RESEARCH MANUSCRIPT

Semantic Uniqueness and the Moduli of Machine Realizations

Truncation, trace semantics, and coherent recursive improvement

Anonymous manuscript

Abstract. An operational semantics into an ∞-topos induces a tower of observations by Postnikov truncation. We model machines equipped with an identification of their observation as a homotopy fiber, whose components and loop spaces record implementation choices and symmetries. We prove criteria for non-uniqueness under hidden decorations, for coherent factorization through traces, and for finality of timelines over groupoidal semantics. We establish when a semantic orbit lifts to a certified machine trajectory and when finite-horizon bounds imply a uniform budget. Deterministic transducers, local systems, and Postnikov extensions illustrate the hypotheses. The paper expands the source note and situates its claims in prior work on coalgebra, operational semantics, and descent.

Keywords. Machine semantics, ∞-topos, Postnikov truncation, homotopy fiber, traces, final functor, recursive improvement.

Prepared for *Higher Categorical Semantics*.

Contents

1 Introduction **5**
1.1 Four questions about uniqueness 5
1.2 Why conservativity alone gives a different answer 6
1.3 Results and organization 7
1.4 Scope, evidence, and conventions 8

2 Classes, markings, and transformation categories **9**
2.1 Equivalences in the marked fiber 10
2.2 Three reflection properties 11
2.3 A group calculation with interface names 12
2.4 Universal representatives and localizations 13

3 Truncation and hidden operational structure **14**
3.1 Examples at the boundary degrees 14
3.2 Hidden decorations and sharp separation 15
3.3 Relative truncation retains the indexing interface 16
3.4 Higher attachments can distinguish machines with the same groups . 17
3.5 All finite stages 17

4 What a finite specification determines **19**
4.1 Determination by soundness and categoricity 19
4.2 Tests can underdetermine even discrete behavior 20
4.3 Finite rules can determine unbounded behavior 21
4.4 Parametrized admissibility and changing specification 21

5 Trace objects, semantic descent, and contextual tests **22**
5.1 A criterion from localization 23
5.2 Tests and the coarsest separating quotient 24
5.3 A branching distinction lost by trace sets 25
5.4 A site of observations 25

6 Deterministic transducers and canonical quotients 26
6.1 Machines, derivatives, and the final behavior object 27
6.2 Reachability and minimality 28
6.3 A family with a single constant behavior 29
6.4 Passing to a behavioral localization 30
6.5 A comparison with nondeterministic semantics 30

7 Families, descent, and local realizability 31
7.1 Descent of marked fibers 31
7.2 A discrete torsor gives local but no global marked realizer 33
7.3 Descent and relative Postnikov data 34

8 Coverage, persistence, and finality of timelines 34
8.1 Hitting and preserving a class 34
8.2 The categorical finality condition 35
8.3 Exact criterion over groupoidal semantics 36
8.4 Finality inside a realization category 37
8.5 Simultaneous requirements and directedness 37

9 Recursive operations and coherent machine dynamics 38
9.1 Typing the operation catalog 38
9.2 Three levels of realization of an orbit 39
9.3 Forward transport and iteration 40
9.4 Commuting update policies and semantic invariance 41
9.5 Periods and equivariance 42
9.6 Equality and improvement 43

10 Feasible improvement, viability, and uniform budgets 43
10.1 Local optimality and feasible continuation 43
10.2 A greatest fixed point for extendible machines 44
10.3 Attainment of a uniform path budget 46
10.4 Computability of paths and observation 47
10.5 An explicit constrained redesign 48

11 Combining observation levels and redesign constraints 49
11.1 A family with two independent machine invariants 49
11.2 Coherent recovery of machine fibers 50
11.3 A joint compactness theorem for observation and time 52
11.4 What a model must supply 53
11.5 Conclusion 54

A Calculus of observation refinements 54
A.1 Decomposing a coarse fiber over its refinements 54
A.2 Interpreting compatible tests by Yoneda 56
A.3 An example with higher marking homotopies 57
A.4 Invariance under a change of model 58

B Further comparisons for local and directed choices 58
B.1 Obstruction by a central extension 58
B.2 A transport criterion in ordinary categories 59
B.3 Two restrictions with different closure properties 60

C Compositional observation and a feedback obstruction 61
C.1 Agreement for wiring along an observed interface 61
C.2 A hidden interface becomes visible after matching 62
C.3 Interface contracts and categories of tests 63

D Source correspondence and notes on prior work 64
D.1 Correspondence with the source note 64
D.2 Higher categories, topoi, and homotopy theory 65
D.3 Localization, descent, and finality 66
D.4 Operational semantics and certified updates 67

References 68

1 Introduction

Two programs can agree on every output exposed by an interface and differ in the state they maintain, the transformations they admit, or the histories through which they produce those outputs. An interface may leave these distinctions invisible deliberately. An engineering equivalence can then be exact on the interface and still leave a substantial choice of implementations. This paper asks how to express that choice when operational objects are higher types and observations discard homotopy above a specified degree. We also ask which extra hypotheses let a family of implementations realize a semantic trajectory.

The starting note, *Semantic Uniqueness and Non-Uniqueness of Machine Realizations* [1], proposes an operational functor $\mathrm{Op} : \mathcal{C} \to \mathcal{X}$ from a category of realizations to an ∞-topos. At degree n it compares machines using $Q_n = \tau_{\leq n} \mathrm{Op}$. It then connects this comparison to specifications, traces, the coverage of timelines, and recursively updated operations. We retain these themes and refine the statements that connect them. The source's numbered sections are discussed individually in Appendix D. We use its text as a point of departure and provide independent definitions and proofs.

1.1 Four questions about uniqueness

Fix a functor $Q : \mathcal{C} \to \mathcal{D}$. There are at least four different questions one could ask about an object $p \in \mathcal{D}$. Does any machine have image equivalent to p? Are any two such machines equivalent? Given specified identifications of their images with p, is there a compatible equivalence between the machines? If such an equivalence exists, how many choices and higher homotopies does it have?

The first two questions concern the map of sets

$$\pi_0(\mathcal{C}^{\simeq}) \longrightarrow \pi_0(\mathcal{D}^{\simeq}).$$

Here the core keeps the objects and equivalences of an ∞-category, so the notation makes the intended isomorphism classes explicit. The last two questions concern the homotopy pullback

$$\mathrm{Real}_Q(p) = \mathcal{C}^{\simeq} \times_{\mathcal{D}^{\simeq}} \{p\}.$$

An object of this space is a machine M with a marking $QM \simeq p$. Retaining a marking is essential when an interface has named ports, indexed actions, or other structure with its own symmetries. It is possible to have a single machine class and several inequivalent markings. It is also possible for every marking to lie in the same component while the realization space has nontrivial loops.

A fifth question enters once one allows irreversible steps. The full category on the realizers can have useful projections, quotients, rewrites, or certified updates. Its terminal objects, when they exist, may provide universal representatives. A terminal representative in that category does not settle the equivalence class of the machines. A two-state toggle that emits zero and a one-state constant-zero transducer have a unique quotient map from the former to the latter. Their reachable state sets have different cardinalities. This example will serve as a concrete check on claims about minimality, contractibility, and directed timelines.

1.2 Why conservativity alone gives a different answer

A conservative functor detects whether a given arrow is an equivalence. If the target supplies an equivalence between two images, there may be no arrow in the source to which conservativity applies. A discrete category of two machines sent to a point provides the simplest example. It is also possible for a nonconservative functor to be injective on machine classes. An endomorphism of the only machine can become invertible while the map on objects has only one possible source class. These observations change the scope of the source note's non-uniqueness claim. We formulate one additional lifting hypothesis under which conservativity does yield object uniqueness, and a stronger criterion on equivalence spaces that detects marked uniqueness.

This distinction persists for truncation itself. The reflector $\tau_{\leq n} : \mathcal{X} \to \mathcal{X}_{\leq n}$ is a localization. Its composite with a general operational functor need not exhibit the target as a localization of the category of machines. For example, send a point to a one-point set and take degree zero. The only arrows available to invert are identities, so the resulting machine category cannot be all sets. Even if the semantic target is restricted to the image on objects, the target can carry morphisms absent from

the localization. We therefore specify whenever we invoke a localization universal property.

1.3 Results and organization

Section 2 compares the unmarked and marked classifications by calculating mapping spaces in a fiber. For a group homomorphism $G \to H$, the fiber of $BG \to BH$ has components indexed by cosets of the image, with a copy of the classifying space of the kernel on every component. This elementary calculation distinguishes lost symmetries from missing changes of marking. It also gives a model for contextual gluing later in the paper.

Section 3 proves a sufficient condition for non-uniqueness with a controlled first observable difference. A category with independent decorations can have pairs with identical observations up to n and a difference at $n+1$. We give a universal construction in spaces, a version for classes of operational objects closed under decoration, and a relative version for families in slices. We also study reconstruction with all finite degrees. Detecting a given arrow, choosing compatible equivalences of finite stages, and making a tower effective in a specified topos require distinct hypotheses.

Sections 4 to 6 move from abstract moduli to semantics defined by tests. A finite presentation can characterize one semantic class through a universal property, or it can leave an entire family of admissible semantic objects. We identify the hypothesis that separates those situations. For traces, we state a localization criterion for coherent factorization and explain the role of contexts and nondeterminism. Deterministic initialized Mealy machines give a complete worked example. Their reachable observable quotient is terminal among all initialized machines of the behavior with strict homomorphisms. The same behavioral class contains many inequivalent reachable machines, and the proof specifies what their replacement maps preserve.

Section 7 treats local families using sheaf descent, and exhibits a specification with local realizers but no global marked realizer. Section 8 distinguishes hitting a class, remaining in it, eventual access to every implementation, and finality of a functor. For a groupoidal semantic target, the classwise comma categories are modeled by marked fibers of the timeline. This computes the actual condition for finality. Directedness

can help promote a single hit to persistence when an invariant is preserved by all subsequent steps.

Sections 9 and 10 separate iteration of operations from iteration of certified machine changes. We give a positive lifting criterion using coCartesian transport and elementary obstructions using directed graphs, monodromy, and cost bounds. For a fixed orbit, viability gives the exact set of states from which a legal continuation exists in the set-based model. Finite branching turns arbitrarily long continuations from a specified machine into an infinite one, although a computable continuation requires more information. We prove an analogous compactness result for attainable uniform budgets along an infinite execution.

1.4 Scope, evidence, and conventions

The machine category and its equivalences are chosen model data. A morphism of initialized machines may be required to preserve transitions, port names, states, or an explicit refinement certificate. When two such choices appear we keep separate notation. An equivalence in a category of spaces does not automatically establish contextual equivalence of programs. Conversely a contextual equivalence relation may be modeled faithfully by discrete observations without introducing higher homotopy. Higher types enter the present analysis through the operational model, the semantic symmetries, or the ambient category of families.

We use ∞-categories in the sense of categories with mapping spaces and composition coherent up to higher homotopy. Results about cores, truncation, localization, descent, and finality use their standard ∞-categorical formulations [17, 18, 26]. Pullbacks and limits in spaces and ∞-categories are homotopy-coherent. In ordinary set-based examples they are ordinary limits unless a classifying space is explicitly taken. Write $[p]$ for a point of $\pi_0(\mathcal{D}^{\simeq})$, and p for the chosen representative. Thus $Q(M) \simeq p$ denotes equivalence of semantic objects. An element of $\mathrm{Eq}_{\mathcal{D}}(QM, p)$ is a specified witness. We write BG for the delooping of a discrete group G, and $K(A, k)$ for an Eilenberg–Mac Lane space when A and k are of the indicated kind. We choose universes large enough for cores of the categories under study, and use small indexing categories in statements about limits and colimits.

The results are mathematical. Constructed finite machines illustrate theorems. The evidence is a proof in each case. There is no empirical claim about the frequency of particular fibers in deployed agents. The term *improvement* is used for an update rule and, when supplied, an order or cost condition. The existence of a semantic endomorphism alone says nothing about performance. The manuscript was prepared with generative AI assistance in drafting, mathematical development, and preparation of references. The submission form provides the place to identify the responsible author or research group.

2 Classes, markings, and transformation categories

Let $Q : \mathcal{C} \to \mathcal{D}$ be a functor. In applications $\mathcal{D} = \mathcal{X}_{\leq n}$ and $Q = Q_n$, although the present results use neither a topos nor a truncation. A core is denoted by a superscript $\simeq$. The following three constructions will all be useful:

$$U_Q(p) = \{[M] \in \pi_0(\mathcal{C}^{\simeq}) : [QM] = [p]\}, \tag{1}$$

$$\mathrm{Real}_Q(p) = \mathcal{C}^{\simeq} \times_{\mathcal{D}^{\simeq}} \{p\}, \tag{2}$$

$$\mathcal{T}_Q(p) = \mathcal{C} \times_{\mathcal{D}} \{p\}. \tag{3}$$

The last expression is a limit in Cat_∞. It includes arrows f with a triangle commuting up to a specified homotopy after applying Q. Every arrow in $\mathcal{T}_Q(p)$ lies over a semantic equivalence. The maximal subgroupoid of $\mathcal{T}_Q(p)$ is $\mathrm{Real}_Q(p)$, since taking the core is a right adjoint to the inclusion of spaces and preserves limits. One can instead form the full subcategory $\mathcal{C}[p] \subset \mathcal{C}$ on all M in (1). Its arrows are unconstrained. For example, a noninvertible semantic endomorphism of p may be the image of an arrow in $\mathcal{C}[p]$ and is then excluded from (3). In a model of machine upgrades it is often appropriate to use only a specified subcategory of $\mathcal{T}_Q(p)$.

2.1 Equivalences in the marked fiber

Proposition 2.1 (Compatible equivalence space)**.** *For $x = (M, \alpha)$, $y = (N, \beta)$ in $\mathrm{Real}_Q(p)$, there is an equivalence*

$$\mathrm{Map}_{\mathrm{Real}_Q(p)}(x, y) \simeq \mathrm{hofib}_{\beta^{-1}\alpha}\Big(\mathrm{Eq}_{\mathcal{C}}(M, N) \longrightarrow \mathrm{Eq}_{\mathcal{D}}(QM, QN)\Big).$$

In particular, the based loop space at x is the homotopy fiber of $\mathrm{Aut}_{\mathcal{C}}(M) \to \mathrm{Aut}_{\mathcal{D}}(QM)$ at the identity.

Proof. The mapping space of a pullback of spaces is the pullback of the mapping spaces at the corresponding endpoints, with the identifications of endpoints specified by the pullback data. A path from M to N in $\mathcal{C}^{\simeq}$ is an equivalence $u : M \simeq N$. For this path to connect (M, α) to (N, β), its image must come with a path from $\beta Q(u)$ to α in $\mathrm{Eq}_{\mathcal{D}}(QM, p)$. Composition with β^{-1} is an equivalence of these path spaces. This identifies the pullback as the stated homotopy fiber. For $M = N$ and $\beta = \alpha$, the distinguished element is the identity up to the canonical cancellation path. □

Corollary 2.2 (Test for marked uniqueness)**.** *Suppose a marked realizer $x = (M, \alpha)$ exists. Then $\mathrm{Real}_Q(p)$ is contractible exactly when it is connected and $\mathrm{Aut}_{\mathcal{C}}(M) \to \mathrm{Aut}_{\mathcal{D}}(QM)$ has contractible homotopy fiber at the identity. Equivalently, contractibility holds if, for every $y \in \mathrm{Real}_Q(p)$, the compatible equivalence space in Proposition 2.1 is contractible.*

Proof. A pointed space is contractible when it is connected and its based loop space is contractible. Apply Proposition 2.1. A contractible space has contractible path space between every two points. Conversely, contractible path spaces from x to every point make the map $x : * \to \mathrm{Real}_Q(p)$ fully faithful and essentially surjective as a map of groupoids. □

The loop-space condition in Corollary 2.2 controls only the homotopy fiber at the identity. In the presence of connectedness of the marked fiber it can be expressed more strongly. Any semantic automorphism γ of QM gives another marking $\alpha\gamma$. A path between (M, α) and $(M, \alpha\gamma)$ gives a source automorphism mapping into the component of γ^{-1}. Thus the automorphism map is surjective on components.

Translation by a source automorphism then identifies the fiber over any component with the identity fiber. Consequently a connected marked fiber is contractible exactly when the entire map of automorphism spaces at the chosen machine is an equivalence. This proof uses the ability to translate within a group-like space, and avoids any assumption that the automorphism spaces are discrete.

More generally, $\pi_0 \operatorname{Real}_Q(p) \to U_Q(p)$ is surjective. Over a fixed class $[M]$, its fiber as a map of sets is the orbit set of the action of $\pi_0 \operatorname{Aut}_{\mathcal{C}}(M)$ on $\pi_0 \operatorname{Eq}_{\mathcal{D}}(QM, p)$ by precomposition through Q. To see this, move any marked machine in the class to M using a machine equivalence. Different choices of move differ by an automorphism, and Proposition 2.1 identifies the resulting orbits. Before taking components, the portion over $[M]$ is the homotopy quotient of $\operatorname{Eq}_{\mathcal{D}}(QM, p)$ by $\operatorname{Aut}_{\mathcal{C}}(M)$. This follows by modeling the connected component of M in $\mathcal{C}^{\simeq}$ as $B \operatorname{Aut}_{\mathcal{C}}(M)$ and pulling back the path fibration with target p. In particular, a fiber can be connected over a given machine class while its automorphisms have higher homotopy.

2.2 Three reflection properties

Definition 2.3. Call Q *injective on equivalence classes* when $\pi_0(\mathcal{C}^{\simeq}) \to \pi_0(\mathcal{D}^{\simeq})$ is injective. Call Q *conservative* when $Q(f)$ an equivalence implies f an equivalence. Call Q *fully faithful on cores* when the map of spaces $Q^{\simeq}$ is fully faithful as a functor of ∞-groupoids.

Proposition 2.4. *The following implications and separations hold.*

1. *Full faithfulness on cores is equivalent to every marked fiber being empty or contractible. It implies injectivity on equivalence classes.*

2. *Conservativity and injectivity on equivalence classes are independent. Full faithfulness on cores does not imply conservativity.*

3. *If every target equivalence $QM \simeq QN$ can be represented up to homotopy by the image of a machine arrow $M \to N$, conservativity implies injectivity on equivalence classes. Fullness on the homotopy categories is sufficient for this lifting premise.*

Proof. A fully faithful map of groupoids is an equivalence onto a union of target components, since any object equivalent to one in its image lies in its essential image, and maps between the image objects agree. Its fibers are thus empty or contractible. Conversely, a map of spaces with empty or contractible fibers is an equivalence onto the union of components on which the fibers are inhabited. In particular it is fully faithful, and its map on components is injective.

The map from a discrete two-object category to the terminal category is conservative and sends distinct classes to the same class. Let $B\mathbb{N}$ here denote the ordinary one-object category of the additive monoid $\mathbb{N}$ of nonnegative integers. Its functor to the terminal category is fully faithful on cores because its only invertible endomorphism is 0, but the arrow 1 becomes invertible. This supplies both remaining counterexamples. Finally, apply conservativity to an arrow whose image is homotopic to the proposed equivalence, using that invertibility is a property of the homotopy class. □

There is a limited conclusion that can always be drawn from a witness of nonconservativity. If $f : M \to N$ becomes an equivalence, the two marked machines obtained using $Q(f)$ and id_{QN} are connected by an arrow in $\mathcal{T}_Q(QN)$. That arrow is noninvertible. One needs a further argument to decide whether some other arrow is a compatible equivalence, or even whether the machines are unmarked-equivalent. For instance, a category with a single object and the monoid $\mathbb{N}$ has a terminal core and a nontrivial transformation category. Nonconservativity witnesses a collapse of arrows, which may occur without any collapse of object classes.

2.3 A group calculation with interface names

Proposition 2.5 (Fiber of a homomorphism of symmetry groups)**.** *For a homomorphism $\phi : G \to H$ of discrete groups, the marked fiber of $B\phi : BG \to BH$ is the action groupoid for G acting on the set H by right multiplication through ϕ. Its components are $H/\phi(G)$, and the component of each marking has homotopy type $B\ker(\phi)$.*

Proof. There is one machine object. A marking is an element of H. A machine arrow g relates h to h' when $h'\phi(g) = h$. This is the action groupoid, using inverse arrows if one writes a right action as $h \mapsto h\phi(g)$. Its orbits are the cosets $h\phi(G)$, and the

automorphisms of h are those g with $\phi(g) = 1$. Choosing a representative in each orbit gives the asserted equivalence component by component. □

If ϕ is injective with proper image, markings can be inequivalent while the underlying machine is unique. If ϕ is onto with nontrivial kernel, there is a unique marked class but the marked symmetry group is the kernel. Both phenomena occur together for the map $C_2 \times C_2 \to C_2 \times C_2$ given by $(a, b) \mapsto (a, 0)$. This map yields two components with symmetry group C_2 on each. The calculation also has a useful interpretation in a machine interface. A semantic symmetry may exchange two indistinguishable ports. If machine morphisms preserve physical connector identities, it may be impossible to lift the exchange. Conversely an internal state symmetry may act as the identity on all ports, and appears as an automorphism of a fixed marking.

2.4 Universal representatives and localizations

Suppose a category $\mathcal{T}_Q(p)$ admits a terminal object t. Then the space obtained by inverting all its arrows is contractible. Indeed the natural transformation from the identity to the constant functor at t induces a homotopy contracting its groupoid completion. This argument uses all arrows, including those that fail to be equivalences in $\mathcal{C}$. The core may still have many components. Even in the simple category $0 \to 1$, the core is a discrete two-point space and the groupoid completion is contractible. An assertion about the core therefore depends on which transformations were declared reversible.

A localization at a class W of arrows formally inverts W and satisfies the corresponding universal property of functors out of the category [8, 9, 24, 3]. It is possible that a semantics Q inverts W and thus factors through $\mathcal{C}[W^{-1}]$ without that factor being fully faithful or essentially surjective. For example, let $\mathcal{C}$ be a point and $\mathcal{D} = BH$ with $H \neq 1$. Every arrow inverted by Q is already invertible. The induced map to the full subcategory of the target on its essential image still lacks the nonidentity automorphisms. The two constructions have the same set of object classes and different equivalence spaces. A useful semantics can deliberately forget or add such changes of coordinates. The localization language is justified only after

specifying the universal property being used.

3 Truncation and hidden operational structure

An object $X \in \mathcal{X}$ is n-truncated if $\mathrm{Map}_{\mathcal{X}}(U, X)$ is n-truncated for every $U \in \mathcal{X}$. The inclusion of such objects has a left adjoint $\tau_{\leq n}$. We use the conventions that (-1)-types are propositions and (-2)-types are contractible spaces, with the terminal object as the only (-2)-truncated object up to equivalence. Truncations compose by taking the minimum of their degrees. They are localizations in their own right [17, Sec. 5.5.6]. For a machine model with operational semantics Op, set $Q_n(M) = \tau_{\leq n} \mathrm{Op}(M)$ and define $M \equiv_n N$ by equivalence of the images. Then $M \equiv_{n+1} N$ implies $M \equiv_n N$. These are equivalence relations on machine classes. They do not specify a category of comparison arrows until such a category has been chosen.

3.1 Examples at the boundary degrees

In $\mathcal{S}$, when $n \geq 0$, S^{n+1} and a point have equivalent n-truncations. They are inequivalent at degree $n+1$ because $\pi_{n+1}(S^{n+1}) \cong \mathbb{Z}$. The canonical projection is a witness of nonconservativity with inequivalent endpoints. For $n = -1$ one can instead use $S^0 \to *$. Both spaces are inhabited, so their (-1)-truncations agree, and their sets of components differ. At $n = -2$ the map $\varnothing \to *$ becomes an equivalence but is detected at -1. This accounts for the special behavior of π_0, which is a set of components rather than a copy of $\mathbb{Z}$ for S^0. These examples prove a statement about spaces and a failure in general. Any restriction of Op to already n-truncated objects loses no additional information at stage n.

Products give a useful way to propagate such examples into machine models. Truncation in an ∞-topos preserves finite products. Indeed, the n-connected maps are stable under pullback and composition, and the product of two units of truncation is a composite of base changes of units. Its target is n-truncated, so it exhibits the required truncation. The same argument works internally in slices. Truncation can fail to preserve other finite limits. In spaces $* \times_{BC_2} * \simeq C_2$, while the pullback of the three zero-truncations is a point. This example will constrain the order of operations

when modeling families.

3.2 Hidden decorations and sharp separation

Theorem 3.1 (Hidden-decoration criterion)**.** *Let $n \geq -2$, let $D \in \mathcal{X}$ satisfy $\tau_{\leq n} D \simeq 1$, and suppose there is an endofunctor $T : \mathcal{C} \to \mathcal{C}$ with a natural equivalence $\mathrm{Op}\, T \simeq \mathrm{Op}(-) \times D$. Then $Q_n T \simeq Q_n$. If M is any machine for which $\tau_{\leq n+1} \mathrm{Op}(M)$ and $\tau_{\leq n+1} \mathrm{Op}(TM)$ are inequivalent, the same observed object at degree n has two inequivalent machine realizations. When T is equipped with a natural map $T \to \mathrm{id}$ modeling projection, its component at M is a specific machine arrow inverted by Q_n and detected by Q_{n+1}.*

Proof. Finite product preservation identifies $Q_n T$ with $Q_n \times \tau_{\leq n} D$. The projection from this product is a natural equivalence. A machine equivalence $M \simeq TM$ would induce an equivalence after applying Q_{n+1}, contrary to the premise. For the last claim, projection becomes the projection off a terminal factor at degree n, while its image has inequivalent endpoints at degree $n + 1$. □

The distinction at the next stage has to be checked. Adding a noncontractible factor to an arbitrary operational space need not change its type. For example a countably infinite discrete space is equivalent to its product with a two-element set. The nullary machine with contractible operational object avoids this absorption. More generally, choosing $D = K(A, n + 1)$ for $n \geq 0$ and A nontrivial abelian gives a sharp difference at degree $n + 1$ on any nonempty M with n-truncated operational space in $\mathcal{S}$. There the first added based homotopy group can be computed in every component, even if the two machines have no preferred base point. For $n = 0$ a nontrivial discrete group also works with its classifying space. The assumption that the machine category permits the decoration matters. A physical machine class with restricted transitions, storage, or process interfaces need not admit T.

There is a different construction in which the marked ambiguity can be calculated without a truncation. Let $\mathcal{B}$ be a category of behaviors, $\mathcal{K}$ a category of internal decorations, and $\mathcal{C} = \mathcal{B} \times \mathcal{K}$. Take $Q = \mathrm{pr}_1$. For each behavior b, the realization space is $\mathcal{K}^{\simeq}$ and the transformation category is $\mathcal{K}$. This follows by taking the pullback of a

product projection. It shows that a single observed class can carry any prescribed space of choices, by taking $\mathcal{K}$ to be that space. The product model assumes independence. If machine updates can couple the decoration to the behavior, the projection may still be a functor but the fibers need not be constant or admit the same transports. The construction is therefore a sharp statement of possibility and a useful baseline for more constrained systems.

3.3 Relative truncation retains the indexing interface

For a context object B, operational families can be modeled in $\mathcal{X}_{/B}$. Write $\tau^B_{\leq n}$ for truncation in this slice. An object $E \to B$ is n-truncated in the slice exactly when the map $E \to B$ is n-truncated. Its relative truncation is the middle object of the factorization of that map into an n-connected map followed by an n-truncated map. This factorization is stable under base change. For $u : A \to B$, pulling back the factorization along u preserves both classes of maps and therefore gives the corresponding factorization for u^*E. By uniqueness,

$$u^*\tau^B_{\leq n}E \simeq \tau^A_{\leq n}(u^*E).$$

The relative form of Theorem 3.1 follows by applying it in the slice, replacing 1 with id_B and products with fiber products over B.

Relative truncation preserves the indexing object as data. As a simple illustration take $E = *$ and $B = BC_2$ in spaces. The fiber of $E \to B$ is the two-element C_2-set with the free transitive action. The map is already 0-truncated, so $\tau^B_{\leq 0}E = E$ over B. Taking absolute zero-truncations first produces $* \to *$, whose base change back to B has a single point in every fiber. The latter forgets the action. This matters when a machine supplies data indexed by a semantic context and changes of context can permute those data. A truncation that discards the context can destroy precisely the symmetries that markings were intended to record.

3.4 Higher attachments can distinguish machines with the same groups

Isomorphism of homotopy groups is weaker than equivalence of Postnikov stages. Consider the generator $x \in H^1(BC_2; \mathbb{F}_2)$, and let $f : BC_2 \to K(\mathbb{F}_2, 3)$ classify x^3. The ring $H^*(BC_2; \mathbb{F}_2)$ is polynomial on x, since BC_2 is the infinite real projective space [11, Chs. 3–4]. Set

$$E_f = \mathrm{hofib}(f), \qquad E_0 = BC_2 \times K(\mathbb{F}_2, 2).$$

The fibration over BC_2 has fiber $K(\mathbb{F}_2, 2)$ in both cases. The long exact sequence gives connected 2-types with $\pi_1 = C_2$, $\pi_2 = \mathbb{F}_2$, and trivial higher groups. The action on π_2 is trivial because the coefficient group has trivial automorphism group. The projections are their 1-truncations. The split fibration $E_0 \to BC_2$ has a section, and a section of $E_f \to BC_2$ would supply a nullhomotopy of f. Since $x^3 \neq 0$, no section exists. Any equivalence $E_f \simeq E_0$ would induce an equivalence on 1-truncations and would carry the property of admitting a section along that equivalence. Thus the two 2-types differ.

This example separates two ways to enrich an observation. Adding separate probes for the homotopy groups can detect new loops and higher loops. Retaining the higher type also records how these groups are assembled. In the example the assembly is expressed by a cohomology class, and the obstruction it creates concerns the existence of a coherent choice over an entire base. This is the same form of obstruction that reappears when local machine certificates must be transported around a context.

3.5 All finite stages

Write $P_\bullet X = (\tau_{\leq n} X)_n$ for the tower with the canonical maps between stages. There are two different assertions about this tower in an ∞-topos. First, a map that becomes an equivalence at every stage is an equivalence after hypercompletion. For each finite m, the unit maps to the m-truncations and the comparison between their targets show that both $Yo\tau_{\leq m}Y$ and $XoYo\tau_{\leq m}Y$ are m-connected. The connectivity estimate for cancellation on the right then gives at least $(m-1)$-connectedness of XoY. Thus

the map is ∞-connected, and hypercompletion localizes at such maps [17, Secs. 6.5.1–6.5.2]. In a hypercomplete topos the given map was already an equivalence. Second, an object may be recovered as the limit of its own truncations. When this holds for the objects in question, compatible equivalences of towers induce an equivalence between them by taking limits. The stronger assertion that every compatible Postnikov tower is effective in the original topos involves essential surjectivity of the tower comparison. We use convergence for the specified objects only where stated. The separate issues of hyperdescent and effectiveness are discussed in [17, Sec. 6.5.3] and [6]. In spaces every object is recovered by its Postnikov tower [11, Ch. 4].

For X and Y whose towers converge, one can make the compatibility test particularly explicit. Adjunction and preservation of limits by mapping out give

$$\begin{aligned} \mathrm{Map}_{\mathcal{X}}(X, Y) &\simeq \lim_n \mathrm{Map}_{\mathcal{X}}(X, \tau_{\leq n} Y) \\ &\simeq \lim_n \mathrm{Map}_{\mathcal{X}_{\leq n}}(\tau_{\leq n} X, \tau_{\leq n} Y). \end{aligned}$$

The transition maps in the last tower are induced by truncation. Restricting to equivalences gives the corresponding formula for $\mathrm{Eq}_{\mathcal{X}}(X, Y)$. One way to verify the restriction is to use convergence of both towers. A point of the limit made of equivalences assembles to a map of the two towers that is an equivalence at every stage, and its limit is an equivalence. The reverse direction follows by applying truncation to a given equivalence. The restriction retains homotopies, since equivalences form unions of components.

Existence of an equivalence separately at every stage supplies inhabitation of each term of the equivalence-space tower. Recovering the equivalence from the displayed formula calls for an element of the homotopy limit. Mere termwise inhabitation is insufficient as a general limit argument. For example the inverse system of nonempty sets $\{n, n+1, \dots\}$ with inclusions has empty limit. A sufficient hypothesis for a countable tower of nonempty spaces is surjectivity on components at every transition. Starting with a point, choose lifts up to path at successive stages, or replace the tower by a tower of fibrations and choose compatible lifts. For finite nonempty sets the same conclusion follows without transition surjectivity by compactness, as the compatibility constraints on finite initial segments have the finite intersection property. These

tests for choosing compatible witnesses add hypotheses beyond equality of stagewise semantic classes.

Finally, reconstruction of $\mathrm{Op}(M)$ is distinct from reconstruction of M. If Op identifies machines already at full precision, no amount of truncation refinement removes that ambiguity. An injectivity or full faithfulness hypothesis on Op must enter any inference from its reconstructed value to a unique machine.

4 What a finite specification determines

The adjective *finite* can refer to the presentation of a syntax, a list of equations, the number of tests, the size of a state space, or the cardinality of a catalog of admissible behaviors. These are different assumptions. To separate them, let P be a presentation and let $\mathcal{A}_P \subseteq \mathcal{D}^{\simeq}$ be the full subspace on semantic objects satisfying it. Assume satisfaction is invariant under equivalence. Let $\mathcal{C}_P \subseteq \mathcal{C}$ be a category of machines admitted by any additional implementation conditions, with a functor $Q_P : \mathcal{C}_P \to \mathcal{D}$ induced by Q. Soundness says that $Q_P^{\simeq}$ factors through $\mathcal{A}_P$. Realizability of a semantic object $p \in \mathcal{A}_P$ says that the marked fiber over p is nonempty.

4.1 Determination by soundness and categoricity

Proposition 4.1. *Assume soundness and let R_P be the image of $\pi_0(\mathcal{C}_P^{\simeq}) \to \pi_0(\mathcal{A}_P)$. The presentation has exactly one realized semantic class if and only if R_P is a singleton. In particular, nonemptiness of $\mathcal{C}_P$ and connectedness of $\mathcal{A}_P$ suffice. If every semantic class is realized, these conditions are also necessary. If $\mathcal{A}_P$ is contractible, the realized semantic object is determined with all equivalence data internal to $\mathcal{A}_P$. The marked machine fiber may still have arbitrary homotopy type.*

Proof. The map on components sends every admitted machine class to its realized semantic class, so its image is exactly the set whose uniqueness is asserted. A connected nonempty space has one component. If every class is realized, the image is all components. For the last assertion, use the product projection $\mathcal{A}_P \times K \to \mathcal{A}_P$ for any space K, as in the preceding section. □

It is useful to separate *semantic categoricity* from uniqueness within the class of implementable systems. If an incomplete theory admits two abstract models but exactly one is implementable on a fixed architecture, the realized class can be unique even when the theory is semantically noncategorical. Conversely, a categorical semantic theory can have no machine realization under a chosen resource constraint. An empty realization category satisfies the pairwise agreement assertion vacuously. For purposes of a journal claim about realization we require a witness, a sound implementation construction, or an independent existence theorem.

There is a parallel distinction between a universal representative and a categorical class. A finite set of generators and equations can specify a theory and thereby specify its free or initial model, provided the relevant colimits or free constructions exist. The theory can have many other models. If one defines P to include the universal property of being the initial model, then the full groupoid of such models is empty or contractible. The property removes all ambiguity at the level of that groupoid because initial objects have contractible mapping spaces to every object, including each other. It does not remove distinct implementations mapping to those models. This reasoning also applies to terminal specifications, such as the universal observable quotient in Section 6, once the class of admissible transformations is fixed.

4.2 Tests can underdetermine even discrete behavior

Take a unary input alphabet and a Boolean output alphabet. A behavior is a sequence in $\{0,1\}^{\mathbb{N}_{>0}}$. A finite test requiring zeros for the first r positions is satisfied by the constant-zero behavior and by the behavior that emits 1 exactly at position $r+1$. Both are implemented by finite initialized transducers. The second uses a counter that saturates one step after the distinguished output. If the semantics regards the entire sequence as a discrete object of a behavioral catalog, these are different semantic classes at every degree at which the catalog is observed as a set of names. Even if the tests inspect finitely many distinct words instead of consecutive prefixes, their maximal length bounds the portion of this unary example they can distinguish.

This construction concerns finite black-box tests without a prior bound that rules out delayed differences. With a uniform bound on the number of reachable states

of two deterministic finite machines, an effective distinguishing depth exists. In the deterministic synchronous model the product of their state sets has at most ab pairs if the sizes are a, b. A shortest distinguishing word has no repeated pair before the edge where the outputs differ. One could otherwise delete the intervening cycle. Its length is therefore at most ab. Thus bounded state cardinality can turn a finite list of word tests into a complete equivalence test for that class. The additional structural hypothesis, rather than finiteness of the presentation alone, provides completeness.

4.3 Finite rules can determine unbounded behavior

A different finite specification says that the initial state is s_0, that every reachable state emits zero on every input, and that the transition is total. There may be infinitely many reachable states and infinitely many input words. For a finite alphabet the output condition is finitely expressible as a universal equation for each letter, or as a single universally quantified rule if letters themselves form a sort. Induction on word length proves that all these machines have the same trace function, even if they are far from isomorphic as transducers.

Here the observation is a family of outputs indexed by words. It can be treated as a point of a set B of behaviors, or as an object in an ambient category of labeled data. It should not be modeled by sending every behavior to an unlabeled one-point space in $\mathcal{S}$. Such a functor collapses every distinction of values before any interesting truncation takes place. A convenient faithful encoding of discrete values into an appropriate slice is to use the maps $\{b\} \hookrightarrow B$, where B is discrete. Two singleton maps are equivalent over B exactly when they specify the same value. The space B may have infinitely many points, and nothing in the encoding restricts the size of the set of tests.

4.4 Parametrized admissibility and changing specification

Suppose P is refined to P' by adding a condition. At the semantic level this is represented by an inclusion $\mathcal{A}_{P'} \hookrightarrow \mathcal{A}_P$. If implementation conditions are fixed, the corresponding space of admitted marked data is obtained by pullback. This gives a

monotonicity statement. A fiber over a retained semantic object is unchanged by the new semantic condition alone, while entire fibers over eliminated objects disappear. An implementation condition can instead cut down a fiber while leaving the set of realized semantic classes unchanged. For example, one may require a cost bound on machines with the same output behavior. It may eliminate slow realizers and retain a fast one without changing the trace function. The distinction anticipates the graph of feasible replacements in Section 10.

There is also a useful restriction concerning empty fibers. A translation of one specification to another may preserve validity of all equations yet lose a realizability guarantee when the target semantics is finer. Lifting along $\tau_{\le n+1}\,\mathrm{Op} \to \tau_{\le n}\,\mathrm{Op}$ requires finding a higher operational object in the specified class, and then a machine realizing it. A semantic Postnikov extension that exists in the topos need not lie in the image of the operational functor. A correct existence argument can proceed in three steps. Construct the extension, show it lies in the admissible semantic subspace, and then use an implementation theorem to supply a machine and its marking.

5 Trace objects, semantic descent, and contextual tests

A single trace is often a path through a machine under one sequence of inputs. It cannot in general determine a machine’s observation class. A trace semantics should first specify whether it keeps a set of traces, a trace tree, a distribution, a space of executions and their homotopies, or an indexed family of observations. It must then specify how machine arrows induce arrows of trace objects. With this convention let $T : \mathcal{C} \to \mathcal{E}$ be a trace functor. A reification at level n is a functor $R_n : \mathcal{E} \to \mathcal{X}_{\le n}$ with a natural equivalence $R_n T \simeq Q_n$. The natural equivalence includes the compatibility with arrows. The equation between maps of classes obtained by taking cores and components is a weaker demand.

5.1 A criterion from localization

Proposition 5.1 (Coherent factorization through a trace quotient)**.** *Suppose $T : \mathcal{C} \to \mathcal{E}$ exhibits the localization at a class W of arrows. Then a coherent factorization $R_nT \simeq Q_n$ exists exactly when $Q_n(w)$ is an equivalence for every $w \in W$. With the equivalence included as part of the data, the space of factorizations is contractible if it is nonempty.*

Proof. By the universal property, restriction along T identifies $\mathrm{Fun}(\mathcal{E}, \mathcal{X}_{\leq n})$ with the full subcategory of $\mathrm{Fun}(\mathcal{C}, \mathcal{X}_{\leq n})$ spanned by functors inverting W. Existence is the essential image condition. Taking cores and the fiber at Q_n gives a contractible space when the target belongs to that full subcategory. □

In particular, for $W_n = Q_n^{-1}(\mathrm{Eq})$ there is a canonical localization factor

$$\mathcal{C} \longrightarrow \mathcal{C}[W_n^{-1}] \xrightarrow{\overline{Q}_n} \mathcal{X}_{\leq n}.$$

If these localizations exist in the chosen size, $W_n \subseteq W_m$ for $m \leq n$ induces functors $\mathcal{C}[W_n^{-1}] \to \mathcal{C}[W_m^{-1}]$. The universal properties supply the coherence of successive comparisons. This yields a second tower, consisting of categories of machines with certain arrows inverted. It is related to the tower of semantic n-types by the factor maps, and can encode additional machine information. For example, independent machines with equal observations and no arrows between them remain in separate components of a discrete machine localization. Equality of their semantic objects does not generate a machine arrow. Even when each quotient is conservative relative to the maps it detects, comparisons between the essential images can fail to be full.

If the trace category is provided by an operational model rather than by a localization, the factorization needs a separate proof. For example let the machine category have one object and the endomorphism monoid freely generated by a. Send it to the terminal trace category. Send the machine to an object p with a chosen nonidentity endomorphism u in a semantic category and take a to u. The two object assignments are compatible through the sole trace object, but no functor from the terminal category can carry a to u, even up to natural equivalence if u is not homotopic to the identity. This example isolates the arrow obstruction independently of whether the trace function on objects was complete.

5.2 Tests and the coarsest separating quotient

Let L be a set of closed programs at a fixed type, let K be a set of typed contexts with an observation operation, and let O be the set of observable outcomes. Define

$$\mathrm{test} : L \longrightarrow O^K, \qquad \mathrm{test}(\ell)(k) = \mathrm{obs}(k[\ell]).$$

Equality in the image of this map agrees with contextual equivalence for these contexts, by definition. The equivalence relation is a congruence for an operation on programs whenever the tests are closed under composing with that operation in each slot and under the required changes of type. In one slot, fix the other inputs and compose a distinguishing context with the term constructor. Equality for this composite context implies equality for the original context after applying the constructor. Varying the slots one at a time proves the congruence for finite arity. This explicit hypothesis is necessary. A set of observations omitting contexts may identify programs that become distinct under composition.

In this setting a semantic model is fully abstract for the given tests when equality of its denotations holds exactly for contextually equivalent programs. Milner's treatment concerns typed lambda calculi and their models, and game models of PCF furnish a richer example of proving both directions [21, 2]. The corresponding claim here, for a chosen machine interface, would require soundness and completeness against a chosen family of tests. An abstract map Q_n supplies neither by itself. One can formulate a higher variant in which testing returns a mapping space of witnesses and composition is coherent. Such a variant needs a specific enriched test model and a comparison of spaces. Equality of observational outcome functions alone only proves the set-level statement above.

A simple universal property makes the observation quotient canonical for these tests. Any map $g : L \to Y$ that is constant on test-equivalence classes factors uniquely through the image of test. Indeed send a realized test vector to the value of g at any preimage, which is independent of the choice. If the tests are enlarged from K to $K' \supset K$, restriction gives a surjection between the two images, because each coarser vector has a program witness whose finer vector restricts to it. This gives an observation hierarchy generated by tests, independent of Postnikov truncation. The

two hierarchies can coexist. For example one can keep all words while truncating the homotopy type of possible executions of each word.

5.3 A branching distinction lost by trace sets

Consider pointed labeled transition systems. The first system has one a-successor and that successor has both a b- and a c-transition to terminal states. The second has two a-successors, one with only a b-transition to a terminal state and the other with only a c-transition to a terminal state. All states can be reached from their respective initial states. Both systems have the same finite trace language $\{\epsilon, a, ab, ac\}$. There is no bisimulation relating the initial states. To match the a-transition to the first system's single successor, one must relate it to one of the two successors in the second. Each such successor lacks a transition required by the bisimulation condition. This argument uses matching of transitions after an intermediate state has been chosen, whereas a trace set quantifies independently over runs for every word.

Consequently a semantic quotient by trace equality can identify systems that a branching-sensitive test distinguishes. The linear-time and branching-time spectrum systematizes such distinctions [10]. The open-map approach to bisimulation also makes precise the role of path extension, by specifying a category of shapes and a lifting condition for morphisms [15]. Our trace functor T can be chosen to retain path-extension data, but Proposition 5.1 applies only if T has the stated universal property. It does not identify open-map bisimulation with equivalence in every possible category of trace objects.

5.4 A site of observations

One model for a structured trace object uses a small category J of contexts or path shapes. A presheaf $E_M : J^{\mathrm{op}} \to \mathcal{S}$ assigns a space of observations to each shape and restriction maps to subshapes. A natural transformation between presheaves preserves the restrictions. In the presheaf ∞-topos, n-truncation is computed pointwise. One can prove this by the objectwise truncation adjunction on diagrams and the fact that a diagram of n-types is an n-truncated object of the functor category. Conversely

evaluation at a representable test object detects if the values are truncated.

For a rooted transition system and a poset of input words ordered by the prefix relation, the value at w can be the set of complete runs of shape w. Restriction deletes later transitions. In the branching example these presheaves are inequivalent because the value at a has different cardinality. Their supports, meaning the presheaves obtained by applying (-1)-truncation pointwise, agree. Thus a support observation has the exact finite trace language here, and a 0-truncated run presheaf can contain more data. Branching equivalence itself is a different comparison, so this observation does not identify it with equality of run presheaves.

Putting a Grothendieck topology on J changes the meaning of observation. Sheafification forces agreement with descent along covers, whereas pointwise truncation controls homotopy in the values of a presheaf. The standard left-exact sheafification functor and truncation are related by universal properties. In a sheaf topos a truncation can be obtained by sheafifying the pointwise truncation. Concretely, sheafification preserves truncated objects as a left-exact functor, and the adjunction against n-truncated sheaves identifies the resulting reflector. Evaluating a sheafified truncation at a context therefore need not be the same as truncating the global observation at that context. The relevant descent data are treated in Section 7.

6 Deterministic transducers and canonical quotients

We now fix a strict model in sets that makes several of the abstract choices explicit. The result is a form of behavioral minimization. Its value here is to identify the relevant functor, the realized behavior, the quotient arrow, and the marked fiber in one model. The construction is closely related to the observational approach to sequential machines [22, 20], to the residual construction for automata, and to the final-coalgebra treatment of deterministic behavior [27, 14].

6.1 Machines, derivatives, and the final behavior object

Fix input and output sets A, B, and write A^* for all finite words with empty word ϵ. An initialized synchronous Mealy machine is a set S, an initial state $s_0 \in S$, and a transition-output function

$$S \times A \longrightarrow B \times S, \qquad (s, a) \longmapsto (\lambda(s, a), \delta(s, a)).$$

A strict morphism $u : S \to S'$ preserves the initial state and satisfies $\lambda'(u(s), a) = \lambda(s, a)$ and $\delta'(u(s), a) = u(\delta(s, a))$. The deterministic dynamics extend δ to words by $\delta(s, \epsilon) = s$ and $\delta(s, wa) = \delta(\delta(s, w), a)$. For a nonempty word $w = va$, let $b_s(w) = \lambda(\delta(s, v), a)$, the last output on w. The entire output sequence on w is then the list of b_s evaluated on its nonempty prefixes.

Let $A^+ = A^* \setminus \{\epsilon\}$, and let $V = B^{A^+}$ be the set of last-output functions. Give V the uninitialized machine structure

$$\lambda_V(b, a) = b(a), \qquad \delta_V(b, a)(v) = b(av), \quad v \in A^+.$$

The derivative after a word w will be denoted by b_w, with $b_w(v) = b(wv)$. The map $s \mapsto b_s$ is a strict uninitialized homomorphism because the formulas for outputs and derivatives agree with the transition-output function. Moreover, any such homomorphism must assign b_s to s. To prove this, output preservation fixes its values on one-letter words and transition preservation inductively fixes its values on longer words. Thus V is a final coalgebra for the functor $S \mapsto (B \times S)^A$. Initializing V at b makes it terminal in the full category of initialized machines with initial behavior b, if unreachable states are allowed. There can be smaller terminal representatives once reachability is required.

The discrete observation functor $S \mapsto b_{s_0} \in V$ is well defined on strict initialized homomorphisms. For such a homomorphism, evaluating the composite into V gives the unique behavior morphism, and the initial states must be sent to the same point of V. We can regard this observation as a functor into the discrete category V and then into the slice $\mathcal{S}_{/V}$ using the singletons over their respective points. The slice embedding makes it an example of Q_n for $n \geq 0$. Working directly with V makes the semantic equivalence spaces discrete and trivial at each behavior.

6.2 Reachability and minimality

Definition 6.1. A machine is *reachable* when every state is $\delta(s_0, w)$ for a word w. It is *observable* when $b_s = b_t$ implies $s = t$.

For $b \in V$, let

$$D_b = \{b_w : w \in A^*\} \subseteq V.$$

This is closed under the derivatives defining the transitions. The initial state is b, and each derivative is reachable by construction. Since its states are literally their behaviors, it is observable.

Theorem 6.2 (Terminal reachable realization and sharp size bound)**.** *In the category of reachable initialized Mealy machines with fixed behavior b, D_b is terminal. Every homomorphism in this category is surjective on states and there is at most one homomorphism between any ordered pair of machines. The unique homomorphism $S \to D_b$ identifies precisely states with the same residual behavior. In particular a finite reachable machine realizes b exactly when D_b is finite. If finite, its size is a lower bound on the sizes of all reachable realizers, attained up to isomorphism only by D_b.*

Proof. The behavior morphism $S \to V$ sends $\delta(s_0, w)$ to b_w, so its image is D_b and it factors as a strict initialized homomorphism to D_b. Uniqueness follows either from the finality of V and the inclusion $D_b \subseteq V$, or directly because a homomorphism must send $\delta(s_0, w)$ to the state reached by w in the target. The latter equation determines every value by reachability and makes the homomorphism surjective when the target is reachable. The fibers of $S \to D_b$ are the equal-behavior classes by definition. Its surjectivity gives the size inequality and the criterion for finite realizability. If the cardinalities coincide and are finite, the map is bijective, and the defining homomorphism equations imply that its inverse also preserves the machine structure. □

The category in the theorem can be described by quotients of A^*. For a reachable machine let $\sim_S$ identify two words when they reach the same state. This is a right congruence under appending letters. It refines $\sim_b$, the relation $v \sim_b w$ when $b_v = b_w$.

Conversely, for any right congruence $\sim$ refining $\sim_b$, form $A^*/\sim$. Its initial state is $[\epsilon]$, its transition is $[w] \mapsto [wa]$ on a, and its output is $b(wa)$. Right congruence ensures well-defined transitions, and $\sim$ refining $\sim_b$ ensures well-defined outputs. A homomorphism exists precisely when $\sim_S \subseteq \sim_{S'}$, since any homomorphism identifies all the equal reached states of its source in its target, and the inclusion defines the quotient map. This proves the following refinement of the classification.

Corollary 6.3 (Behavioral fiber as a poset of congruences)**.** *The transformation category of all reachable realizers of b is equivalent to the poset of right congruences on A^* refining $\sim_b$, ordered by inclusion. It has a greatest element $\sim_b$ and a least element, the equality congruence. Its core is equivalent to the discrete set of these congruences.*

Proof. The quotient construction gives representatives for all objects up to isomorphism. The preceding argument gives full faithfulness with the poset. The stated extremal relations are right congruences refining $\sim_b$. In a poset the only equivalences are identity arrows, so taking the core gives the discrete set. The least element may yield an infinite machine and can be excluded by a finiteness bound. □

6.3 A family with a single constant behavior

Take $A = \{a\}$, $B = \{0, 1\}$, and $b(a^k) = 0$ for $k \geq 1$. Each initialized m-cycle S_m on states $\mathbb{Z}/m\mathbb{Z}$, with successor given by adding 1 and initial state 0, is reachable and has behavior b. Its congruence identifies lengths exactly modulo m. An initialized homomorphism $S_m \to S_d$ exists exactly when $d \mid m$. Indeed the candidate is forced to be reduction of lengths, and its well-definedness is the divisibility condition. Thus these machines are pairwise nonisomorphic for different m, and the one-state machine is terminal. The same observation class contains every cycle length, even when unreachable state is forbidden. The length encodes a machine invariant that the interface does not expose.

One may enrich the observation to retain the phase. For example add a port that emits 1 exactly when the current state is 0. If two lengths $m < d$ produce the same new behavior, their outputs at the $(m + 1)$st transition would agree, yet the m-cycle has returned to the initial state while the d-cycle has not. Hence the cycles become

distinguishable. This refinement changes the observation functor on morphisms. A quotient from S_m to S_d preserving the original constant output may no longer preserve the phase port. A refined observation can therefore require restricting the category of certified comparisons, not only increasing the amount of data attached to each object.

6.4 Passing to a behavioral localization

Let $\mathcal{M}_{\mathrm{r}}$ be the category of reachable initialized machines and V_{r} its image of realized behaviors. Different behaviors have no strict initialized arrows between them. Let $Q : \mathcal{M}_{\mathrm{r}} \to V_{\mathrm{r}}$ be the behavior functor. Inverting all strict morphisms exhibits V_{r} as the groupoid completion of $\mathcal{M}_{\mathrm{r}}$. For each b, $\mathcal{M}_{\mathrm{r}}[b]$ is contractible after groupoid completion by terminality of D_b. Taking the coproduct over behaviors gives the result. The argument still works in the finite reachable category because D_b is finite for every finitely realized b. Thus here the localization claim holds with a concrete proof and a specified machine category.

The proof illustrates why semantic uniqueness and machine uniqueness have different meanings in the same model. Before inversion the core records each congruence, and the strict arrows record which machines can be quotiented to which other machines. After inversion all reachable machines with fixed behavior are equivalent. Adding costs or restricting allowable arrows can change the localization. The map to behaviors remains available when every retained arrow preserves behavior, but each behavior category may cease to have a terminal object or contractible groupoid completion.

6.5 A comparison with nondeterministic semantics

The proof of Theorem 6.2 uses the unique continuation from a state for each letter. In the nondeterministic example of Section 5, the trace language loses branching information, and a derivative of the language may correspond to a set of reachable states rather than a single state. Determinization can construct such a subset state, yet the resulting arrow, if desired as a comparison, has to be described in a category allowing the appropriate relations. A strict state map preserving individual transitions is a different morphism. Thus the Mealy construction supplies a terminal

representative under its own hypotheses. Its proof cannot be used to declare a terminal trace representative in an arbitrary transition category.

7 Families, descent, and local realizability

A machine may be specified separately in several overlapping contexts. A context can describe an environment, a set of input names, or an available set of permissions. If the specification must work across changes of context, existence of a local machine at each context does not alone settle global existence. This section gives two precise models. The first models the core of machine choices as a sheaf of spaces. The second models semantic outputs as objects of slices and compares local equivalence, global equivalence, and relative truncation.

7.1 Descent of marked fibers

Let $(J, \mathcal{J})$ be a small site, let $\mathcal{X} = \mathrm{Shv}_{\mathcal{J}}(J; \mathcal{S})$, and let $q : \mathscr{M} \to \mathscr{B}$ be a map of sheaves of spaces. The value of $\mathscr{M}$ at U is the space of machine data available in context U. The object $\mathscr{B}$ is the space of semantic data. For a specification $p : U \to \mathscr{B}$ at an object $U \in \mathcal{X}$, put

$$\mathscr{R}_p = U \times_{\mathscr{B}} \mathscr{M} \longrightarrow U.$$

The space of marked realizers in a context $V \to U$ is $\mathrm{Map}_{\mathcal{X}_{/U}}(V, \mathscr{R}_p)$. By preservation of limits in the target by representables this is the homotopy fiber at $p|_V$ of $\mathrm{Map}_{\mathcal{X}}(V, \mathscr{M}) \to \mathrm{Map}_{\mathcal{X}}(V, \mathscr{B})$. In particular it has the same notion of marking as (2).

Proposition 7.1 (Descent for marked machine choices)**.** *Let $e : V \to U$ be an effective epimorphism and $V_\bullet$ its Čech nerve. Then*

$$\mathrm{Map}_{\mathcal{X}_{/U}}(U, \mathscr{R}_p) \simeq \lim_{[k] \in \Delta} \mathrm{Map}_{\mathcal{X}_{/U}}(V_k, \mathscr{R}_p).$$

Local inhabitation is the condition that $\mathscr{R}_p \to U$ be an effective epimorphism. If $\mathscr{R}_p \to U$ is also a monomorphism, local inhabitation implies that the space of global marked realizers over U is contractible.

Proof. The Čech nerve of an effective epimorphism has colimit U, also in the slice. Mapping out of this colimit into $\mathscr{R}_p$ gives the limit. The (-1)-truncation of $\mathscr{R}_p$ in the slice is its image subobject of U, so local inhabitation says that this image is the entire base. If $\mathscr{R}_p \to U$ is a monomorphism it already is (-1)-truncated. Its relative (-1)-truncation is then itself, so the map is an equivalence and its section space is contractible. □

The totalization has concrete content. Its degree-zero entries are local machines with markings. Degree-one entries compare the two restrictions of a local machine to overlaps by compatible equivalences. Degree-two entries give homotopies between compositions on triple overlaps. Further entries encode the remaining coherences. When values are discrete the degree-one condition reduces to equality of restrictions, and higher compatibility is determined. For groupoid-valued data, degree one contains genuine isomorphisms, and the triple-overlap condition is a cocycle equation. For higher data one must retain the homotopy limit. This is ordinary Čech descent for an effective cover. Descent for every hypercover is a stronger ambient condition in general. The standard comparison with simplicial presheaves is developed in [7, 17, 12].

One may apply the same calculation to a relative equivalence object. Given $E, F \in \mathcal{X}_{/U}$, the internal object of equivalences over U is the subobject of the internal mapping object whose generalized elements are equivalences. The slice is locally cartesian closed, and the equivalence condition can be imposed by asking separately for a left inverse and a right inverse, each with its corresponding homotopy. The space of such data over an equivalence is contractible because composition with an equivalence induces equivalences of mapping spaces. Thus sections of the equivalence object classify equivalences between E and F over U. Existence on each member of a cover says the equivalence object is locally inhabited. It supplies a global equivalence if the local sections can be chosen with coherent overlap data. This rephrases reconstruction with compatible witnesses in a spatial rather than temporal diagram.

7.2 A discrete torsor gives local but no global marked realizer

Let $G = C_2$ and work in $\mathcal{X} = \mathrm{Fun}(BG, \mathcal{S})$, a presheaf ∞-topos after identifying BG with its opposite. Let $\mathscr{B} = 1$, and let $\mathscr{M} = G$ be the discrete regular left G-set. For the unique global p, $\mathscr{R}_p = G \to 1$. Effective epimorphisms in this functor topos are computed pointwise, and the unique map from the nonempty underlying set to a point is surjective on components. Thus this is an effective epimorphism. The pullback to the cover $G \to 1$ has the diagonal section. A global section would be a G-fixed point of the regular action. There is no such point. Since the action is on a discrete set, a homotopy-fixed point amounts to an ordinary fixed point, so passing to the homotopy-coherent definition does not change this conclusion.

The diagonal section after base change does not satisfy equality of the two restrictions on the overlap $G \times G$: those restrictions select its two coordinates. Since the target is discrete, there is no homotopy that supplies the required compatibility when the coordinates differ. Thus this is a descent failure for a particular local choice, with an actual obstruction to any global choice. It is not a failure of the sheaf condition. The sheaf condition determines the space of global sections from the entire descent diagram, and that space is empty. In a machine interpretation, local naming permits one to choose a physical connector, while a global choice would have to be invariant under exchanging the connectors.

A related phenomenon is local uniqueness of semantic objects without a global semantic equivalence. In the same G-topos take the regular two-element G-set and the trivial two-element G-set. They become isomorphic over the regular cover $G \to 1$ because the action on the cover freely chooses a coordinate frame. They cannot be equivariantly isomorphic globally, as their fixed-point sets have cardinalities zero and two. Equivalences over the cover form a torsor under the group of port permutations. Identifying only their underlying cardinalities hides the distinction between the two semantic families.

7.3 Descent and relative Postnikov data

Relative n-truncation commutes with pulling back to each level of an effective cover. If $E, F \to U$ are operational families, a compatible collection of equivalences between $\tau_{\leq n}^{V_k} E|_{V_k}$ and $\tau_{\leq n}^{V_k} F|_{V_k}$ therefore gives an equivalence between $\tau_{\leq n}^{U} E$ and $\tau_{\leq n}^{U} F$. This follows either by descent of objects in slices or by applying Proposition 7.1 to their internal equivalence object. The construction has two indices, one for the cover and one for the truncation level. If convergence is known and the equivalences are compatible in both indices, the two homotopy limits can be interchanged because limits commute with limits. This yields an equivalence of the original families in the slice. Checking agreement of local classes separately in both indices omits the choices needed for this conclusion.

The same care applies to recursive rules in a sheaf. If F is given only locally, its local formulas must commute with restrictions and agree coherently over overlaps before there is a global map $F : \mathscr{B} \to \mathscr{B}$. After that one may form semantic iterates internally. Even then, a lift to $\mathscr{M}$ can fail by the same torsor mechanism. The obstruction can lie in selecting a marked machine or in selecting a compatible family of certified transitions between machines already given.

8 Coverage, persistence, and finality of timelines

A timeline is a functor $M : I \to \mathcal{C}$. Usually I is a directed poset of stages or the ordinary category $(\mathbb{N}, \leq)$. We first define the coverage condition using only the induced function on object classes. We then compare it with finality, which depends on entire comma categories. This addresses the source note's use of *n-cofinal* for an objectwise condition.

8.1 Hitting and preserving a class

Let $R \subseteq \pi_0(\mathcal{D}^{\simeq})$ be a set of realized classes. Say that M *covers* R through Q if R is contained in the image of $i \mapsto [QM_i]$. For a single presentation with unique realized class $[p]$, the condition reduces to existence of a stage with $QM_i \simeq p$. This explains

why the coverage result of the source note does not use directedness. One can find the stage by choosing any actual realizer and using the defining coverage hypothesis.

For a directed poset I, the stronger statement *eventually always in the class* means that there is some i for which every $j \geq i$ satisfies $QM_j \simeq p$. A sufficient condition is a hit at i together with the requirement that $Q(M_i \to M_j)$ be an equivalence for every $j \geq i$. On $\mathbb{N}$ it suffices that the consecutive arrows on that tail be sent to equivalences, since arbitrary tail arrows are composites. A single hit does not imply this hypothesis. A diagram in the ordered category $0 < 1$ can start in 0 and remain in 1 forever after its first step. A diagram in a groupoidal semantic target automatically stays in the same component on any connected indexing category. A timeline through several different semantic classes therefore needs a base that allows transitions between components of its core, or an indexing category with several components.

8.2 The categorical finality condition

For $u : I \to \mathcal{D}$ and $d \in \mathcal{D}$, define

$$d \downarrow u = \mathcal{D}_{d/} \times_{\mathcal{D}} I.$$

An object consists of a stage i and an arrow $d \to u(i)$. A morphism includes a morphism of stages and the commuting triangle up to homotopy. Write $|K|$ for the groupoid completion of an ∞-category K, namely the space obtained by inverting all its arrows. The finality criterion states that u is final if and only if $|d \downarrow u|$ is contractible for every d [17, 18, 23, Sec. 4.1]. A final functor preserves the value of colimits by restriction. For a diagram $X : \mathcal{D} \to \mathcal{E}$, the colimit over I of Xu agrees with the colimit over $\mathcal{D}$ when computed. We use *final* for the colimit orientation. Some literature calls the same property *cofinal*.

Nonemptiness of $d \downarrow u$ guarantees a comparison arrow. Contractibility asks in addition for connectedness and the triviality of all higher homotopy in the comparison category after inversion. Coverage of classes is different still, since it asks for an equivalence $d \simeq u(i)$. When $\mathcal{D}$ has noninvertible arrows, a functor can be final and miss equivalence classes. For example the inclusion of the terminal point 1 in the ordinary category $0 \to 1$ is final, because each object has one arrow to 1. It misses

the class of 0. Conversely the inclusion of the discrete set of objects of any ordinary category covers its object classes and may fail finality. In $0 \to 1$, the comma category at 0 has the two objects given by the arrows $0 \to 0$ and $0 \to 1$, with no arrow between them in the discrete indexing category.

8.3 Exact criterion over groupoidal semantics

Theorem 8.1 (Timeline finality over a semantic space)**.** *Let B be a space viewed as an ∞-category and $u : I \to B$. Let $|u| : |I| \to B$ be the induced map. For each $b \in B$ there is an equivalence*

$$|b \downarrow u| \simeq \mathrm{hofib}_b(|u|).$$

Thus coverage of the classes of B is equivalent to inhabitation of all these fibers. Finality is equivalent to contractibility of all these fibers, or to $|u|$ being an equivalence.

Proof. The category $b \downarrow u \to I$ is the left fibration for the functor $i \mapsto \mathrm{Map}_B(b, u(i))$. For a functor from I to spaces the groupoid completion of its Grothendieck construction is its colimit. Mapping that completion into a space Y is the space of compatible maps from each of its fibers to Y, which is the mapping property of the colimit. On the other hand the diagram of points $* \to B$, one at each $u(i)$, has colimit $|I| \xrightarrow{|u|} B$ in $\mathcal{S}_{/B}$. This uses the colimit of the constant diagram of points in $\mathcal{S}$, and the given compatible maps to B. Pulling back this colimit along b preserves colimits in spaces, and the pullback of the point at $u(i)$ is the path space $\mathrm{Map}_B(b, u(i))$. This gives the formula. For coverage, the path space is inhabited for some i exactly when b is in a represented component. The finality criterion and detection of equivalences of spaces by contractible homotopy fibers give the last assertions. □

For a nonempty directed poset I, $|I|$ is contractible. Every finite diagram has a cone at a common upper bound, which contracts the finite data appearing in maps from finite simplicial spheres into the nerve. Equivalently, the nerve of a filtered category is weakly contractible. If B is a connected semantic groupoid with a nontrivial symmetry group H, then a constant diagram $I \to BH$ covers its sole class. The formula identifies the comma-space type with $\Omega BH \simeq H$. Thus even with a directed indexing category and a single semantic class, coverage misses finality when markings have inequivalent

symmetries. A nonconstant diagram with contractible $|I|$ into BH yields the same conclusion, since its induced map is homotopic to a constant map.

For contrast, if the target is a discrete behavioral set, coverage of each component reduces to nonemptiness of the corresponding subdiagram of stages. Finality requires each such subdiagram to have contractible groupoid completion. For a nonempty directed diagram into a discrete set the only covered class is its one constant value, and the functor is final onto that singleton. In that case finality onto the class makes no statement about coverage of distinct machine realizers in the fiber.

8.4 Finality inside a realization category

Suppose $M : I \to \mathcal{T}_Q(p)$ takes values in certified realizers of one behavior. For any realizer R, inhabitation of $R \downarrow M$ means some later representative is accessible from R by an allowed transformation with compatible marking. If these comma categories are all contractible after inversion, the timeline is final within this realization category. This compares arrows and their coherence, in contrast with coverage in $\mathcal{D}$. If the transformation category has a terminal object t, the singleton diagram at t is final. It can omit all other machine classes, even if there are infinitely many of them.

In the unary Mealy example the quotient of any cycle to the one-state machine gives just such a terminal comparison. Changing the allowable arrows to isomorphisms alone changes the answer. The transformation category becomes its discrete core of congruences, so a diagram that only includes the one-state machine misses every other machine class. In practical terms the quotient uses a one-way loss of internal phase. If such a loss is forbidden by a change policy, neither finality nor the set of reachable upgrades can be inferred from the earlier terminality argument.

8.5 Simultaneous requirements and directedness

Directedness is useful when combining a finite list of requirements that become persistent. Suppose $R_1, \dots, R_k$ are subsets of stages, each containing a tail $\{j : j \geq i_r\}$. A common upper bound i of $i_1, \dots, i_k$ lies in every tail, and all later stages do as well. The conclusion uses persistence of each requirement, which can arise from independent

invariants along arrows. Mere coverage of each requirement at some stage supplies no such joint witness. For instance, a diagram can alternate between two mutually exclusive observations using noninvertible transition arrows, hitting each requirement infinitely often and never satisfying both. Countably many persistent requirements need not share a common stage. The requirements $j \geq r$ for $r \in \mathbb{N}$ on the natural timeline each persist eventually and have empty intersection. A limit stage, if the chosen timeline includes one, requires a separate semantic continuity argument.

9 Recursive operations and coherent machine dynamics

The source note describes recursive self-improvement as iteration of a self-map on a truncated space of operations. The idea is meaningful with an explicit distinction between a space classifying operations and an individual operational type. We make this distinction first, then ask for a machine lift of a semantic orbit. The lift problem depends on the comparison arrows and the domain of the update rule. It is already nontrivial when every operation is discrete.

9.1 Typing the operation catalog

Let $\mathscr{O}$ be an object in an ∞-topos, interpreted as a catalog of operations, and let $\mathscr{O}_n = \tau_{\leq n}\mathscr{O}$. Suppose $F : \mathscr{O}_n \to \mathscr{O}_n$ and $o_0 : 1 \to \mathscr{O}_n$. Put $o_k = F^k o_0$. If desired, the external sequence can be assembled as a map from the coproduct $\coprod_{k\in\mathbb{N}} 1$ to $\mathscr{O}_n$, with the usual successor shift. Alternatively F may come from $\widetilde{F} : \mathscr{O} \to \mathscr{O}$ by functoriality of truncation. In that case naturality of the truncation unit ensures that the iterates obtained by truncating the original iterates are the same as those of F up to the supplied coherence.

The expression $Q_n(M) \simeq o_k$ requires interpretation. The left side denotes an object of $\mathcal{X}_{\leq n}$. The right side as written denotes a point of a catalog. For a small external catalog O, one can specify a classifier $j : O \to (\mathcal{X}_{\leq n})^{\simeq}$. Then the comparison is $Q_n(M) \simeq j(o_k)$. For an internal catalog, a family $E \to \mathscr{O}_n$ in the slice can be part

of the semantic data, with $j(o_k) = o_k^* E$. An n-truncated family gives n-truncated operational objects at each point. One can also work with a topos-valued functor of contexts and the analogous relative pullbacks. These are ways of making the link from a catalog to the semantic category explicit. They do not assert that $\tau_{\leq n}$ of an arbitrary universe of operations classifies all n-truncated operational objects. In general the universe of n-types has nontrivial automorphisms and different truncation behavior from its elements.

An endomorphism F assigns operations, but it need not supply a path $o \to F(o)$ in the core of the catalog. If F changes components no such path exists. A directed model of revisions uses a category $\mathcal{B}$ with objects representing operations and arrows representing allowed semantic steps, plus chosen arrows $b_k \to b_{k+1}$ along the orbit. This defines a functor $b : \mathbb{N} \to \mathcal{B}$ once composites and their coherence are chosen. One can freely generate a category from a graph of operation steps if the steps have no imposed relations, or impose further relations when different routes are required to agree. The semantic catalog is part of the input to this construction.

9.2 Three levels of realization of an orbit

For a functor $q : \mathcal{E} \to \mathcal{B}$ of allowed machine updates, there are three increasingly structured requirements. Stagewise realization requires that every b_k has a marked object in the fiber of q. A functorial realization consists of a functor $m : \mathbb{N} \to \mathcal{E}$ with a natural equivalence $qm \simeq b$. The third is a policy that can lift the orbit from every permitted initial machine, or a machine endofunctor lifting a semantic endofunctor on an entire class of operations. These demands have different quantifiers. A functorial realization is a section of the pullback $b^*\mathcal{E} \to \mathbb{N}$, interpreted in the homotopy-coherent sense. Choosing one point in each fiber need not give a section.

For example, let $\mathcal{B}$ be the category on $0 \to 1$, let $\mathcal{E}$ consist of two isolated objects over 0 and 1, and let the semantic path use the single nonidentity arrow. Both fibers are nonempty and the path has no lift. There is an obstruction even when each consecutive semantic arrow has some lift. Take a base $0 \to 1 \to 2$, with fibers consisting of a over 0, b, b' over 1, and c over 2. Let the only nonidentity edges upstairs be $a \to b$ and $b' \to c$, mapping to the two consecutive base arrows. Neither choice

over 1 can connect both edges. The failure concerns the shared intermediate object. Additional edges or homotopies that identify b with b' could remove the obstruction, but none are present in this example.

The opposite separation is also possible. One can have a functorial realization from some initial machines and no realization from others. With the same two-stage base, add another isolated object over 0 to a lifting pair. It realizes the first operation and cannot take the required step. An assumption that a semantics is realized at every stage gives no uniform replacement theorem at the level of machines.

9.3 Forward transport and iteration

Theorem 9.1 (Lifting an orbit by forward transport)**.** *Let $q : \mathcal{E} \to \mathcal{B}$ be a coCartesian fibration of ∞-categories. For every $b : \mathbb{N} \to \mathcal{B}$ and every chosen object over b_0, there exists a lift $m : \mathbb{N} \to \mathcal{E}$ starting at that object, with all arrows coCartesian over their images. The successive objects are obtained by the forward transport functors associated to the consecutive arrows of b.*

Proof. Straighten the pullback coCartesian fibration over $\mathbb{N}$ to a functor $K : \mathbb{N} \to \mathrm{Cat}_\infty$. A chosen initial object is a functor $* \to K(0)$. The constant diagram at $*$ on $\mathbb{N}$ is the left Kan extension of $*$ along $\{0\} \hookrightarrow \mathbb{N}$, since the category of arrows from 0 to each stage is contractible. The given functor thus extends to a natural transformation from the constant diagram at $*$ to K, with its initial value identified with the chosen object. Under unstraightening this is a section of the pullback taking every arrow to a coCartesian arrow. The value at each successor stage is the transport of the preceding value, by the defining property of coCartesian lifts. □

The proof uses the standard theory of coCartesian fibrations and straightening [17, 18, 26, Secs. 2.4 and 3.2]. The theorem is a sufficient condition for certified change along an orbit. It makes no requirement that the lift be computable. It also makes no assertion that the coCartesian choice is optimal for a cost functional. A coCartesian arrow is universal for factoring arrows over compatible base maps. Cost constraints on future edges may invalidate that arrow as a feasible policy if the

admissible subcategory after imposing those constraints no longer has such forward transports.

There is a useful refinement for the space of choices. Let $\mathcal{E}^{\mathrm{tr}}$ denote the category over $\mathcal{B}$ obtained by keeping only coCartesian arrows, so its fibers are groupoids. Straightening now gives a functor $\mathcal{B} \to \mathcal{S}$ of cores of the original fibers with their transports. CoCartesian lifts of b, including equivalences of lifts over the identity of b, are sections of the resulting left fibration over $\mathbb{N}$. Their space is the limit of the diagram of fiber spaces. Because 0 is initial in $\mathbb{N}$, the limit is equivalent to the value at 0. This can be seen directly by propagating the initial point and the required coherences along its unique maps to later stages. The space of such lifts with a specified marking of the initial object is thus contractible. Arbitrary lifts using non-coCartesian arrows can still vary widely. This result distinguishes coherent deterministic transport up to contractible choice from the collection of all legal paths.

A Cartesian fibration supplies transport in the other direction. It gives canonical lift choices from a prescribed final object for finite intervals. It need not produce a compatible infinite forward choice from an initial machine. Consider the right fibration over $\mathbb{N}$ for the inverse system of discrete sets $X_k = \{k, k+1, \ldots\}$ with inclusions. Every fiber is nonempty and every finite initial interval has a section. A section of the full timeline would give an integer in their intersection, which is empty. This example mirrors the difference between compatibility and stagewise existence in the Postnikov analysis.

9.4 Commuting update policies and semantic invariance

A global machine update rule is a functor $G : \mathcal{E} \to \mathcal{E}$. Suppose the semantic update is $H : \mathcal{B} \to \mathcal{B}$ and a natural equivalence $qG \simeq Hq$ is part of the data. Iterating it yields equivalences $qG^k \simeq H^k q$ coherently by composition. In particular $G^k(M_0)$ realizes $H^k(qM_0)$ as objects. To supply a directed machine trajectory one may additionally give an update arrow or a natural transformation $\eta : \mathrm{id}_{\mathcal{E}} \to G$ mapping to the intended semantic step. The functor G alone is an object update that acts on comparison morphisms. It does not canonically give $M \to G(M)$. When a natural transformation exists, its iterates generate a functorial trajectory. A semantic update $\theta : \mathrm{id}_{\mathcal{B}} \to H$

can be lifted canonically up to the appropriate coherence by coCartesian transport in the setting of Theorem 9.1. One constructs $G(M)$ as the transported object along θ_{qM}, and defines its effect on arrows using naturality of θ and the coCartesian universal property.

Conversely, if a machine functor G is given first, its action can descend to an observational quotient. If q is a localization at W, then qG factors through q exactly when $qG(w)$ is invertible for all $w \in W$, by Proposition 5.1. When q is instead just a projection to a set of behaviors, the analogous condition is constancy of qG on the fibers of q. The distinction matters because the first criterion includes effects on comparison arrows, while the second is only about values. For example, a rewrite could consult an internal counter while preserving the initial semantics on a handful of machines and fail to do so uniformly on an entire behavioral fiber. Existence of an induced semantic update requires uniformity.

9.5 Periods and equivariance

An infinite lift of a repeating semantic sequence need not repeat with the same period. Let the base be BC_2, and consider the left fibration classified by the regular C_2-set C_2. Its monodromy around the nonidentity loop swaps the two machines in the fiber. An orbit in which that loop is used at every step has a lift along the natural timeline by transport, alternating between the two machines. A lift respecting the loop relation as a single globally chosen machine would be a section over BC_2, which is equivalent to a fixed point of the action. No fixed point exists. Traversing the generator twice returns to the starting machine. Thus one can distinguish a lift on the linear sequence of visits, a lift invariant under one period, and a lift invariant after passing to a larger period.

More generally for a discrete set of machines with an action of a group G, a section over BG corresponds to an invariant machine. If the fiber is a space, the same section condition uses homotopy fixed points. A candidate fixed class in π_0 may still require higher coherence for a homotopy-fixed point. The torsor of Section 7 shows failure already at π_0. Higher twists can create further obstructions even if all the local machine choices lie in a single component. The 2-type in Section 3, viewed as a

fibration over BC_2, has a connected fiber and no section. It exhibits an obstruction to a coherent choice despite the absence of a set-level obstruction.

9.6 Equality and improvement

A semantic fixed point $H(b) \simeq b$ says that the new operation lies in the same class, with no numerical comparison. A machine trajectory may continue through inequivalent realizers in the fiber. An order or a functional on the machines can regard some transitions as improvements while observations are invariant. Conversely a semantic step may lead to a new output behavior and a worse cost. To reason about improvement, the next section adds quantitative predicates and distinguishes pointwise optima, feasible execution, and uniform performance bounds.

10 Feasible improvement, viability, and uniform budgets

We now use ordinary sets and directed graphs. This specialization makes the choice and computability issues explicit. Fix a semantic orbit $b = (b_k)$, and let X_k be the set of admissible marked machines at stage k. For each consecutive step let $E_k \subseteq X_k \times X_{k+1}$ be the relation of certified updates. A path in this graph gives an ordinary functorial lift if the updates are taken as generating arrows over the orbit and their composites are allowed. In a category with several distinct certificates for the same step one may replace E_k by a set of edges with source and target maps. The existence arguments below then apply to that multigraph. They do not identify different certificates as higher homotopies.

10.1 Local optimality and feasible continuation

A scalar cost $c_k : X_k \to [0, \infty)$ might measure state count, a declared resource metric, or a cost model that has been justified independently. Two states can have the same observed semantics and different costs. The pointwise minimizer of c_k need

not be compatible with any minimizing state of c_{k+1}. For example, let $X_0 = \{a, a'\}$, $X_1 = \{b, b'\}$, with edges $a \to b'$, $a' \to b$. Give a, b cost 0 and a', b' cost 1. Both pointwise minima are unique and cannot be composed into an execution. At least one unit of cost is unavoidable somewhere on every lift. Additional reachable states at later times can make the first, locally cheaper choice incompatible with every infinite lift. Thus an optimization criterion needs to refer to feasible paths or to a transport theorem ensuring compatibility.

It is also possible for a semantic improvement measured at one layer to violate an invariant at another. For example, changing a Mealy machine from S_m to S_1 preserves constant output and reduces its state count. Adding the phase port makes the change distinguishable. An evaluation policy declaring the quotient admissible under the original interface does not imply that it remains admissible when the interface changes. One can impose a joint observation consisting of the original behavior and the retained phase behavior. The allowed transformation graph is then the one computed for that joint observation. No universal relationship among implementation costs follows from an equivalence of observable behavior.

10.2 A greatest fixed point for extendible machines

Let $X = \coprod_{k\in\mathbb{N}} X_k$, remembering the stage even when machine presentations repeat. For $U \subseteq X$ define

$$\Gamma(U) = \{x \in X : \exists\, y \in U \text{ with an allowed edge } x \to y\}.$$

It is monotone under inclusion. A subset with $U \subseteq \Gamma(U)$ supplies at least one successor within U for every state it contains.

Theorem 10.1 (Viability of a certified update)**.** *In classical set theory with dependent choice, a state x admits an infinite certified continuation if and only if it belongs to the greatest fixed point of Γ. If each state has finitely many outgoing edges, this fixed point is*

$$V_\omega = \bigcap_{j<\omega} \Gamma^j(X).$$

Consequently, in the finitely branching case arbitrarily long legal continuations from a fixed *state imply an infinite one.*

Proof. The union V of all U with $U \subseteq \Gamma(U)$ is itself postfixed by monotonicity. Since $\Gamma(V)$ is then also postfixed, $\Gamma(V) \subseteq V$, and hence $V = \Gamma(V)$. This is the usual greatest-fixed-point argument [28]. Every infinite continuation has a set of states that is postfixed and contains its starting state. Conversely dependent choice applied to the successor relation on V yields an infinite continuation from any $x \in V$.

For the second claim the sets $V_j = \Gamma^j(X)$ decrease, and V_j is exactly the set of states with a continuation of length j, by induction. Fix $x \in V_\omega$. For every j, x has a successor in V_j because $x \in V_{j+1}$. The sets of its successors in V_j are a decreasing sequence of nonempty subsets of a fixed finite set. Their intersection is nonempty. Hence x has a successor in V_ω, making V_ω postfixed. Every postfixed set is contained in each V_j, again by induction. The equality follows. □

The same proof of the finite branching statement can be phrased as Kőnig's lemma on the tree of finite legal paths from x. This form makes it clear why the state must be fixed. One may have an arbitrarily large finite duration somewhere in the union of the initial states while each initial state separately has bounded duration. It is also important that finite branching is a hypothesis about outgoing alternatives from a state, including the admissibility predicates. Finite presentation of the rule set need not imply finite branching, for example when a rule may synthesize an integer parameter of arbitrary size.

For a counterexample with infinite branching, let a root have one edge to the start of a chain of length m for each $m \in \mathbb{N}$. Every continuation commits to one such chain and terminates. The root lies in V_ω but has no successor in V_ω. The root is removed in the next pruning step. For a set of states in general, transfinite pruning computes the greatest fixed point. Start at X, apply Γ at successors, and intersect at limits. The process is decreasing. Before reaching a fixed point each successor step removes a state not removed before, so stabilization occurs before the process could produce distinct removals indexed by a cardinal strictly larger than $|X|$. Every postfixed set is contained in every stage by transfinite induction. The stable stage is therefore the

greatest fixed point. The argument concerns existence and does not prescribe an effective stabilization procedure.

10.3 Attainment of a uniform path budget

Let $x \in X_0$ be fixed and suppose its legal path tree is finitely branching. Each edge e has a cost $c(e) \geq 0$. For a finite path s, write $C_\Sigma(s) = \sum_{e \in s} c(e)$ for cumulative cost and $C_{\max}(s) = \max_{e \in s} c(e)$ for bottleneck cost, with empty value 0 in both cases. Fix one of these objectives and write $C(s)$. For an infinite path s, set $C(s) = \sup_m C(s|_m)$. Let b_m be the minimum cost of a legal path of length m, with $b_m = +\infty$ when none exists. If there is an edge of cost $+\infty$, the same definitions can be used in the extended nonnegative reals.

Theorem 10.2 (Uniform budget from finite horizons)**.** *If there are legal paths of every finite length and the path tree is finitely branching, there is an infinite legal path and*

$$\min_{s \text{ infinite}} C(s) = \sup_{m \in \mathbb{N}} b_m.$$

In particular, if $B < \infty$ and each horizon has a legal path with cost at most B, then a single infinite path stays within B.

Proof. There are finitely many paths of any fixed length, so finite-horizon minima are attained whenever such paths exist. Costs are monotone under extending a prefix. Every infinite path thus has cost at least $L = \sup_m b_m$. If $L < \infty$, keep only prefixes whose cost is at most L. They form a prefix-closed finitely branching tree containing a path at each depth because the minima b_m are attained. By the finite branching conclusion of Theorem 10.1, there is an infinite branch. Its prefixes have cost at most L, so it has cost L. If $L = +\infty$, apply the same existence argument to the whole tree. Every infinite branch then has cost $+\infty$ by the lower bound. □

For the bottleneck objective, the budget condition can also be imposed by discarding all edges above B. For cumulative cost the remaining budget is part of the relevant state, since the cost of an edge that can be taken depends on the already traversed prefix. One may view the tree of budget-feasible prefixes as the new layered graph.

The theorem uses nonnegative increments for the cumulative objective. More generally it applies to any prefix-monotone extended cost on finite paths. For a prefix-monotone vector cost and a fixed componentwise bound one can use the same pruning argument. A claim of a single componentwise optimal path is stronger. Distinct paths may optimize different components, and the componentwise infimum may be unattainable even at a finite horizon.

Finite branching also supports a useful description of the set of feasible budgets. At any fixed bound it is sufficient to examine whether the corresponding prefix tree has nodes at all depths. The closed feasible budget interval begins at L when $L < \infty$. There need not be a finite horizon m witnessing equality $b_m = L$. Along the unique path with edge costs 2^{-k-1} at stage k, the cumulative optimum is 1 and every finite optimum is strictly smaller. Thus finite-horizon optimization approximates the uniform optimum monotonically, while certification of the exact threshold can require information about the whole tail.

10.4 Computability of paths and observation

Even a computable finitely branching tree need not have a computable infinite branch. Here is a direct diagonal construction. Fix an effective enumeration (φ_e) of partial computable functions on $\mathbb{N}$ and a step-counted universal simulation. Keep a binary string s of length m under the following condition. For every $e < m$, if $\varphi_e(e)$ halts within m steps with a bit d, then $s(e)$ must equal $1 - d$. Membership is decidable by bounded simulation. The tree is prefix closed because a computation halting within a shorter bound would also be detected at a longer bound, and constraints at one depth concern distinct coordinates and therefore are consistent. There are strings at every depth. An infinite branch exists by finite branching. If a branch $a : \mathbb{N} \to \{0, 1\}$ were computable with index e, then $\varphi_e(e) = a(e)$ would eventually halt, forcing $a(e) = 1 - a(e)$. This is impossible. The tree gives a computable update graph with viable initial state and no computable certified continuation. It is the usual kind of obstruction distinguished in work on weak Kőnig's lemma and computable trees [5].

There is no uniform algorithm deciding viability of the root for computable binary trees. For a program with index e, keep the zero string of length m if the program

on input e has not halted in the first m steps, and keep no other string of length m. The resulting tree is uniformly decidable and has an infinite path exactly when the computation never halts. A decision algorithm for viability would decide nonhalting and hence halting. This reduction holds even for a graph with at most one outgoing edge per state. An oracle for the viability predicate would make the construction of a path effective when successors can be effectively listed, by selecting the first viable successor at each step. Thus a proof of viability, a decision procedure for viability, and a computable realization policy are distinct resources.

Observation equality is a separate computability question. For total computable unary output sequences, construct a program that emits 1 on input m exactly when a specified computation has halted within m steps and otherwise emits 0. It agrees everywhere with the constant-zero sequence if and only if the computation never halts. Hence equality of the chosen total trace functions can already be undecidable even though each requested output can be computed. This is an explicit reduction for the interface considered here. Rice's theorem concerns nontrivial extensional properties of recursively enumerable sets and gives broader context for such separations [25]. Neither reduction rules out decidable equivalence for the bounded deterministic finite-state class treated in Section 4.

10.5 An explicit constrained redesign

Return to the constant-output cycles S_m. If only strict quotient changes are allowed, a path of redesigns is a chain $m_{k+1} \mid m_k$. Under the cost $c(S_m) = m$, a strict improvement is possible only when the divisor is proper, and any strict chain starting at m_0 has length at most $\lfloor \log_2 m_0 \rfloor$. This follows from $m_{k+1} \leq m_k/2$ at each proper step. The trajectory can continue indefinitely only if nonstrict steps are allowed, in which case it eventually stays at a constant size. The trace is already fixed throughout. If instead the policy requires retaining every phase in an added port, even the step to S_1 can be forbidden by the change in behavior. The improvement claim therefore depends on the declared behavioral contract, arrow policy, and cost.

One can also give a cumulative budget on the allowed changes, for example by attaching cost 1 to any proper quotient and 0 to a no-op. A uniform budget then bounds

the number of proper changes. Imposing a requirement that the cycle eventually reach S_1 is an additional liveness condition. With a no-op allowed, the viability of an infinite path with a small budget may be witnessed by a path that never reaches S_1. For a liveness claim one can impose a fixed deadline in the admissible prefixes, or prove a progress condition that forces target reachability. The compactness theorem controls prefix-closed feasibility. Arbitrary liveness properties need a separate proof.

11 Combining observation levels and redesign constraints

The preceding results describe different projections of the same problem. We now combine them in a model where the distinctions can be seen simultaneously. We then prove two comparison results. The first identifies coherent lifts along all observation degrees. The second gives a finite-branching compactness condition that combines the horizon and observation degrees without assuming existence of compatible witnesses in advance.

11.1 A family with two independent machine invariants

Fix $G = C_2$, $U = BG$, and the unary constant-zero behavior b of Section 6. Let L be the category of positive cycle lengths with one arrow $m \to d$ when $d \mid m$. Regard the arrow as the strict quotient $S_m \to S_d$ of initialized machines. Form the machine category

$$\mathcal{C}_b = L \times \mathcal{S}_{/U}.$$

A machine is a pair $(m, E \to U)$ whose first factor is its reachable transducer and whose second factor specifies auxiliary data parametrized by contexts. For a variable initial behavior, the corresponding ambient semantic slice can be taken over $V \times U$, with V the discrete set of behavior values. In the fixed-behavior model we can omit the constant b and set $\mathrm{Op}(m, E) = E$ in $\mathcal{S}_{/U}$. The slice equivalences retain the base map, so changing context remains part of the operational data.

Use E_f, E_0 from Section 3. Both are relatively 2-truncated, and their relative 1-truncations are the terminal object over U. Therefore all pairs (m, E_f), (d, E_0) agree at relative degree 1. At degree 2, E_f and E_0 are distinguished by the existence of a section. This separates machines with different auxiliary families even if the cycle length is the same. Changing the length alone remains invisible at all degrees of this operational functor, although the transducers can be distinguished by the additional phase port. Thus refinement of truncation and enlargement of the family of tests eliminate different ambiguities.

One can explicitly determine the fiber at the degree-one semantic object $p = \mathrm{id}_U$. Let $K \subset \mathcal{S}_{/U}$ be the full subcategory on families $E \to U$ with relatively contractible 1-truncation. The space of equivalences from each such truncation to p is contractible, since p is terminal. Every arrow between such truncations is an equivalence. Consequently the transformation fiber of $(m, E) \mapsto \tau^U_{\leq 1} E$ is equivalent to $L \times K$, and its marked core to $L^{\simeq} \times K^{\simeq}$. It has a terminal transformation $(1, p)$. However its marked core has distinct components for every cycle length and for both auxiliary families above. Terminality of the coarse transformation model therefore coexists with two independent ways to distinguish machines.

Now require a certificate in the form of a global section of $E \to U$, together with compatibility with the section for all allowed changes. This restriction replaces the category of families by a category of pointed families. The E_f-machine is excluded because its section space is empty. Nevertheless it is eligible in each context of the cover $* \to BG$ before descent. The pullback has fiber $K(\mathbb{F}_2, 2)$, which has points. Thus the additional admissibility condition changes the fiber by using coherent global witnesses even though the original degree-one semantics was constant. If one separately requires an update to preserve the degree-two family, its projection to p cannot be certified as degree-two preserving unless the family was already equivalent to p. Each of these restrictions changes a different mathematical ingredient.

11.2 Coherent recovery of machine fibers

Let $Q_\infty = \mathrm{Op} : \mathcal{C} \to \mathcal{X}$, and suppose a semantic object $p \in \mathcal{X}$ is fixed. Put $p_n = \tau_{\leq n} p$, with the coherent transition identifications. There are natural maps of

marked realization spaces

$$\mathrm{Real}_{Q_\infty}(p) \longrightarrow \mathrm{Real}_{Q_n}(p_n), \qquad \mathrm{Real}_{Q_{n+1}}(p_{n+1}) \longrightarrow \mathrm{Real}_{Q_n}(p_n).$$

The transition applies truncation to the marking and retains the machine. One can compare the full fiber with the entire inverse system.

Theorem 11.1 (Recovery of compatible marked realizers)**.** *Suppose the Postnikov towers of p and every object $\mathrm{Op}(M)$ converge in $\mathcal{X}$. Then the canonical map*

$$\mathrm{Real}_{Q_\infty}(p) \longrightarrow \lim_{n\geq -2} \mathrm{Real}_{Q_n}(p_n)$$

is an equivalence of spaces.

Proof. The limit of the diagram of marked fibers is the pullback of the limits of the three diagrams that define them. The machine-core diagram is constant and its limit is $\mathcal{C}^{\simeq}$, since the index has contractible groupoid completion. The point diagram has limit a point with image the tower of p. The remaining term is $\lim_n(\mathcal{X}_{\leq n})^{\simeq}$. Thus the right-hand side is the space of machines with an equivalence between the truncation tower of $\mathrm{Op}(M)$ and the given tower of p. The comparison is a map over $\mathcal{C}^{\simeq}$. On the homotopy fiber over any M it is the comparison $\mathrm{Eq}_{\mathcal{X}}(\mathrm{Op}(M), p) \to \lim_n \mathrm{Eq}_{\mathcal{X}_{\leq n}}(Q_n M, p_n)$, an equivalence by the mapping-space calculation in Section 3. Equivalences of spaces over a base are detected on their fibers. □

This theorem can be read as a precise assertion about machine identity across degrees. A point on the right is a compatible family of marked realizers, with machine equivalences as part of the transition compatibilities. It does not permit choosing an unrelated machine at every degree. A proof that every $\mathrm{Real}_{Q_n}(p_n)$ is nonempty only establishes existence at finite degrees. For example, in spaces let $p = \prod_{k\geq 1} K(\mathbb{F}_2, k)$, and let the allowed machines be the discrete collection indexed by $r \geq 0$ with operational object $P_r p$. At any finite degree n with $n \geq 0$, every machine $r \geq n$ is a realizer of p_n. No machine in the collection realizes p, since each machine's operational object has a largest possible nonzero homotopy degree and p has nonzero groups in arbitrarily high degrees. The tower of realization spaces is therefore termwise

nonempty with empty limit. This example uses convergent Postnikov towers in spaces and an explicit failure to choose a fixed admissible machine. It places the obstruction in the allowed machine class, without invoking any failure of Postnikov convergence.

If the machine category includes a machine realizing p, further conclusions need hypotheses about full-precision equivalence. In the independent decoration model of Section 3, Op can be a projection with a large fiber even at infinite precision. By contrast, if Op is fully faithful on cores, every inhabited full-precision marked fiber is contractible by Proposition 2.4. In that case the limit in Theorem 11.1 recovers a unique marked machine up to coherent equivalence whenever it is nonempty. Neither conclusion follows solely from convergence of operational objects.

11.3 A joint compactness theorem for observation and time

The budget argument can be adapted to handle many observation levels at once. Consider a finitely branching tree T of legal execution prefixes from a fixed initial machine. For each $n \geq 0$, let $T_n \subseteq T$ be a prefix-closed subtree. Membership in T_n means that the entire finite prefix is certified at observation degree n, according to the chosen test on the prefix. Assume that the conditions are nested, so $T_{n+1} \subseteq T_n$. This applies to constraints preserved by lowering the degree, including compatibility tests on certified replacement arrows when those tests are imposed consistently on prefixes.

Theorem 11.2 (Joint finite-horizon and finite-degree feasibility)**.** *Suppose for every pair n, m there exists a path of length m in T_n. Then T has an infinite branch every finite prefix of which belongs to every T_n. If T_n also enforces a uniform prefix-monotone budget, the same branch satisfies that budget at every horizon.*

Proof. For fixed m, the set of all paths of length m in T is finite. The paths of length m lying in T_n form a decreasing family of nonempty subsets of it. The intersection is nonempty. Hence $\bigcap_n T_n$ is a prefix-closed finitely branching tree with a node at every depth. It has an infinite branch by Theorem 10.1. The additional budget condition holds on every prefix by membership in the intersection. □

The theorem concerns paths whose finite prefixes satisfy all the predicates being imposed. If the predicates say that a particular displayed arrow is sent to an equivalence

at degree n, the branch uses the same arrows for all n. When the operational topos is hypercomplete, these arrows then induce equivalences of full operational objects. If the predicates say only that endpoints can be matched by some equivalence, higher compatibility may require more work, as Theorem 11.1 shows. One can sometimes include equivalence certificates in the path tree to impose coherence, but finite branching must then be checked for the enlarged tree, including those certificates.

Failure of finite branching can invalidate the result in a simple model. From a root choose an integer r, then follow an infinite path marked r on every edge. Define T_n to retain the root and precisely the branches with $r \geq n$. Every T_n has paths of all lengths, while the intersection has no edge out of the root. The same issue arises if the first stage can select any one finite approximation to a desired operation but the policy never permits replacing that approximation. Some additional compactness condition, distinct from finiteness of each test, is what permits combining the constraints.

11.4 What a model must supply

The examples support a concise division of assumptions. An operational functor supplies the object and arrow observations. A specification determines a predicate on semantic objects, with a separate assertion of realization. A machine equivalence relation determines the core whose fibers classify implementations. A certified update model specifies which arrows actually connect those implementations, possibly across different semantic objects. A cost functional and a feasibility condition specify the mathematical sense in which those arrows can be used to improve a machine. The theorems make different demands on each component.

In particular, there is no common notion of *uniqueness* that subsumes the computations in the paper without discarding useful information. A constant trace can have a universal strict quotient and many inequivalent machines. The unique object class of BH can have several marked choices even for a one-object machine category. A local family can be trivial on every chosen chart and obstructed globally. A cost threshold can be met at every finite horizon and require a noncomputable policy on the full timeline. Each example points to a distinct comparison that the formalism can express.

11.5 Conclusion

The source note's proposal becomes precise when semantic equivalence is paired with a model of the choices it forgets. The marked fiber records machines and their identifications with an observed type. The associated transformation category records permissible one-way comparisons. Postnikov towers control which higher operational distinctions become visible, and local descent controls whether choices can be made across contexts. The finality computation and transport theorem then specify when a timeline covers a class, when it computes colimits, and when it can be lifted as a certified sequence.

An application to a concrete architecture would need an operational semantics and a proof that its trace and certification layers represent the intended tests and changes. The mathematical results here apply once those data are supplied. The set-based examples show why existence of a semantic iteration alone gives no conclusion about convergence, reachability of a preferred machine, or resource improvement, and the positive results state conditions under which particular such claims can be proved.

A Calculus of observation refinements

This appendix collects comparison results for changing observation models. Its constructions apply to any composable functors, and specialize to consecutive Postnikov degrees. They can also be used to compare a semantic quotient obtained by a new test with one obtained by discarding higher data.

A.1 Decomposing a coarse fiber over its refinements

Let $Q : \mathcal{C} \to \mathcal{D}$, $T : \mathcal{D} \to \mathcal{B}$, and fix $b \in \mathcal{B}$. Define

$$\mathcal{D}_b = \mathcal{D}^{\simeq} \times_{\mathcal{B}^{\simeq}} \{b\}.$$

Its points are candidate fine observations d equipped with a marking $\delta : Td \simeq b$. There is a canonical forgetful map from $\mathrm{Real}_{TQ}(b)$ to $\mathcal{D}_b$, sending (M, α) to (QM, α). It lets one recover the coarse realization space from the fine ones, with monodromy as part of the answer.

Proposition A.1 (Fiber pasting for semantic refinement)**.** *The homotopy fiber of* $\mathrm{Real}_{TQ}(b) \to \mathcal{D}_b$ *at* (d, δ) *is* $\mathrm{Real}_Q(d)$*. These fibers, with the transports induced by paths in* $\mathcal{D}_b$*, have colimit* $\mathrm{Real}_{TQ}(b)$ *in spaces. In particular, the forgetful map is an equivalence onto the union of realized components of* $\mathcal{D}_b$ *if every inhabited fine marked fiber is contractible.*

Proof. The square

$$\begin{array}{ccc} \mathrm{Real}_{TQ}(b) & \longrightarrow & \mathcal{C}^{\simeq} \\ \downarrow & & \downarrow{\scriptstyle Q^{\simeq}} \\ \mathcal{D}_b & \longrightarrow & \mathcal{D}^{\simeq} \end{array}$$

is a pullback by associativity of pullbacks. Pulling it back along the specified point (d, δ) gives the fiber formula. A map of spaces over $\mathcal{D}_b$ is a local system of spaces, and its total space is the colimit of that local system. This follows by straightening for spaces, or by the same groupoid-completion argument as in Theorem 8.1. Finally a map of spaces with empty or contractible fibers is an equivalence onto its image of components. □

This formula preserves gluing. For instance, if $\mathcal{D}_b = BH$ is connected and the fine fiber is K, the coarse space is equivalent to the homotopy quotient of K by the induced coherent H-action. It agrees with $BH \times K$ when the local system is equivalent to the constant one. Even if all fine fibers are mutually equivalent, that equivalence alone does not trivialize their transport. The groupoid calculation Proposition 2.5 is an instance. There the homotopy quotient of the set of markings sees both orbits and stabilizers. At the level of components, loops in the base can exchange different fine realizer components, so simply taking a disjoint union over representatives of $\pi_0(\mathcal{D}_b)$ would miss identifications.

For consecutive degrees take $\mathcal{D} = \mathcal{X}_{\leq n+1}$, $\mathcal{B} = \mathcal{X}_{\leq n}$, and $T = \tau_{\leq n}$. The base $\mathcal{D}_b$ is the space of possible next-stage types with a chosen equivalence of their truncation with b. Restricting to realizable fine observations removes its empty-fiber components. A next-stage type, even when realized, can have many inequivalent implementations, which remain in the fine fiber. The result explains why the phrase *fiber of truncation*

and the phrase *fiber of machine semantics* refer to different maps. The latter can be analyzed through the former together with an independent realization problem.

A.2 Interpreting compatible tests by Yoneda

There is a test characterization of the relation $M \equiv_n N$ that does not require choosing coordinates for the homotopy groups. For each operational object X define the covariant functor

$$\mathcal{P}_n(X) : \mathcal{X}_{\leq n} \longrightarrow \mathcal{S}, \qquad A \longmapsto \mathrm{Map}_{\mathcal{X}}(X, A).$$

This tests maps *out of* X into every n-truncated object. The adjunction identifies this with the corepresentable functor of $\tau_{\leq n}X$. The covariant Yoneda embedding is fully faithful on the opposite category, so a natural equivalence of the two test functors is equivalent to an equivalence of their corepresenting objects. In particular,

$$X \equiv_n Y \quad \Longleftrightarrow \quad \mathcal{P}_n(X) \simeq \mathcal{P}_n(Y),$$

where the notation on the left abbreviates equivalence of n-truncations in the operational topos. This statement uses all n-truncated test objects and their maps. Arbitrary componentwise bijections of the sets of test results are not the supplied natural equivalence. Selecting a small family of probes can be useful, but a conservativity or density statement about that family is then needed to recover the same relation.

For ordinary sets and the case $n = 0$, a labeled output can be treated by the same test logic in a slice. If V is the catalog of values, an object p_v is the inclusion of the singleton v. In $\mathcal{S}_{/V}$, $\mathrm{Map}(p_v, p_w)$ is empty for different labels and contractible for equal labels. The labels therefore give a complete family of probes for this class of singleton objects. Forgetting the slice replaces all these probes by maps of unlabeled singletons. A statement about the sensitivity of the tests must refer to the category in which the maps are computed.

In the other direction, an equivalence of two n-truncated objects can be detected by mapping *into* them from every object of the topos, naturally in the domain. On a site presentation it often suffices to take the designated representables after sheafification, because the induced functors of sections on the site determine the sheaf. The direction

of the first test family is adapted to the truncation adjunction. For a map into a general operational object there is no rule replacing $\mathrm{Map}(A, X)$ with $\mathrm{Map}(A, \tau_{\leq n} X)$ as an equivalence for arbitrary A. In fact these mapping spaces can have different components when A has nontrivial homotopy.

A.3 An example with higher marking homotopies

The group calculation in Proposition 2.5 has a genuinely higher counterpart even with a single machine. Let $\mathcal{C} = *$, take $\mathcal{X} = \mathcal{S}$, and send the sole object to $p = K(A, n)$ for a nonzero abelian group A and $n \geq 1$. At level n there is exactly one unmarked machine class. The space of marked realizers is

$$\mathrm{Real}_{Q_n}(p) \simeq \mathrm{Aut}_{\mathcal{S}}(p).$$

Evaluation at the base point gives a map $\mathrm{Aut}(p) \to p$ whose fiber consists of based self-equivalences. The space of based maps $K(A, n) \to K(A, n)$ has components $H^n(K(A, n); A) \cong \mathrm{Hom}(A, A)$. For $j > 0$, its jth homotopy groups are reduced cohomology in degree $n - j$ of $K(A, n)$, with zero groups for negative degrees. This follows from representability of cohomology and the based suspension adjunction. Translation in the mapping space uses the abelian group model of $K(A, n)$, so the calculation applies at every component. Reduced cohomology vanishes below degree n, and consequently the based self-equivalence space is discrete with components $\mathrm{Aut}(A)$. The unbased self-equivalence space has a component equivalent to $K(A, n)$ for each automorphism of A. Indeed translation by a point of the topological or simplicial abelian group model of $K(A, n)$ carries a based map to an unbased map taking the base point to that point, and subtracting the value at the base point reverses the construction.

For $A = \mathbb{F}_2$ there is only one automorphism. The marked realization space is then connected and has nontrivial π_n, despite the underlying category having a single object and only its identity. This extra structure belongs to the space of choices of semantic identifications. It is not a hidden automorphism of the machine. After applying degree $n - 1$, the target becomes a point and the marked space becomes contractible. This example makes the distinction visible in an actual truncation

comparison, complementing the discrete group examples.

A.4 Invariance under a change of model

Suppose two presentations of the machine category are related by an equivalence $\Phi : \mathcal{C} \to \mathcal{C}'$ and there is a natural equivalence $Q'\Phi \simeq Q$ with the semantic functors. Taking cores and pullbacks gives equivalences of all their marked fibers. The same result for their transformation fibers follows because the square of functors commutes up to equivalence and Φ is an equivalence of categories. If Φ only induces an equivalence after localizing at selected rewrite arrows, the statement holds for the corresponding localized models when both semantic functors descend, and can fail for the unlocalized cores. The Mealy quotient of Section 6 is a concrete example of why the two comparisons differ.

When the change of presentation is only an equivalence on cores, the marked fiber classification survives but the transformation classification need not. For example $B\mathbb{N} \to *$ has equivalent cores, while their transformation categories over the point differ. Thus a proof that a representation change preserves observable object classes is weaker than a proof that it preserves reversible machine comparisons, and the latter is weaker than a proof that it preserves certified update arrows. This hierarchy is one reason to keep the choice of category explicit in computational applications.

B Further comparisons for local and directed choices

B.1 Obstruction by a central extension

The torsor example has empty global section space because there is no invariant point even at the set level. A central extension provides a different obstruction with a connected fiber. Consider the quotient $C_4 \to C_2$ with kernel C_2. The induced map $BC_4 \to BC_2$ has fiber BC_2 by Proposition 2.5. In particular its fiber is connected. A section would induce a homomorphism $C_2 \to C_4$ whose composite with the quotient is

the identity on C_2. A nonidentity element of C_2 can only be lifted to an odd residue modulo 4, and each odd residue has order 4. Hence the homomorphism cannot exist. The section obstruction is present even though every point of the base has a connected, nonempty fiber and the action on π_1 of that fiber is trivial.

The obstruction can be computed on the two group elements. Write $C_2 = \{0, 1\}$ additively and choose the set-theoretic lift $s(0) = 0, s(1) = 1$ in C_4. The defect

$$c(g, h) = \frac{s(g) + s(h) - s(g + h)}{2} \mod 2$$

is 1 on $(1, 1)$ and 0 on the other pairs. The numerator is interpreted in the subgroup $\{0, 2\} \subset C_4$, with division by 2 its isomorphism to C_2. Associativity of addition in C_4 gives the cocycle equation. Changing the lift by a function $2t$ changes c by the coboundary of t. A normalized t has coboundary zero at $(1, 1)$, because $t(0) = 0$ and $t(1) + t(1) = 0$ in C_2. Thus c represents a nonzero class in $H^2(C_2; C_2)$, in agreement with the classification of central extensions [4, Ch. IV].

In the language of machine updates, a change of context can be lifted individually for each group element by selecting $s(g)$. Composition of those changes is incompatible with the specified order-two relation. Along a linear timeline using the generator, one can lift each step by $1 \in C_4$ and retain every composite. At the second step the total machine symmetry is 2, which acts trivially on the observation and nontrivially in the fiber. To impose the relation after two steps would require coherently trivializing that symmetry. This is a different obstruction from an inaccessible successor in an update graph.

B.2 A transport criterion in ordinary categories

For clarity, the forward transport hypothesis admits a direct description in ordinary categories. Let $q : \mathcal{E} \to \mathcal{B}$. For $x \in \mathcal{E}$ and $f : qx \to b$, an arrow $u : x \to y$ over f is coCartesian when for every $v : x \to z$ and every $g : b \to qz$ with $qv = gf$ there exists a unique $w : y \to z$ over g satisfying $wu = v$. The arrow selects a target with a universal factorization. Choosing any merely admissible successor supplies existence of a single step but need not supply this universal property.

Suppose now q is constructed from an ordinary functor $K : \mathcal{B} \to \mathbf{Cat}$. A machine over b is $x \in K(b)$, and an arrow over $f : b \to b'$ is a morphism $K(f)x \to y$ in $K(b')$. The arrow over f using the identity of $K(f)x$ is coCartesian. To factor an arrow $K(gf)x \to z$ after it, use the same arrow as a morphism from $K(g)(K(f)x)$. Uniqueness is immediate. Thus a family of update categories with compositional forward transport supplies the desired fibration. If the chosen rule does not transport every object, one can restrict the source fibers to objects with viable continuations along a particular timeline. This can restore existence of successors along that timeline, but it does not by itself give a universal factorization for arrows over all base maps.

B.3 Two restrictions with different closure properties

One natural restriction retains only arrows with a semantic certificate, such as a proof that the displayed comparison is an equivalence after truncation. This is closed under composition because equivalences compose. Another restriction retains only edges under a prescribed one-step latency bound. Closure under composition depends on the cost rule. Adding two durations each below the bound can exceed the bound. One can instead take paths as composites with explicit accumulated budgets or use a filtered collection of arrow sets $E_{\leq B}$ with

$$E_{\leq B_2} \circ E_{\leq B_1} \;\subseteq\; E_{\leq B_1 + B_2}.$$

These choices affect whether the admissible comparisons form a category over the semantic steps or are better treated as a graph of immediate actions. The viability result was stated for such a graph to avoid assuming an invalid closure under composition. The categorical transport result assumes a category, so if applied to budgeted redesign it requires an explicit composition rule. In either setting preservation of a semantic condition can be checked separately from the cost bound.

A third possibility is to require progress without a uniform cost bound, for example by assigning a rank in a well-founded strict order and decreasing it at each proper step. Then an infinite trajectory can contain no infinite sequence of proper steps. If the rank is a nonnegative integer, the number of proper steps is bounded by the initial rank. This condition is compatible with an infinite timeline containing identity

or cost-preserving steps. If the rank lies in the nonnegative real numbers instead, strictly decreasing sequences can be infinite, such as 2^{-k}. A statement about eventual stabilization should therefore specify the order in addition to the existence of an improvement metric.

C Compositional observation and a feedback obstruction

This appendix asks whether one can truncate each component of a composite operation before assembling the components. The answer depends on the limits used for assembly. It is positive for a product and for a pullback over an already truncated interface. A homotopically richer interface can carry hidden information into an observable degree.

C.1 Agreement for wiring along an observed interface

Let $A \to B \leftarrow C$ be a diagram in $\mathcal{X}$. One can form the composite $A \times_B C$ and then observe it at degree n, or apply $\tau_{\leq n}$ to the diagram and take the resulting pullback. There is a natural comparison map from the former to the latter. Indeed the target is n-truncated because the truncated subcategory is closed under limits, and there is a map from the untruncated pullback to it. The adjunction gives the comparison. It need not be invertible.

Proposition C.1 (Truncation across an already observed interface)**.** *If B is n-truncated, the comparison*

$$\tau_{\leq n}(A \times_B C) \longrightarrow \tau_{\leq n}A \times_B \tau_{\leq n}C$$

is an equivalence. Here the structure maps on the right are induced by the adjunction.

Proof. For any $X \to B$, its relative n-truncation has an n-truncated structure map and thus has n-truncated total object because $B \to 1$ is n-truncated. The unit from X to that total object is n-connected, as in the relative factorization of Section 3. Therefore the relative truncation has underlying object $\tau_{\leq n}X$. Relative truncation

preserves products in the slice by the product argument of Section 3. Apply it to the product of $A \to B$ and $C \to B$. □

It follows by induction that an iterated matching of spans along n-truncated interfaces can be performed using the truncated apices at each step. In each step, take the pullback over the designated common endpoint and retain the remaining structural maps. If the remaining endpoints are n-truncated too, the resulting maps can be obtained by the truncation adjunction. Associativity of pullbacks makes the comparison independent of the chosen bracketing up to its canonical coherence. This is a restricted compositionality statement, suited to a model in which interfaces and span apices have been defined explicitly. It supplies the comparisons for those compositions.

In the case of products alone no interface data beyond a terminal object is needed. For example, in a model where independent parallel composition has operational object $\mathrm{Op}(M) \times \mathrm{Op}(N)$, the equality $\tau_{\leq n}(\mathrm{Op}(M) \times \mathrm{Op}(N)) \simeq Q_n(M) \times Q_n(N)$ allows equivalences of component observations to be composed. This helps distinguish an observation equivalence compatible with an explicitly modeled constructor from a contextual equivalence for an entire machine language. If the language includes other constructors, their effects must be considered as well.

C.2 A hidden interface becomes visible after matching

Take $\mathcal{X} = \mathcal{S}$, choose $n \geq 0$, and let $B = K(A, n+1)$ for a nonzero abelian group A. Take both legs of the diagram $* \to B \leftarrow *$ to be the base point. Their homotopy pullback is $\Omega B \simeq K(A, n)$. In particular the degree-n observation of the composite is $K(A, n)$, where for $n = 0$ this is the discrete set A. Truncating the three input objects first yields the diagram $* \to * \leftarrow *$, whose pullback is a point. This displays a failure at precisely the degree of observation. The mechanism is that matching two marked points in an interface asks for paths between them. A path space can lower the degree of a hidden homotopy group into the range of the observation.

The same effect occurs in a pullback interpreted as a constraint that two interface values agree up to a chosen witness. Keeping that witness retains a path space. If

the composition instead takes a mere existence condition, it can discard the number or higher homotopy of such witnesses. The relevant mathematical question is then whether the model uses a homotopy pullback, a truncated pullback, or an image of a projection. These operations can have different observable consequences even when all are interpreted informally as matching an interface.

One can characterize a degree loss under repeated loop formation. For a pointed space B and $j \geq 1$, the iterated loop space $\Omega^j B$ is computed by repeated pointed pullback. With $B = K(A, n+j)$, the degree-n observation of $\Omega^j B$ is $K(A, n)$, while the degree-n truncation of B is contractible. Thus for a composite that takes j successive loop spaces, data arbitrarily far above the nominal observation degree can enter the answer as j varies. A bound for a class of circuits or feedback systems would depend on how their composition uses such operations.

C.3 Interface contracts and categories of tests

The preceding example suggests a formal criterion for when an observational quotient is substitutive under a constructor. Let $C : \mathcal{C}^r \to \mathcal{C}$ be a functor, and let $Q : \mathcal{C} \to \mathcal{D}$ be an observation. If there is a functor $\overline{C} : \mathcal{D}^r \to \mathcal{D}$ and an equivalence $QC \simeq \overline{C}Q^r$, replacing an input by another with equivalent observation preserves the composite observation. The proof applies $\overline{C}$ to the component equivalences. In product models one can exhibit $\overline{C}$ by the product calculation. In the span example an induced pullback constructor cannot agree on all diagrams if the only retained data are their degree-n truncations. The two diagrams $* \to K(A, n+1) \leftarrow *$ and $* \to * \leftarrow *$ have the same truncated diagrams and observably different composites.

If a language supplies several constructors, the testing method of Section 5 can form an observational relation closed under those constructors by permitting tests after composition. This is a way to refine the observation until it is sensitive to the contexts under consideration. It may require observing more than the truncation of each component considered in isolation. Conversely, if the intended contexts only wire through n-truncated interfaces and the observational semantics respects the span model, Proposition C.1 gives a reason why the chosen coarse observation is already adequate for that part of the language.

D Source correspondence and notes on prior work

D.1 Correspondence with the source note

The following map records how the numbered parts of the source [1] enter the manuscript. The earlier text is a web note. The statements and proofs here are presented independently, with changes to the hypotheses made explicit.

Source sections 1–2. The observed n-type and equality of semantic classes are retained. Section 2 makes the core explicit and adds markings and transformation arrows. The source's implication from nonconservativity to non-uniqueness of objects requires inequivalent endpoints or an independent proof of object collapse. The one-object monoid and discrete two-object category give counterexamples to opposite implications. Full faithfulness on cores is the appropriate criterion for contractible marked fibers, and lifting of target equivalences by source arrows makes conservativity useful for the object question.

Source section 3. The sphere example proves failure of reflection in spaces for $n \geq 0$. The negative degrees require separate treatment. Section 3 extends the example to operational models with hidden products, distinguishes relative and absolute truncation, and gives a cohomological example that cannot be detected by listing homotopy groups. The compatibility needed to reconstruct observations across the tower is used again in Theorem 11.1.

Source section 4. A finite specification can have a unique realized semantic class if it is sound, has a realizer, and all realizers have the same observed type. Section 4 separates the assumed observed type from the constraints that could actually determine it. Finite tests and finite rules can have different force. Section 6 proves a uniqueness claim for a concrete universal representative under an explicit homomorphism policy.

Source section 5. The factorization through traces is retained as a hypothesis and proved from a universal property when the trace category is a localization. Section 5 distinguishes a map on objects from coherent factorization of functors

and contrasts linear tests with branching observations. Appendix C shows why truncating a component and then placing it in a context need not preserve the behavior of the composite.

Source section 6. The source's definition of n-cofinality is coverage of equivalence classes. Its theorem gives a stage at which the type appears. Section 8 states the additional hypothesis for staying in that class and identifies actual finality through comma categories. Theorem 8.1 gives a formula for the comma-space type when the target is groupoidal.

Source section 7. A semantic update is typed using a catalog and a classifier. Section 9 proves a positive criterion for lifting a semantic orbit with arrows and gives explicit failures of stagewise realization to lift. Section 10 adds a cost model and proves viability and budget statements with their finiteness and choice assumptions.

Source section 8. The conclusion that semantic equivalence can coexist with many inequivalent realizations is made constructive in several models. The composite of an operational semantics with truncation is called a localization of the machine category in examples where a universal property establishes that claim. This separates the always available operation of semantic truncation from a property of a particular machine model.

D.2 Higher categories, topoi, and homotopy theory

Lurie's *Higher Topos Theory* supplies the truncation adjunction, the connected–truncated factorization, coCartesian fibrations, and the standard finality criterion [17]. Kerodon and Riehl–Verity give alternative treatments of many of the ∞-categorical constructions [18, 26]. The use of a homotopy fiber over p is standard for objects equipped with an identification with p. The contribution here is its use in distinguishing machine uniqueness from interface markings, including the successive-fiber comparison and the connection to timelines and recursive redesign.

The homotopy type theory book gives a type-theoretic account of truncations,

equivalences, and contractibility [13, Chs. 4, 7]. Under that interpretation a marked realizer is a dependent pair whose second coordinate is an identification of its observed type with the given one. The distinction between unique inhabitant up to path and a contractible type includes the role of higher paths in that contractibility test. The present paper works externally in ∞-categories and uses internal objects for the constructions of Sections 7 and 9.

Hatcher's chapters on cohomology and homotopy theory supply the ingredients for the Postnikov example [11], and Brown treats group extensions through cohomology [4]. Those tools explain why two operational objects can share homotopy groups and have different section spaces. Our examples instantiate two degrees of such an obstruction with explicit groups and a specified base. Clausen and Mathew's discussion of hyperdescent provides context for the difference between detecting maps by homotopy sheaves and reconstructing objects from truncations [6]. We impose tower convergence where reconstruction requires it.

D.3 Localization, descent, and finality

Dwyer–Kan localization retains information in mapping spaces when arrows are inverted [8, 9]. Rezk's complete Segal spaces and Barwick–Kan's relative categories provide models for the relevant homotopy theories [24, 3]. Those works motivate taking the universal property as the meaning of localization. Here the finite reachable Mealy category provides a positive example after inversion of all strict comparison arrows. The monoid and disconnected examples isolate ways that a general observation can fail to have the same property.

For descent, Mac Lane and Moerdijk supply the ordinary sheaf and topos background [19], and the higher variants used here are related to the work of Hirschowitz–Simpson and Dugger–Hollander–Isaksen [12, 7]. We use the Čech nerve of an effective epimorphism in an ∞-topos and the universal property of its colimit. Higher descent is needed when the ambient comparison is by hypercovers. The C_2-torsor example already works in an ordinary sheaf topos on a classifying site, so the absence of global choices does not require higher homotopy.

Quillen's Theorem A is the classical ancestor of the comma-category test for finality

[23, 17, 18]. In the groupoidal target case Theorem 8.1 reduces that test to fibers in spaces using the groupoid completion of the indexing category. This specialization is what makes it directly comparable with the source's claim about semantic coverage. For semantic arrows that change the observed object through irreversible transitions, the general comma test applies.

D.4 Operational semantics and certified updates

Moore studies what experiments can infer about sequential machines, and Mealy describes the synchronous transition-output model [22, 20]. Rutten and Jacobs put behavioral maps and state minimization in a coalgebraic setting [27, 14]. Our deterministic quotient is an instance of those methods. The right-congruence description emphasizes the full strict transformation category in a fixed fiber and explains why its groupoid completion and its core have different classifications. The unary cycles give an explicit infinite family of inequivalent finite machines of one trace behavior. The proof relies on deterministic transitions and initialization.

Milner's full abstraction and the game semantics of PCF distinguish semantic equality and contextual observations through explicit adequacy and definability arguments [21, 2]. Joyal, Nielsen, and Winskel relate bisimulation to lifting along paths, while van Glabbeek maps the range of linear and branching equivalences [15, 10]. These are prior comparisons for our trace section. They also explain why an unqualified assertion that a trace is complete has no fixed meaning across models. The certified arrow layer is distinct from this observational adequacy problem. As an example from compilation, Leroy's compiler correctness theorem proves a particular semantic preservation statement with an explicit source language, target, and observable behavior [16]. Our lift theorems are abstract conditions on an already supplied category of certified updates. A concrete compiler requires a separate proof of preservation for its operational model.

The fixed-point argument in Theorem 10.1 uses the lattice-theoretic method of Tarski [28]. The finite-branching arguments are compactness arguments on trees, and the computable-tree example belongs to the classical distinction between a path and a computable path [5]. The role of these results here is to state exactly what follows from

finite-horizon evidence for a fixed semantic orbit. In particular, the uniform-budget formula and joint-constraint theorem are applications of compactness with the full prefix cost and observation predicates specified. The reductions in Section 10 concern explicit computational interfaces. The broader link to undecidability of extensional properties is provided by Rice [25].

References

[1] *Semantic Uniqueness and Non-Uniqueness of Machine Realizations*. Web note, sections 1–8, accessed November 4, 2026. Source note.

[2] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Information and Computation* 163(2), 409–470, 2000. Oxford Research Archive.

[3] Clark Barwick and Daniel M. Kan. Relative categories. Another model for the homotopy theory of homotopy theories. *Indagationes Mathematicae* 23(1–2), 42–68, 2012. arXiv:1011.1691.

[4] Kenneth S. Brown. *Cohomology of Groups*. Graduate Texts in Mathematics 87, Springer, 1982. Springer.

[5] Douglas Cenzer and Jeffrey B. Remmel. *Effectively Closed Sets*. Author-hosted manuscript, 2017 version. See Sec. 6.1 on weak Kőnig's lemma. Author's copy.

[6] Dustin Clausen and Akhil Mathew. Hyperdescent and étale K-theory. *Inventiones Mathematicae* 225, 981–1076, 2021. arXiv:1905.06611.

[7] Daniel Dugger, Sharon Hollander, and Daniel C. Isaksen. Hypercovers and simplicial presheaves. *Mathematical Proceedings of the Cambridge Philosophical Society* 136(1), 9–51, 2004. arXiv:math/0205027.

[8] William G. Dwyer and Daniel M. Kan. Simplicial localizations of categories. *Journal of Pure and Applied Algebra* 17(3), 267–284, 1980. DOI record.

[9] William G. Dwyer and Daniel M. Kan. Calculating simplicial localizations. *Journal of Pure and Applied Algebra* 18(1), 17–35, 1980. University of Rochester copy.

[10] Rob J. van Glabbeek. The linear time – branching time spectrum I. The semantics of concrete, sequential processes. In J. A. Bergstra, A. Ponse, and S. A. Smolka, eds., *Handbook of Process Algebra*, 3–99, Elsevier, 2001. Institutional record.

[11] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002. Author's chapter downloads.

[12] André Hirschowitz and Carlos Simpson. Descente pour les n-champs (Descent for n-stacks). Preprint, 1998, revised 2001. arXiv:math/9807049.

[13] The Univalent Foundations Program. *Homotopy Type Theory. Univalent Foundations of Mathematics*. Institute for Advanced Study, 2013. Book site.

[14] Bart Jacobs. *Introduction to Coalgebra. Towards Mathematics of States and Observation*. Cambridge University Press, 2016. Publisher record.

[15] André Joyal, Mogens Nielsen, and Glynn Winskel. Bisimulation from open maps. *Information and Computation* 127(2), 164–185, 1996. BRICS precursor. Journal DOI.

[16] Xavier Leroy. Formal verification of a realistic compiler. *Communications of the ACM* 52(7), 107–115, 2009. Article copy.

[17] Jacob Lurie. *Higher Topos Theory*. Annals of Mathematics Studies 170, Princeton University Press, 2009. Author's electronic edition. arXiv record.

[18] Jacob Lurie. *Kerodon*. Online text, accessed November 4, 2026. Text. Section 7.2.3.

[19] Saunders Mac Lane and Ieke Moerdijk. *Sheaves in Geometry and Logic. A First Introduction to Topos Theory*. Universitext, Springer, 1992. Springer.

[20] George H. Mealy. A method for synthesizing sequential circuits. *Bell System Technical Journal* 34(5), 1045–1079, 1955. Publisher record.

[21] Robin Milner. Fully abstract models of typed λ-calculi. *Theoretical Computer Science* 4(1), 1–22, 1977. Publisher record.

[22] Edward F. Moore. Gedanken-experiments on sequential machines. In C. E. Shannon and J. McCarthy, eds., *Automata Studies*, 129–153. Princeton University Press, 1956. Carnegie Mellon copy.

[23] Daniel Quillen. Higher algebraic K-theory I. In H. Bass, ed., *Algebraic K-Theory I*, Lecture Notes in Mathematics 341, 85–147, Springer, 1973. Bibliographic record. Chapter DOI.

[24] Charles Rezk. A model for the homotopy theory of homotopy theory. *Transactions of the American Mathematical Society* 353(3), 973–1007, 2001. arXiv:math/9811037.

[25] Henry Gordon Rice. Classes of recursively enumerable sets and their decision problems. *Transactions of the American Mathematical Society* 74(2), 358–366, 1953. University of Beira Interior copy.

[26] Emily Riehl and Dominic Verity. *Elements of ∞-Category Theory*. Cambridge Studies in Advanced Mathematics 194, Cambridge University Press, 2022. Publisher record. arXiv:1808.09860.

[27] Jan J. M. M. Rutten. Universal coalgebra. A theory of systems. *Theoretical Computer Science* 249(1), 3–80, 2000. CWI repository.

[28] Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics* 5(2), 285–309, 1955. DOI record.