states. In a higher categorical model this family also contains paths, automorphisms, and higher coherence. We call the study of this residual structure **fiber observability**. The term specifies a framework in this paper; no claim is made that it denotes a previously standardized theory.

Classical nonlinear observability distinguishes states through input-output behavior [1]. Coalgebra organizes behavioral equivalence and observable quotients [2]. Higher topos theory and homotopy type theory provide fibers, truncation, and descent [3,4]. Statistical experiments distinguish states through output laws and decision risks [5-7]. We combine these ingredients into one typed account, with explicit proofs and finite examples.

The contribution is a synthesis: a relative definition of coherent observability; a probe-resolution obstruction; tests using generators and descent; finite-horizon reconstruction; and statistical impossibility and confidence bounds. Several ingredients are standard facts or elementary consequences. Their roles and assumptions are kept explicit; the framework does not supply a general algorithm for arbitrary higher objects.

2. Observations and their coherent fibers

Fix an infinity-topos $\mathcal{E}$ and work in a universe permitting the small limits used below. A known context is an object B . A hidden-state object $p : X \rightarrow B$ and a record object $q : Y \rightarrow B$ belong to the slice $\mathcal{E}_{/B}$. An **observer** is a morphism $o : X \rightarrow Y$ in that slice. Products and equalities below are homotopy-coherent unless we explicitly work with sets.

Definition 2.1 (observation fiber). For a generalized record $y : T \rightarrow Y$, define

$$F_y = X \times_Y T \quad \text{in } \mathcal{E}_{/T}.$$

Internally, over a record y , the fiber has the type

$$F_y \simeq \sum_{x:X} (o(x) =_Y y).$$

Thus a realization includes a state and a witness that its observation matches the record. An inhabited fiber expresses realizability; a subterminal fiber expresses at most one coherent realization; a contractible fiber expresses exactly one. A subterminal object has its map to the terminal object a monomorphism. It need not be globally inhabited.

Definition 2.2 (full fiber observability). The observer is fully fiber-observable if o is a monomorphism, equivalently a (-1) -truncated map. Its image factorization is

$$X \xrightarrow{e} I \xrightarrow{j} Y,$$

where e is an effective epimorphism and j is a monomorphism. The object I records supported observations. An effective epimorphism is a map whose augmented Čech nerve has its codomain as colimit.

Proposition 2.3 (unique reconstruction on the image). The following are equivalent: (i) o is fully fiber-observable; (ii) its diagonal $X \rightarrow X \times_Y X$ is an equivalence; (iii) every generalized fiber is subterminal; (iv) $e : X \rightarrow I$ is an equivalence.

Proof. The diagonal criterion is the definition of a monomorphism in an infinity-category with pullbacks. Base change gives the equivalence with subterminal fibers, including the test $T = Y$. If $o = je$ is mono, then e is mono; an effective epimorphism that is mono is an equivalence. Conversely an equivalence followed by j is mono. $\square$

Consequently reconstruction from I is unique up to a contractible space of choices. Records outside the image need not have realizations. For sets this is ordinary injectivity. For spaces it is stronger than injectivity on connected components: compatibility witnesses are part of the chosen semantics.

Example 2.4 (hidden symmetry). Let G be a nontrivial discrete group and BG its one-object classifying groupoid. The map $BG \rightarrow *$ has one connected component in its fiber, but its fiber has loop group G and is not subterminal. Conversely $* \rightarrow BG$ has one underlying state, yet its homotopy fiber at the basepoint is G : there are multiple compatibility witnesses. These examples prevent interpreting a homotopy fiber as only a set of state labels.

3. Probe refinement and a truncation barrier

An additional probe $a : X \rightarrow Z$ over B gives the joint observer

$$Q = (o, a) : X \longrightarrow Y \times_B Z.$$

Theorem 3.1 (residual fiber under a new probe). For compatible records $y : T \rightarrow Y$ and $z : T \rightarrow Z$, let $Z_T = Z \times_B T$ and let $a_y : F_y \rightarrow Z_T$ be induced by a . Then

$$\text{fib}_{(y,z)}(Q) \simeq F_y \times_{Z_T} T_z.$$

Here $T_z \rightarrow Z_T$ is the section given by z . Joint observation is fully fiber-observable exactly when these residual fibers are subterminal for all generalized records.

Proof. Both sides are limits of the diagram imposing $o(x) = y$ and $a(x) = z$ over their common context. Associativity of pullbacks identifies the limits, with all compatibility data retained. The final assertion is Proposition 2.3. $\square$

The theorem describes which ambiguity a probe removes. Adding a probe can resolve states, preserve a residual symmetry, or reveal that a proposed joint record has no realization.

We next bound probe resolution. For $n \geq 0$, an object is **n -truncated** when every mapping space into it has no homotopy groups above degree n . A morphism is n -truncated when it is an n -truncated object in its codomain slice. Relative truncation $\tau_{\leq n}^Y$ is the left adjoint to inclusion of these objects in $\mathcal{E}_Y$ [3, Proposition 5.5.6.18].

Theorem 3.2 (probe-resolution barrier). Regard $o : X \rightarrow Y$ as an object in $\mathcal{E}_Y$. Let $r_i : X \rightarrow Z_i$ be a small family of probes over Y , where every $Z_i \rightarrow Y$ is n -truncated. The joint probe factors, uniquely up to contractible choice, as

$$X \longrightarrow \tau_{\leq n}^Y X \longrightarrow \prod_{i,Y} Z_i.$$

If the joint probe is a monomorphism, then $X \rightarrow Y$ is n -truncated. Conversely, if $X \rightarrow Y$ is n -truncated, the unrestricted class of such probes contains a fully observing probe.

Proof. The product is n -truncated, since truncated objects are closed under limits. The truncation adjunction gives the factorization and its uniqueness. If $r : X \rightarrow \prod_{i,Y} Z_i$ is mono, its composite with the product's structure map is n -truncated: composition of a (-1) -truncated map and an n -truncated map is n -truncated. This composite is o . For the converse choose $Z = X$ and $r = \text{id}_X$. $\square$

The converse is an expressiveness statement, not an implementable sensor construction. In spaces, $\tau_{\leq 0} BG \simeq *$, so every set-valued probe on a fixed BG is constant on its internal symmetry. More generally $K(A, n+1)$, for a nonzero abelian group A , is invisible beyond its terminal n -truncation to such probes. A scalar invariant defined on a different moduli space may distinguish components representing different groups; that does not contradict the theorem about a specified domain.

4. Adequate tests, locality, and descent

Global states alone need not detect observability in a topos. Call a family $\mathcal{G}$ of test objects **conservative** when the functors $\text{Map}(G, -)$, $G \in \mathcal{G}$, jointly detect equivalences.

Theorem 4.1 (test adequacy). For a conservative family, $o : X \rightarrow Y$ is mono if and only if

$$\text{Map}(G, X) \longrightarrow \text{Map}(G, Y)$$

is a monomorphism of spaces for every $G \in \mathcal{G}$.

Proof. Mapping out of G preserves limits, so a mono gives a mono. Conversely these conditions make the images of the diagonal $X \rightarrow X \times_Y X$ equivalences. Conservativity makes the diagonal an equivalence. $\square$

The test uses whole mapping spaces. Checking injectivity of their π_0 alone fails even in spaces, as $BG \rightarrow *$ shows.

Theorem 4.2 (locality). If $u : U \rightarrow Y$ is an effective epimorphism, then $o : X \rightarrow Y$ is mono if and only if $o_U : X \times_Y U \rightarrow U$ is mono.

Proof. One direction is stability under pullback. For the other, the pullback of the diagonal of o is the diagonal of o_U , hence an equivalence. Pullback along an effective epimorphism is conservative: equivalences descend along its Čech nerve [3, Lemma 6.2.3.16]. Therefore the original diagonal is an equivalence. $\square$

Locality of this property must be distinguished from the existence of chosen global states. In the infinity-topos of sheaves of spaces on S^1 , let P be the sheaf of local sections of the double cover $z \mapsto z^2$. The map $P \rightarrow 1$ is locally a two-element set over 1, hence effective epi and not mono. It has no global section: such a section $s : S^1 \rightarrow S^1$ would satisfy $2 \deg(s) = 1$. Nevertheless the global-sections map $\emptyset \rightarrow \{*\}$ is injective. Global sections therefore miss the failure of observability and local choices need not glue.

5. Dynamic observation and finite horizons

Let X be a nonempty finite set of size N , A a finite input alphabet, $\delta_a : X \rightarrow X$ transitions, and $h : X \rightarrow O$ an output function. For a word $w = a_1 \cdots a_k$, put $\delta_w = \delta_{a_k} \cdots \delta_{a_1}$, with $\delta_\emptyset = \text{id}$. Define

$$Q_L(x) = (h(\delta_w x))_{|w| \leq L}, \quad R_L = \ker(Q_L).$$

The kernel here is an ordinary equivalence relation. The protocol assumes access to the indicated responses from the same initial state, for example by resetting the system. It is not automatically a single unresettable experiment.

Theorem 5.1 (finite stabilization). The relations satisfy

$$R_{L+1} = R_0 \cap \bigcap_{a \in A} (\delta_a \times \delta_a)^{-1}(R_L).$$

They stabilize by $L = N - 1$. The limit R_∞ is the largest transition congruence contained in $\ker(h)$. All states are distinguishable by finite words exactly when Q_{N-1} is injective.

Proof. Every nonempty word of length at most $L + 1$ is a first input followed by a word of length at most L , giving the recurrence. The sequence refines partitions. Each strict refinement increases the number of blocks, which is at most N . If one step is equal, the recurrence makes all later steps equal; thus there cannot be N strict steps. At stability the relation lies in $\ker(h)$ and is preserved by every transition. Conversely any such congruence identifies states with equal output after every word, by induction, so is contained in every R_L . The last assertion follows. $\square$

This is the finite deterministic specialization of behavioral observability [2]. For $X = \{0, 1\}^2$, $h(a, b) = a$, and swap transition $\delta(a, b) = (b, a)$, Q_0 has two states per fiber while Q_1 distinguishes all four. With the identity transition, the second bit remains hidden forever.

Finite-state stabilization has no unrestricted infinite-state analogue. On $X = \{0, 1\}^\mathbb{N}$, all coordinate probes jointly determine a sequence, while any finite subfamily leaves infinitely many coordinates free. These probes are continuous for the compact Cantor topology; compactness of the state space does not force a finite separating subfamily.

6. Noisy probes and identification from laws

Let X, Y be finite sets. A stochastic observer is a channel $K : X \rightsquigarrow Y$, assigning a probability vector K_x to each state. Its **law observer** is $\kappa : X \rightarrow \Delta(Y)$, $\kappa(x) = K_x$, where $\Delta(Y)$ is regarded as a set of laws, not a topological simplex. Equal-law fibers differ from fibers of a realized sample. Markov categories describe channel composition [6,7]; the quantitative assumptions below are additional.

Write $\text{TV}(P, Q) = \frac{1}{2} \sum_y |P(y) - Q(y)|$.

Theorem 6.1 (testing obstruction and garbling). For two states with laws P, Q , any randomized test based on m conditionally independent samples has equal-prior error at least

$$\frac{1 - \text{TV}(P^{\otimes m}, Q^{\otimes m})}{2}.$$

Equality is achieved by a likelihood-ratio test. Consequently equal-law states cannot be uniformly identified with error below $1/2$ by any finite number of these samples. Postprocessing by another channel cannot remove this obstruction.

Proof. At a sample string s , deciding for either hypothesis contributes at least $\frac{1}{2} \min(P^{\otimes m}(s), Q^{\otimes m}(s))$ to the error; randomized decisions cannot improve a pointwise minimum. Summing uses $\sum_s \min(P(s), Q(s)) = 1 - \text{TV}(P, Q)$. Choose the more likely hypothesis at each string to attain equality. If $P = Q$, their powers coincide. For a postprocessing channel M , the triangle inequality gives

$$\frac{1}{2} \sum_z \left| \sum_y (P(y) - Q(y)) M(z | y) \right| \leq \frac{1}{2} \sum_y |P(y) - Q(y)|.$$

Thus total variation cannot increase, and equal laws remain equal. $\square$

This elementary obstruction is consistent with the comparison-of-experiments viewpoint [5,7]. It concerns the declared experiment; a new intervention or a new channel may distinguish states that this channel does not.

Theorem 6.2 (finite confidence bound). Suppose $|X| \geq 2$, $M = |Y| < \infty$, the channel is known, and

$$\Delta = \min_{x \neq x'} \|K_x - K_{x'}\|_\infty > 0.$$

From m independent samples, let $\hat{p}$ be the empirical distribution and choose a state minimizing $\|\hat{p} - K_x\|_\infty$, with a fixed rule for ties. Uniformly over the true state,

$$\Pr(\hat{x} \neq x) \leq 2M \exp(-m\Delta^2/2).$$

Hence $m \geq (2/\Delta^2) \log(2M/\delta)$ suffices for confidence $1 - \delta$.

Proof. If $\|\hat{p} - K_x\|_\infty < \Delta/2$, the triangle inequality makes every other law farther from $\hat{p}$. A union bound over the M coordinates, using Hoeffding's inequality [8] for their indicator variables, bounds the complement by $2Me^{-2m(\Delta/2)^2}$. $\square$

This is a finite identification guarantee for known, separated laws. For Bernoulli parameters $1/4$ and $3/4$, $M = 2$ and $\Delta = 1/2$; $m = 36$ suffices for $\delta = 0.05$ by this bound. Homotopical observability alone does not give a separation margin, independent samples, or a statistical rate.

7. Task recovery and coherent descent

Reconstructing a whole state may exceed what a task requires. For sets and $I = \text{im}(o)$, a task $g : X \rightarrow Z$ is recoverable from records if $g = \bar{g}e$ for some $\bar{g} : I \rightarrow Z$.

Proposition 7.1 (set-valued task criterion). Such a $\bar{g}$ exists exactly when g is constant on every fiber of e ; if it exists it is unique.

Proof. A factorization is constant on fibers. Conversely define $\bar{g}(y)$ as the common value of g on the nonempty fiber $e^{-1}(y)$. This value is uniquely determined, so no arbitrary representative choice enters the definition. Surjectivity gives uniqueness. $\square$

In an infinity-topos, let $X_\bullet$ be the Čech nerve of the effective epi $e : X \rightarrow I$, so $X_k = X \times_I \cdots \times_I X$ with $k + 1$ factors.

Theorem 7.2 (coherent task recovery). For every Z ,

$$\text{Map}(I, Z) \simeq \lim_{[k] \in \Delta} \text{Map}(X_k, Z).$$

Here Δ is the simplex category. A task $g : X \rightarrow Z$ extends precisely when it admits the descent data on the right with degree-zero component g . Its reconstruction space is

$$\text{Rec}_e(g) = \text{fib}_g(\text{Map}(I, Z) \rightarrow \text{Map}(X, Z)).$$

Proof. Effective epimorphy identifies I with the colimit of $X_\bullet$. Mapping this colimit into Z gives the stated limit. Taking the fiber over the prescribed degree-zero component gives exactly the extensions with their compatibility witness. $\square$

The descent data include a homotopy between the two pullbacks of g to X_1 and compatible higher homotopies. Once complete descent data are specified, the corresponding realization is unique up to contractible choice. Specifying g alone need not specify those data.

For example, $e : * \rightarrow BG$ is an effective epi when G is a discrete group. The basepoint task $g : * \rightarrow BG$ admits both the identity extension $BG \rightarrow BG$ and the constant extension. For nontrivial G they differ on π_1 . Thus effective epimorphy alone does not make a higher-valued extension unique. Conversely the task id_{BG} does not factor through $BG \rightarrow *$: such a factorization would kill its nontrivial fundamental group.

For a binary task on finite stochastic states, if an equal-law fiber contains different task labels, Theorem 6.1 gives the same $1/2$ worst-case error obstruction for that task. Exact state reconstruction can therefore fail while a task constant on the law fibers remains identifiable.

8. Interpretation and limits

The framework specifies four obligations: a context and observation type; residual coherent fibers; an admissible probe or intervention family; and, for statistical claims, an experiment with quantitative assumptions. Full fiber observability is relative to those choices. Truncating the chosen equality can

legitimately weaken the reconstruction demand, but the discarded symmetries then lie outside that demand.

For neural representations or distributed agents, an observation might record predictions, traces, or local compatibility tests. The results apply after these are represented by the stated maps or finite channels. They do not imply that architecture, state dimension, or a list of scalar scores determines an observation fiber without constructing the observation map. Finite simulations illustrate the discrete statements; they do not verify the general infinity-topos results.

Further work can add costs, adaptive interventions, effective presentations, or quantitative coherence bounds to this specification.

Appendix A. Exact witnesses and reproducibility

The companion program `verify_examples.py` uses only the Python 3 standard library. Run `python verify_examples.py` to regenerate `example-results.json`. It checks finite instances of the dynamical theorem using two independent descriptions: explicit finite-word output signatures and reachability in an automaton of pairs of states. The latter searches all reachable pairs for an output disagreement, without assuming a finite word horizon.

The exhaustive enumeration covers all binary-output systems on one, two, or three states with two input letters, and all four-state binary-output systems with one input letter: **9994 machines and 118,282 ordered state pairs**. Finite-word distinguishability at horizon $N - 1$ agreed with the pair-search oracle in every case. The swap and identity-transition systems reproduce the hidden-bit examples in Section 5.

Deterministic observer in spaces or sets	Residual observation fiber
$\{0, 1\}^2 \rightarrow \{0, 1\}, (a, b) \mapsto a$	Two states for each supported record
Swap system, $(a, b) \mapsto (h(x), h(\delta x))$	A singleton for each supported record
$BG \rightarrow *$ with nontrivial G	BG , including its loop group G
Basepoint $* \rightarrow BG$	$\Omega BG \simeq G$, recording compatibility witnesses

For $G = C_n$, $2 \leq n \leq 5$, the program enumerates group endomorphisms and distinguishes the identity from the trivial homomorphism. This is a finite algebraic witness for the different extensions in Section 7; it is not a computation of a general infinity-categorical mapping space.

For Bernoulli observations, exact rational binomial sums give the optimal equal-prior error with random tie breaking. For sample counts one through ten, a separate enumeration of all binary sample strings computes total variation of the product laws and confirms the testing identity in Theorem 6.1. Equal laws give error 1/2 in all these cases. At $m = 36$, the exact error for parameters 1/4 and 3/4 is

$$\frac{205333278708616879}{295147905179352825856} \approx 0.0006956962,$$

while the sufficient confidence bound is $4e^{-36/8} \approx 0.044436 < 0.05$. The confidence bound is conservative; exact finite risk and an upper bound have different purposes.

All finite probability sums are rational. The displayed exponential comparison and confidence sample count use floating-point elementary functions. The computations provide inspectable examples and implementation checks. The general truncation and descent arguments are justified by the mathematical proofs and cited foundations, rather than by finite enumeration.

References

- [1] Robert Hermann and Arthur J. Krener. *Nonlinear controllability and observability*. IEEE Transactions on Automatic Control 22(5):728-740, 1977. [Author-hosted paper](#). DOI: 10.1109/TAC.1977.1101601.
- [2] J. J. M. M. Rutten. *Universal coalgebra: a theory of systems*. Theoretical Computer Science 249(1):3-80, 2000. [CWI publication and full text](#).
- [3] Jacob Lurie. *Higher Topos Theory*. Annals of Mathematics Studies 170, Princeton University Press, 2009; maintained manuscript. [Author's text](#). See Sections 5.5.6, 6.1, and 6.2.3 for truncation, descent, and effective epimorphisms.
- [4] The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study, 2013. [Official book and maintained editions](#), Chapters 4 and 7.
- [5] David Blackwell. *Equivalent comparisons of experiments*. Annals of Mathematical Statistics 24(2):265-272, 1953. [Publication](#). DOI: 10.1214/aoms/1177729032.
- [6] Tobias Fritz. *A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics*. Advances in Mathematics 370:107239, 2020. [Primary preprint and publication metadata](#).
- [7] Tobias Fritz, Tomas Gonda, Paolo Perrone, and Eigil Fjeldgren Rischel. *Representable Markov Categories and Comparison of Statistical Experiments in Categorical Probability*. Theoretical Computer Science 961:113896, 2023. [Primary preprint](#). DOI: 10.1016/j.tcs.2023.113896.
- [8] Wassily Hoeffding. *Probability inequalities for sums of bounded random variables*. Journal of the American Statistical Association 58(301):13-30, 1963. [Publisher record](#). DOI: 10.1080/01621459.1963.10500830.

Authorship and evidence. Anonymous author. Prepared with AI assistance. Proofs are mathematical arguments from the stated assumptions and cited standard foundations; they are not proof-assistant certified. The accompanying code checks finite examples with exact arithmetic. No empirical validation or priority claim for the standard constituent theorems is asserted.