Quantification and Kan Extension in the Internal Logic of an ∞-Topos

Anonymous

11 November 2026

Abstract

For a map of contexts f in an ∞-topos, the adjoints to substitution are the dependent sum and dependent product on types. On propositions, the left adjoint is the image of the dependent sum, and the right adjoint is the dependent product. We prove that these constructions are the groupoidal case of internal left and right Kan extension. For arbitrary internal ∞-categories, the pointwise formulas use comma categories. We give equivalent descriptions by weighted coends and ends, and express their propositional reflections using existential and universal quantification over the corresponding arrow types. We characterize when the canonical base change maps are equivalences for all type-valued or all proposition-valued diagrams. The two criteria differ by pointwise propositional truncation, and finite directed examples separate them. We then specify an ontology learner for a swarm observing a finite signal type. Left extension pools keyed observations, right extension describes compatible marginal predictions, and a finite likelihood optimization selects global models. A three-agent example on a 64-bit signal type exhibits the value of paired observations.

Keywords. ∞-topos, internal logic, dependent quantification, Kan extension, Beck–Chevalley, ontology fusion.

1 Contexts, types, and truth values

Lawvere's account of quantification identifies the existential and universal quantifiers as the left and right adjoints to substitution in a predicate doctrine [1]. In an ∞-topos, substitution also acts on families whose fibers have higher homotopy. Its adjoints at this level encode witnesses and coherent sections. The internal theory of Kan extensions [4] puts these operations in a wider setting by allowing directed arrows in the indexing categories. The relation between the two settings is reflected in the following adjoint triples.

Coefficient and index	Left adjoint to restriction	Right adjoint to restriction
Types over a map f	Σ_f	Π_f
Propositions over f	$\exists_f = \tau_{-1}^B \Sigma_f$	$\forall_f = \Pi_f\vert_{\mathsf{Prop}}$
Types over an internal functor u	Lan_u by comma colimits	Ran_u by comma limits
Propositions over u	$\tau_{-1}\,\mathrm{Lan}_u$	$\mathrm{Ran}_u\,\vert_{\mathsf{Prop}}$

Let $\mathcal{E}$ be a Grothendieck ∞-topos. For a context $T \in \mathcal{E}$, write $\mathsf{Type}(T) = \mathcal{E}_{/T}$ and $\mathsf{Prop}(T) = \mathrm{Sub}_{\mathcal{E}}(T)$. A type over T is a map $X \to T$. A proposition over T is a (-1)-truncated map, or equivalently a subobject of T. For subobjects $P, Q \hookrightarrow T$, write $P \leq Q$ for entailment. The mapping space over T between P and Q is contractible when $P \leq Q$ and empty otherwise. The subobject classifier Ω represents the underlying truth values [2, §6.1.6]. The corresponding internal poset $\mathsf{Prop}_{\mathcal{E}}$ has implications as its arrows.

We use $\mathsf{Type}_{\mathcal{E}}$ for the large internal ∞-category of types and functions whose category of sections over T is $\mathcal{E}_{/T}$. Its indexed description follows from descent for slices [2, §6.1.3]. The universe for internal groupoids in [4, §2.10] has this role. We place the slices and functor categories in a larger universe and take indexing internal categories to be small. All limits, colimits, and total spaces of types below are formed in the resulting slices. Generalized objects in a context T may be restricted along any $T' \to T$. Formulas in internal type notation mean identities of types in every such context. The dependent type notation follows the usual interpretation of homotopy type theory in an ∞-topos [7, 6].

2 Quantification along a map of contexts

2.1 The adjoints on types and propositions

Let $f : A \to B$. Substitution is the pullback functor $f^* : \mathcal{E}_{/B} \to \mathcal{E}_{/A}$. Composition with f defines its left adjoint Σ_f. Pullback preserves colimits by universality of colimits in $\mathcal{E}$. The slices are presentable, so f^* also has a right adjoint Π_f [2, Remark 6.1.3.11]. Explicitly, for $X \to A$ and $Y \to B$ the defining equivalences are

$$\mathrm{Map}_{/B}(\Sigma_f X, Y) \simeq \mathrm{Map}_{/A}(X, f^*Y), \tag{1}$$

$$\mathrm{Map}_{/A}(f^*Y, X) \simeq \mathrm{Map}_{/B}(Y, \Pi_f X). \tag{2}$$

The first follows from the pullback universal property. The second identifies a map to $\Pi_f X$ over B with a section of X over $Y \times_B A$.

Write $\iota_T : \mathsf{Prop}(T) \hookrightarrow \mathsf{Type}(T)$ for the inclusion and τ_{-1}^T for its left adjoint. The latter takes the image of $X \to T$ in the effective epimorphism–monomorphism factorization. Both ι_T and τ_{-1}^T commute with restriction of context, by pullback stability of this factorization [2, §6.2.3].

Theorem 2.1 (Quantifier adjunctions)**.** *On propositions there is an adjoint triple*

$$\exists_f \dashv f^* \dashv \forall_f, \qquad \exists_f P = \tau_{-1}^B \Sigma_f \iota_A P, \qquad \iota_B \forall_f P \simeq \Pi_f \iota_A P.$$

If $P \hookrightarrow A$ is a predicate and b is a generalized element of B, set

$$A_b = \sum_{a:A} \mathrm{Id}_B(f(a), b).$$

In the context of b the formulas are

$$(\exists_f P)(b) \simeq \| \sum_{(a,p):A_b} P(a) \|_{-1}, \qquad (\forall_f P)(b) \simeq \prod_{(a,p):A_b} P(a). \tag{3}$$

The corresponding type-valued formulas are the untruncated sum and the same dependent product with any type family in place of P.

Proof. Pullback preserves monomorphisms. For every $Y \to B$, the right side of (2) is (-1)-truncated when X is (-1)-truncated, since the left side is then (-1)-truncated. This verifies that Π_f takes propositions to propositions. For $Q \hookrightarrow B$, the reflection adjunction and (1) give

$$\mathrm{Map}_{/B}(\iota_B \exists_f P, \iota_B Q) \simeq \mathrm{Map}_{/A}(\iota_A P, f^* \iota_B Q).$$

These mapping spaces encode exactly the corresponding entailments. For $\forall_f$, use (2) on the same full subcategories.

In any context $b : T \to B$, the pullback of $A \to B$ is the homotopy fiber A_b in $\mathcal{E}_{/T}$. The fiber of $\Sigma_f X$ is the total space of the pulled back family over A_b. For any $Z \to T$, maps $Z \to b^*\Pi_f X$ over T correspond by (2) to sections of that family over $Z \times_T A_b$. This is the dependent product universal property in $\mathcal{E}_{/T}$. Pullback stability of images gives the relative truncation in (3). □

For a projection $A = B \times K \to B$, the formula becomes the usual $\exists_{k:K} P(b,k)$ and $\forall_{k:K} P(b,k)$. For the projection of a dependent type $A \to B$, the bound type is the family A_b. For a general map, the identity path in A_b records the constraint that the witness lies over b. Propositional truncation has the elimination rule into propositions [7, §3.7]. This explains the difference between constructing an existential proof and producing a witness of a specified type.

2.2 Substitution, conjunction, and implication

Proposition 2.2 (Beck–Chevalley and Frobenius)**.** *For a pullback of contexts*

$$\begin{array}{ccc} A' & \xrightarrow{h} & A \\ {\scriptstyle f'}\downarrow & \lrcorner & \downarrow{\scriptstyle f} \\ B' & \xrightarrow[g]{} & B, \end{array}$$

the canonical maps $\Sigma_{f'} h^* \to g^* \Sigma_f$ and $g^* \Pi_f \to \Pi_{f'} h^*$ are equivalences. The same holds for $\exists$ and $\forall$. There is also a natural equivalence

$$\Sigma_f (X \times_A f^* Y) \simeq \Sigma_f X \times_B Y. \tag{4}$$

On propositions it gives $\exists_f (P \wedge f^* Q) \simeq (\exists_f P) \wedge Q$.

Proof. The assertion for Σ and (4) are both the associativity equivalence for pullback. Applying the same assertion to the other pair of sides gives $f^* \Sigma_g \simeq \Sigma_h f'^*$. For $Z \to B'$ and $X \to A$ we obtain

$$\begin{aligned} \mathrm{Map}_{/B'}(Z, g^* \Pi_f X) &\simeq \mathrm{Map}_{/B}(\Sigma_g Z, \Pi_f X) \\ &\simeq \mathrm{Map}_{/A}(f^* \Sigma_g Z, X) \\ &\simeq \mathrm{Map}_{/A}(\Sigma_h f'^* Z, X) \\ &\simeq \mathrm{Map}_{/A'}(f'^* Z, h^* X) \simeq \mathrm{Map}_{/B'}(Z, \Pi_{f'} h^* X). \end{aligned}$$

By construction this identifies the right mate. Relative images commute with base change. Taking the image of the left side of (4) for a subobject $Q \hookrightarrow B$ gives the intersection of the image of $\Sigma_f P$ with Q. Indeed the effective epimorphism onto the image pulls back along Q and the resulting subobject of Q is also a subobject of B. □

The unit $P \le f^* \exists_f P$ introduces an existential proof. The counit $f^* \forall_f P \le P$ specializes a universal proof. The other counit and unit encode the corresponding elimination and introduction rules by transposition. Composition of maps composes the quantifiers in the order forced by the adjunctions. For example, if $g : B \to C$, then $\exists_{gf} \simeq \exists_g \exists_f$ and $\forall_{gf} \simeq \forall_g \forall_f$.

The Heyting implication of subobjects $P, Q \hookrightarrow B$ is $P \Rightarrow Q = \forall_m(m^*Q)$, where $m : P \hookrightarrow B$. Indeed $R \leq \forall_m m^*Q$ iff $R \wedge P \leq Q$. Frobenius and the adjunctions then give

$$\forall_f(f^*Q \Rightarrow P) \simeq (Q \Rightarrow \forall_f P).$$

For example, for each $R \hookrightarrow B$, entailment into the left side is equivalent to $f^*R \wedge f^*Q \leq P$, which is equivalent to $R \wedge Q \leq \forall_f P$.

2.3 Truth and local witnesses

Corollary 2.3 (Generalized truth conditions)**.** *Let $b : T \to B$. The proposition $(\exists_f P)(b)$ holds over T if and only if $T \times_B P \to T$ is an effective epimorphism. Equivalently, there is an effective epimorphism $T' \to T$ over which there is a witness. The proposition $(\forall_f P)(b)$ holds over T if and only if $T \times_B P \to T \times_B A$ has a section over $T \times_B A$.*

Proof. Existential truth means that the image subobject is all of T. When $T \times_B P \to T$ is effective epic, take $T' = T \times_B P$ with its tautological witness. Conversely, any cover admitting a witness factors through the image. An effective epimorphism that factors through a subobject forces that subobject to be the entire context. Universal truth follows from (2) applied to the terminal object of $\mathcal{E}_{/T}$. □

Example 2.4 (Sheaves and local existence)**.** Let $\mathcal{E} = \mathrm{Shv}(\mathbb{R}; \mathcal{S})$ and let U, V be the subterminal sheaves for $(-\infty, 1)$ and $(0, \infty)$. The projection $f : U \amalg V \to \mathbf{1}$ is effective epic, so $\exists_f \top = \top$. A global element of $\Sigma_f \top$ would yield a partition of $\mathbb{R}$ into disjoint open subsets subordinate to U, V, by disjointness and stability of coproducts. Connectedness forces one part to be all of $\mathbb{R}$, contrary to the chosen opens. The existential holds by descent while the space of global witnesses is empty.

3 Internal Kan extensions of type-valued diagrams

3.1 Internal diagrams and their (co)limits

Use complete Segal objects for internal ∞-categories. If $I \in \mathrm{Cat}(\mathcal{E})$ and T is a context, I_T denotes the pullback of I to $\mathcal{E}_{/T}$. For i, j in context T, the mapping object $\underline{\mathrm{Map}}_I(i, j)$ lies in $\mathcal{E}_{/T}$. The Segal structure supplies composition and its higher coherences. An internal functor $H : I_T \to \mathsf{Type}_{\mathcal{E}_{/T}}$ is equivalently a left fibration $p : L \to I_T$ [3, Theorem 4.5.1]. Its fiber at i is $H(i)$.

Proposition 3.1 (Realization and sections)**.** *The colimit of H in $\mathcal{E}_{/T}$ is the geometric realization $|L|$ of the simplicial object L. Its limit is the object of sections of p, characterized, for each $Z \to T$, by*

$$\mathrm{Map}_{/T}(Z, \lim_{I_T} H) \simeq \mathrm{Map}_{\mathrm{Cat}(\mathcal{E}_{/T})_{/I_T}}(I_T \times Z, L), \tag{5}$$

where Z is regarded as a constant simplicial object. Both constructions commute with restriction of ambient context.

Proof. An internal functor $L \to Z$ is equivalent to a map $|L| \to Z$ by the realization adjunction. It is also equivalent to a map of left fibrations $L \to I_T \times Z$ over I_T, or a natural transformation $H \to \Delta Z$ by internal straightening. This proves the colimit formula. Realization commutes with context pullback because colimits are universal.

For the limit, the functor $Z \mapsto I_T \times Z$ into the slice of simplicial objects over I_T preserves colimits, degree by degree. It therefore has a right adjoint by presentability. Its value at L represents (5), since the full inclusion of complete Segal objects preserves mapping spaces. Straightening identifies the right side with the space of cones $\Delta Z \to H$. For $v : T' \to T$, postcomposition and pullback are adjoint on objects and degreewise on simplicial objects. The canonical equivalence $\Sigma_v(v^* I_T \times Z) \simeq I_T \times \Sigma_v Z$ gives, by testing (5) against $\Sigma_v Z$, the same representation after restriction to T'. This verifies the internal universal properties in every context. □

3.2 Comma objects, coherent witnesses, and weights

Let $u : C \to D$ be an internal functor and $F : C \to \mathsf{Type}_{\mathcal{E}}$ a covariant diagram. For $d : \mathbf{1}_T \to D_T$, form the internal comma categories

$$(u \downarrow d) = C_T \times_{D_T} (D_T)_{/d}, \qquad (d \downarrow u) = C_T \times_{D_T} (D_T)_{d/}, \tag{6}$$

using source for the first pullback and target for the second. Their projections to C_T are denoted by π_- and π_+, respectively. Their objects are pairs $(c, \alpha : u(c) \to d)$ and $(c, \beta : d \to u(c))$. A morphism $m : c \to c'$ in the first comma category includes a coherent triangle $\alpha' \circ u(m) \simeq \alpha$. In the second it includes $u(m) \circ \beta \simeq \beta'$.

For an internal I, $\underline{\mathrm{Nat}}_I(H, K)$ denotes the internal mapping object of natural transformations, including all higher coherences. An end below is the internal end, and a coend is the internal coend. In particular $\underline{\mathrm{Nat}}_I(H, K) = \int_{i \in I} K(i)^{H(i)}$ for type-valued diagrams. For a presheaf $W : I^{\mathrm{op}} \to \mathsf{Type}_{\mathcal{E}}$, the weighted colimit of $H : I \to \mathsf{Type}_{\mathcal{E}}$ is $\int^{i \in I} W(i) \times H(i)$. The exponential and product are those in the ambient slice. The end and coend use both variables of the corresponding bifunctor, so they impose the relations and coherences of all simplices.

Theorem 3.2 (Internal pointwise extension)**.** *The restriction u^* between internal functor categories into $\mathsf{Type}_{\mathcal{E}}$ has left and right adjoints. In each context T,*

$$(\mathrm{Lan}_u F)(d) \simeq \operatorname*{colim}_{(u \downarrow d)} F_T \pi_- \simeq \int^{c \in C_T} \underline{\mathrm{Map}}_{D_T}(u(c), d) \times F_T(c), \tag{7}$$

$$(\mathrm{Ran}_u F)(d) \simeq \lim_{(d \downarrow u)} F_T \pi_+ \simeq \int_{c \in C_T} F_T(c)^{\underline{\mathrm{Map}}_{D_T}(d, u(c))}. \tag{8}$$

In particular, $\underline{\mathrm{Nat}}_D(\mathrm{Lan}_u F, G) \simeq \underline{\mathrm{Nat}}_C(F, u^* G)$ *and* $\underline{\mathrm{Nat}}_C(u^* G, F) \simeq \underline{\mathrm{Nat}}_D(G, \mathrm{Ran}_u F)$.

Proof. Work in an arbitrary context and suppress the subscript T. By internal straightening, the comma projection π_- is the right fibration classified by $W_d = \underline{\mathrm{Map}}_D(u(-), d)$, and π_+ is the left fibration classified by $V_d = \underline{\mathrm{Map}}_D(d, u(-))$. For a type Z, cocones on $F\pi_-$ with value Z correspond to maps of presheaves $W_d \to Z^{F(-)}$. This is straightening over C^{op}: the map on the fiber of W_d at c assigns to each $\alpha : u(c) \to d$ a map $F(c) \to Z$, and its naturality is precisely cocone compatibility along the comma triangles. The exponential adjunction gives the same mapping object out of the coend in (7). Dually, cones from Z to $F\pi_+$ correspond to maps of covariant diagrams $V_d \to F(-)^Z$. Their mapping object is represented by the end in (8). These identifications hold in all contexts. Existence also follows from Proposition 3.1. Functoriality in d is induced by postcomposition on W_d and precomposition on V_d.

For the global adjunction, ends compute natural transformations. Using the cartesian tensor–hom adjunction and the Yoneda lemma internally gives

$$\begin{aligned}\underline{\mathrm{Nat}}_D(\mathrm{Lan}_u F, G) &\simeq \int_{c\in C}\left(\int_{d\in D} G(d)^{\underline{\mathrm{Map}}_D(u(c),d)}\right)^{F(c)}\\ &\simeq \int_{c\in C} G(u(c))^{F(c)} \simeq \underline{\mathrm{Nat}}_C(F, u^*G).\end{aligned}$$

The first equivalence interchanges the end over D with the end over C obtained by mapping out of the coend. For the right adjunction one has

$$\begin{aligned}\underline{\mathrm{Nat}}_D(G, \mathrm{Ran}_u F) &\simeq \int_{c\in C} F(c)^{\int^{d\in D}\underline{\mathrm{Map}}_D(d,u(c))\times G(d)}\\ &\simeq \int_{c\in C} F(c)^{G(u(c))} \simeq \underline{\mathrm{Nat}}_C(u^*G, F)\end{aligned}$$

by co-Yoneda. The internal Yoneda and co-Yoneda lemmas apply to these small diagrams [3, 4]. This proves the adjunctions and identifies the comma formulas with the internal pointwise Kan extensions of [4, §6.3]. □

Remark 3.3 (Other coefficient categories). For an internal target category $\mathcal{K}$, the comma formulas still apply when the indicated colimits and limits exist internally. The internal pointwise theorem constructs the extensions from these colimits and limits in all contexts [4, Theorem 6.3.5 and its dual]. In a tensored and cotensored target, the weighted coend and end descriptions use the target's tensors and cotensors by types in place of products and exponentials.

Terms and computation. The left formula has a canonical map from $F(c)$ for each arrow $\alpha : u(c) \to d$. Its eliminator requires such maps into the target together with their coherent compatibility along the comma category. The right formula consists of coherent assignments $x_{c,\beta} : F(c)$ for $\beta : d \to u(c)$. A comma morphism as in (6) imposes a path from $F(m)(x_{c,\beta})$ to $x_{c',\beta'}$, with coherences on higher simplices. Thus realization and sections give explicit constructions for the two internal extension operators.

Example 3.4 (The walking arrow). Let $u : \mathbf{1} \to [1]$ select 0 in the category $0 \to 1$. For a type X at 0, $\mathrm{Lan}_u X$ is the identity diagram $X \to X$ and $\mathrm{Ran}_u X$ is the terminal map $X \to \mathbf{1}$. At 1 the left comma category has the arrow $0 \to 1$ as its sole object, while the right comma category is empty. The directed arrow therefore contributes to the extension at 1 even though the homotopy fiber of u there is empty.

4 Propositional Kan extension and its logical form

Theorem 4.1 (Reflection to truth values). *For a predicate diagram $P : C \to \mathsf{Prop}_{\mathcal{E}}$, let $\iota P : C \to \mathsf{Type}_{\mathcal{E}}$ be its inclusion. The adjoints to restriction with proposition coefficients are*

$$\mathrm{Lan}_u^{\mathsf{Prop}} P \simeq \tau_{-1}\,\mathrm{Lan}_u(\iota P), \qquad \iota\,\mathrm{Ran}_u^{\mathsf{Prop}} P \simeq \mathrm{Ran}_u(\iota P),$$

where the reflection acts pointwise in every context. With $R_u(c,d) = \|\underline{\mathrm{Map}}_D(u(c),d)\|_{-1}$ and $S_u(d,c) = \|\underline{\mathrm{Map}}_D(d,u(c))\|_{-1}$ their internal formulas are

$$(\mathrm{Lan}_u^{\mathsf{Prop}} P)(d) \simeq \exists_{c:C_0}\big(R_u(c,d) \wedge P(c)\big), \tag{9}$$

$$(\mathrm{Ran}_u^{\mathsf{Prop}} P)(d) \simeq \forall_{c:C_0}\big(S_u(d,c) \Rightarrow P(c)\big). \tag{10}$$

Here C_0 is the type of objects, and $P(c)$ denotes the underlying predicate over it.

Proof. The reflection $\tau_{-1} \dashv \iota$ in the indexed coefficient categories is internal, because it is natural under context pullback. It induces the pointwise adjunction on diagram categories. The left statement follows by composing adjunctions. The right statement follows because the limit of a proposition-valued diagram is a proposition and the inclusion preserves this limit.

Here are the corresponding pointwise calculations. In a fixed context, let $H : I \to \mathsf{Prop}_{\mathcal{E}}$ and let $L \to I$ be its left fibration as a type-valued diagram. By Proposition 3.1, the left value in propositions is $\|\|L\|\|_{-1}$. The map from L_0 induces an equivalence

$$\|\|L\|\|_{-1} \simeq \|L_0\|_{-1} \simeq \|\sum_{i:I_0} H(i)\|_{-1}.$$

Indeed the reflection takes realization to a join of the supports of the L_n. This join equals the support of L_0 because every L_n maps to L_0 by taking a vertex, and the term $n = 0$ occurs.

For the right value, the section object is equivalent to $\prod_{i:I_0} H(i)$. To see this, test in each extended context. A section in degree zero of the monomorphism $L_0 \to I_0$ makes it an equivalence. The defining left fibration identities $L_n \simeq I_n \times_{I_0} L_0$, using the initial vertex, then make $L \to I$ an equivalence and supply its unique coherent section. Restriction to degree zero gives the converse. Now apply these two calculations to the comma categories. Their object types are $\sum_{c:C_0} \underline{\mathrm{Map}}_D(u(c), d)$ and $\sum_{c:C_0} \underline{\mathrm{Map}}_D(d, u(c))$, respectively. The resulting formulas are

$$\|\sum_{c:C_0} \sum_{\alpha:u(c)\to d} P(c)\|_{-1} \quad \text{and} \quad \prod_{c:C_0} \prod_{\beta:d\to u(c)} P(c).$$

Truncation eliminates into propositions. In the left expression it replaces the arrow witness by its support inside the existential. In the right expression it replaces an arrow-indexed family of proofs of $P(c)$ by an implication from the support of the arrow type. This gives (9) and (10). □

For any type-valued $F : C \to \mathsf{Type}_{\mathcal{E}}$, there is also a canonical equivalence

$$\tau_{-1} \operatorname{Lan}_u F \simeq \operatorname{Lan}_u^{\mathsf{Prop}}(\tau_{-1} F).$$

Both sides are left adjoint to the same composite $u^* \iota = \iota u^*$ on proposition-valued diagrams over D. For right extension there is a comparison $\tau_{-1} \operatorname{Ran}_u F \to \operatorname{Ran}_u^{\mathsf{Prop}}(\tau_{-1} F)$, induced by the unit $F \to \iota \tau_{-1} F$. Example 5.2 gives a groupoidal instance where this comparison fails.

The logical rules can also be read directly. If $P(c)$ entails $Q(u(c))$, then a witness $\alpha : u(c) \to d$ transports this proof to $Q(d)$. Elimination from truncation gives $\operatorname{Lan}_u^{\mathsf{Prop}} P \leq Q$. Conversely, specialize at $d = u(c)$ and use the identity arrow. For the right rule, from $Q(u(c)) \leq P(c)$ a proof of $Q(d)$ can be transported along any $\beta : d \to u(c)$ to prove $P(c)$. This is exactly the product in (10). The converse again uses identity arrows. Covariance supplies the stated transports, and transitivity of the two arrow relations gives the transports of the extended predicates.

5 The groupoidal specialization

Theorem 5.1 (Kan extension over types)**.** *View $A, B \in \mathcal{E}$ as internal groupoids via constant complete Segal objects, and $f : A \to B$ as an internal functor. Under straightening,* $\mathrm{Fun}_{\mathcal{E}}(A, \mathsf{Type}_{\mathcal{E}})$ *is the indexed category of families over A. Under this identification there are canonical equivalences*

$$\operatorname{Lan}_f \simeq \Sigma_f, \qquad \operatorname{Ran}_f \simeq \Pi_f, \qquad \operatorname{Lan}_f^{\mathsf{Prop}} \simeq \exists_f, \qquad \operatorname{Ran}_f^{\mathsf{Prop}} \simeq \forall_f.$$

The equivalences respect the units, counits, and restriction of ambient context.

Proof. A left fibration over the constant simplicial object A has its n-simplices equivalent to $L_0 \times_A A \simeq L_0$ by the left fibration identities. The degeneracies give the equivalence with the constant family $L_0 \to A$. This holds over every context T with A replaced by $T \times A \to T$. Restriction along f is pullback of this family. Uniqueness of adjoints gives the asserted equivalences with their units and counits.

Pointwise, a slice arrow in the groupoid B is a path. The comma category $(f \downarrow b)$ is the homotopy fiber $\sum_{a:A} \mathrm{Id}_B(f(a), b)$, regarded as a groupoid. Path inversion identifies $(b \downarrow f)$ with the same fiber. A colimit of a type family indexed by a groupoid is its total space, by Proposition 3.1, and the limit is its type of dependent sections. Equations (7) and (8) therefore give the sum and product of Theorem 2.1. Applying Theorem 4.1 gives the proposition formulas. □

Example 5.2 (Group actions and uniform choices)**.** In $\mathcal{S}$, let G be a nontrivial discrete group and let $f : BG \to \mathbf{1}$. A type family over BG is a homotopy-coherent G-action on its fiber X. The dependent sum is the homotopy orbit space X_{hG} and the dependent product is the homotopy fixed point space X^{hG}. For the trivial action on a point, $\Sigma_f \top \simeq BG$ and $\exists_f \top = \top$.

For the family $E = (\mathbf{1} \to BG)$ classifying the regular G-action on the fiber G, the relative support $\tau_{-1}^{BG} E$ is terminal over BG. Yet E has no section, since a homotopy section of $\mathbf{1} \to BG$ would make the identity of BG factor up to homotopy through a point, contradicting its fundamental group. Consequently

$$\tau_{-1}^{\mathbf{1}} \Pi_f E \simeq \mathbf{0}, \qquad \Pi_f \tau_{-1}^{BG} E \simeq \mathbf{1}.$$

Truncating the fibers asserts local inhabitation. Taking a dependent product first demands a coherent choice across all paths in BG.

The path type itself also belongs to this distinction. In context $(a, a') : A \times A$, the diagonal $A \to A \times A$ has fiber $\mathrm{Id}_A(a, a')$, whereas the image of the diagonal represents $\|\mathrm{Id}_A(a, a')\|_{-1}$. At the basepoint of BG the former has loop space G and the latter is true. Equality used as a proposition is the support of the identity type.

6 Base change for directed indexing

Let the following be a square of internal categories, equipped with an equivalence $uh \simeq gu'$:

$$\begin{array}{ccc} C' & \xrightarrow{\ h\ } & C \\ {\scriptstyle u'}\downarrow & & \downarrow{\scriptstyle u} \\ D' & \xrightarrow[\ g\ ]{} & D. \end{array} \tag{11}$$

The left and right mates of the restriction equivalence are

$$\mathrm{Lan}_{u'} h^* \longrightarrow g^* \mathrm{Lan}_u, qquad g^* \mathrm{Ran}_u \longrightarrow \mathrm{Ran}_{u'} h^*. \tag{12}$$

They are induced on the comma formulas by the functors from the primed comma categories to the unprimed comma categories. Finality of the left comma functors and initiality of the right comma functors suffice for the respective comparisons, in each context. Exact squares and their relation to pointwise Kan extension are studied via modules by Riehl and Verity [5, §5].

Here is a coefficient-specific criterion for (12). In a context containing $d' \in D'$, abbreviate

$$W' = \underline{\mathrm{Map}}_{D'}(u'(-), d'), \quad W = \underline{\mathrm{Map}}_{D}(u(-), g(d')),$$
$$V' = \underline{\mathrm{Map}}_{D'}(d', u'(-)), \quad V = \underline{\mathrm{Map}}_{D}(g(d'), u(-)).$$

The maps on arrow types have adjoint transposes

$$\lambda : \mathrm{Lan}_{h^{\mathrm{op}}} W' \longrightarrow W, qquad\rho : \mathrm{Lan}_h V' \longrightarrow V. \tag{13}$$

For instance, $W'(c') \to W(h(c'))$ takes an arrow to its image under g composed with the structural equivalence of the square.

Theorem 6.1 (Tests on weights)**.** *Fix a generalized d' and allow all diagrams in its context and further restrictions. The left mate in (12) is an equivalence at d' for all type-valued diagrams precisely when λ is an equivalence. The right mate is an equivalence for all such diagrams precisely when ρ is an equivalence. For proposition-valued diagrams the corresponding conditions are equivalence of $\tau_{-1}\lambda$ and of $\tau_{-1}\rho$.*

Proof. For a presheaf K on C, write $K \otimes_C F = \int^{c \in C} K(c) \times F(c)$. Adjunction of extension and restriction on presheaves, tested against $Z^{F(-)}$, gives

$$W' \otimes_{C'} h^* F \simeq (\mathrm{Lan}_{h^{\mathrm{op}}} W') \otimes_C F.$$

Thus the left mate at d' is $\lambda \otimes_C F$. Equivalence of λ suffices. Conversely, take the covariant representables $F = \underline{\mathrm{Map}}_C(c, -)$ in every context. Co-Yoneda identifies the comparison with $\lambda(c)$.

The right mate at d' identifies with precomposition

$$\underline{\mathrm{Nat}}_C(V, F) \longrightarrow \underline{\mathrm{Nat}}_C(\mathrm{Lan}_h V', F).$$

This follows from (8) and the left Kan adjunction along h. It is an equivalence for all F exactly when ρ is an equivalence, by Yoneda in the diagram category.

For the left proposition test, (-1)-truncation preserves finite products in each slice. Indeed, the product of the effective epimorphisms to the two image objects is effective epic by base change and composition. By the colimit formula over the right fibration of K, and the realization argument in Theorem 4.1, the support of $K \otimes_C F$ is the support of the sum over $c : C_0$ of $K(c) \times F(c)$. Product compatibility rewrites this as $\exists_{c:C_0}(\tau_{-1}K(c) \wedge \tau_{-1}F(c))$. Thus the reflected weight comparison suffices for all predicates. For necessity, take $P = \tau_{-1}\underline{\mathrm{Map}}_C(c, -)$. Product compatibility and reflection show that this tests the truncation of the same representable calculation as before, giving $\tau_{-1}\lambda(c)$.

For the right proposition test, the reflection adjunction gives $\underline{\mathrm{Nat}}_C(V, \iota P) \simeq \underline{\mathrm{Nat}}_C(\tau_{-1}V, P)$, and the same identification for $\mathrm{Lan}_h V'$. It identifies precomposition with the map induced by $\tau_{-1}\rho$. Testing all P in the proposition-valued diagram category again applies Yoneda. □

Corollary 6.2 (Logical base change criteria)**.** *In the context of d', the left proposition mate is an equivalence for all predicates if and only if the following implication holds internally for every $c : C_0$:*

$$\|u(c) \to g(d')\|_{-1} \quad \Longrightarrow \quad \exists_{c':C'_0}(\|c \to h(c')\|_{-1} \wedge \|u'(c') \to d'\|_{-1}). \tag{14}$$

For the right proposition mate the criterion is

$$\|g(d') \to u(c)\|_{-1} \quad \Longrightarrow \quad \exists_{c':C'_0}(\|d' \to u'(c')\|_{-1} \wedge \|h(c') \to c\|_{-1}). \tag{15}$$

The arrow types refer to their respective internal categories.

Proof. For the left weight, apply the comma formula for $\mathrm{Lan}_{h^{\mathrm{op}}} W'$ at c and take its support. The object type of its indexing comma category consists of $(c', c \to h(c'))$. The support of a colimit of types over an internal category is the support of the total type over the object type, by the realization argument used above. This gives the right side of (14). The map $\tau_{-1}\lambda(c)$ already supplies the implication from right to left by composition. Its being an equivalence is equivalent to the other implication. The same calculation with $\mathrm{Lan}_h V'$ uses $(c', h(c') \to c)$ and gives (15). Apply Theorem 6.1. □

Example 6.3 (A pullback that fails base change)**.** In Cat_∞, take $u : \mathbf{1} \to [1]$ to select 0 and $g : \mathbf{1} \to [1]$ to select 1. Their pullback C' is empty, since the two objects are inequivalent. At the unique object of D' the left mate with $F = \mathbf{1}$ compares $\mathbf{0}$ with $\mathbf{1}$. Reversing the roles of 0 and 1 gives a pullback for which the right mate with $F = \mathbf{0}$ compares $\mathbf{0}$ with $\mathbf{1}$. These examples already fail with proposition coefficients. The groupoidal case recovers Proposition 2.2, because its comma categories are homotopy fibers.

Example 6.4 (Truth preserves base change in a larger class)**.** Let D be the walking parallel pair with two arrows $a, b : 0 \to 1$. Let $D' = [1]$ and let $g : D' \to D$ keep the two objects and select a. Let $C' = C = \mathbf{1}$, $h = \mathrm{id}$, and let u', u select 0. This is a pullback square. At 1 the left weight comparison is the inclusion of one point into a discrete pair. For a type X the left mate at 1 is $X \to X \amalg X$, into the a summand. For every proposition P its reflected comparison is an equivalence, since the two weights have the same support. At 0 the weights are already equivalent. If instead u', u select 1, the right mate at 0 is the a projection $X \times X \to X$. Again it is an equivalence on all propositions.

7 Ambient pullback and geometric inverse image

The preceding substitution laws take place within a fixed ∞-topos. There is also change of topos. For a geometric morphism $q : \mathcal{F} \to \mathcal{E}$, the inverse image q^* preserves finite limits and colimits, hence dependent sums and relative images. It preserves $\exists$ in the sense of the comparison $q^*(\exists_f P) \simeq \exists_{q^* f}(q^* P)$. A dependent product has a canonical comparison $q^*(\Pi_f X) \to \Pi_{q^* f}(q^* X)$ induced by evaluation. It can fail even on propositions.

Example 7.1 (A stalk and a countable universal quantifier)**.** Take the stalk at 0 on $\mathrm{Shv}(\mathbb{R}; \mathcal{S})$. Let $A = \coprod_{n \geq 1} \mathbf{1}$ and $f : A \to \mathbf{1}$. Define a predicate on the nth component by the open $(-1/n, 1/n)$. Its universal quantification is the meet of these opens in the lattice of subterminals, which is the interior of their intersection and hence empty. Every component predicate has terminal stalk at 0. After applying the stalk the universal quantification is therefore true. This gives an empty-to-terminal comparison for $\forall$.

8 A finite observation model for an agent swarm

The adjunctions above specify three operations in a distributed ontology learner. A left Kan extension collects observations according to reachability in the communication history. A right Kan extension assembles the matching families of marginal predictions. Quantification along a map of candidates to predictions expresses realizability and agreement among optimal explanations. The following construction gives these operations finite data representations and an objective function.

8.1 The environment and the observation interface

Fix a finite coordinate set J, a nonempty finite type X_j for each $j \in J$, and a finite ordered set of environmental events $[N] = \{1, \dots, N\}$. For a scope $V \subseteq J$, put

$$X_V = \prod_{j \in V} X_j, \qquad X = X_J, \qquad \pi_{V,W} : X_V \longrightarrow X_W \quad (W \subseteq V).$$

The environment produces a trace $(x_e)_{e \in [N]} \in X^N$ according to an unknown law $Q_\star$. The definition also accepts a fixed trace. The event identifier e is assigned at the observation interface and is distinct from the time a message is delivered. An agent i has a declared scope V_i and can emit a report (e, i, y) with $y = \pi_{J,V_i} x_e$. The same event identifier allows agents to combine views of the same signal. A record key (e, i) identifies a single observation, regardless of how many agents relay it. The finite type of possible records is

$$\mathsf{Rec} = \sum_{e:[N]} \sum_{i:I} X_{V_i},$$

where I is the finite set of agents. This specification uses exact projections first. A stochastic observation interface is given below.

We use ordinary finite sets for the worked calculation. With the same finite types and their structure maps interpreted in a slice of $\mathcal{E}$, the constructions by images, finite limits, and dependent types also have an internal interpretation. The statistical assertions below concern the external sampling law of the finite model.

8.2 Transport and deduplication by left extension

Let D be the finite directed category of information flow. Its vertices include each agent's successive states and one source vertex for each batch of newly observed records. Sources have no incoming edges. Other edges represent retention of a state or a message delivery from a sending state to a receiving state. A finite time-ordered execution gives an acyclic graph. We can use its path category or the preorder of reachability. The source vertices form a discrete subcategory S, with inclusion $u : S \to D$. Let $P(s, r)$ be the predicate that source s produced record r. Apply the left Kan extension to the S-diagram of subobjects of Rec, equivalently apply Theorem 4.1 in context $r : \mathsf{Rec}$.

Proposition 8.1 (Reachable evidence)**.** *At state $d \in D$, the available ledger is the subobject L_d of* Rec *given by*

$$r \in L_d \quad \Longleftrightarrow \quad (\mathrm{Lan}_u^{\mathsf{Prop}} P)(d, r) \quad \Longleftrightarrow \quad \exists_{s:S_0} \big(\|s \to d\|_{-1} \wedge P(s, r) \big). \tag{16}$$

For a finite execution this is computed by taking the set union of new observations and all ledgers delivered along incoming edges. In a static strongly connected communication graph with reliable synchronous rounds and no further observations, every agent has the global union after at most the directed diameter of the graph.

Proof. Equation (16) is the left comma formula on propositions applied recordwise. A path from a source to a state either enters directly from a source or has a last delivery or retention edge. Induction in a topological order therefore gives the stated union recurrence. A record can be tested by its full typed content, so relaying the same record along several paths leaves its membership

unchanged. The union recurrence is a grow-only set in the usual sense of replicated data types [8]. After k synchronous rounds, each agent has every initial record whose source agent is at directed distance at most k, by induction on path length. Strong connectivity gives the diameter bound. □

One may retain the source and path as provenance in a separate type-valued ledger. The score below uses the predicate $r \in L_d$. For a restriction of a communication history that keeps the same source vertices, equality of these ledgers for every source predicate is equivalent to preserving reachability from each source to the chosen state. Sufficiency follows from (16). For necessity, test a record produced at only the source in question. This is a recordwise instance of the proposition-valued base change test in Theorem 6.1.

8.3 Event reconstruction and overlap

For a ledger L, let U_e be the union of the scopes of its records with event identifier e. Define

$$\mathsf{Glue}_e(L) = \sum_{y : X_{U_e}} \prod_{(e,i,z) \in L} \mathrm{Id}_{X_{V_i}}(\pi_{U_e,V_i} y, z). \tag{17}$$

This type is contractible if the reports agree on all overlaps, and empty if two values disagree at a shared coordinate. Agreement lets each coordinate of U_e take the unique reported value there. This also proves uniqueness, since the report scopes cover U_e. With no reports at an event, $U_e = \varnothing$ and the glued value is the terminal tuple. The existential support of $\mathsf{Glue}_e(L)$ tests whether the reports admit a common signal on the observed coordinates. Two contradictory reports under the same key give an empty gluing type in this exact sensor model.

More generally, for any downward-closed family of scopes $\mathcal{A} \subseteq \mathcal{P}(J)$ with union V, the comparison $X_V \to \lim_{\mathcal{A}^{\mathrm{op}}} X_W$ is an equivalence. Its inverse is the coordinate assignment just constructed. The right Kan extension formula identifies the limit as a value of the extension of $W \mapsto X_W$ from observed subscopes to V. This is the gluing law for signal values. Marginal distributions have a different gluing law.

9 Ontology types and coherent predictions

9.1 A finite grammar for explanations

An ontology o in the candidate type $\mathcal{O}$ specifies a finite latent type Z_o, a finite vocabulary of predicates on Z_o, and logical axioms for the vocabulary. Axioms can impose inclusion, exclusion, or exhaustiveness of concepts. For example, an implication $p \Rightarrow q$ is enforced by $\prod_{z : Z_o}(p(z) \Rightarrow q(z))$. The ontology also has a rational initial law w_o on Z_o, a rational transition kernel $K_o(z' \mid z)$, and an emission law $E_o(x \mid z)$ on X. The emission law ties concepts to observable signals. One can impose restrictions on its support, such as requiring a concept to emit only vectors satisfying a specified relation. These data define the trace law

$$Q_o(x_{1:N}) = \sum_{z_{1:N} \in Z_o^N} w_o(z_1) \left(\prod_{e=1}^{N} E_o(x_e \mid z_e) \right) \left(\prod_{e=1}^{N-1} K_o(z_{e+1} \mid z_e) \right). \tag{18}$$

An independent model uses identical rows in K_o equal to w_o. Bounds on description length, latent state count, and permitted rational tables or circuits give a finite grammar. We use $\ell(o)$ for the

integer bit length of a prefix-free description. The common grammar fixes how the agents interpret each candidate.

A local vocabulary with a declared map $b_i : X_{V_i} \to B_i$ has predictive law $(b_i)_* q_o^{V_i}$ at that scope, where $q_o^{V_i}$ is the relevant marginal. Thus an agent's finite claim about the distribution of its concepts is a predicate of the global candidate. Claims admitted as hard constraints may be included by taking the dependent subtype of $\mathcal{O}$ satisfying their product. Local models may also propose candidate descriptions for the shared grammar. Each proposed candidate is then scored against the canonical event evidence in L.

9.2 The right extension and the realizability fiber

For any fixed event, choose M so that the model marginals of interest have denominator dividing M. Write $\mathcal{M}(V) = \Delta_M(X_V)$ for the finite type of distributions of the form $(k_x/M)_{x \in X_V}$, with nonnegative integers k_x summing to M. Marginalization makes $\mathcal{M}$ a presheaf on the scope poset $\mathcal{P}(J)$. One can also apply the same construction to the coordinates of a trace by replacing J with $[N] \times J$.

Let $a : \mathcal{A} \hookrightarrow \mathcal{P}(J)$ be the inclusion of the downward-closed family of sampled scopes. From the right comma formula in Theorem 3.2,

$$(\mathrm{Ran}_{a^{\mathrm{op}}} \mathcal{M}|_{\mathcal{A}^{\mathrm{op}}})(V) \simeq \lim_{\substack{W \in \mathcal{A}^{\mathrm{op}} \\ W \subseteq V}} \mathcal{M}(W). \tag{19}$$

An arrow from V to W in the opposite scope category is precisely an inclusion $W \subseteq V$. Thus a term of the limit is a family of distributions that agree under all the listed marginalization maps. A global distribution gives such a term by restriction. Let R be the right side of (19) at J, and let $\mathrm{res} : \mathcal{O} \to R$ send a candidate to its family of marginals. For $m : R$ the type of extensions in the model grammar is the fiber

$$\mathsf{Real}(m) = \sum_{o:\mathcal{O}} \mathrm{Id}_R(\mathrm{res}(o), m).$$

Its propositional support gives the existential quantifier along res applied to truth. Its elements retain the possible global explanations. Exact compatibility here is a condition on the model predictions. Empirical frequencies at different scopes can differ by sampling error, so the optimizer below compares their likelihoods using the actual keyed reports.

Example 9.1 (Compatible marginals with an empty fiber)**.** Take three bits and all their pair scopes. Put the uniform law on the two unequal pairs at each scope. All three induced one-bit laws are uniform, so these laws define a term of the right extension. Any joint law extending them would have probability one on each of $x_1 \neq x_2$, $x_2 \neq x_3$, and $x_1 \neq x_3$. The intersection of the three events is empty. Thus even an unrestricted joint probability law cannot realize this term. This is a finite instance of the marginal extension problem [9].

Proposition 9.2 (A sufficient condition for distribution gluing)**.** *Suppose the maximal scopes are vertices of a tree with the running intersection property: the scopes containing any fixed coordinate form a connected subtree. If rational marginal laws agree on the overlaps of neighboring scopes, they extend to a rational law on the union of the scopes.*

Proof. Root the tree and attach scopes in parent-first order. Let V be the next scope, U the union already attached, and S its intersection with the parent. The running intersection property gives $U \cap V = S$, since the path to any earlier scope containing a shared coordinate passes through

the parent. Let q be the already constructed law on U, m_V the new law, and m_S their common marginal on S. Define

$$q'(x_{U\cup V}) = \begin{cases} q(x_U)m_V(x_V)/m_S(x_S), & m_S(x_S) > 0, \\ 0, & m_S(x_S) = 0. \end{cases}$$

Summing over coordinates of $V \setminus S$ gives q, since the conditional factor sums to one when $m_S(x_S) > 0$, and $q(x_U) = 0$ otherwise. Summing over $U \setminus S$ gives m_V by the same calculation. This proves normalization and the induction step. The resulting rational law can require a finer denominator grid or a larger grammar than the one used for the individual marginals. □

10 The optimization problem and its evaluation

10.1 A score on event evidence

Consider a ledger L whose event gluing types are inhabited. Write $y_e \in X_{U_e}$ for the glued tuple. At a fixed observation schedule, the trace evidence is the cylinder

$$B_L = \{x_{1:N} : x_e|_{U_e} = y_e \text{ for all } e \in [N]\}.$$

The observed-data likelihood of o is $Q_o(B_L)$. A two-part description-length objective [10] is

$$\mathcal{J}_L(o) = \ell(o) - \log_2 Q_o(B_L), \tag{20}$$

$$\mathcal{W}_L(o) = 2^{-\ell(o)} Q_o(B_L). \tag{21}$$

Positive-weight maximizers of $\mathcal{W}_L$ are the finite-cost minimizers of $\mathcal{J}_L$. Zero likelihood has infinite cost. One can compare the weights exactly by rational arithmetic. A full-support candidate ensures a positive weight whenever all event gluing types are inhabited and the candidate is admitted.

In the independent case with one-event law q_o, the likelihood factors as

$$Q_o(B_L) = \prod_{e=1}^{N} q_o^{U_e}(y_e), \qquad q_o^{U_e}(y_e) = \sum_{x|_{U_e}=y_e} q_o(x). \tag{22}$$

The factor is indexed by the event. If two agents view the same event, their reports determine a joint constraint on a single x_e. The product of the two marginal likelihoods in general differs from the mass of that constraint. When an incoming report extends an existing event scope, the contribution of that event is recomputed using the larger scope.

For noisy observations, let $k_e(r_e \mid x)$ denote the joint conditional law for the canonical tuple r_e of reports of event e. If the reports factor across events conditional on the signal trace, the likelihood uses $\sum_x E_o(x \mid z)k_e(r_e \mid x)$ as the emission factor for that event. A known sampling mechanism that selects the scope based on the hidden signal can be included in this channel. The displayed cylinder likelihood applies to a fixed design or to an exogenous schedule conditioned upon. The population result below assumes independence of scope and signal explicitly.

10.2 A finite solver and swarm agreement

For a fixed candidate, define $b_e(z) = \sum_{x|_{U_e}=y_e} E_o(x \mid z)$. The usual forward recursion for a hidden Markov model [11] gives

$$\begin{aligned} \alpha_1(z) &= w_o(z) b_1(z), \\ \alpha_{e+1}(z') &= b_{e+1}(z') \sum_z \alpha_e(z) K_o(z' \mid z), \\ Q_o(B_L) &= \sum_z \alpha_N(z). \end{aligned} \tag{23}$$

By induction $\alpha_e(z)$ is the joint probability of the reports up to e and latent state z at e. Summation over the next latent state gives the indicated transition step, and final summation gives the likelihood. The cost for a candidate with k states is $O(Nk^2)$ arithmetic operations after computing the marginal emission factors. Computing those factors can dominate for large X. The worked example uses factored emissions for which the marginals are explicit.

Proposition 10.1 (Finite optimization and agreement)**.** *For each ledger the following computation returns exactly the positive-weight optimizer type. Merge incoming ledgers by union, group the canonical records by event, construct each* $\mathsf{Glue}_e(L)$*, and evaluate* (23) *and* (21) *for each candidate admitted by the grammar. Return the entire set of maximal positive weights with their candidate descriptions. A conflicting exact event produces its overlap obstruction, and an empty feasible set produces an empty optimizer type. When all agents use the same grammar and the final ledgers agree, the optimizer types agree. In the synchronous setting of Proposition 8.1, agreement therefore holds by the directed diameter after the last injection of observations.*

Proof. The recurrence computes the exact probability of the canonical event evidence. It is insensitive to message order because the grouped event tuples are determined by set union. Exhausting the finite grammar and retaining all maximal weights returns the argmax, which equals the argmin of $\mathcal{J}_L$ on the positive weights. Equal ledgers give equal likelihoods for each candidate. The diameter bound follows from Proposition 8.1. □

After fixing a common ledger, the swarm can partition the candidate type as $\mathcal{O} = \coprod_i \mathcal{O}_i$ and evaluate one part per agent. Each agent returns its maximal positive weight and the candidates attaining it, or the pair $(0, \varnothing)$ if its part has no positive weight. To combine two results, retain the larger weight and its candidates, or take the union of their candidate sets at equal weight. The combined value is the maximum over the union of the two parts and contains all ties. Induction gives the same $\mathsf{Opt}(L)$ as a single exhaustive evaluation. This rule requires that the workers score the same snapshot.

Formally, with decidable rational comparisons, the type returned on a consistent ledger is

$$\mathsf{Opt}(L) = \sum_{o:\mathcal{O}} \left((\mathcal{W}_L(o) > 0) \ \wedge \ \prod_{o':\mathcal{O}} (\mathcal{W}_L(o') \le \mathcal{W}_L(o)) \right).$$

Every condition after the dependent sum is a proposition. For a semantic claim φ about a candidate, the proposition that some optimum validates it and the proposition that every optimum validates it are

$$\| \sum_{o:\mathsf{Opt}(L)} \varphi(o) \|_{-1}, \qquad \prod_{o:\mathsf{Opt}(L)} \varphi(o).$$

These are the quantifiers of Theorem 2.1 along the optimizer projection to the current ledger context. A universal claim intended as evidence of agreement can be conjoined with $\|\mathsf{Opt}(L)\|_{-1}$ to require existence of a fit.

10.3 What the observation design identifies

Consider the independent-model specialization. Suppose events are independent with common signal law $p_\star$ and observation scopes are sampled independently of signals with law w. For one event, the log loss of o on (V, y) is $-\log_2 q_o^V(y)$. Write $R(o)$ for its expectation and $R_\star$ for the expectation when using the true marginal. Then

$$R(o) - R_\star = \sum_V w(V) D_{\mathrm{KL}}(p_\star^V \| q_o^V). \tag{24}$$

This follows by conditioning on V and expanding the expectation. Consequently a candidate reaches $R_\star$ precisely when its marginals agree with the true ones at every positive-probability scope. Candidates with the same marginals at all such scopes have the same observation law. Jointly observing complementary scopes of the same event can enlarge the set of positive-probability scopes, since each glued scope is a union of reported scopes. Equation (24) concerns the environment as seen through that observation design.

For a finite-sample bound, assume the resulting pairs (V_e, y_e) are independent and identically distributed and that each candidate assigns every possible observed value on the sampling support mass at least $\varepsilon > 0$. Put $B = \log_2(1/\varepsilon)$ and let R_n be the average empirical loss over n events. Hoeffding's inequality [12] and a union bound give, with probability at least $1 - \delta$,

$$\sup_{o \in \mathcal{O}} |R_n(o) - R(o)| \le B \sqrt{\frac{\ln(2|\mathcal{O}|/\delta)}{2n}} =: \eta_n.$$

Indeed each loss lies in $[0, B]$. If $o^\star$ minimizes the population risk over the grammar and $\hat{o}$ minimizes $R_n(o) + \ell(o)/n$, optimality and the uniform bound give

$$R(\hat{o}) - R(o^\star) \le 2\eta_n + \frac{\ell(o^\star) - \min_{o \in \mathcal{O}} \ell(o)}{n}.$$

The conclusion controls predictive risk for the observed scopes. The induced observation law determines an equivalence class of ontologies. Even full observations give equal likelihoods to latent descriptions inducing the same trace law.

10.4 Choosing a measurement by information gain

For a finite set of feasible acquisition actions, let $p_o^a(y \mid L)$ be the predictive law for the event report returned by action a. An action may schedule multiple views of the same event. When $Z_L = \sum_o \mathcal{W}_L(o) > 0$, let $\pi_L(o) = \mathcal{W}_L(o)/Z_L$ on the positive-weight support and $m_a(y) = \sum_o \pi_L(o) p_o^a(y \mid L)$. A cost-adjusted choice is to maximize

$$\sum_o \pi_L(o) D_{\mathrm{KL}}(p_o^a(- \mid L) \| m_a) - c(a). \tag{25}$$

The first term is the conditional mutual information between the candidate and the new report [13]. By the chain rule it equals the expected decrease in entropy of the normalized weights after updating on the report. The cost $c(a)$ is expressed in equivalent bits. This uses the same finite probability model for acquisition and for evaluating the fused ledger.

11 Worked swarm on a 64-bit signal type

Let $X = \{0,1\}^{64}$. The first three bits of a signal have parity $x_1 \oplus x_2 \oplus x_3$. Each candidate below gives the last 61 bits an independent fair law, also independent of the first three bits. Agents A, B, C have the following scopes:

$$V_A = \{1,2\} \cup \{4,\dots,23\},$$
$$V_B = \{2,3\} \cup \{24,\dots,43\},$$
$$V_C = \{1,3\} \cup \{44,\dots,64\}.$$

An external generator produces signals and supplies the event identifiers. Each agent receives only its declared projection when it samples an event. The agents relay ledgers along the bidirectional chain $A \leftrightarrow C \leftrightarrow B$.

The grammar has three candidates. The baseline o_0 uses independent fair bits. The candidates o_+ and o_- have two mutually exclusive latent concepts Even and Odd. They emit each first-three-bit configuration of the corresponding parity uniformly within that concept. The independent latent draws assign Even probability $7/8$ for o_+ and $1/8$ for o_-. Use prefix-free descriptions $0, 100, 101$ with lengths $1, 3, 3$, respectively. Each candidate assigns every two-bit pattern among the first three probability $1/4$. Therefore all three have the same exact matching family on the individual sensor scopes and their intersections.

For a worked trace, let A and B observe events 1 through 8. Let C observe events $1, 4, 9$. The first eight first-three-bit tuples, which A and B jointly determine after sharing, are

$$(000,\ 011,\ 101,\ 110,\ 000,\ 011,\ 101,\ 001).$$

The first-and-third-bit observation at event 9 is $(1, 0)$. All choices of the last 61 bits give the same score comparisons in this grammar. This trace is possible under an environment with either biased parity law, and the learner is given the reports and the grammar rather than the generator's law.

Before communication, the three ledgers contain $8, 8, 3$ records. After one synchronous round the sizes are $11, 11, 19$. After two rounds every ledger contains the same 19 records. Grouping them by event gives eight jointly observed triples and one observed pair among the first three bits. The repeated view of bit 2 on each paired event is checked for equality. The additional reports from C on events 1 and 4 also agree on their shared coordinates. There are $8 \cdot 40 + 3 \cdot 21 = 383$ observed suffix bits in total, contributing 383 bits to the negative log likelihood of every candidate.

Candidate	Even triple mass	Odd triple mass	Code length	$\mathcal{J}_L - 383$
o_0	$1/8$	$1/8$	1	27
o_+	$7/32$	$1/32$	3	$45 - 7\log_2 7$
o_-	$1/32$	$7/32$	3	$45 - \log_2 7$

The masses refer to individual triples of the indicated parity. There are seven even triples and one odd triple in the fused ledger, with multiplicity across distinct event identifiers. The last pair contributes two common bits. The reduced scores are approximately 27, 25.3485, 42.1926, and o_+ is the unique optimizer. The inequality $45 - 7\log_2 7 < 27$ is equivalent to the exact integer comparison $7^7 > 2^{18}$. Repeated deliveries leave this inequality and the entire objective unchanged.

If the agents share only their separate marginal histograms, the per-agent marginal log losses are equal across the grammar and the shorter baseline code wins for that criterion. Keeping

event identifiers allows the union scope of two views at the same event to contain the three parity coordinates, which changes the observed law. Prior to any reports the normalized code weights are $(2/3, 1/6, 1/6)$. A fresh single-sensor view has zero information gain because its distribution is common to all three candidates. Scheduling A and B at the same fresh event has information gain $(1 - h_2(7/8))/3 \simeq 0.1521$ bits, where h_2 is binary entropy. The predictive mixture has fair parity, while the two biased candidates have parity entropy $h_2(7/8)$. The suffix coordinates contribute zero to the comparison. This gives a concrete acquisition decision by (25).

Conclusion. The adjoints to substitution on families recover the adjoints to substitution on predicates by reflecting the left adjoint and restricting the right. Internal Kan extensions implement the same pattern with directed comma categories in place of homotopy fibers. Their weighted formulas show which arrow data can be forgotten when checking base change for truth values. The weight comparison must be an equivalence to transport the full types, and equivalence of supports suffices for the logic of propositions.

In the finite swarm, these distinctions determine the evidence to pool, the local predictions to match, and the global explanation type to optimize. Event identifiers govern which partial views contribute to a joint likelihood. The observation design then determines which differences among candidate laws can be learned from that likelihood.

Disclosure. AI assistance was used in drafting and typesetting this manuscript. Exact rational arithmetic was used to verify the worked scores and compare the forward recursion with enumeration on a small masked trace.

References

[1] F. W. Lawvere. Adjointness in foundations. *Dialectica* 23 (1969), 281–296. Reprinted with commentary in *Reprints in Theory and Applications of Categories* 16 (2006), 1–16. Publisher reprint.

[2] J. Lurie. *Higher Topos Theory*. Annals of Mathematics Studies 170, Princeton University Press, 2009. Author's PDF.

[3] L. Martini. Yoneda's lemma for internal higher categories. 2021. arXiv:2103.17141.

[4] L. Martini and S. Wolf. Colimits and cocompletions in internal higher category theory. *Higher Structures* 8(1) (2024), 97–192. doi:10.21136/HS.2024.03.

[5] E. Riehl and D. Verity. Kan extensions and the calculus of modules for ∞-categories. *Algebraic & Geometric Topology* 17(1) (2017), 189–271. arXiv:1507.01460.

[6] M. Shulman. All $(\infty, 1)$-toposes have strict univalent universes. 2019. arXiv:1904.07004.

[7] The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study, 2013. Open-access book.

[8] M. Shapiro, N. Preguiça, C. Baquero, and M. Zawirski. A comprehensive study of convergent and commutative replicated data types. INRIA Research Report RR-7506, 2011. Research report.

[9] N. N. Vorob'ev. Consistent families of measures and their extensions. *Theory of Probability and Its Applications* 7(2) (1962), 147–163. doi:10.1137/1107014.

[10] P. D. Grünwald. *The Minimum Description Length Principle*. MIT Press, 2007. Publisher's edition.

[11] L. R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE* 77(2) (1989), 257–286. doi:10.1109/5.18626.

[12] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58(301) (1963), 13–30. doi:10.1080/01621459.1963.10500830.

[13] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. 2nd edition, Wiley, 2006. doi:10.1002/047174882X.