Fraïssé–Nominal Probability and Commutative Monads for Exchangeable Relational Interfaces

Liz Lemma   Future Detective

January 16, 2026

Abstract

We give a systematic semantics for probabilistic programming over exchangeable relational data in settings where relation queries must be deterministic and memoized. Starting from a Fraïssé limit M of a class of finite structures $\mathcal{K}$, we form the nominal topos $\mathrm{Nom}(M)$ of finitely supported $\mathrm{Aut}(M)$-sets, generalizing the Rado-nominal sets used to model Erdős–Rényi graphons. We then construct an internal probability monad P_M on $\mathrm{Nom}(M)$ whose kernels compose by internal integration on finitely supported events. As in the source paper, full Fubini need not hold; we therefore extract commutative affine submonads generated by specified self-commuting 'generic measures' using a least-submonad construction. This yields Bernoulli-based distributive Markov equational theories for relational programming interfaces in which `new` samples fresh atoms and relation tests are deterministic, yet the induced finite distributions exhibit genuinely 'gray' exchangeable behavior. We work out the random tournament as the first non-graph example and then generalize using ω-categoricity (finite orbit/type partitions). The result is a uniform semantic toolkit connecting program equations, symmetry, and exchangeable limit laws beyond classical measure-theoretic models.

Table of Contents

4. 4. Definability vs support: finitely supported subsets 2^X, parameter-definable sets, and orbit partitions; explicit characterization for $X = |M|^n$.

5. 5. Internal probability measures and integration: support-bounded countable additivity; definition of $\mathsf{P}_M(X)$; internal Lebesgue integral on finitely supported $f : X \to [0,1]$; monotone convergence in $\mathrm{Nom}(M)$.

6. 6. The probability monad P_M: unit (Dirac), bind (integration), kernels and composition; when Fubini fails; determinism and commutative-let laws.

7. 7. Commutative affine submonads from generators: least strong submonad generated by commuting maps; specialization to $\nu : 1 \to \mathsf{P}_M(|M|)$ and bern : $[0,1] \to \mathsf{P}_M(2)$; Bernoulli base on numeral types.

8. 8. Semantics for relational interfaces: interpret `atom` as $|M|$, interpret each relation test by the corresponding deterministic predicate of M, interpret `new` by ν; induced exchangeable finite Σ-models; statement of soundness (program equations $\to$ exchangeability).

9. 9. Worked example: random tournament. Explicit construction of ν by averaging over one-point extension patterns; self-commutation check; resulting exchangeable law and its 'constant-gray' characterization.

10. 10. Generalization schema: free amalgamation Fraïssé classes. Orbit/type-based formula for ν_θ; commutation criteria reduced to finite two-point extension checks; discussion of when uniqueness holds (e.g. tournaments force symmetry).

11. 11. Limitations and future directions: non-free amalgamation (e.g. linear orders), mixtures, conditioning, combining with real-valued primitives, and prospects for computer-checked orbit/type case analyses.

1 Introduction

We study a simple but rigid programming interface: a base type of *atoms* together with an operation new for allocating a fresh atom and, for each relation symbol R in a fixed finite relational signature Σ, a deterministic test

$$R : \mathsf{atom}^{\mathrm{ar}(R)} \to \mathsf{bool}.$$

From these primitives we build terms using products, sums, and sequencing (Kleisli composition). The central semantic constraint is that relational queries are *deterministic* and therefore implicitly *memoized*: once an atom has been allocated, every subsequent evaluation of a relation test on tuples involving that atom must return the same truth value as before. In other words, the probabilistic effects come only from allocation and base random choice (e.g. Bernoulli), while the relational structure itself is not resampled; it is a single ambient object that the program can interrogate repeatedly.

This memoization constraint is easy to state operationally but it cuts across many standard denotational treatments of probabilistic computation. If one interprets atom as an arbitrary measurable space A and new as an ordinary probability measure on A, then deterministic tests $R : A^k \to 2$ are measurable predicates and one can form ordinary Markov kernels. However, such a model does not by itself capture the requirement that the program cannot distinguish atoms except through finitely many comparisons and queries, and that it should be invariant under renaming of atoms that preserves all previously observed information. Put differently, the operational behavior is constrained by symmetries: if we permute the unseen portion of the ambient relational universe, the program should have no observation capable of detecting that permutation. The semantics must therefore encode both probability and symmetry in a way that makes invariance and finite dependence manifest.

The second pressure point is that the event space relevant to such programs is not naturally presented as an external σ-algebra. A program can branch on relation tests involving atoms it has sampled, store atoms in data structures, and later reuse them; the predicates it can form are therefore *definable from finitely many atoms* and are stable under automorphisms fixing those atoms. If we insist on an external measurable structure on A, we must choose which subsets are measurable, and the choice must be compatible with the action of renamings and with the finite-information nature of programs. In many canonical "random structure" constructions, the most natural measurable sets are those generated by cylinder events, but the memoized query interface does not range over arbitrary cylinders: it ranges over definable sets determined by finitely many parameters. We therefore seek a semantics in which the correct event space is not imposed from outside, but arises internally from the same finite-support principle that governs definability and program observability.

These considerations lead us to nominal semantics. The basic move is to fix a countable structure M (in fact the Fraïssé limit of a class of finite Σ-structures) and to take its underlying set $V = |M|$ as the domain of atoms. The automorphism group $G = \mathrm{Aut}(M)$ acts on V, and more generally on all sets built from V. We interpret program denotations not as arbitrary functions, but as G-equivariant maps between G-sets that satisfy a finite-support condition: every element of such a G-set depends only on finitely many atoms in the sense that any automorphism fixing those atoms pointwise fixes the element. This nominal condition expresses precisely the "finite information" intuition: a term of the language can only mention finitely many atoms, hence its behavior must be invariant under automorphisms that leave those atoms unchanged.

Once we adopt nominal objects, deterministic relational queries become canonical. Each relation symbol $R \in \Sigma$ is interpreted by the actual relation $R^M \subseteq V^{\mathrm{ar}(R)}$; the test map $V^{\mathrm{ar}(R)} \to 2$ is G-equivariant by construction. The memoization constraint is then automatic: the relations are properties of the fixed structure M, so the outcome of a query is a deterministic function of the atoms supplied. What remains is to equip this deterministic nominal world with probabilistic choice in a way that respects finite support, yields a useful notion of integration, and validates the equational principles expected of probabilistic programming (notably those arising from commutative sequencing).

The key technical choice is how to represent events and measures internally. For a nominal object X, we take as the "measurable subsets" precisely the finitely supported subsets $S \subseteq X$; these form the internal powerobject 2^X. This reflects the idea that an event is observable only if it is invariant under renamings that fix some finite set of atoms: equivalently, it is definable from finitely many parameters. We then define a probability measure on X as a finitely supported function $\mu : 2^X \to [0, 1]$ satisfying normalization and countable additivity for disjoint families that are themselves controlled by a single finite support. The restriction to *support-bounded* countable families is not an ad hoc weakening; it is the correct internal analogue of countable additivity in a setting where one cannot quantify over an arbitrary external σ-algebra without smuggling in choice principles incompatible with finite support. It is also exactly what is needed to support an internal Lebesgue integration theory for finitely supported functions $X \to [0, 1]$, built by simple-function approximation using finitely supported predicates.

With this event space and integration in hand, we obtain a probability monad on the nominal category: the unit is the Dirac measure, and bind is given by internal integration against finitely supported kernels. This monad plays the role that the Giry monad plays over measurable spaces, but it is tuned to the definability and symmetry constraints imposed by atom allocation and memoized queries. Importantly, the semantics is not merely a rephrasing of classical probability in equivariant language: because we insist on finite support at every stage, the resulting measures and integrals are intrinsically ω-categorical/definability-driven objects rather than arbitrary countably additive measures on an externally fixed σ-algebra.

A further reason to work internally is that commutativity properties cannot be taken for granted. In standard measure-theoretic settings one often relies on Fubini/Tonelli to justify program equations such as commutative sequencing of independent random draws. In our nominal setting, not every measure will commute with itself in the sense required to validate those equations, and even when commutation holds, it must be proven using the finite-support event structure rather than by appealing to external product σ-algebras. Accordingly, we isolate a pragmatic principle: instead of demanding that the full nominal probability monad be commutative, we construct the least strong submonad generated by the specific probabilistic primitives of the language, namely Bernoulli choice on ground booleans and atom allocation by a fixed measure ν on V. When ν satisfies an internal self-commutation (a Fubini symmetry condition against finitely supported test functions), the generated submonad is commutative and affine, and therefore supports the familiar equational theory of commutative probabilistic computation (including weakening/projectivity properties that correspond to deleting unused samples).

The interaction between finite support and probability becomes especially tractable when M is ω-categorical. Oligomorphicity of G implies that over any finite support set $A \subseteq V$, there are only finitely many orbits of n-tuples; equivalently, only finitely many n-types over A. This orbit-finiteness has two consequences that we exploit throughout. First, it identifies finitely supported subsets of V^n with first-order definable subsets over finite parameter sets, so the internal powerobject 2^{V^n} coincides with definable events. Second, it turns many ostensibly analytic statements into finite combinatorics: integrals of definable indicator functions reduce to finite sums over orbit partitions, and self-commutation of a measure ν can be checked by verifying equality on finitely many definable cases over a chosen support. Thus, the nominal-probabilistic semantics is not only conceptually aligned with memoized queries; it is also computationally amenable because definability yields finite partitions.

At a higher level, this semantics is designed to produce exchangeable and projective families of finite Σ-structures. Given ν as the interpretation of new, we may sample n atoms, evaluate all relation tests on the resulting n-tuple, and thereby obtain a random finite Σ-structure on $\{1, \ldots, n\}$. Equivariance of the semantics yields exchangeability under permutations of the sampled atoms, while affineness yields projectivity under deletion of unused samples. In free-amalgamation settings, one expects more: the law should be "local" in the sense that one-point extension information determines the distribution by independent per-relation choices. The nominal framework isolates precisely where such locality comes from (namely, from explicit ν constructed by averaging over one-point extension types) and where it can fail (namely, when amalgamation imposes global constraints that break the multiplicative extension pattern).

We therefore proceed as follows. After recalling in the next section the Fraïssé and oligomorphic background needed to control orbits and definability, we develop the nominal topos $\mathrm{Nom}(M)$ and identify its powerobjects with finitely supported subsets. We then define internal measures and integration and prove that they assemble into a strong probability monad. Finally, we study commutative affine submonads generated by designated probabilistic primitives and give explicit constructions of self-commuting atom-allocation measures in free-amalgamation classes, recovering the canonical "constant-gray" exchangeable laws in examples such as the random graph and the random tournament.

2 Fraïssé limits and oligomorphic groups

We fix throughout a finite relational signature Σ, i.e. a finite set of relation symbols R each equipped with an arity $\mathrm{ar}(R) \geq 1$. A (finite) Σ-structure A consists of a finite set $|A|$ together with interpretations $R^A \subseteq |A|^{\mathrm{ar}(R)}$. All embeddings are assumed to be Σ-embeddings (injective homomorphisms preserving and reflecting all relations).

2.1 Fraïssé classes and their limits

A class $\mathcal{K}$ of finite Σ-structures is a *Fraïssé class* if it satisfies the following standard conditions.

- *Hereditary property (HP):* if $B \in \mathcal{K}$ and A embeds into B, then $A \in \mathcal{K}$.
- *Joint embedding property (JEP):* for all $A, B \in \mathcal{K}$ there exists $C \in \mathcal{K}$ into which both A and B embed.
- *Amalgamation property (AP):* for all embeddings $f_i : A \hookrightarrow B_i$ $(i = 1, 2)$ with $A, B_1, B_2 \in \mathcal{K}$, there exist $C \in \mathcal{K}$ and embeddings $g_i : B_i \hookrightarrow C$ such that $g_1 \circ f_1 = g_2 \circ f_2$.
- *Countability and unboundedness:* there are only countably many isomorphism types in $\mathcal{K}$, and for every n there exists $A \in \mathcal{K}$ with $|A| \geq n$.

Fraïssé's theorem then yields a distinguished countable Σ-structure M characterized uniquely up to isomorphism by two properties: $\mathrm{Age}(M) = \mathcal{K}$ (every finite substructure of M lies in $\mathcal{K}$, and every $A \in \mathcal{K}$ embeds into M), and M is *ultrahomogeneous* (every isomorphism between finite substructures extends to an automorphism of M). We refer to M as the *Fraïssé limit* of $\mathcal{K}$, and we write $V := |M|$ for its underlying set.

We shall use repeatedly a convenient reformulation of the above characterization, namely the *one-point extension property*: if $A \subseteq M$ is finite and $B \in \mathcal{K}$ is a finite structure extending A by one new element, then every embedding $A \hookrightarrow M$ extends to an embedding $B \hookrightarrow M$. In practice, when $\mathcal{K}$ is presented by local constraints (e.g. forbidding finitely many finite patterns) this property is often easy to check directly and serves as a combinatorial proxy for ultrahomogeneity.

We also fix the automorphism group

$$G := \mathrm{Aut}(M),$$

acting on V and diagonally on all finite powers V^n. For a finite set $A \subseteq V$, we write

$$G_A := \{\sigma \in G : \forall a \in A,\ \sigma(a) = a\}$$

for the pointwise stabilizer of A. The pair (M, G) is the symmetry datum governing all later constructions: orbit decompositions under G and under G_A will serve as finite partitions on which definable events and integrals are computed.

2.2 Standard examples

We list the examples that will serve as running test cases.

The random (Rado) graph. Let $\Sigma = \{E\}$ with a single binary relation symbol, interpreted as an irreflexive symmetric edge relation. Let $\mathcal{K}$ be the class of all finite graphs. This class has AP, in fact *free amalgamation*: given B_1, B_2 over A, we may amalgamate by taking the disjoint union over A and adding no edges between the new points of $B_1 \setminus A$ and $B_2 \setminus A$. The Fraïssé limit M is the Rado graph, characterized by the extension property: for every finite disjoint $U, W \subseteq V$ there exists $v \in V$ adjacent to all of U and to none of W.

The countable random tournament. Let $\Sigma = \{\to\}$ with a single binary relation symbol, interpreted as a tournament orientation: for distinct x, y exactly one of $x \to y$ or $y \to x$ holds, and $\to$ is irreflexive. Let $\mathcal{K}$ be the class of all finite tournaments. This class again has free amalgamation in the appropriate oriented sense (the amalgam places no constraints on orientations between new points beyond those already imposed), and the limit M is the random tournament, characterized by the extension property for finite prescribed in/out neighborhoods.

Random k-uniform hypergraphs. Fix $k \geq 2$ and let $\Sigma = \{R\}$ with one k-ary relation symbol, interpreted as an irreflexive symmetric k-uniform hyperedge predicate. Taking $\mathcal{K}$ to be all finite k-uniform hypergraphs, we again obtain a free-amalgamation Fraïssé class and a homogeneous limit M with the expected extension property: given a finite set $A \subseteq V$, every pattern of membership/non-membership for k-tuples involving one new point and $k-1$ points from A is realized by some vertex in V.

Forbidden-substructure generics. Many natural Fraïssé classes are obtained by forbidding a finite configuration; for instance, the Henson graphs are the K_n-free finite graphs. Such classes typically still have AP but not free amalgamation, and their limits exhibit extension properties constrained by the forbidden patterns. They provide useful nontrivial examples where orbit partitions remain finite but one-point extension types are no longer independent across different relations.

2.3 ω-categoricity and orbit finiteness

A central assumption in our development is that the limit M is ω-categorical. We recall the relevant equivalences in the present setting. Let $\mathrm{Th}(M)$ denote the complete first-order theory of M in the language Σ. Then M is ω-*categorical* if $\mathrm{Th}(M)$ has, up to isomorphism, exactly one countable model. By the Engeler–Ryll-Nardzewski–Svenonius theorem, for a countable structure M this is equivalent to the automorphism group $G = \mathrm{Aut}(M)$ being *oligomorphic*, meaning that for every $n \geq 1$ the induced action of G on V^n has only finitely many orbits.

We emphasize the parameterized form of this finiteness, which we will use systematically. Fix a finite parameter set $A \subseteq V$. The group G_A acts on V^n, and ω-categoricity implies that there are only finitely many G_A-orbits on V^n for each n. Equivalently, there are only finitely many complete n-types over A. In more concrete terms: the information carried by an n-tuple $\bar{v} \in V^n$, as far as formulas with parameters from A can distinguish it, takes only finitely many values.

We shall also need the definability of these orbits. If M is ω-categorical, then for each finite $A \subseteq V$ and each n, every G_A-orbit $\mathcal{O} \subseteq V^n$ is definable over A: there exists a first-order formula $\varphi(\bar{x}; \bar{a})$ with parameters $\bar{a}$ enumerating A such that

$$\mathcal{O} = \{\bar{v} \in V^n : M \models \varphi(\bar{v}; \bar{a})\}.$$

One convenient way to see this is to note that the orbit of $\bar{v}$ over A is exactly its type over A, and in an ω-categorical theory every type over a finite set is isolated by a single formula. We will not require a specific normal form (e.g. quantifier elimination), only the consequence that orbit partitions are finite and definable.

We record the orbit-finiteness principle as a standing lemma schema.

Orbit finiteness over finite parameters. *Assume M is ω-categorical. For every finite $A \subseteq V$ and every $n \geq 1$, the action of G_A on V^n has* *finitely many orbits, and each orbit is definable by a first-order formula with parameters from A.*

This principle underlies two later reductions. First, any set or predicate on V^n that is invariant under G_A can be decomposed into a finite union of G_A-orbits, and hence into a finite Boolean combination of definable pieces over A. Second, any equality between functions built from such predicates can be checked orbitwise, reducing analytic-looking statements to finite combinatorics.

2.4 Types, one-point extensions, and free amalgamation

For free-amalgamation Fraïssé classes, it is useful to think in terms of *one-point extension types.* Fix a finite $A \subseteq V$. A one-point extension of A is determined by declaring, for each relation symbol $R \in \Sigma$, which tuples involving one new element x and $\mathrm{ar}(R) - 1$ elements of A satisfy R. When $\mathcal{K}$ has free amalgamation, essentially any such local specification consistent with the basic relation axioms (e.g. symmetry/irreflexivity, tournament antisymmetry) is realized by some $x \in V$. Moreover, the ultrahomogeneity of M implies that the G_A-orbit of x is determined exactly by this one-point pattern relative to A: two points $x, y \in V$ lie in the same orbit under G_A iff they induce the same extension type over A.

We illustrate this in the random graph case. Given $A \subseteq V$ finite and a subset $U \subseteq A$, the pattern "x is adjacent precisely to U" determines a one-point extension type, and by the extension property it is realized. Thus the G_A-orbits on V correspond to subsets $U \subseteq A$, hence there are exactly $2^{|A|}$ such orbits. Similarly, for the random tournament the G_A-orbits on V correspond to orientation patterns, again $2^{|A|}$ possibilities. For a k-uniform hypergraph, the orbits correspond to all choices of which $(k-1)$-tuples from A form a hyperedge together with the new point, yielding $2^{\binom{|A|}{k-1}}$ patterns.

We will later exploit this concrete orbit description to define measures on V by *averaging over one-point extension types.* The key point for the present section is that in the free-amalgamation setting these extension types are purely local data over A, and orbit finiteness is not merely abstract but comes with an explicit finite index set parameterizing the orbits. In non-free settings (e.g. forbidden configurations), extension types may be restricted and may interact globally, but ω-categoricity still guarantees that over each finite A the space of possible types/orbits remains finite.

2.5 Consequences for later constructions

We summarize the two consequences of ω-categoricity that will be used repeatedly.

Finite partitions. For each finite parameter set A and arity n, there exists a finite partition

$$V^n = \bigsqcup_{i=1}^{m} \mathcal{O}_i$$

into G_A-orbits $\mathcal{O}_i$, each definable over A. Any G_A-invariant subset $S \subseteq V^n$ is a union of some subcollection of the $\mathcal{O}_i$. Thus, whenever an object or construction depends only on finitely many parameters A, it is controlled by a finite amount of orbit data.

Checking equalities orbitwise. Suppose $f, g : V^n \to Y$ are functions into a set Y such that f and g are G_A-invariant (equivalently: constant on each $\mathcal{O}_i$). Then $f = g$ holds iff f and g agree on one representative from each orbit. In later measure-theoretic arguments, the relevant f and g will be probabilities obtained by integrating definable indicator functions; orbitwise constancy will allow us to reduce statements such as commutation of iterated integrals to the verification of finitely many cases.

These are the structural inputs we require from Fraïssé theory and oligomorphic group actions. In the next section we pass from the structure M and its symmetry group G to the nominal category built from the G-action, where "dependence on finitely many atoms" is expressed as finite support and will serve as the internal notion of context.

3 The nominal topos $\mathrm{Nom}(M)$

We now fix the Fraïssé limit M and its automorphism group $G = \mathrm{Aut}(M)$, acting on the underlying set $V = |M|$ and diagonally on all finite powers V^n. The category $\mathrm{Nom}(M)$ is the categorical environment in which "dependence on finitely many atoms" is expressed intrinsically, without choosing external names or a global enumeration of V.

3.1 Finitely supported G-sets

A G-set is a set X equipped with an action map $G \times X \to X$, $(\sigma, x) \mapsto \sigma \cdot x$. For a finite set $A \subseteq V$ we write $G_A \leq G$ for the pointwise stabilizer of A. We say that a finite set $A \subseteq V$ *supports* an element $x \in X$ if

$$\forall \sigma \in G_A, \qquad \sigma \cdot x = x.$$

Equivalently, the stabilizer subgroup $\mathrm{Stab}(x) = \{\sigma \in G : \sigma \cdot x = x\}$ contains G_A. A G-set X is *finitely supported* if every element $x \in X$ has some finite support $A \subseteq V$. The objects of $\mathrm{Nom}(M)$ are finitely supported G-sets, and morphisms are *equivariant* functions $f : X \to Y$ (i.e. $f(\sigma \cdot x) = \sigma \cdot f(x)$ for all $\sigma \in G$).

It is often convenient to keep in mind the topological reformulation. The group G carries the pointwise convergence topology (equivalently, the permutation group topology whose basic open neighborhoods of the identity are the subgroups G_A for finite $A \subseteq V$). If X is a G-set with the discrete topology, then the action $G \times X \to X$ is continuous exactly when each point stabilizer $\mathrm{Stab}(x)$ is open, i.e. contains some G_A. Thus "continuous G-action" and "finite support" coincide. We will freely switch between these viewpoints: support calculations are combinatorial, while topos-theoretic properties are inherited from the category of continuous actions of a topological group.

3.2 Least supports

Because our action group arises from an ultrahomogeneous structure, the support relation has a strong closure property under intersection. This yields a canonical notion of least support for each element, which we will use to speak of *the* parameters on which an element depends.

Lemma (intersection of supports). *Let $X \in \mathrm{Nom}(M)$ and $x \in X$. If $A, B \subseteq V$ are finite supports of x, then $A \cap B$ is also a support of x.*

Proof. Let $\sigma \in G_{A \cap B}$. Consider the partial map $h : A \cup B \to A \cup \sigma(B)$ defined as the identity on A and as σ on B. This is well-defined because σ fixes $A \cap B$ pointwise, and it is an isomorphism between the induced finite substructures on $A \cup B$ and $A \cup \sigma(B)$ because it is the identity on A and agrees with the automorphism σ on B. By ultrahomogeneity of M, h extends to some $\tau \in G$. By construction, $\tau \in G_A$ and $\tau(b) = \sigma(b)$ for all $b \in B$, hence $\tau^{-1}\sigma \in G_B$. Since A supports x, we have $\tau \cdot x = x$; since B supports x, we have $(\tau^{-1}\sigma) \cdot x = x$. Therefore $\sigma \cdot x = \tau \cdot ((\tau^{-1}\sigma) \cdot x) = x$, so $A \cap B$ supports x. □

In particular, the family of finite supports of a given x is closed under finite intersections. Since the intersection of any collection of finite sets is again finite, we may define

$$\mathrm{supp}(x) \; := \; \bigcap \{A \subseteq V : A \text{ finite and supports } x\}.$$

By iterating the previous lemma, $\mathrm{supp}(x)$ is itself a support of x, and it is contained in every other support. Hence $\mathrm{supp}(x)$ is the *least* support of x. In later sections we will use least supports to make parameter dependence explicit: whenever a construction is equivariant, it cannot introduce new dependencies on atoms, and its output support is bounded by the supports of its inputs.

Two immediate consequences will be used repeatedly. First, if $f : X \to Y$ is an equivariant map, then $\mathrm{supp}(f(x)) \subseteq \mathrm{supp}(x)$ for all $x \in X$, since any automorphism fixing $\mathrm{supp}(x)$ fixes x and hence fixes $f(x)$. Second, if

$(x_1, \dots, x_n) \in X_1 \times \cdots \times X_n$, then

$$\mathrm{supp}(x_1, \dots, x_n) \;=\; \bigcup_{i=1}^{n} \mathrm{supp}(x_i),$$

as is clear from the componentwise action and the definition of support.

3.3 Limits, colimits, and ground objects

The category $\mathrm{Nom}(M)$ has the expected finite limits and colimits, inherited from the category of all G-sets, with the additional observation that finite support is preserved by these constructions.

Products are given by Cartesian products with componentwise action:

$$\sigma \cdot (x, y) \;:=\; (\sigma \cdot x, \sigma \cdot y).$$

If x is supported by A and y is supported by B, then (x, y) is supported by $A \cup B$. Coproducts are disjoint unions with action defined by $\sigma \cdot \iota_i(x) := \iota_i(\sigma \cdot x)$. Equalizers and coequalizers are formed as the corresponding subsets/quotients in **Set** with the induced action; equivariance of the defining maps ensures stability under the group action, and finite support follows because the underlying action is still continuous.

A distinguished role is played by *ground* objects, by which we mean finite sets equipped with the trivial G-action. If n is a finite set with trivial action, then every element has empty support, $\mathrm{supp}(x) = \varnothing$. Such objects serve as semantics for non-atom data types (booleans, finite sums, finite products of ground data, and so on). In particular, any equivariant map out of a ground object is uniquely determined pointwise, without any hidden dependence on V.

The terminal object 1 is a singleton with trivial action, and the initial object 0 is the empty set. More generally, any finite limit/colimit of ground objects remains ground. This is the categorical reflection of the distinction between parameters of type `atom` (carrying G-action) and purely discrete parameters (with trivial action).

3.4 Exponentials and internal function spaces

The category $\mathrm{Nom}(M)$ is cartesian closed. Given $X, Y \in \mathrm{Nom}(M)$, the exponential Y^X may be described explicitly as follows: its underlying set consists of *finitely supported functions* $f : X \to Y$ (not assumed equivariant), with the conjugation action

$$(\sigma \cdot f)(x) \;:=\; \sigma \cdot f(\sigma^{-1} \cdot x).$$

A finite set $A \subseteq V$ supports f precisely when $\sigma \cdot f = f$ for all $\sigma \in G_A$, i.e. when f is invariant under renamings that fix A. Evaluation $\mathrm{ev} : Y^X \times X \to Y$

is the usual $\mathrm{ev}(f, x) = f(x)$, and the transpose correspondence $Z \times X \to Y \leftrightarrow Z \to Y^X$ is given by currying/uncurrying as in **Set**, with equivariance ensured by the conjugation action.

For later use we note the support bound implicit in this construction: if $f \in Y^X$ has support A and $x \in X$ has support B, then $f(x)$ is supported by $A \cup B$. In other words, applying a finitely supported function does not introduce fresh atoms; it can only combine the atoms already present in the function and the argument.

3.5 Contexts as supports

The intended operational intuition is that elements of V are atoms, and finite supports are contexts of atoms on which an object may depend. This becomes precise in two complementary ways.

First, if $A \subseteq V$ is finite, then invariance under the subgroup G_A expresses "no dependence outside A". For $x \in X$, the statement "A supports x" can be read as: any automorphism of M that fixes A pointwise is a permissible renaming of the rest of the universe, and such renamings do not change x. Thus $\mathrm{supp}(x)$ is the smallest context of atoms that must be fixed to keep x unchanged.

Second, equivariant maps are precisely the maps that respect this notion of context. If $f : X \to Y$ is equivariant, then f commutes with renaming, and hence cannot create additional dependencies: $\mathrm{supp}(f(x)) \subseteq \mathrm{supp}(x)$. This monotonicity is the categorical form of parametricity with respect to atoms.

In the programming interpretation developed later, environments carrying atom variables are represented by products V^n, and a term in context $\Gamma = (a_1 : \mathsf{atom}, \ldots, a_n : \mathsf{atom})$ denotes a G-equivariant map out of V^n (or, in the probabilistic case, a finitely supported kernel out of V^n). The least-support operation then corresponds to extracting the minimal finite set of atoms that a semantic value actually depends upon, which is crucial when we restrict attention to constructions that are stable under the symmetries of M.

This concludes the structural description of $\mathrm{Nom}(M)$. In the next section we connect finite support to first-order definability over M by identifying finitely supported subsets of V^n with unions of orbits over finite parameter sets, and we make this correspondence explicit via orbit partitions.

3.6 Definability, support, and orbit partitions

The internal powerobject of a nominal object X is not the full set-theoretic powerset, but the set 2^X of *finitely supported* subsets of X. Concretely, a subset $S \subseteq X$ is regarded as an element of 2^X when it is stable under renaming outside a finite context: a finite $A \subseteq V$ supports S if

$$\forall \sigma \in G_A, \qquad \sigma[S] = S,$$

where $\sigma[S] = \{\sigma \cdot x : x \in S\}$ is the transported subset under the G-action on X. In this case we also say that S is G_A-invariant. By the least-support lemma from the previous section (applied in the nominal object 2^X), every $S \in 2^X$ has a least finite support $\mathrm{supp}(S) \subseteq V$. Thus 2^X exactly captures those events whose truth is determined by finitely many atoms.

For arbitrary $X \in \mathrm{Nom}(M)$ this is already the correct internal notion of an event space. However, the objects that arise from typing contexts in our programming interpretation are finite powers V^n, equipped with the diagonal G-action. In that case, finite support admits an explicit model-theoretic characterization in terms of first-order definability over the Fraïssé limit M. We now make this correspondence precise, because it will reduce many nominal measurability and commutation questions to finite orbit calculations.

Fix $n \geq 1$. For a finite parameter set $A \subseteq V$, the pointwise stabilizer G_A acts on V^n, and we write $\bar{v} \sim_A \bar{w}$ if $\bar{w} = \sigma \cdot \bar{v}$ for some $\sigma \in G_A$. The equivalence classes are the G_A-orbits in V^n. We call the induced partition the *A-orbit partition* of V^n. A subset $S \subseteq V^n$ is supported by A exactly when it is a union of $\sim_A$-classes: if $\bar{v} \in S$ and $\bar{v} \sim_A \bar{w}$, then $\bar{w} \in S$. In other words, support by A is the same as invariance under the renamings that fix A.

The relevance of ω-categoricity is that these orbit partitions are *finite*. More precisely, since M is ω-categorical, the action of G on V^n has finitely many orbits, and likewise the action of G_A on V^n has finitely many orbits for each finite A. Equivalently, for fixed A there are only finitely many complete n-types over A, and each such type is isolated by a formula over A. We record this in the form we will use.

Lemma (orbit finiteness and definability over parameters). *Let $A \subseteq V$ be finite. Then the G_A-action on V^n has finitely many orbits. Moreover, for each orbit $O \subseteq V^n$ there is a first-order formula $\chi_O(\bar{x}; \bar{a})$ with parameters $\bar{a}$ enumerating A such that*

$$O = \{\bar{v} \in V^n : M \models \chi_O(\bar{v}; \bar{a})\}.$$

We regard this as a standard consequence of the Ryll–Nardzewski theorem together with the usual identification of G_A-orbits with n-types over A: two tuples $\bar{v}, \bar{w} \in V^n$ are in the same G_A-orbit iff they satisfy the same first-order formulas with parameters from A. Since there are only finitely many such types in the ω-categorical setting, each orbit is definable (indeed, definable by a disjunction of formulas isolating the corresponding type).

With this in hand, finitely supported subsets of V^n become exactly the parameter-definable subsets of V^n. We state this identification explicitly.

Proposition (definability = finite support on V^n). *Let $S \subseteq V^n$. Then $S \in 2^{V^n}$ if and only if there exists a first-order formula $\varphi(\bar{x}; \bar{a})$ in the language Σ, with parameters $\bar{a}$ from some finite $A \subseteq V$, such that*

$$S = \{\bar{v} \in V^n : M \models \varphi(\bar{v}; \bar{a})\}.$$

Moreover, when $S \in 2^{V^n}$ and $A = \operatorname{supp}(S)$, the set S is a union of G_A-orbits, hence is definable over A by a finite disjunction of orbit formulas χ_O.

Proof. For the forward direction, assume $S \in 2^{V^n}$ and let $A = \operatorname{supp}(S)$. Then S is G_A-invariant, so it is a union of G_A-orbits. By the lemma, there are only finitely many such orbits, each definable by some formula $\chi_O(\bar{x}; \bar{a})$ over A. Hence S is definable by the finite disjunction

$$\varphi(\bar{x}; \bar{a}) := \bigvee_{O \subseteq S} \chi_O(\bar{x}; \bar{a}),$$

where the disjunction ranges over those G_A-orbits O contained in S.

For the reverse direction, suppose S is defined by a formula $\varphi(\bar{x}; \bar{a})$ with parameters from a finite A. If $\sigma \in G_A$, then σ fixes $\bar{a}$ pointwise and is an automorphism of M, so for all $\bar{v} \in V^n$ we have

$$M \models \varphi(\bar{v}; \bar{a}) \iff M \models \varphi(\sigma \cdot \bar{v}; \sigma \cdot \bar{a}) \iff M \models \varphi(\sigma \cdot \bar{v}; \bar{a}).$$

Thus S is G_A-invariant, i.e. supported by A, hence $S \in 2^{V^n}$. □

This proposition gives an effective translation principle: whenever we encounter an element of 2^{V^n}, we may treat it as a definable predicate on n-tuples with finitely many parameters, and conversely any such definable predicate is a valid internal subset. In particular, the least support $\operatorname{supp}(S)$ can be read as a canonical "minimal parameter set" for the definable set S. We emphasize that the least support is intrinsic, whereas a defining formula is not unique; nevertheless, in the ω-categorical setting every supported set admits a definition over $\operatorname{supp}(S)$ obtained by orbit decomposition.

It is useful to state the orbit-decomposition viewpoint in a form that exposes finiteness. Fix finite $A \subseteq V$. Let $\mathcal{O}_{A,n}$ denote the finite set of G_A-orbits on V^n. Then every A-supported subset $S \subseteq V^n$ corresponds uniquely to a subset $I \subseteq \mathcal{O}_{A,n}$ via

$$S = \bigcup_{O \in I} O.$$

Thus, relative to a fixed support A, the collection of events on V^n is (externally) just a finite boolean algebra. Internally, this finiteness is what makes countable additivity tractable later: for a fixed support, there are only finitely many disjoint definable pieces to consider at once.

The same definability principle extends from subsets to predicates and functions. An element of 2^{V^n} is equivalently a finitely supported map $V^n \to$

2 (its characteristic function). More generally, a finitely supported function $f : V^n \to Y$ into a ground object Y is constant on each G_A-orbit for some finite A, hence determined by finitely many values indexed by $\mathcal{O}_{A,n}$. In the special case $Y = [0,1]$ with trivial action, any finitely supported $f : V^n \to [0,1]$ is again constant on G_A-orbits for some A, so it may be viewed as an A-definable $[0,1]$-valued predicate, constant on types over A. This observation will later allow us to reduce integrals of supported functions to finite sums after partitioning into orbits.

Finally, we note the compatibility with the relational interface. Each relation symbol $R \in \Sigma$ is interpreted in M as a subset $R^M \subseteq V^{\mathrm{ar}(R)}$. Since R^M is definable without parameters (by the atomic formula $R(\bar{x})$), it is supported by the empty set, i.e. $R^M \in 2^{V^{\mathrm{ar}(R)}}$ with $\mathrm{supp}(R^M) = \varnothing$. More generally, any first-order definable construction on M with finitely many parameters yields a finitely supported subset or function in $\mathrm{Nom}(M)$. Thus, when we interpret programs in contexts V^n, the internal events they test and the internal predicates they build remain within the finitely supported universe precisely because they are definable from finitely many atoms.

In summary, on the fundamental objects V^n we may freely pass between three equivalent descriptions: finite support (invariance under some G_A), orbit unions (unions of finitely many G_A-orbits), and first-order definability over finite parameter sets. This equivalence is the point at which ω-categoricity enters the semantics in an essential way: it converts nominal "measurability" questions into finite combinatorics of types and orbits.

3.7 Internal probability measures and integration

To interpret probabilistic programs in $\mathrm{Nom}(M)$ we require a notion of probability measure and integration that (i) lives entirely inside the finitely supported universe, and (ii) is sufficiently countably additive to support sequential composition of probabilistic kernels. The guiding constraint is that in a nominal topos, an arbitrary countable family of events need not be meaningfully "given" without additional choice; accordingly, the correct substitute for σ-additivity is countable additivity along families that are controlled by a *single* finite context.

Support-bounded families. Let $X \in \mathrm{Nom}(M)$. Recall that the internal powerobject 2^X consists of finitely supported subsets $S \subseteq X$. A sequence $(S_i)_{i\in\mathbb{N}} \subseteq 2^X$ is called *support-bounded* if there exists a finite $A \subseteq V$ such that A supports every S_i, i.e. each S_i is G_A-invariant. Equivalently, the entire sequence is a map $\mathbb{N} \to 2^X$ that is itself supported by A when $\mathbb{N}$ carries the trivial action. This boundedness condition is the internal analogue of "measurable sequence": it guarantees that the countable decomposition is not changing its defining parameters along the sequence.

We will only ask for countable additivity along disjoint support-bounded families. This is sufficient for the integration theory needed for the monad, and it becomes concrete on the basic objects V^n because, over any fixed finite support A, there are only finitely many G_A-orbits (hence only finitely many disjoint definable pieces) to which any supported event can reduce.

Finitely supported probability measures. A *(finitely supported) probability measure* on X is a finitely supported function

$$\mu : 2^X \to [0,1]$$

such that

1. $\mu(X) = 1$, and
2. for every disjoint support-bounded sequence $(S_i)_{i\in\mathbb{N}} \subseteq 2^X$ we have

$$\mu\Big(\bigsqcup_{i\in\mathbb{N}} S_i\Big) = \sum_{i\in\mathbb{N}} \mu(S_i),$$

   where $\bigsqcup$ denotes disjoint union in 2^X and the sum is the usual sum in $[0,1]$.

We denote by $\mathsf{P}_M(X)$ the set of all such μ. The group G acts on $\mathsf{P}_M(X)$ by transport of structure:

$$(\sigma\cdot\mu)(S) \ := \ \mu(\sigma^{-1}[S]), \qquad \sigma\in G,\ S\in 2^X.$$

With this action, $\mathsf{P}_M(X)$ is again a nominal object: the support condition on $\mu : 2^X \to [0,1]$ is exactly the requirement that there exists a finite $A \subseteq V$ such that $\mu(\sigma[S]) = \mu(S)$ for all $\sigma \in G_A$ and all $S \in 2^X$. Intuitively, $\mathrm{supp}(\mu)$ is the finite set of atoms relative to which the law μ breaks symmetry.

The restriction to support-bounded families is essential. Without it, countable additivity would quantify over arbitrary countable collections of supported subsets with unbounded parameters, which is incompatible with finite support in general. In contrast, for a fixed support A, the collection of A-supported events in 2^X forms (externally) a set with a finitary description: on $X = V^n$ it is a finite Boolean algebra generated by the finitely many G_A-orbits. Thus the above additivity axiom is a genuine strengthening beyond finite additivity, but remains stable under the nominal finiteness constraints.

Simple functions and their integrals. Let $f : X \to [0,1]$ be finitely supported (here $[0,1]$ has the trivial G-action). We call $s : X \to [0,1]$ a *finitely supported simple function* if there exist finitely supported, pairwise disjoint subsets $S_1,\dots,S_m \in 2^X$ with $\bigcup_{j=1}^m S_j = X$ and scalars $r_1,\dots,r_m \in [0,1]$ such that

$$s(x) = \sum_{j=1}^{m} r_j\,[x\in S_j].$$

Equivalently, s has finite image and each fiber $s^{-1}(r)$ is finitely supported. Any such s is supported by the union of the supports of the S_j's, and thus by a finite subset of V.

Given $\mu \in \mathsf{P}_M(X)$ and a simple s as above, we define its integral by

$$\int s \, d\mu \;:=\; \sum_{j=1}^{m} r_j \, \mu(S_j).$$

This is well-defined: if s admits two simple presentations, then passing to a common refinement (given by intersections of the supporting partitions) yields a presentation with disjoint pieces on which both descriptions agree, and finite additivity of μ along that finite partition implies that the resulting value of $\int s \, d\mu$ is independent of the chosen presentation. In particular, for any $S \in 2^X$ we have

$$\int [x \in S] \, d\mu = \mu(S).$$

Internal Lebesgue integral. For a general finitely supported $f : X \to [0, 1]$, we define the integral by approximation from below by simple functions, exactly as in the classical construction:

$$\int f \, d\mu \;:=\; \sup\Big\{ \int s \, d\mu : s \leq f, \;\; s \text{ a finitely supported simple function} \Big\}.$$

The supremum is taken in the usual complete lattice $[0, 1]$; the expression is meaningful internally because the collection of finitely supported simple functions below f is itself determined by finitely much data (in particular, the support of f). Concretely, if A supports f, then f is constant on each G_A-orbit in X; for $X = V^n$, orbit finiteness over A implies that f is determined by finitely many values, and the above supremum reduces to an ordinary supremum over a directed set of finite approximants.

The integral so defined satisfies the expected basic properties whenever the relevant algebraic operations preserve finite support:

1. *Monotonicity:* if $f \leq g$ pointwise then $\int f \, d\mu \leq \int g \, d\mu$.

2. *Normalization:* $\int 1 \, d\mu = 1$.

3. *Finite linearity:* if f, g are finitely supported and $\alpha, \beta \in [0, 1]$ then $\int (\alpha f + \beta g) \, d\mu = \alpha \int f \, d\mu + \beta \int g \, d\mu$ whenever $\alpha f + \beta g \leq 1$ (or more generally for bounded nonnegative combinations, with the obvious codomain adjustment).

Each property is proved by checking it for simple functions (where it follows from the defining axioms of μ on finite partitions) and then passing to suprema using directedness of the approximation order $s \leq f$.

Monotone convergence under support-boundedness. The key analytical tool for monadic composition is monotone convergence. In our setting the correct hypothesis is again a support-boundedness condition, now on the approximating predicates.

Lemma (internal monotone convergence). *Let $\mu \in \mathsf{P}_M(X)$. Let $(f_i)_{i\in\mathbb{N}}$ be an increasing sequence of finitely supported functions $f_i : X \to [0,1]$ which is support-bounded, i.e. supported by a common finite $A \subseteq V$. Then*

$$\int \Big(\sup_{i\in\mathbb{N}} f_i\Big)\, d\mu \;=\; \sup_{i\in\mathbb{N}} \int f_i \, d\mu.$$

Proof. Write $f = \sup_i f_i$. The inequality $\sup_i \int f_i \, d\mu \;\le\; \int f \, d\mu$ follows immediately from monotonicity of the integral.

For the reverse inequality, fix $\varepsilon > 0$. By definition of $\int f \, d\mu$, there exists a finitely supported simple $s \le f$ such that

$$\int f \, d\mu \;\le\; \int s \, d\mu + \varepsilon.$$

Let $r_1, \dots, r_m$ be the (finite) set of values taken by s. For each j define the supported set $S_j := \{x \in X : s(x) = r_j\}$, so $s = \sum_{j=1}^{m} r_j \, [x \in S_j]$. Since the sequence (f_i) is supported by A, the pointwise supremum f and the inequality $s \le f$ are all witnessed in the A-invariant world: for each $x \in S_j$ we have $r_j \le f(x) = \sup_i f_i(x)$, hence there exists $i = i(x)$ such that $f_{i(x)}(x) \ge r_j$. Consider the sets

$$T_{j,i} := \{x \in S_j : f_i(x) \ge r_j\} \in 2^X.$$

Each $T_{j,i}$ is supported by $A \cup \mathrm{supp}(s)$, uniformly in i and j, and for fixed j the family $(T_{j,i})_{i\in\mathbb{N}}$ is increasing with $\bigcup_i T_{j,i} = S_j$. By countable additivity of μ along support-bounded families (applied to the disjoint decomposition of S_j into successive differences $T_{j,i+1} \setminus T_{j,i}$), we obtain

$$\mu(S_j) = \sup_{i\in\mathbb{N}} \mu(T_{j,i}).$$

Choose i_j such that $\mu(S_j) \le \mu(T_{j,i_j}) + \varepsilon/m$, and let $N := \max_j i_j$. Since $f_N \ge f_{i_j}$ pointwise, we have $T_{j,i_j} \subseteq \{x : f_N(x) \ge r_j\}$, hence

$$\mu(S_j) \le \mu(\{x : f_N(x) \ge r_j\}) + \varepsilon/m.$$

Define the simple function $t(x) := \sum_{j=1}^{m} r_j \, [f_N(x) \ge r_j]$. Then $t \le f_N$ pointwise, and

$$\int s \, d\mu = \sum_{j=1}^{m} r_j \, \mu(S_j) \le \sum_{j=1}^{m} r_j \, \mu(\{x : f_N(x) \ge r_j\}) + \varepsilon = \int t \, d\mu + \varepsilon \le \int f_N \, d\mu + \varepsilon.$$

Combining inequalities,

$$\int f \, d\mu \le \int s \, d\mu + \varepsilon \le \int f_N \, d\mu + 2\varepsilon \le \sup_i \int f_i \, d\mu + 2\varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, we conclude $\int f \, d\mu \le \sup_i \int f_i \, d\mu$, as required. □

This lemma is the point at which the support-boundedness restriction is used critically: it ensures that the relevant increasing unions of events are measured by μ via its countable additivity axiom. In practice, our supported predicates and kernels will always arise from program terms in finite contexts, hence automatically satisfy such uniform support bounds.

We have thus obtained, for each $X \in \mathrm{Nom}(M)$, a well-behaved nominal object $\mathsf{P}_M(X)$ of finitely supported probability measures together with a corresponding integral $\int f \, d\mu$ defined for finitely supported $f : X \to [0,1]$ and enjoying the monotone convergence property in the internal (support-bounded) sense. In the next section we use this integral to define the Kleisli bind operation and thereby assemble P_M into a strong probability monad on $\mathrm{Nom}(M)$.

3.8 The probability monad P_M

We now use the integral constructed above to assemble the assignment $X \mapsto \mathsf{P}_M(X)$ into a probability monad on $\mathrm{Nom}(M)$. The intended interpretation is standard: $\mathsf{P}_M(X)$ is the object of (finitely supported) probability laws on X, a Kleisli morphism $X \to \mathsf{P}_M(Y)$ is a finitely supported Markov kernel from X to Y, and sequential composition is given by integration.

Functoriality (pushforward). For an equivariant map $f : X \to Y$ and $\mu \in \mathsf{P}_M(X)$, we define the pushforward measure $\mathsf{P}_M(f)(\mu) \in \mathsf{P}_M(Y)$ by

$$\mathsf{P}_M(f)(\mu)(T) \;:=\; \mu\big(f^{-1}[T]\big), \qquad T \in 2^Y.$$

Equivariance of f implies that $f^{-1}[T] \in 2^X$ whenever $T \in 2^Y$, and preservation of support-bounded disjoint unions ensures that $\mathsf{P}_M(f)(\mu)$ satisfies the additivity axiom. Moreover, if A supports both f and μ, then A supports $\mathsf{P}_M(f)(\mu)$, so $\mathsf{P}_M(f)$ is a morphism in $\mathrm{Nom}(M)$. This defines P_M on morphisms and makes $\mathsf{P}_M : \mathrm{Nom}(M) \to \mathrm{Nom}(M)$ a functor.

Unit (Dirac measures). For each $X \in \mathrm{Nom}(M)$ we define the unit $\eta_X : X \to \mathsf{P}_M(X)$ by Dirac measures:

$$\eta_X(x)(S) \;:=\; [x \in S], \qquad x \in X,\; S \in 2^X.$$

The map η_X is equivariant and supported by $\mathrm{supp}(x)$. Countable additivity along support-bounded disjoint families is immediate, since $[x \in \bigsqcup_i S_i] = \sum_i [x \in S_i]$ for disjoint families. Thus $\eta_X(x) \in \mathsf{P}_M(X)$ for all x.

Bind (Kleisli extension by integration). Let $\mu \in \mathsf{P}_M(X)$ and let $k : X \to \mathsf{P}_M(Y)$ be a finitely supported map (an equivariant kernel). We define the composite measure $\mu >\!>\!= k \in \mathsf{P}_M(Y)$ by

$$(\mu >\!>\!= k)(T) \ := \int k(x)(T)\,\mu(dx), \qquad T \in 2^Y, \tag{1}$$

where $x \mapsto k(x)(T)$ is a finitely supported function $X \to [0,1]$ and the integral is the internal Lebesgue integral from the previous subsection. To check that $\mu >\!>\!= k$ is indeed a probability measure, we note:

1. Normalization holds since $k(x)(Y) = 1$ for all x, hence $(\mu >\!>\!= k)(Y) = \int 1\,d\mu = 1$.
2. For countable additivity, let $(T_i)_{i\in\mathbb{N}} \subseteq 2^Y$ be disjoint and support-bounded by some finite $A \subseteq V$. Then $x \mapsto k(x)(T_i)$ is supported by $A \cup \operatorname{supp}(k)$, uniformly in i, and the sequence $\big(k(-)(T_i)\big)_{i\in\mathbb{N}}$ is support-bounded in the same sense. By additivity of each $k(x)$ and the monotone convergence lemma applied to partial sums, we obtain

$$(\mu >\!>\!= k)\Big(\bigsqcup_i T_i\Big) = \int \sum_i k(x)(T_i)\,\mu(dx) = \sum_i \int k(x)(T_i)\,\mu(dx) = \sum_i (\mu >\!>\!= k)(T_i).$$

Equivariance of (1) follows from equivariance of k and μ together with invariance of the integral under the G-action on finitely supported functions. Finally, if A supports both μ and k, then A supports $\mu >\!>\!= k$, since for each T the integrand $x \mapsto k(x)(T)$ is A-invariant.

Monad laws. The data $(\mathsf{P}_M, \eta, >\!>\!=)$ satisfy the usual laws.

1. *Left unit.* For $x \in X$ and $k : X \to \mathsf{P}_M(Y)$,

$$(\eta_X(x) >\!>\!= k)(T) = \int k(x')(T)\,\eta_X(x)(dx') = k(x)(T),$$

   since integration against the Dirac measure evaluates at x.

2. *Right unit.* For $\mu \in \mathsf{P}_M(X)$,

$$(\mu >\!>\!= \eta_X)(S) = \int \eta_X(x)(S)\,\mu(dx) = \int [x \in S]\,\mu(dx) = \mu(S).$$

3. *Associativity.* For $\mu \in \mathsf{P}_M(X)$, $k : X \to \mathsf{P}_M(Y)$, and $\ell : Y \to \mathsf{P}_M(Z)$, we must show $(\mu >\!>\!= k) >\!>\!= \ell = \mu >\!>\!= (x \mapsto k(x) >\!>\!= \ell)$. Evaluated at $U \in 2^Z$, both sides reduce to an equality of iterated integrals:

$$\int \Big(\int \ell(y)(U)\,k(x)(dy)\Big)\mu(dx) \ = \ \int \ell(y)(U)\,(\mu >\!>\!= k)(dy).$$

> This is the defining property of bind, obtained by unwinding (1) and using the characterization of the integral as a supremum over simple approximants together with internal monotone convergence to justify passage to directed suprema. Concretely, one first proves the identity for simple functions of y, where it is an instance of finite linearity, and then extends to general finitely supported predicates $y \mapsto \ell(y)(U)$ by approximation.

Thus P_M is a monad on $\mathrm{Nom}(M)$.

Kernels and Kleisli composition. A *(finitely supported) probability kernel* from X to Y is precisely a morphism $k : X \to \mathsf{P}_M(Y)$ in $\mathrm{Nom}(M)$. Given $k : X \to \mathsf{P}_M(Y)$ and $\ell : Y \to \mathsf{P}_M(Z)$, their Kleisli composite $\ell \circ_{\mathrm{Kl}} k : X \to \mathsf{P}_M(Z)$ is defined pointwise by

$$(\ell \circ_{\mathrm{Kl}} k)(x) \;:=\; k(x) >\!\!>\!= \ell.$$

Associativity of $\circ_{\mathrm{Kl}}$ is exactly the associativity law of bind. The unit for Kleisli composition is the Dirac embedding η. This is the categorical form of sequential composition of probabilistic programs: sampling according to $k(x)$ and then sampling according to ℓ from the intermediate result.

Strength and pairing. To interpret probabilistic programs in context (i.e. with parameters that are passed through computations), we use the canonical strength of P_M. Define

$$\mathrm{st}_{X,Y} : X \times \mathsf{P}_M(Y) \to \mathsf{P}_M(X \times Y)$$

by

$$\mathrm{st}_{X,Y}(x,\mu)(S) \;:=\; \int [(x,y) \in S]\, \mu(dy), \qquad S \in 2^{X \times Y}.$$

Equivariance is immediate, and the additivity axiom follows from additivity of μ together with finite additivity of indicators and monotone convergence for support-bounded unions. This strength is compatible with η and $>\!\!>\!=$ in the usual sense, so P_M is a strong monad. In particular, given $\mu \in \mathsf{P}_M(X)$ and $k : X \to \mathsf{P}_M(Y)$, we may form a joint law on $X \times Y$ by the standard construction

$$\mathrm{joint}(\mu,k) \;:=\; \mu >\!\!>\!= \big(x \mapsto \mathrm{st}_{X,Y}(x,k(x))\big) \in \mathsf{P}_M(X \times Y),$$

which internalizes the familiar notion of "sample $x \sim \mu$, then $y \sim k(x)$".

Deterministic maps as kernels. Every equivariant function $f : X \to Y$ induces a deterministic kernel $\eta_Y \circ f : X \to \mathsf{P}_M(Y)$. Under Kleisli composition, these deterministic kernels compose as ordinary functions:

$$(\eta_Z \circ g) \circ_{\mathrm{Kl}} (\eta_Y \circ f) = \eta_Z \circ (g \circ f),$$

and pushing forward a measure $\mu \in \mathsf{P}_M(X)$ along f may be expressed via bind:

$$\mathsf{P}_M(f)(\mu) = \mu >>= (\eta_Y \circ f).$$

This is the formal reason that deterministic constructs (products, coproduct case analysis, and the relation tests interpreted by $[\![R]\!]$) behave as expected in the probabilistic semantics: they are interpreted by equivariant maps and hence embed into the Kleisli category without introducing additional probabilistic effects.

When Fubini fails. Although bind is defined by integration, P_M need not be a commutative monad. Commutativity would assert that for all $\mu \in \mathsf{P}_M(X)$, $\nu \in \mathsf{P}_M(Y)$, and all finitely supported $f : X \times Y \to [0, 1]$, the two iterated integrals agree:

$$\int\!\!\int f(x, y)\, \nu(dy)\, \mu(dx) \;=\; \int\!\!\int f(x, y)\, \mu(dx)\, \nu(dy),$$

equivalently that the two ways of forming a joint law on $X \times Y$ coincide. In classical measure theory this is ensured by σ-additivity and standard Fubini–Tonelli hypotheses. In our nominal setting the additivity axiom is deliberately restricted to support-bounded families, and this restriction is not stable under arbitrary constructions of product measures. Intuitively, even if f is finitely supported, the decompositions of $X \times Y$ that arise when approximating f by simple functions may require countable unions indexed by data that is not uniformly supported after one marginal has been integrated out. The monotone convergence lemma gives exactly the amount of countable additivity needed for associativity of bind, but it does not force symmetry under exchange of integration order.

This noncommutativity has a direct programming interpretation. In the equational theory of probabilistic let-binding, commutativity corresponds to the *commutative-let* equation

$$\mathsf{let}\ x = t\ \mathsf{in}\ \mathsf{let}\ y = u\ \mathsf{in}\ e \;=\; \mathsf{let}\ y = u\ \mathsf{in}\ \mathsf{let}\ x = t\ \mathsf{in}\ e,$$

when x is not free in u and y is not free in t. In the Kleisli semantics this is precisely the assertion that the induced joint law on pairs (x, y) is symmetric under swapping the order of sampling. Since P_M is defined for all finitely supported measures and kernels, but without an a priori Fubini theorem, we do not in general validate such commutations.

Consequence: the role of commutative submonads. We therefore distinguish two layers. The full monad P_M provides a sound semantics for sequential probabilistic computation in $\mathrm{Nom}(M)$, with deterministic constructs embedded by η and composition given by bind. However, when one wishes to validate commutative-let equations (and, later, weakening/affineness principles relevant to projectivity), one must restrict to a commutative affine submonad in which the chosen primitive sampling operations commute. The next section makes this restriction precise by constructing the least strong submonad generated by specified commuting maps, and we will apply it to the two primitives relevant for our relational interface: the Bernoulli family on bool and the atom generator $\nu : 1 \to \mathsf{P}_M(V)$.

3.9 Commutative affine submonads generated by primitives

We now formalize the restriction sketched at the end of the previous subsection. Fix the strong monad P_M on $\mathrm{Nom}(M)$. By a *strong submonad* $\mathsf{T} \subseteq \mathsf{P}_M$ we mean a subfunctor $\mathsf{T}(X) \subseteq \mathsf{P}_M(X)$ on objects, natural in X, such that (i) $\eta_X(x) \in \mathsf{T}(X)$ for all x, (ii) $\mu \in \mathsf{T}(X)$ and $k : X \to \mathsf{T}(Y)$ imply $\mu >>= k \in \mathsf{T}(Y)$, and (iii) $\mathrm{st}_{X,Y}(x, \mu) \in \mathsf{T}(X \times Y)$ whenever $\mu \in \mathsf{T}(Y)$. Equivalently, T is closed under the Kleisli operations needed to interpret probabilistic programs in context.

Commutation of Kleisli arrows. Let $a : A \to \mathsf{P}_M(X)$ and $b : B \to \mathsf{P}_M(Y)$ be Kleisli arrows. We say that a and b *commute* if for every finitely supported predicate $f : X \times Y \to [0, 1]$ the two ways of sampling independently and evaluating f agree, i.e. the following equality holds as a finitely supported map $A \times B \to [0, 1]$:

$$\iint f(x, y)\, b(b_0)(dy)\, a(a_0)(dx) \;=\; \iint f(x, y)\, a(a_0)(dx)\, b(b_0)(dy), \qquad (a_0, b_0) \in A \times B.$$

In the internal language this is the familiar Fubini symmetry for the induced joint law on $X \times Y$; in monadic terms it is the equality of the two composites obtained by using strength to form a joint distribution and then pushing forward along the swap $X \times Y \to Y \times X$. When P_M is commutative, all arrows commute; here we only require this for the chosen primitives.

Least strong submonad generated by a family of arrows. Let $\mathcal{G}$ be a set of Kleisli arrows $g_i : A_i \to \mathsf{P}_M(B_i)$ in $\mathrm{Nom}(M)$. We define the least strong submonad $\langle \mathcal{G} \rangle \subseteq \mathsf{P}_M$ generated by $\mathcal{G}$ by closure under the monadic operations, implemented via a transfinite construction. Concretely, we build classes $\mathsf{T}_\alpha(X) \subseteq \mathsf{P}_M(X)$ by recursion on ordinals:

- $\mathsf{T}_0(X)$ is the set of all Dirac measures $\eta_X(x)$ together with all measures obtainable from the generators by deterministic postprocessing, i.e. $\mathsf{P}_M(f)(\mu)$ where f ranges over equivariant maps and μ is in the image of some g_i.

- Given T_α, we let $\mathsf{T}_{\alpha+1}(X)$ be the closure of $\mathsf{T}_\alpha(X)$ under bind and strength in the following sense:

  $$\mathsf{T}_{\alpha+1}(Y) \;\supseteq\; \{\mu >>= k : \mu \in \mathsf{T}_\alpha(X),\ k : X \to \mathsf{T}_\alpha(Y)\},$$

  and similarly $\mathsf{T}_{\alpha+1}(X \times Y) \supseteq \{\mathrm{st}_{X,Y}(x, \mu) : x \in X,\ \mu \in \mathsf{T}_\alpha(Y)\}$, and we close under equivariant pushforward.

- At limit stages λ we set $\mathsf{T}_\lambda(X) = \bigcup_{\alpha<\lambda} \mathsf{T}_\alpha(X)$.

Since each stage is defined by operations available in $\mathrm{Nom}(M)$ and preserves finite support, the union $\mathsf{T}_\infty(X) := \bigcup_\alpha \mathsf{T}_\alpha(X)$ (for α ranging over all ordinals) yields a well-defined subfunctor of P_M. Standard closure arguments show that T_∞ is a strong submonad, contains every generator g_i, and is least with these properties. We denote it by

$$\langle \mathcal{G} \rangle \;:=\; \mathsf{T}_\infty.$$

Propagation of commutativity from commuting generators. Assume now that the family $\mathcal{G}$ is *pairwise commuting* in the sense above. We claim that $\langle \mathcal{G} \rangle$ is a *commutative* strong monad. The proof is a closure argument: consider the property $\mathcal{C}(\mathsf{T})$ that *all* Kleisli arrows $a : A \to \mathsf{T}(X)$ and $b : B \to \mathsf{T}(Y)$ commute. One checks:

1. $\mathcal{C}$ holds at the base stage for the arrows generated by η and by the given g_i, because (i) Dirac measures commute with everything by definition of integration, and (ii) the g_i commute pairwise by hypothesis and commutation is preserved by deterministic postprocessing.

2. If $\mathcal{C}$ holds for T_α, then it holds for $\mathsf{T}_{\alpha+1}$. Indeed, closure under pushforward does not affect commutation (it amounts to substituting $f \circ (f_X \times f_Y)$ for f). For bind, if a commutes with b, and if a' is obtained from a by sequencing with a kernel k, then the two iterated integrals witnessing commutation of a' and b reduce to the commutation of a and b together with associativity of bind; formally, one expands using $(\mu >>= k)(T) = \int k(x)(T)\,\mu(dx)$ and applies Fubini symmetry at the previous stage to the integrand. Strength is handled similarly: st merely packages a parameter together with a draw, and commutation reduces to the symmetry of the underlying draws.

3. At limits, $\mathcal{C}$ is preserved because any two arrows land in some common stage T_α by finite support considerations, hence commute there.

Thus every stage T_α is commutative in the above sense, and so is $\langle \mathcal{G} \rangle$. In particular, within $\langle \mathcal{G} \rangle$ we validate the commutative-let equations for programs built using only the generators in $\mathcal{G}$ together with deterministic constructs.

Affineness. A strong monad T on a cartesian category is *affine* if there is a canonical map $\mathsf{T}(1) \to 1$ that is an isomorphism (equivalently: discard is a homomorphism for the effect). In our setting, $\mathsf{P}_M(1) \cong 1$ because the only finitely supported probability measure on the terminal object is the unique one. Hence every strong submonad $\mathsf{T} \subseteq \mathsf{P}_M$ inherits $\mathsf{T}(1) \cong 1$, and is therefore affine. Operationally, this affineness underpins weakening/projection principles: unused probabilistic computations may be discarded without affecting the rest of the semantics.

Specialization to ν and bern**.** We now fix the two primitives relevant for our relational interface. Let

$$\nu : 1 \longrightarrow \mathsf{P}_M(V) \qquad \text{and} \qquad \mathrm{bern} : [0,1] \longrightarrow \mathsf{P}_M(2), \quad r \longmapsto r\delta_\top + (1-r)\delta_\bot,$$

where $[0,1]$ carries the trivial G-action. We write

$$\mathsf{P}_{M,\nu} \;:=\; \langle \{\nu, \mathrm{bern}\} \rangle \;\subseteq\; \mathsf{P}_M.$$

By construction, $\mathsf{P}_{M,\nu}$ is a strong submonad containing both primitives and closed under bind and strength, hence sufficient to interpret programs that use new and Bernoulli sampling. If ν commutes with itself (in the sense that ν and ν commute as Kleisli arrows $1 \to \mathsf{P}_M(V)$), then the generating family $\{\nu, \mathrm{bern}\}$ is pairwise commuting: bern commutes with itself because it lives over a trivial G-object and corresponds to an ordinary commutative probability choice, and ν commutes with bern because sampling a fresh atom and sampling an independent boolean do not interact (formally: the two iterated integrals factor through the product measure on $V \times 2$, whose construction uses only strength and bind and is symmetric under swap when one component is a ground distribution). Consequently, under the self-commutation hypothesis on ν, the monad $\mathsf{P}_{M,\nu}$ is commutative and affine.

Bernoulli base on numeral types. Finally, we record the fact that on ground finite types, $\mathsf{P}_{M,\nu}$ reduces to ordinary finite probability. Let n be a numeral object in $\mathrm{Nom}(M)$, i.e. a finite set with trivial G-action. Then every subset of n is finitely supported, so $2^n = \mathcal{P}(n)$, and the additivity axiom for measures reduces to finite additivity. Hence $\mathsf{P}_M(n)$ identifies with the standard simplex of probability distributions on the finite set n, and Kleisli maps $m \to \mathsf{P}_M(n)$ are exactly stochastic matrices $m \rightsquigarrow n$.

The point is that this identification already holds for the generated submonad: $\mathsf{P}_{M,\nu}(n) = \mathsf{P}_M(n)$ for all numerals n. Indeed, the Bernoulli generator bern provides arbitrary binary probabilistic choice on 2, and closure under coproducts, bind, and deterministic maps allows us to build arbitrary convex combinations on any finite set. Concretely, given a distribution $(p_i)_{i=1}^k$ on k, we can implement it by iterated binary choices (e.g. using a decision tree that first selects between $\{1\}$ and $\{2, \ldots, k\}$ with probability p_1, then recurses), and this construction uses only bern, bind, and deterministic case analysis. Since $\mathsf{P}_{M,\nu}$ is closed under precisely these operations, it contains every such finite distribution. Therefore, when we restrict the Kleisli category of $\mathsf{P}_{M,\nu}$ to numeral objects, we recover the classical category FinStoch of finite sets and stochastic matrices. This gives a faithful "observation layer": programs of ground type denote ordinary finite probability distributions, even though the internal semantics lives in the nominal world and is constrained by finite support.

3.10 Semantics for relational interfaces

We interpret the typed language Σ-Prog in the Kleisli category of the commutative affine monad $\mathsf{P}_{M,\nu}$ on $\mathrm{Nom}(M)$. Types are interpreted as nominal objects by the usual clauses: $[\![\mathsf{atom}]\!] := V$, $[\![\mathsf{bool}]\!] := 2$, and products/coproducts by the cartesian and cocartesian structure of $\mathrm{Nom}(M)$. Ground numeral types n are interpreted as finite sets with trivial G-action. Terms in context are interpreted as Kleisli arrows

$$[\![\Gamma \vdash t : \tau]\!] : [\![\Gamma]\!] \longrightarrow \mathsf{P}_{M,\nu}([\![\tau]\!]),$$

defined by the standard monadic interpretation of sequencing (let) via bind, and deterministic constructs via the unit η and functoriality. Since $\mathsf{P}_{M,\nu} \subseteq \mathsf{P}_M$ is a strong submonad, all operations needed for the usual call-by-value interpretation are available.

The relational interface is interpreted by fixing, for each relation symbol $R \in \Sigma$ of arity $k = \mathrm{ar}(R)$, the deterministic predicate induced by the Fraïssé limit M:

$$[\![R]\!] : V^k \longrightarrow 2, \qquad [\![R]\!](\bar{v}) = \top \iff M \models R(\bar{v}).$$

Equivariance is immediate because $G = \mathrm{Aut}(M)$ preserves all relations. The allocation primitive $\mathsf{new} : 1 \to \mathsf{atom}$ is interpreted as the given finitely supported measure $\nu : 1 \to \mathsf{P}_{M,\nu}(V)$. Thus, a program may sample atoms using ν, form tuples, and interrogate M through the deterministic tests $[\![R]\!]$; any remaining probabilistic choice at ground types is provided by bern, and therefore also lies in $\mathsf{P}_{M,\nu}$.

A convenient way to package the observable output of a relational program is to regard a finite Σ-structure on a fixed labeled domain $n = \{1, \ldots, n\}$ as a tuple of truth tables, one for each relation symbol. Accordingly we define, for each numeral object n, the finite set

$$\mathrm{Str}_\Sigma(n) \;:=\; \prod_{R \in \Sigma} 2^{n^{\mathrm{ar}(R)}},$$

where the exponentiation is ordinary set-theoretic since n and 2 carry trivial G-action, hence every subset is finitely supported. Elements of $\mathrm{Str}_\Sigma(n)$ may be identified with Σ-structures with underlying set n (with labeled vertices), and the assignment $n \mapsto \mathrm{Str}_\Sigma(n)$ is functorial for functions between finite sets via pullback/reindexing of relation tables.

Given an n-tuple $\bar{v} = (v_1, \ldots, v_n) \in V^n$, we obtain an induced labeled Σ-structure on n by interpreting the i-th label as v_i and reading off all relation facts from M. This yields an equivariant map

$$\mathrm{ev}_n : V^n \longrightarrow \mathrm{Str}_\Sigma(n),$$

defined componentwise by

$$\big(\mathrm{ev}_n(\bar{v})\big)_R(\bar{i}) = \top \iff M \models R(v_{i_1}, \dots, v_{i_k}), \qquad \bar{i} = (i_1, \dots, i_k) \in n^k.$$

Equivariance follows from automorphism-invariance of M, and triviality of the G-action on $\mathrm{Str}_\Sigma(n)$. Thus, any probability measure on V^n induces, by pushforward along ev_n, an ordinary distribution on the finite set $\mathrm{Str}_\Sigma(n)$.

We now define the canonical n-vertex sampler associated to ν. Consider the Kleisli arrow $\mathsf{new}^n : 1 \to \mathsf{P}_{M,\nu}(V^n)$ obtained by iterating ν using bind and pairing. Concretely, in internal language this is the program

$$\mathsf{let}\ x_1 = \mathsf{new}\ \mathsf{in}\ \cdots\ \mathsf{let}\ x_n = \mathsf{new}\ \mathsf{in}\ (x_1, \dots, x_n),$$

whose denotation we also write as $\nu^{(n)} \in \mathsf{P}_{M,\nu}(V^n)$. We then set

$$\mu_n \ := \ \mathsf{P}_{M,\nu}(\mathrm{ev}_n)\big(\nu^{(n)}\big) \ \in \ \mathsf{P}_{M,\nu}(\mathrm{Str}_\Sigma(n)).$$

Since $\mathrm{Str}_\Sigma(n)$ is a finite object with trivial action, μ_n may be read as an ordinary probability distribution on labeled n-vertex Σ-structures (equivalently: on $\mathrm{Str}_\Sigma(n)$ as a finite set). This is the basic "finite-model law" induced by ν and the interface Σ.

The commutativity of $\mathsf{P}_{M,\nu}$ supplies exchangeability of these finite-model laws. Let $\pi \in S_n$ be a permutation. It acts on V^n by permuting coordinates, via the deterministic map $\pi_* : V^n \to V^n$, $(v_1, \dots, v_n) \mapsto (v_{\pi^{-1}(1)}, \dots, v_{\pi^{-1}(n)})$. On $\mathrm{Str}_\Sigma(n)$ it acts by relabeling the domain, via $\pi^\sharp : \mathrm{Str}_\Sigma(n) \to \mathrm{Str}_\Sigma(n)$ given by precomposition on each relation table:

$$\big(\pi^\sharp(s)\big)_R(\bar{i}) = s_R(\pi^{-1}(\bar{i})).$$

By construction, $\mathrm{ev}_n \circ \pi_* = \pi^\sharp \circ \mathrm{ev}_n$. Moreover, commutativity of $\mathsf{P}_{M,\nu}$ implies that the n-fold sampler $\nu^{(n)}$ is invariant under reordering of the n independent uses of new, i.e.

$$\mathsf{P}_{M,\nu}(\pi_*)\big(\nu^{(n)}\big) = \nu^{(n)}.$$

Indeed, this is precisely the content of the commutative-let equation specialized to repeated sampling from ν and deterministic tuple formation. Pushing forward along ev_n and using naturality, we obtain

$$\mathsf{P}_{M,\nu}(\pi^\sharp)(\mu_n) = \mathsf{P}_{M,\nu}(\mathrm{ev}_n)\big(\mathsf{P}_{M,\nu}(\pi_*)(\nu^{(n)})\big) = \mathsf{P}_{M,\nu}(\mathrm{ev}_n)\big(\nu^{(n)}\big) = \mu_n.$$

Thus μ_n is invariant under relabeling by S_n, i.e. the induced random labeled Σ-structure on n vertices is exchangeable.

Affineness gives the expected projectivity with respect to deletion (or, more generally, restriction along injections of finite sets). Let $\iota : m \hookrightarrow n$ be an injection between numeral objects. There is an induced restriction map

$\iota^\sharp : \mathrm{Str}_\Sigma(n) \to \mathrm{Str}_\Sigma(m)$ obtained by pulling back each relation table along ι. On tuples, we have the deterministic projection map $\iota_* : V^n \to V^m$ selecting the coordinates in the image of ι. Again, $\mathrm{ev}_m \circ \iota_* = \iota^\sharp \circ \mathrm{ev}_n$. Consider the program that samples n atoms, then discards the unused coordinates and keeps only the m selected ones; in the Kleisli semantics, discarding is interpreted by the unique map to 1, and affineness ensures that discarding commutes with probabilistic effects. Consequently the marginal of $\nu^{(n)}$ along ι_* coincides with $\nu^{(m)}$:

$$\mathsf{P}_{M,\nu}(\iota_*)\big(\nu^{(n)}\big) = \nu^{(m)}.$$

Pushing forward through ev_m yields the projectivity equation

$$\mathsf{P}_{M,\nu}(\iota^\sharp)(\mu_n) = \mathsf{P}_{M,\nu}(\mathrm{ev}_m)\big(\mathsf{P}_{M,\nu}(\iota_*)(\nu^{(n)})\big) = \mathsf{P}_{M,\nu}(\mathrm{ev}_m)\big(\nu^{(m)}\big) = \mu_m.$$

In particular, the family $(\mu_n)_{n\in\mathbb{N}}$ forms a consistent system of finite-dimensional distributions: restricting an n-vertex sample to any m labels yields the m-vertex sample.

These two properties (exchangeability and projectivity) may be viewed as soundness consequences of the equational theory validated by a commutative affine monad: commutativity of let yields invariance under permutations of independent draws, and affineness yields stability under weakening (discarding unused probabilistic computations). In our setting, this can be stated purely denotationally, without appealing to any external probability space. For any closed term t of ground type (e.g. $t : 1 \to n$ or $t : 1 \to \mathrm{Str}_\Sigma(n)$), its denotation is a measure in $\mathsf{P}_{M,\nu}(n)$ or $\mathsf{P}_{M,\nu}(\mathrm{Str}_\Sigma(n))$, hence an ordinary finite distribution. Program equations derivable from the commutative affine axioms (together with $\beta\eta$ for the deterministic fragment) are sound because they hold in the Kleisli category of $\mathsf{P}_{M,\nu}$; in particular, any two syntactic ways of expressing "sample n atoms and then compute the induced Σ-structure" denote the same μ_n, and any two orderings of the n samples are provably equal and therefore yield the same distribution.

Finally, we record the operational reading relevant for relational modeling. A program that first allocates finitely many atoms and then returns a value of numeral type (or, more generally, any finite coproduct of numerals) is observed as an ordinary stochastic computation, because the codomain carries trivial action and $\mathsf{P}_{M,\nu}$ agrees there with finite probability. When the program returns the full relational table $\mathrm{Str}_\Sigma(n)$, we obtain precisely an exchangeable random finite Σ-structure on n labeled vertices, and the projectivity equations above ensure compatibility across n. Thus the semantics of the Σ-interface produces, from the single primitive measure ν on V, a canonical exchangeable projective family of finite Σ-models. In the subsequent worked examples we will identify these families explicitly for particular Fraïssé limits and particular choices of ν, showing that the abstract nominal-monadic construction recovers the expected "constant-gray" laws in free-amalgamation settings.

3.11 Worked example: the random tournament

We specialize to the Fraïssé class of finite tournaments. Thus $\Sigma = \{T\}$ where T is a binary relation intended to satisfy, for all distinct x, y, exactly one of $T(x, y)$ and $T(y, x)$, and never $T(x, x)$. Let M be the countable random tournament (the Fraïssé limit), $V = |M|$, and $G = \mathrm{Aut}(M)$. In this setting there is essentially a unique "constant-gray" choice: antisymmetry forces the edge bias to be $1/2$. We now give an explicit finitely supported measure $\nu_{1/2} \in \mathsf{P}_M(V)$, prove that it self-commutes, and identify the induced finite-model laws μ_n as the classical exchangeable random tournament on n labeled vertices.

One-point extension patterns and unary orbit partitions. Fix a finite parameter set $A \subseteq V$. For each function $\epsilon : A \to 2$, consider the set

$$O_\epsilon^A \;:=\; \Big\{ v \in V \setminus A \;:\; \forall a \in A,\; \big(T(v, a) \leftrightarrow \epsilon(a) = 1\big) \Big\}.$$

By ultrahomogeneity and the one-point extension property of the random tournament, each O_ϵ^A is nonempty; moreover, each O_ϵ^A is a single orbit for the pointwise stabilizer G_A, and the family $\{O_\epsilon^A\}_{\epsilon \in 2^A}$ partitions $V \setminus A$. Together with the singleton orbits $\{a\}$ for $a \in A$, we obtain a finite G_A-orbit decomposition of V.

By Theorem B, every A-supported event $S \in 2^V$ is definable over A, hence a union of G_A-orbits. In particular, $S \cap (V \setminus A)$ is a union of some of the O_ϵ^A, and membership of S on $V \setminus A$ is constant on each O_ϵ^A.

Definition of $\nu_{1/2}$ by uniform pattern averaging. For an event $S \in 2^V$, let $A := \mathrm{supp}(S)$ be its least support (Lemma 1). We define

$$\nu_{1/2}(S) \;:=\; \frac{1}{2^{|A|}} \sum_{\epsilon \in 2^A} \big[\, O_\epsilon^A \subseteq S \,\big]. \tag{2}$$

Equivalently (using that $S \cap (V \setminus A)$ is a union of some O_ϵ^A), we may write

$$\nu_{1/2}(S) \;=\; \frac{|\{\epsilon \in 2^A : O_\epsilon^A \subseteq S\}|}{2^{|A|}},$$

and we stipulate that points of A themselves receive no mass (indeed, they do not appear in the sum). Intuitively: relative to parameters A, a "fresh" vertex is specified exactly by the orientation pattern it makes with A, and $\nu_{1/2}$ chooses this pattern uniformly.

We must check that (2) is well-defined (i.e. independent of the chosen supporting set). Suppose $B \supseteq A$ is another finite set supporting S. Then each O_ϵ^A refines into $2^{|B \setminus A|}$ many G_B-orbits $O_{\epsilon'}^B$, where $\epsilon' : B \to 2$ extends ϵ

arbitrarily on $B \setminus A$. Because S is B-supported, it is a union of G_B-orbits, so for a fixed ϵ either all extensions ϵ' give $O^B_{\epsilon'} \subseteq S$, or none do. Consequently,

$$\frac{1}{2^{|B|}} \sum_{\epsilon' \in 2^B} [O^B_{\epsilon'} \subseteq S] \;=\; \frac{1}{2^{|A|}} \sum_{\epsilon \in 2^A} [O^A_{\epsilon} \subseteq S],$$

as required. Thus $\nu_{1/2}$ is a well-defined function $2^V \to [0,1]$.

Measure axioms and finite support. For fixed support A, the algebra of A-supported events is finite: it is generated by the finitely many orbits $\{a\}$ ($a \in A$) and O^A_ϵ ($\epsilon \in 2^A$). On this finite Boolean algebra, $\nu_{1/2}$ is plainly a probability measure: $\nu_{1/2}(V) = 1$, and finite additivity holds because (2) is just uniform counting over $\epsilon \in 2^A$. Countable additivity for support-bounded disjoint families reduces to finite additivity since there are only finitely many A-orbits available to partition $V \setminus A$. Hence $\nu_{1/2} \in \mathsf{P}_M(V)$.

Moreover $\nu_{1/2}$ is G-invariant, hence supported by $\emptyset$. Indeed, for any $\sigma \in G$ and event S, we have $\mathrm{supp}(\sigma S) = \sigma(\mathrm{supp}(S))$, and σ bijects the pattern-orbits O^A_ϵ with $O^{\sigma A}_{\epsilon \circ \sigma^{-1}}$. Therefore the uniform average in (2) is preserved:

$$\nu_{1/2}(\sigma S) = \nu_{1/2}(S).$$

Two-point extension patterns. We now verify that $\nu_{1/2}$ self-commutes. By Lemma 4 it suffices to check the Fubini symmetry on indicator functions of definable subsets $D \subseteq V^2$ supported by a fixed finite $A \subseteq V$. Since $\nu_{1/2}$ assigns mass 0 to any finite set, we may ignore the A-diagonals and equalities (i.e. tuples where $x \in A$ or $y \in A$ or $x = y$); these contribute 0 to both iterated integrals. Thus we restrict attention to $(V \setminus A)^2$ with $x \neq y$.

For $\epsilon, \epsilon' \in 2^A$ and $t \in 2$, define the $(A, \epsilon, \epsilon', t)$-cell

$$O^A_{\epsilon,\epsilon',t} := \Big\{(x,y) \in (V\setminus A)^2 : x \neq y,\ \forall a \in A,\ T(x,a) \leftrightarrow \epsilon(a) = 1,\ T(y,a) \leftrightarrow \epsilon'(a) = 1,\ T(x,y) \leftrightarrow t = 1$$

As before, by ultrahomogeneity and free choice of orientations in a tournament, each $O^A_{\epsilon,\epsilon',t}$ is nonempty and is a single G_A-orbit in $(V \setminus A)^2$. The family $\{O^A_{\epsilon,\epsilon',t}\}$ yields a finite orbit partition of the generic part of V^2, and any A-supported definable $D \subseteq V^2$ is a union of some of these orbits (up to the measure-zero exceptional set where $x \in A$ or $y \in A$ or $x = y$).

Computation of iterated integrals and symmetry. Fix such a definable D supported by A, and write 1_D for its indicator. Consider the iterated integral $\int \big(\int 1_D(x,y)\, \nu_{1/2}(dy)\big)\, \nu_{1/2}(dx)$. For a given $x \in V \setminus A$, its A-type is the unique $\epsilon \in 2^A$ such that $x \in O^A_\epsilon$. The inner event $\{y : (x,y) \in D\}$ is definable over $A \cup \{x\}$. By the definition of $\nu_{1/2}$, $\nu_{1/2}\big(\{y : (x,y) \in D\}\big)$ is obtained by uniformly averaging over the $2^{|A|+1}$ one-point patterns of y

relative to $A \cup \{x\}$; these patterns consist of an $\epsilon' \in 2^A$ together with the bit $t \in 2$ describing the direction between x and y (equivalently, $T(x, y)$). Since D is a union of the cells $O^A_{\epsilon,\epsilon',t}$, the value of the inner integral depends on x only through ϵ, and equals

$$\frac{1}{2^{|A|+1}} \sum_{\epsilon' \in 2^A} \sum_{t \in 2} [\, O^A_{\epsilon,\epsilon',t} \subseteq D \,].$$

Averaging once more over x means averaging uniformly over $\epsilon \in 2^A$. Hence

$$\iint 1_D(x, y)\, \nu_{1/2}(dy)\, \nu_{1/2}(dx) \;=\; \frac{1}{2^{2|A|+1}} \sum_{\epsilon \in 2^A} \sum_{\epsilon' \in 2^A} \sum_{t \in 2} [\, O^A_{\epsilon,\epsilon',t} \subseteq D \,]. \quad (3)$$

If we reverse the order of integration, the same argument yields

$$\iint 1_D(x, y)\, \nu_{1/2}(dx)\, \nu_{1/2}(dy) \;=\; \frac{1}{2^{2|A|+1}} \sum_{\epsilon' \in 2^A} \sum_{\epsilon \in 2^A} \sum_{t \in 2} [\, O^A_{\epsilon,\epsilon',t} \subseteq D \,],$$

which is equal to (3) by commutativity of finite sums. This establishes the required Fubini symmetry on definable indicators, hence $\nu_{1/2}$ self-commutes.

Induced finite-model law and "constant-gray". Let $\nu^{(n)} \in \mathsf{P}_{M,\nu_{1/2}}(V^n)$ be the n-fold sampler obtained by iterating new, and let μ_n be the pushforward of $\nu^{(n)}$ along $\mathrm{ev}_n : V^n \to \mathrm{Str}_\Sigma(n)$. We claim that μ_n is the uniform distribution on labeled tournaments on n vertices, i.e. each orientation of the $\binom{n}{2}$ unordered pairs occurs with probability $2^{-\binom{n}{2}}$.

This follows by an induction that mirrors the pattern-averaging definition. Suppose we have already sampled $x_1, \dots, x_m$, and condition on their realized tuple. The conditional distribution of x_{m+1} relative to the finite set $A = \{x_1, \dots, x_m\}$ is, by construction of $\nu_{1/2}$, uniform over the 2^m one-point extension patterns $\epsilon : A \to 2$. But such a pattern is exactly a choice, independently for each $i \le m$, of whether $T(x_{m+1}, x_i)$ holds. Thus, conditional on the past, the orientations $\big(T(x_{m+1}, x_i)\big)_{i \le m}$ are independent fair coins. Iterating this step shows that for the induced labeled tournament on $\{1, \dots, n\}$, all edge directions are independent and unbiased. In other words, μ_n is the classical random tournament law.

In the language of limit objects, this is precisely the "constant-gray" tournamenton: the measurable kernel W on $[0, 1]^2$ is (up to null sets) constantly $1/2$ on the off-diagonal, and antisymmetry forces $W(x, y) = 1 - W(y, x)$, hence forces the constant to be $1/2$. Consequently, unlike the graph case where an Erdős–Rényi parameter α yields a family ν_α, the tournament signature admits no nontrivial one-parameter constant-gray deformation compatible with antisymmetry. Our explicit $\nu_{1/2}$ therefore realizes the canonical exchangeable projective family of finite random tournaments arising from the random tournament Fraïssé limit and the uniform one-point extension mechanism.

3.12 Generalization schema: free amalgamation Fraïssé classes

We explain a uniform construction of "constant-gray" measures for Fraïssé classes with free amalgamation, and how the commutation requirement reduces to a finite, orbit-wise calculation. Throughout, Σ is a finite relational signature and $\mathcal{K}$ is a Fraïssé class of finite Σ-structures whose Fraïssé limit M is ω-categorical. We additionally assume that $\mathcal{K}$ has free amalgamation in the usual relational sense: whenever $B_1, B_2 \in \mathcal{K}$ intersect in a common substructure A, there is an amalgam $C \in \mathcal{K}$ whose underlying set is the disjoint union of B_1 and B_2 over A, and such that no new relations are imposed between $B_1 \setminus A$ and $B_2 \setminus A$ beyond those forced by the axioms of the class (e.g. irreflexivity, symmetry/antisymmetry conventions, etc.). In this regime, one-point extensions over a finite parameter set are combinatorially unconstrained, and this is precisely what allows a simple multiplicative specification of ν.

One-point extension orbits and their parametrization. Fix a finite parameter set $A \subseteq V$. Consider the action of the pointwise stabilizer G_A on $V \setminus A$. By ω-categoricity (Lemma 2), there are finitely many G_A-orbits, and by Theorem B these orbits are exactly the realized 1-types over A (equivalently, quantifier-free one-point extension patterns over A, since M is ultrahomogeneous and $\mathcal{K}$ is relational). We denote the orbit set by

$$\mathrm{Orb}_1(A) \;:=\; (V \setminus A)/G_A,$$

and for $p \in \mathrm{Orb}_1(A)$ we write $O_p^A \subseteq V \setminus A$ for the corresponding orbit. In free amalgamation classes, an orbit p can be concretely described by the truth-values of all atomic Σ-facts in which the new variable x appears together with a tuple of parameters from A, modulo the equational constraints on $\mathcal{K}$ (e.g. prohibiting repetitions, enforcing symmetry, etc.). In particular, $\mathrm{Orb}_1(A)$ is finite and effectively enumerable from $|A|$ and Σ once the axioms of $\mathcal{K}$ are fixed.

Orbit/type weights and the formula for ν_θ. We choose a parameter vector θ assigning a bias to each atomic relation template in which a fresh point participates. Concretely, for each $R \in \Sigma$ of arity $r = \mathrm{ar}(R)$, we pick a real number $\theta_R \in [0,1]$, with the understanding that θ_R is interpreted as the intended probability that $R(x, \bar{a})$ holds for a fresh x and a parameter tuple $\bar{a} \in A^{r-1}$ satisfying whatever distinctness conditions are imposed by $\mathcal{K}$. (If $\mathcal{K}$ enforces identifications between R-facts, e.g. symmetry, one first fixes a canonical representative scheme and assigns θ_R to those representatives; we suppress this bookkeeping since it is signature- and axiom-dependent but finite.)

Given A and an orbit $p \in \mathrm{Orb}_1(A)$, define its (unnormalized) weight by the multiplicative rule

$$w_A^\theta(p) \;:=\; \prod_{R\in\Sigma} \; \prod_{\bar a \in \mathrm{Adm}_R(A)} \theta_R^{\,\mathbf{1}[R(x,\bar a)\in p]} \,(1-\theta_R)^{\,\mathbf{1}[\neg R(x,\bar a)\in p]}, \tag{4}$$

where $\mathrm{Adm}_R(A)$ ranges over those $(r-1)$-tuples $\bar a$ of parameters to which the class allows relating a new point x via R (e.g. excluding tuples with repetitions if the class is irreflexive in the relevant coordinates), and where the exponents refer to whether the corresponding atomic fact is forced true/false by the type p. Free amalgamation ensures that every such assignment of atomic facts consistent with the axioms is realized in M, so each admissible pattern contributes an orbit p with $w_A^\theta(p) \geq 0$, and at least one orbit has positive weight provided $\theta_R \notin \{0,1\}$ are chosen compatibly with the axioms.

Let $Z_A^\theta := \sum_{p\in\mathrm{Orb}_1(A)} w_A^\theta(p)$. Define the normalized distribution $\pi_A^\theta(p) := w_A^\theta(p)/Z_A^\theta$. For an event $S \in 2^V$ with least support $A = \mathrm{supp}(S)$, we then set

$$\nu_\theta(S) \;:=\; \sum_{p\in\mathrm{Orb}_1(A)} \pi_A^\theta(p)\cdot \big[\, O_p^A \subseteq S \,\big]. \tag{5}$$

Thus $\nu_\theta(S)$ is determined by which one-point orbits over A are included in S, weighted according to θ. This is the direct analogue of uniform pattern averaging, with uniformity replaced by the product-form bias (4).

Well-definedness under enlargement of supports. To verify that (5) does not depend on the particular supporting set used to describe S, we must compare the distributions π_A^θ and π_B^θ when $B \supseteq A$. The crucial property is a finite marginalization identity: each orbit $q \in \mathrm{Orb}_1(B)$ restricts to a unique orbit $\mathrm{res}_{B\to A}(q) \in \mathrm{Orb}_1(A)$, and in free amalgamation classes the weight factors for atomic facts involving parameters in $B \setminus A$ separate from those involving only A. Concretely, for fixed $p \in \mathrm{Orb}_1(A)$, we have

$$\sum_{q:\ \mathrm{res}_{B\to A}(q)=p} w_B^\theta(q) \;=\; w_A^\theta(p)\cdot C_{A,B}^\theta,$$

where $C_{A,B}^\theta$ is a constant independent of p (it is the total weight contributed by choices of atomic facts between x and $B \setminus A$, summed over all consistent such choices; in the product-form case, it is a finite product of terms of the form $\theta_R + (1-\theta_R) = 1$, possibly multiplied by class-dependent constants arising from canonical-representative conventions). After normalization, this yields

$$\pi_A^\theta(p) \;=\; \sum_{q:\ \mathrm{res}_{B\to A}(q)=p} \pi_B^\theta(q).$$

Since S is B-supported iff it is a union of G_B-orbits, the indicator $[O_q^B \subseteq S]$ is constant on all q restricting to the same p, and the above marginalization implies that the right-hand side of (5) computed with A agrees with the same expression computed with B. Hence ν_θ is well-defined as a function $2^V \to [0,1]$.

Measure axioms and equivariance. For a fixed finite A, the Boolean algebra of A-supported events in 2^V is finite, being generated by the finitely many singleton orbits in A together with the finitely many orbits O_p^A in $V \setminus A$. On this finite algebra, ν_θ is a probability measure by construction: $\nu_\theta(V) = 1$ and additivity on disjoint unions follows from additivity of the sum in (5). Countable additivity for support-bounded disjoint families reduces to finite additivity because there are only finitely many relevant orbits over the common support. Thus $\nu_\theta \in \mathsf{P}_M(V)$.

Equivariance is immediate from the orbit-based definition. For $\sigma \in G$, the map σ sends $\mathrm{Orb}_1(A)$ bijectively to $\mathrm{Orb}_1(\sigma A)$ and preserves the truth-values of atomic facts, hence preserves the weights w_A^θ and therefore the normalized probabilities π_A^θ. Consequently $\nu_\theta(\sigma S) = \nu_\theta(S)$ for all S, so ν_θ is G-invariant (supported by $\emptyset$).

Self-commutation via two-point extension checks. Assuming ν_θ is defined as above, we now indicate the general commutation argument. By Lemma 4, it suffices to check Fubini symmetry on indicators 1_D where $D \subseteq V^2$ is definable over a fixed finite A. As in the tournament computation, we may ignore a finite exceptional set (coordinates in A, or diagonals) since ν_θ gives mass 0 to finite subsets.

On the generic part of $(V \setminus A)^2$ with $x \neq y$, the stabilizer G_A has finitely many orbits $\mathrm{Orb}_2(A)$, corresponding to realized 2-types over A. Each $q \in \mathrm{Orb}_2(A)$ specifies simultaneously: the one-point orbit $p_x \in \mathrm{Orb}_1(A)$ of x, the one-point orbit $p_y \in \mathrm{Orb}_1(A)$ of y, and the atomic Σ-facts relating x and y (and, in higher arity, facts involving both x and y together with some parameters from A). Free amalgamation again ensures that all consistent such specifications occur and form single orbits.

The key observation is that sequential sampling according to ν_θ induces a product-form weight on $\mathrm{Orb}_2(A)$ that is symmetric in the two variables. More precisely, if we define an unnormalized two-point weight $W_A^\theta(q)$ by the same multiplicative recipe as (4), but now ranging over all admissible atomic instances in which at least one of x, y appears (including those connecting x to y), then $W_A^\theta(q)$ is invariant under swapping x and y. Moreover, the iterated integral $\iint 1_D(x,y)\,\nu_\theta(dy)\,\nu_\theta(dx)$ reduces, by orbit constancy, to a finite sum of the form

$$\frac{1}{\widetilde{Z}_A^\theta} \sum_{q \in \mathrm{Orb}_2(A)} W_A^\theta(q) \cdot \big[\, O_q^A \subseteq D \,\big],$$

for an appropriate normalization constant $\widetilde{Z}_A^\theta$. Reversing the order of integration yields the same expression, because the summand depends only on the (unordered) 2-type and the set of included orbits of D, and because finite sums commute. Thus ν_θ self-commutes. In practice, this argument can be implemented as a finite case analysis over $\mathrm{Orb}_2(A)$, whose size is controlled by oligomorphicity; for many concrete $\mathcal{K}$, one can compute these orbit partitions explicitly from the one- and two-point extension axioms.

On uniqueness of parameters. The extent to which θ yields a genuine family depends on algebraic constraints in $\mathcal{K}$. In some signatures, the axioms force θ to a unique value. The tournament case is the prototypical example: for distinct x, y, exactly one of $T(x, y)$ and $T(y, x)$ holds, and swapping the roles of x and y forces any constant bias p to satisfy $p = 1 - p$, hence $p = 1/2$. By contrast, for the graph signature (an irreflexive symmetric binary relation), the same scheme yields the Erdős–Rényi family $\theta = \alpha \in [0, 1]$. For higher-arity free amalgamation classes (e.g. k-uniform hypergraphs) one obtains analogous constant-gray families with parameters θ_R for each relation symbol (subject only to the signature-imposed symmetries). In general, any definable dependency among atomic facts involving a fresh point constrains the admissible θ, and uniqueness is exactly the phenomenon that the only self-consistent exchangeable "local" bias is forced by the axioms.

3.13 Limitations and future directions

Beyond free amalgamation. Our explicit construction of ν_θ and the accompanying commutation argument rely on a particularly strong form of local independence: one-point extension data over a finite parameter set A can be chosen essentially atom-by-atom, and restrictions along $B \supseteq A$ factor through a marginalization identity whose constant does not depend on the A-type. This is exactly what free amalgamation provides. Once amalgamation introduces genuine compatibility constraints (forbidden configurations, algebraic closure, definable equivalence relations, order, metric inequalities, *etc.*), the product-form weight (4) can fail to be consistent under support extension, or can become consistent only for a thin set of parameter values (often forcing uniqueness, or degeneracy).

A particularly instructive non-example is the Fraïssé limit $(\mathbb{Q}, <)$ of finite linear orders. Over a finite parameter set $A = \{a_1 < \cdots < a_k\}$, the realized 1-types are the $k + 1$ cuts, hence $\mathrm{Orb}_1(A)$ grows with $|A|$. Any attempt to define ν by assigning a fixed probability to each cut immediately encounters coherence constraints: if we enlarge the support by inserting a new parameter b into an interval, then the mass assigned to that interval must split additively into the masses of the two subintervals. Iterating this refinement along longer and longer finite chains forces the mass of every nonempty open interval to be arbitrarily small, and one quickly runs into tension with normalization $\nu(V) = 1$ when we try to maintain countable additivity on support-bounded families. In other words, the sort of finitely supported, orbit-wise specification we used in free amalgamation classes does not interact well with the order-induced refinement structure. We view this not as a defect of the formalism, but as a signal that new in ordered contexts should not be expected to behave like "independent generic choice" without importing additional analytic structure (for instance, sampling from an ambient continuous distribution and then passing to induced order statistics). Such analytic structure is deliberately absent from the present nominal setting.

More generally, for non-free amalgamation classes one may still hope to obtain self-commuting ν by an orbit/type averaging principle, but the averaging will have to be carried out against the genuine one-point extension counts in $\mathcal{K}$ rather than a simple multiplicative bias. This suggests two concrete research tasks: first, to characterize when the projective system of finite-dimensional distributions arising from a candidate ν exists at all (a "Kolmogorov extension" problem internal to $\mathrm{Nom}(M)$); second, to determine when such a ν can be chosen self-commuting. Both questions appear to depend sensitively on combinatorial properties of $\mathcal{K}$ (e.g. strong amalgamation variants, elimination of algebraic closure, or structural Ramsey features) and on group-theoretic properties of $G = \mathrm{Aut}(M)$ (amenability and its strengthenings). We have not attempted to systematize these dependencies here.

Mixtures and hierarchical randomness. The measures ν_θ obtained in the free amalgamation regime are "extremal" in the sense that the induced finite-model laws are constant-gray with fixed parameters θ. Many natural exchangeable laws are mixtures of such extremals, and in applications one typically wants hierarchical models where θ itself is random and learned from data. Semantically, this asks for an additional source of randomness on a parameter object (e.g. $[0,1]$ or a finite-dimensional cube $[0,1]^m$) together with the ability to *bind* that parameter into the subsequent uses of new.

At the level of our generated monads, the obstruction is clear: $\mathsf{P}_{M,\nu}$ is generated by ν on V and bern on 2, so while we can express finite probabilistic choices and hence randomize among finitely many parameter values, we cannot express genuinely continuous priors on θ. Extending the language with a primitive unif : $1 \to [0,1]$ (or more generally, with a class of sampling primitives on trivial-action objects) would enable such mixtures, but would also force us to confront the status of probability measures on objects with trivial G-action beyond the finite ones (cf. the discussion below on real-valued primitives). A workable compromise is to add *finitely supported* priors (finite mixtures), which are already definable using Bernoulli choices and sums; this suffices for many algorithmic uses (model selection among finitely many regimes), but does not capture de Finetti-style integral representations.

Conditioning and posterior support. Our semantics is purely generative: programs denote kernels built from unit, bind, ν, and deterministic tests. Conditioning—either hard conditioning on events or soft conditioning via likelihood weights—is not primitive. One can attempt to define conditioning internally by

$$\mathrm{cond}_S(\mu)(T) \;:=\; \frac{\mu(T \cap S)}{\mu(S)} \quad \text{when } \mu(S) > 0,$$

for finitely supported $S \in 2^X$, but this is necessarily partial, and the map $(\mu, S) \mapsto \mathrm{cond}_S(\mu)$ is not everywhere defined nor obviously stable under bind in a way that preserves commutativity. Even when defined, conditioning typically increases support: if S is supported by A, then $\mathrm{cond}_S(\mu)$ should be supported by $A \cup \mathrm{supp}(\mu)$, reflecting the familiar fact that observations introduce dependencies on the observed parameters. This is compatible with nominal reasoning, but it means that equational principles such as commutative-let should not be expected to survive unrestricted conditioning, since reordering two conditionings can change intermediate normalizing constants.

A more robust direction is to treat conditioning as a separate effect (for instance, via a "measure transformer" or "scoring" monad) and to relate it to P_M only after normalization. Doing so in a nominal topos raises further issues: the relevant notion of σ-algebra is replaced by 2^X, and countable additivity is only required for support-bounded families, so standard disintegration theorems do not apply verbatim. Establishing a usable theory of Bayesian inversion in $\mathrm{Nom}(M)$ (even restricted to definable observations) remains open in our setting.

Real-valued primitives and the role of trivial-action objects. Although $[0, 1]$ appears in our development, it appears only as a *codomain* for predicates and expectations, equipped with the trivial G-action. This choice is intentional: it allows us to speak about probabilities and integrals without committing to a rich supply of probability measures on $[0, 1]$ itself. Indeed, because the action on $[0, 1]$ is trivial, every subset is finitely supported, so $2^{[0,1]}$ is simply the full powerset. A "probability measure" in the sense of our definition would therefore be a countably additive measure defined on all subsets of $[0, 1]$, which is far stronger than standard Borel measurability and is typically unavailable in ordinary foundations. Consequently, $\mathsf{P}_M([0, 1])$ should be expected to contain very few measures (e.g. Dirac measures), and this makes it unreasonable to add a primitive continuous sampler $[0, 1]$ without further modifying the notion of event.

There are several plausible remedies. One can enrich the base topos so that objects come equipped with a chosen σ-algebra (or an internal analogue) rather than taking 2^X as the full finitely supported powerset; alternatively, one can restrict attention to a designated subobject of "measurable" subsets for trivial-action objects while retaining 2^X for nominal ones. Another possibility is to work with quasi-Borel or synthetic-measure-theoretic structures internal to $\mathrm{Nom}(M)$, so that $[0,1]$ carries a canonical measurable structure not equal to the full powerset. Each of these approaches amounts to adding analytic structure orthogonal to the nominal symmetry, and we have not pursued it here. The present framework is therefore best viewed as a semantics for probabilistic computation whose observable randomness is ultimately *finitary* (via numeral types and Bernoulli), even though the hidden nominal state V is infinite.

Algorithmic and machine-checked orbit/type calculations. Even in the free amalgamation regime, verifying self-commutation by hand can become tedious as the signature grows, because one must reason about orbit partitions $\mathrm{Orb}_n(A)$ and the effect of bind on finitely supported definable sets. At the same time, our reduction lemmas make these verifications intrinsically finite: over any fixed finite support A, there are finitely many relevant n-orbits, and integrals reduce to finite sums of orbit indicators. This finiteness strongly suggests computer assistance.

A concrete goal is a toolchain which, given a relational signature Σ and a finite axiom scheme describing $\mathcal{K}$ (e.g. free amalgamation together with irreflexivity/symmetry conventions and finitely many forbidden patterns), produces (i) an explicit description of $\mathrm{Orb}_1(A)$ and $\mathrm{Orb}_2(A)$ up to isomorphism over A, (ii) the corresponding weight expressions w_A^θ and the marginalization constants $C_{A,B}^\theta$, and (iii) a mechanically checked proof that the induced ν_θ is well-defined and self-commuting. For many classes, orbit representatives can be encoded as finite relational tables with marked parameters, and consistency can be discharged by SAT/SMT solving. One can then export the finite orbit-sum computations as proof certificates checked in a proof assistant (Lean/Coq/Isabelle), thereby separating the combinatorial enumeration from the semantic argument.

We emphasize that such mechanization is not merely an implementation detail: it would provide a systematic way to explore the boundary between classes where commutative `new`-semantics exists and classes where it fails, and it would enable rapid prototyping of new interfaces (higher-arity relations, multiple interacting sorts of atoms, or signatures with constrained symmetries). Ultimately, one would like to turn the abstract statement "commutation reduces to finitely many orbit checks" into a practical method that either constructs ν and proves commutation, or produces a minimal counterexample orbit configuration witnessing failure.

Summary. Our results are strongest when $\mathcal{K}$ presents genuinely local extension freedom, in which case ν_θ and commutation follow a uniform orbit-averaging pattern. Extending beyond this regime raises substantive questions: whether suitable ν exist, how to support mixtures and conditioning without sacrificing commutative-let reasoning, how to incorporate real-valued sampling without collapsing into set-theoretic pathologies on trivial-action objects, and how to mechanize the finite orbit/type arguments that underlie both definability and commutation. We expect that progress on these questions will require a tighter integration of Fraïssé-theoretic combinatorics, topological dynamics of G, and synthetic or internalized measure theory within nominal toposes.