Aldous–Hoover by Program Equations: k-Uniform Hypergraphs, Markov Categories, and Hypergraphons

Liz Lemma Future Detective

January 16, 2026

Abstract

We extend the semantic correspondence between probabilistic programming interfaces for random graphs and graphons to higher-arity relational data. Fix $k \geq 2$ and consider a minimal probabilistic programming interface providing a type of vertices, a fresh-vertex sampler new, and a deterministic symmetric predicate $\mathsf{hyperedge}_k$ deciding k-hyperedges. For any Bernoulli-based equational theory (modeled in a distributive Markov category with an observation map into finite stochastic matrices), we build programs that output finite k-uniform hypergraphs as Boolean incidence tensors. We prove that the resulting sequence of finite random hypergraphs is exchangeable, projective, and dissociated (induced substructures on disjoint vertex sets are independent). By the dissociated Aldous–Hoover/Kallenberg representation theorem, every such theory therefore determines a k-uniform hypergraphon. Conversely, we construct a universal category from finite k-uniform hypergraphs via coproduct completion and a monoidal-indeterminate (para) construction adjoining new; Bernoulli-based quotients of this universal model are in bijection with hypergraphons. This work generalizes the graphon-by-equations phenomenon to hypergraphs and exchangeable arrays, providing a semantics-first route to designing modern probabilistic APIs for multi-relational and tensor-valued data.

Table of Contents

1. 1. Introduction and motivation: exchangeable interfaces for hypergraphs/arrays; summary of contributions; relationship to graphon correspondence and Aldous–Hoover.

2. 2. Language and axioms: typed core with products/sums/let; interface ($\mathsf{vertex}, \mathsf{new}, \mathsf{hyperedge}_k$); determinism and symmetry/irreflexivity equations; Bernoulli base as observation semantics.

3. 3. Background on k-uniform hypergraphs and hypergraphons: finite incidence tensors; exchangeability and projectivity; dissociation/k-locality; statement of the dissociated Aldous–Hoover/Kallenberg representation theorem in the form needed.

4. 4. From program equations to random hypergraphs: define the incidence-tensor program $t_{n,k}$; prove it yields valid k-uniform hypergraphs; establish exchangeability, projectivity, and k-locality from the let-laws and determinism.

5. 5. From random hypergraphs to hypergraphons: apply representation theorem; discuss equivalence of hypergraphons (measure-preserving transformations) and what is identifiable from the program theory.

6. 6. A universal semantics for the interface: define the category Hyp_k of finite k-uniform hypergraphs; build $\mathsf{Fam}(\mathsf{Hyp}_k^{\mathrm{op}})$; adjoin new via a monoidal-indeterminate construction to form $\mathcal{U}_k$; define the universal interpretation of $\mathsf{hyperedge}_k$.

7. 7. Hypergraphons as Bernoulli-based quotients: characterize distributive Markov functors $(\mathcal{U}_k)_{\mathbb{N}} \to \mathsf{FinStoch}$ as exactly exchangeable/projective/k-local models; then use a contextual-equivalence quotient construction to obtain a Bernoulli-based theory realizing any given hypergraphon.

8. 8. Examples and variants: $k = 2$ recovers graphons; $k = 3$ and higher; discussion of directed/colored hyperedges; array interfaces as bipartite/higher-sorted variants.

9. 9. Discussion and future directions: sparse hypergraphs, approximate/metric observation, computability, combination with other effects (real-valued primitives), and prospects for mechanization.

1 1. Introduction and motivation: exchangeable interfaces for hypergraphs/arrays; summary of contributions; relationship to graphon correspondence and Aldous–Hoover.

We study a simple but rigid phenomenon: once a probabilistic program is allowed to allocate *fresh* abstract names (vertices) and to query a deterministic, symmetric, irreflexive k-ary relation on those names, the distributions produced by such programs are forced into the well-known exchangeable world of hypergraph limits. Our aim is to make this constraint precise in a way that is simultaneously (i) syntactic, by an equational theory of programs, (ii) semantic, by interpretation in a suitable probabilistic categorical model, and (iii) structural, by identifying the induced laws with the standard representation theory of exchangeable arrays.

The guiding example is the graphon correspondence for $k = 2$. There, one may view a graphon W as a limit object encoding all finite-dimensional distributions of an exchangeable random graph, and conversely every exchangeable, projective, and dissociated random graph law is generated by some W (uniquely up to measure-preserving rearrangement). The customary formulation starts from random adjacency matrices and invariance under relabeling; our point of entry is instead an *interface* in a probabilistic programming language: a constructor for new vertices and a predicate testing whether an edge is present. The fact that a program never receives vertex labels as data, but only obtains them by repeated allocation, is precisely what enforces exchangeability. The fact that the predicate is deterministic, combined with the usual laws of sequencing, is what enforces dissociation/locality.

For hypergraphs the same conceptual picture holds, but the correct limit object is more intricate. A k-uniform hypergraph is a $\{0, 1\}$-valued symmetric k-array with the additional constraint forbidding repeated indices, and the appropriate exchangeability notion is invariance under finite permutations of the vertex set. When $k \geq 3$, the Aldous–Hoover/Kallenberg representation theorem describes general exchangeable k-arrays via a hierarchy of latent random variables indexed by all nonempty proper subsets of $[k]$. In particular, one does not in general obtain a model from a naive kernel $[0, 1]^k \to [0, 1]$ depending only on vertex-level randomness; rather, one requires dependence on the full collection of lower-dimensional coordinates (e.g. pair-coordinates for $k = 3$), and this dependence is exactly what the hypergraphon formalism records.

Our contribution is to show that this representation theory is already latent in a minimal programming interface for k-uniform hypergraphs. On the one hand, from any interpretation of the interface in a probabilistic semantics satisfying standard "let" equations, we obtain for each n a canonical closed program producing the entire incidence tensor of a random k-uniform hypergraph on n sampled vertices. This yields a sequence of distributions $(p_n)_{n\geq 1}$ on finite hypergraphs. We prove that these distributions are necessarily exchangeable, consistent under restriction to subsets of vertices, and k-local in the sense that induced sub-hypergraphs on disjoint vertex blocks are independent. The argument is not probabilistically deep; rather, it is an exercise in extracting probabilistic invariances from program equations. Commutativity of independent "let"-bindings corresponds to relabeling invariance; weakening corresponds to marginalization; and deterministic substitution for relation tests, together with the monoidal structure implicit in pairing, yields dissociation.

On the other hand, once these three properties are established, the classical representation theorem applies: every exchangeable, projective, dissociated k-uniform hypergraph model arises from a hypergraphon W, and the hypergraphon is unique up to the standard measure-preserving equivalence. Thus our syntactic and categorical assumptions isolate exactly the hypotheses required by the dissociated Aldous–Hoover/Kallenberg theorem, and we may pass from program denotations to hypergraphons without additional modeling choices. In this sense, hypergraphons appear not as an externally imposed analytical gadget, but as the unavoidable completion of the equational theory governing name allocation and deterministic relational queries.

A key methodological choice is to phrase semantics in the language of *distributive Markov categories*. This framework separates (i) the deterministic structural fragment (products, coproducts, and their distributive interaction) from (ii) the probabilistic fragment (Markov kernels and sequencing), and it packages the equational reasoning about programs into categorical equalities. However, to compare program denotations with classical probability distributions we require an observation map on a suitable "numeral" fragment. We therefore assume a *Bernoulli base*, i.e. a faithful functor into $\mathsf{FinStoch}$ on finite objects, ensuring that closed programs of finite type determine honest finite probability distributions. This step is conceptually small but technically clarifying: it lets us state results directly in terms of equality of distributions, while keeping the ambient semantics abstract enough to accommodate different models.

Besides extracting hypergraphons from program theories, we also pursue a converse direction: every hypergraphon should be realizable as the semantics of some Bernoulli-based equational theory for the interface. To do so we construct a universal semantic object, a distributive Markov category freely generated by deterministic k-uniform hypergraph structure together with a single probabilistic generator representing "fresh vertex creation". This universal category plays the role of a syntax-free initial model for the interface. By analyzing distributive Markov functors from its numeral fragment into $\mathsf{FinStoch}$, we obtain a classification in terms of exchangeable, projective, and k-local hypergraph models; composing with the representation theorem yields the promised correspondence with hypergraphons. In particular, given a concrete hypergraphon W we build a functor encoding its finite-dimensional sampling scheme and then take a suitable quotient of the universal category to obtain a semantic model whose induced (p_n) agrees with the standard $p_{W,n}$ for all n.

Several points deserve emphasis. First, the results are not restricted to i.i.d. edge models; indeed, the hypergraphon formalism permits rich dependencies among hyperedges, mediated by the shared latent variables indexed by lower-dimensional subsets. From the programming perspective, this is precisely what one expects: a program may allocate vertices and then make correlated relational decisions through shared hidden state in the semantic model, even if individual relation queries are deterministic. Second, the constraint of determinism for $\mathsf{hyperedge}_k$ is essential for dissociation: if relation tests themselves were randomized, then additional sources of randomness indexed by k-tuples would appear, and one would recover the full (not necessarily dissociated) Aldous–Hoover form. Our setting thus isolates the "randomness comes from names, not from predicates" regime and identifies it with the dissociated part of the theory.

Third, although we phrase the main development for a single-sorted vertex type and a single k-ary predicate, the organizing idea is more general: one may treat probabilistic interfaces as generators of exchangeable structures whenever fresh-name creation is the only way to obtain indistinguishable individuals. This viewpoint connects directly with standard probabilistic symmetries (exchangeability, separate exchangeability, partial exchangeability) and suggests a systematic way to read off representation-theoretic consequences from program equations. Hypergraphs are a natural test case because they already exhibit the higher-order latent-variable hierarchy, making clear why the correct limit object is genuinely higher-dimensional when $k \geq 3$.

Finally, we regard this work as part of a broader program: to relate operationally meaningful, equationally specified probabilistic programming interfaces to the invariant objects studied in probability, ergodic theory, and combinatorial limits. In the present case, the bridge runs from program equations, through categorical semantics with a faithful finite observation principle, to dissociated exchangeable hypergraph laws, and thence to hypergraphons. The resulting correspondence recovers the graphon story at $k = 2$, while providing a uniform account for all $k \geq 2$ that is compatible with the classical Aldous–Hoover/Kallenberg framework for $\{0, 1\}$-valued symmetric arrays.

2 Language, interface, and Bernoulli-based observation

We work in a typed, call-by-value core calculus with finite products, finite coproducts, and an explicit sequencing construct. Types are generated from a terminal type $\mathbf{1}$, binary products $A \times B$, binary coproducts $A + B$, and a distinguished boolean type

$$\mathsf{bool} \;:=\; \mathbf{1} + \mathbf{1},$$

together with a type constant vertex. We use standard derived notations: A^n for the n-fold product, and we write $\langle t, u \rangle$ for pairing, π_i for projections, inl and inr for injections, and case for coproduct elimination. Contexts Γ are finite lists of typed variables and are interpreted as products of types; a judgment $\Gamma \vdash t : A$ is read as a program which, given an input tuple of values in Γ, produces an output of type A.

The probabilistic structure is expressed by the *let*-construct. Formally, if $\Gamma \vdash t : A$ and $\Gamma, x{:}A \vdash u : B$, then $\Gamma \vdash \mathsf{let}\ x = t\ \mathsf{in}\ u : B$. Operationally (in call-by-value style), we evaluate t to a random value a and then continue as u with x bound to a. This construct is the only source of sequencing; in particular, we do not assume a primitive monadic syntax, but we treat let as the syntactic reflection of Kleisli composition in a probabilistic semantics.

The hypergraph interface consists of two term constants:

$$\mathsf{new} : \mathbf{1} \to \mathsf{vertex}, \qquad \mathsf{hyperedge}_k : (\mathsf{vertex})^k \to \mathsf{bool},$$

where $k \geq 2$ is fixed once and for all. We emphasize that vertex is abstract: programs cannot inspect vertices except by feeding them into $\mathsf{hyperedge}_k$ (and by duplicating or discarding them using product structure). The intention is that new allocates a fresh vertex and $\mathsf{hyperedge}_k$ queries the presence of a k-uniform hyperedge among its arguments.

We reason about programs up to an equational theory $\equiv$ extending the usual β/η-equalities for products and coproducts with a collection of *probabilistic let-laws.* We shall only use the following schemes, each understood as an equation in any context where the terms are well-typed, with the usual side conditions to avoid variable capture. First, *associativity of sequencing*:

$$\mathsf{let}\ x = t\ \mathsf{in}\ (\mathsf{let}\ y = u\ \mathsf{in}\ v) \;\equiv\; \mathsf{let}\ y = (\mathsf{let}\ x = t\ \mathsf{in}\ u)\ \mathsf{in}\ v.$$

Second, *commutativity of independent lets*: if $x \notin \mathrm{FV}(u)$ and $y \notin \mathrm{FV}(t)$ then

$$\mathsf{let}\ x = t\ \mathsf{in}\ (\mathsf{let}\ y = u\ \mathsf{in}\ v) \;\equiv\; \mathsf{let}\ y = u\ \mathsf{in}\ (\mathsf{let}\ x = t\ \mathsf{in}\ v).$$

Third, *pairing/strength* for products, expressing that sequencing distributes over forming tuples:

$$\mathsf{let}\ x = t\ \mathsf{in}\ \langle u, v \rangle \;\equiv\; \Big\langle \mathsf{let}\ x = t\ \mathsf{in}\ u,\ \mathsf{let}\ x = t\ \mathsf{in}\ v \Big\rangle.$$

Fourth, *weakening* (discard): if $x \notin \mathrm{FV}(u)$ then

$$\mathsf{let}\ x = t\ \mathsf{in}\ u \quad \equiv \quad u.$$

Finally, we require a *substitution principle for deterministic terms.* Concretely, we distinguish a syntactic class of deterministic terms (those built without let and without probabilistic primitives, hence interpreting as deterministic maps in the semantics below). For deterministic d we impose that sequencing into d is ordinary substitution:

$$\mathsf{let}\ x = t\ \mathsf{in}\ d \quad \equiv \quad d[t/x].$$

In particular, deterministic coproduct elimination and deterministic product operations commute with probabilistic binding in the expected way. These laws are the sole source of probabilistic reasoning we assume at the syntactic level; all subsequent invariance and independence properties are extracted from repeated use of these equalities.

The interface is constrained by equations enforcing that $\mathsf{hyperedge}_k$ behaves as the adjacency predicate of a simple undirected k-uniform hypergraph. First, *irreflexivity on repeated vertices*: in any context Γ and for any k-tuple of variables or terms $(x_1, \dots, x_k)$ in which some x_i and x_j are definitionally equal (or, more generally, provably equal in $\equiv$), we impose

$$\Gamma \vdash \mathsf{hyperedge}_k(x_1, \dots, x_k) \quad \equiv \quad \mathsf{false} : \mathsf{bool}.$$

Second, *symmetry*: for each permutation $\pi \in S_k$ and any $\Gamma \vdash (x_1, \dots, x_k) : (\mathsf{vertex})^k$ we impose

$$\Gamma \vdash \mathsf{hyperedge}_k(x_1, \dots, x_k) \quad \equiv \quad \mathsf{hyperedge}_k(x_{\pi(1)}, \dots, x_{\pi(k)}) : \mathsf{bool}.$$

These equations ensure that hyperedges are sets rather than ordered tuples and that loops are forbidden. Third, and crucially for dissociation, we require *determinism* of $\mathsf{hyperedge}_k$: it is treated as a deterministic constant, so that the deterministic substitution law above applies when d contains occurrences of $\mathsf{hyperedge}_k$. Intuitively, the only randomness in the interface comes from allocating vertices; relational queries are measurements of the underlying (possibly random) structure but are not themselves randomized.

The intended semantics is given in a *distributive Markov category* $\mathcal{C}$. We assume the reader is familiar with the basic structure: morphisms are Markov kernels; deterministic morphisms form a wide subcategory closed under products; every object carries copy and discard maps, and $\mathcal{C}$ has finite coproducts distributing over products. An interpretation assigns to each type A an object $[\![A]\!]$ of $\mathcal{C}$, with $[\![A \times B]\!] \cong [\![A]\!] \otimes [\![B]\!]$ and $[\![A + B]\!]$ a coproduct, and to each term $\Gamma \vdash t : A$ a morphism

$$[\![t]\!] : [\![\Gamma]\!] \longrightarrow [\![A]\!].$$

The defining clause is for let: if $\Gamma \vdash t : A$ and $\Gamma, x{:}A \vdash u : B$, then $[\![\mathsf{let}\ x = t\ \mathsf{in}\ u]\!]$ is the Kleisli composite obtained by first applying $[\![t]\!]$ and then $[\![u]\!]$, using the monoidal structure to thread the ambient context through. The stated let-laws are precisely the equalities valid in any such semantics; conversely, we may view them as axiomatizing the fragment of Markov categorical reasoning we will need.

The interface constants are interpreted by chosen morphisms

$$[\![\mathsf{new}]\!] : \mathbf{1} \to [\![\mathsf{vertex}]\!], \qquad [\![\mathsf{hyperedge}_k]\!] : [\![\mathsf{vertex}]\!]^{\otimes k} \to [\![\mathsf{bool}]\!],$$

with $[\![\mathsf{hyperedge}_k]\!]$ required to lie in the deterministic subcategory. The irreflexivity and symmetry axioms are imposed as equalities between the corresponding morphisms, using the symmetry isomorphisms of the monoidal product to interpret permutations of arguments.

To connect denotations with ordinary finite probability distributions, we assume a *Bernoulli base* $\Psi : \mathcal{C}_{\mathbb{N}} \hookrightarrow \mathsf{FinStoch}$ on the numeral fragment. Here $\mathcal{C}_{\mathbb{N}}$ denotes the full subcategory of $\mathcal{C}$ generated by finite coproducts of $\mathbf{1}$ (equivalently, by the usual finite objects $n := \mathbf{1} + \cdots + \mathbf{1}$), and $\mathsf{FinStoch}$ is the category of finite sets of cardinality n and stochastic matrices between them. The functor Ψ is assumed faithful and distributive Markov, so it preserves the probabilistic and distributive structure and reflects equality of morphisms on finite objects. In particular, if $t : \mathbf{1} \to \mathsf{bool}^m$ is closed, then $\Psi([\![t]\!])$ is an honest probability distribution on $\{0,1\}^m$, and two such programs have the same observed distribution if and only if their denotations are equal in $\mathcal{C}$.

We will apply this observation principle to the canonical closed programs producing incidence tensors. Fix $n \geq 1$ and choose once and for all an enumeration of $\binom{[n]}{k}$; this identifies $\mathsf{bool}^{\binom{n}{k}}$ with a finite product bool^m where $m = \binom{n}{k}$. We define $t_{n,k} : \mathbf{1} \to \mathsf{bool}^{\binom{n}{k}}$ by sampling $x_1, \ldots, x_n$ using new and returning the tuple of booleans $\big(\mathsf{hyperedge}_k(x_{i_1}, \ldots, x_{i_k})\big)_{\{i_1 < \cdots < i_k\} \in \binom{[n]}{k}}$. Applying Ψ to $[\![t_{n,k}]\!]$ yields a distribution p_n on $\{0,1\}^{\binom{[n]}{k}}$, which we identify with a distribution on labeled k-uniform hypergraphs on vertex set $[n]$ via the usual incidence-function encoding.

The remainder of the development uses only the equational principles listed above, together with the interface axioms, to establish structural properties of the family $(p_n)_{n \geq 1}$ and to connect those properties to the standard representation theory of exchangeable k-uniform hypergraphs.

3 Background on k-uniform hypergraphs and hypergraphons

We briefly recall the combinatorial and probabilistic objects that will mediate between programs and limit representations. Throughout, $k \geq 2$ is fixed. For

$n \in \mathbb{N}$ we write $[n] = \{1, \ldots, n\}$ and $\binom{[n]}{k}$ for the set of k-element subsets of $[n]$.

Finite k-uniform hypergraphs and incidence tensors. A (simple, undirected) k-uniform hypergraph on vertex set $[n]$ is a pair $H = ([n], E)$ with $E \subseteq \binom{[n]}{k}$. Equivalently, it is an incidence function

$$A_H : \binom{[n]}{k} \longrightarrow \{0, 1\}, \qquad A_H(e) = \mathbf{1}_{\{e \in E\}}.$$

We will freely identify H with A_H. Fixing an enumeration $\binom{[n]}{k} \cong [m]$ with $m = \binom{n}{k}$ identifies the space of incidence functions with $\{0, 1\}^m$. This is the finite space on which our observed distributions p_n will live.

It is often convenient to pass between the set-indexed representation $A_H(e)$, $e \in \binom{[n]}{k}$, and an ordered k-array representation. Given an incidence function $A : \binom{[n]}{k} \to \{0, 1\}$, define a k-array $(X_{i_1,\ldots,i_k})_{(i_1,\ldots,i_k) \in [n]^k}$ by

$$X_{i_1,\ldots,i_k} := \begin{cases} A(\{i_1, \ldots, i_k\}) & \text{if } i_1, \ldots, i_k \text{ are pairwise distinct,} \\ 0 & \text{otherwise.} \end{cases}$$

Then $(X_{i_1,\ldots,i_k})$ is symmetric under permutations of coordinates and vanishes on the diagonal (repeated indices), and conversely any such symmetric, diagonal-free $\{0, 1\}$-valued k-array corresponds to a unique incidence function on $\binom{[n]}{k}$. We will use whichever presentation is more convenient, but our primary viewpoint is the incidence function on k-subsets.

Random k-uniform hypergraph models; exchangeability and projectivity. A random k-uniform hypergraph on $[n]$ is a random variable taking values in $\{0, 1\}^{\binom{[n]}{k}}$, i.e. a probability measure p_n on that finite space. A random k-uniform hypergraph model is a family $(p_n)_{n \geq 1}$ of such measures, one for each n. In later sections p_n will arise as the observed output distribution of a closed term producing the incidence tensor on n sampled vertices.

Two structural properties will be central.

First, *exchangeability* expresses invariance under relabeling of the sampled vertices. If $\sigma : [n] \to [n]$ is a bijection, it induces a bijection $\binom{[n]}{k} \to \binom{[n]}{k}$ by $e \mapsto \sigma(e) := \{\sigma(i) \mid i \in e\}$. For an incidence function A define $(\sigma \cdot A)(e) := A(\sigma^{-1}(e))$, and for a measure p_n define its pushforward $\sigma_* p_n$ along $A \mapsto \sigma \cdot A$. We say that (p_n) is exchangeable if $\sigma_* p_n = p_n$ for all n and all bijections σ.

Second, *projectivity* (or consistency) expresses compatibility under restriction to fewer vertices. Let $\iota : [n] \hookrightarrow [n + 1]$ be the standard inclusion.

Restriction of incidence functions along $\binom{[n]}{k} \subseteq \binom{[n+1]}{k}$ gives a marginal map

$$\mathrm{res}_{n+1\to n} : \{0,1\}^{\binom{[n+1]}{k}} \longrightarrow \{0,1\}^{\binom{[n]}{k}}, \qquad A \longmapsto A|_{\binom{[n]}{k}}.$$

We say that (p_n) is projective if $(\mathrm{res}_{n+1\to n})_* p_{n+1} = p_n$ for all n.

Exchangeability and projectivity together say that the family (p_n) behaves like the finite-dimensional marginals of an infinite exchangeable random hypergraph on vertex set $\mathbb{N}$. Indeed, by Kolmogorov extension one can often package such a family into a probability measure on $\{0,1\}^{\binom{\mathbb{N}}{k}}$; we will not rely on that construction explicitly, but it provides intuition for the representation theorem stated below.

Dissociation and k-locality. Beyond exchangeability and projectivity, we require an independence property. Let $A \subseteq [n]$ and write $\binom{A}{k}$ for the k-subsets of A. Restricting incidence functions defines a random induced sub-hypergraph on A (as a random variable valued in $\{0,1\}^{\binom{A}{k}}$). If $A, B \subseteq [n]$ are disjoint, we can consider the pair of induced sub-hypergraphs

$$\Big(A_H|_{\binom{A}{k}},\ A_H|_{\binom{B}{k}} \Big) \in \{0,1\}^{\binom{A}{k}} \times \{0,1\}^{\binom{B}{k}}.$$

We say that a model (p_n) is *dissociated* (or *k-local* in the terminology we will use for program reasoning) if for every n and all disjoint $A, B \subseteq [n]$ the induced sub-hypergraphs on A and on B are independent under p_n. This is strictly stronger than mere conditional independence given latent vertex labels; it asserts unconditional independence between disjoint induced substructures. In exchangeability theory this property is called dissociation and is precisely the hypothesis under which the Aldous–Hoover representation simplifies to a single "hypergraphon" function without an additional global random parameter.

Hypergraphons in dissociated array form. For $k = 2$, a graphon is a measurable symmetric function $W : [0,1]^2 \to [0,1]$ generating exchangeable random graphs by sampling i.i.d. latent labels $U_i \sim \mathrm{Unif}[0,1]$ and then including each edge $\{i,j\}$ with probability $W(U_i, U_j)$ independently over pairs conditional on the labels. For $k \geq 3$, this naive vertex-only parameterization is insufficient: the most general dissociated exchangeable k-uniform hypergraph requires auxiliary randomness attached to lower-dimensional faces.

A *k-uniform hypergraphon* (in dissociated Aldous–Hoover form) is a measurable function

$$W : [0,1]^{2^k - 2} \longrightarrow [0,1]$$

whose coordinates are indexed by the nonempty proper subsets $S \subsetneq [k]$, and which is invariant under the natural action of S_k permuting indices: for each permutation $\pi \in S_k$ we require

$$W\big((u_S)_{\emptyset \neq S \subsetneq [k]}\big) \;=\; W\big((u_{\pi(S)})_{\emptyset \neq S \subsetneq [k]}\big) \quad \text{for a.e. } (u_S).$$

Given such a W, we define for each n a probability distribution $p_{W,n}$ on $\{0,1\}^{\binom{[n]}{k}}$ by the following sampling scheme.

1. For every nonempty $J \subseteq [n]$ with $|J| < k$, sample an independent uniform random variable $U_J \sim \mathrm{Unif}[0,1]$.

2. For each hyperedge $e \in \binom{[n]}{k}$, form the tuple of lower-face labels

$$\big(U_J\big)_{\emptyset \neq J \subsetneq e} \in [0,1]^{2^k-2},$$

   using the identification $e \cong [k]$ up to permutation; the S_k-invariance of W ensures the resulting quantity is well-defined almost surely (i.e. does not depend on how we order the elements of e).

3. Conditional on the entire family $(U_J)_{|J|<k}$, include e as a hyperedge with probability $W((U_J)_{\emptyset \neq J \subsetneq e})$, independently across distinct $e \in \binom{[n]}{k}$.

The resulting random incidence function has distribution $p_{W,n}$. By construction, $(p_{W,n})_{n \geq 1}$ is exchangeable and projective, and it is dissociated because disjoint vertex sets depend on disjoint families of latent variables (U_J).

Two hypergraphons can induce the same model. The appropriate notion of equivalence is the usual measure-preserving one: if one applies measure-preserving transformations to the underlying probability space $[0,1]$ in a way compatible with the face-indexing (equivalently, replaces the i.i.d. family (U_J) by another i.i.d. family with the same law), the induced distributions $p_{W,n}$ do not change. We will therefore treat W as determined only up to this standard a.e. equivalence.

Representation theorem (dissociated Aldous–Hoover/Kallenberg). We now record the form of the representation theorem that we will use. There are several equivalent formulations in the literature (infinite arrays indexed by $\mathbb{N}$, projective families of finite marginals, or random measures); we state it directly at the level of the family (p_n).

Theorem 3.1 (Dissociated exchangeable representation for k-uniform hypergraphs)**.** *Let $(p_n)_{n\geq 1}$ be a family of probability measures on $\{0,1\}^{\binom{[n]}{k}}$ that is exchangeable and projective. Assume moreover that the family is dissociated: induced sub-hypergraphs on disjoint vertex sets are independent. Then there exists a k-uniform hypergraphon $W : [0,1]^{2^k-2} \to [0,1]$, invariant under the S_k-action, such that*

$$p_n \;=\; p_{W,n} \qquad \text{for all } n \geq 1.$$

The representing hypergraphon is unique up to the standard measure-preserving equivalence (i.e. equality a.e. after a suitable measure-preserving reparameterization of the underlying i.i.d. family (U_J)).

For our purposes, the content of Theorem 3.1 is twofold. First, it identifies dissociated exchangeable random k-uniform hypergraphs with the hypergraphon sampling scheme above; this will justify viewing a program-induced family (p_n) as determined by an abstract measurable object W. Second, it provides the correct ambient notion of "limit object" for k-ary relation models: for $k \geq 3$ the domain dimension $2^k - 2$ is forced by general exchangeability considerations, and restricting to a naive $[0,1]^k$ kernel would exclude valid dissociated models. In the next section we will show that the program equations enforce exactly the hypotheses of Theorem 3.1 for the family induced by the canonical incidence-tensor program.

4 From program equations to random hypergraphs

Fix $n \geq 1$. Our first task is to define, purely syntactically in Lang_k, the closed incidence-tensor term

$$t_{n,k} : \mathbf{1} \longrightarrow \mathsf{bool}^{\binom{n}{k}}$$

which samples n fresh vertices and then queries the deterministic predicate $\mathsf{hyperedge}_k$ on every k-subset of the sampled vertices. The precise arrangement of products needed to realize the exponent $\mathsf{bool}^{\binom{n}{k}}$ is immaterial, since we work in a cartesian fragment and hence may choose any fixed enumeration $\binom{[n]}{k} \cong [m]$ with $m = \binom{n}{k}$; different enumerations yield canonically isomorphic types and, under Ψ, canonically isomorphic finite sample spaces. We therefore tacitly fix such enumerations once and for all.

The incidence-tensor program. Let $\vec{x} = (x_1, \ldots, x_n)$ be n variables of type vertex. For each $e = \{i_1 < \cdots < i_k\} \in \binom{[n]}{k}$ write

$$h_e(\vec{x}) \;:=\; \mathsf{hyperedge}_k(x_{i_1}, \ldots, x_{i_k}) : \mathsf{bool}.$$

Using pairing and the chosen enumeration $\binom{[n]}{k} \cong [m]$, we may form the bool^m-valued term

$$\mathrm{inc}_{n,k}(\vec{x}) \;:=\; \big(h_e(\vec{x})\big)_{e \in \binom{[n]}{k}} : \mathsf{bool}^{\binom{n}{k}}.$$

The closed term $t_{n,k}$ is then defined by sampling vertices sequentially and returning $\mathrm{inc}_{n,k}$:

$$t_{n,k} \;:=\; \mathsf{let}\; x_1 = \mathsf{new}(\star)\;\mathsf{in}\;\mathsf{let}\; x_2 = \mathsf{new}(\star)\;\mathsf{in}\;\cdots\;\mathsf{let}\; x_n = \mathsf{new}(\star)\;\mathsf{in}\;\mathrm{inc}_{n,k}(x_1, \ldots, x_n),$$

where $\star : \mathbf{1}$ denotes the unique term of unit type. This term is well-typed by construction and uses only the core fragment together with the two interface constants.

Observed distributions. Interpreting Lang_k in a distributive Markov category $\mathcal{C}$, we obtain a morphism

$$[\![t_{n,k}]\!]_{\mathcal{C}} : \mathbf{1} \longrightarrow [\![\mathsf{bool}^{\binom{n}{k}}]\!]_{\mathcal{C}}.$$

By the Bernoulli base, the object $[\![\mathsf{bool}^{\binom{n}{k}}]\!]_{\mathcal{C}}$ lies in the numeral fragment and hence is observed as a finite set of cardinality 2^m with $m = \binom{n}{k}$. Concretely, writing $\Psi([\![t_{n,k}]\!]_{\mathcal{C}})$ for the induced stochastic map in $\mathsf{FinStoch}$, we obtain a probability distribution

$$p_n \;\in\; \Delta\big(\{0,1\}^{\binom{[n]}{k}}\big)$$

by identifying $\mathsf{bool}^{\binom{n}{k}}$ with $\{0,1\}^{\binom{[n]}{k}}$ via the chosen enumeration. All structural properties we prove below are equalities in $\mathsf{FinStoch}$ transported back through Ψ; faithfulness of Ψ on numerals ensures that program equations suffice to establish the corresponding probabilistic equalities.

Well-formedness: support on simple k-uniform hypergraphs. The first point is that, although $t_{n,k}$ returns an m-tuple of booleans, it is forced by the interface axioms to behave as the incidence function of a *simple, undirected* k-uniform hypergraph.

Indeed, consider the k-array presentation induced from the tuple by setting

$$X_{i_1,\ldots,i_k} \;:=\; \mathsf{hyperedge}_k(x_{i_1},\ldots,x_{i_k}) \qquad ((i_1,\ldots,i_k) \in [n]^k).$$

If $(i_1,\ldots,i_k)$ has a repeated index, then by irreflexivity we have the program equation

$$\mathsf{hyperedge}_k(x_{i_1},\ldots,x_{i_k}) \;\equiv\; \mathsf{false}.$$

If $\pi \in S_k$ is a permutation, then by symmetry we have

$$\mathsf{hyperedge}_k(x_{i_1},\ldots,x_{i_k}) \;\equiv\; \mathsf{hyperedge}_k(x_{i_{\pi(1)}},\ldots,x_{i_{\pi(k)}}).$$

Since these are equations in the program theory, they hold under interpretation in $\mathcal{C}$, and hence under Ψ they hold almost surely with respect to p_n. It follows that p_n is supported on those boolean tensors which are diagonal-free and symmetric, equivalently on incidence functions $A : \binom{[n]}{k} \to \{0,1\}$ of simple k-uniform hypergraphs.

Exchangeability from commutative let. Let $\sigma \in S_n$ be a permutation of $[n]$. Consider the term obtained by sampling vertices in the permuted order:

$$t^{\sigma}_{n,k} \;:=\; \mathsf{let}\ y_1 = \mathsf{new}(\star)\ \mathsf{in}\ \cdots\ \mathsf{let}\ y_n = \mathsf{new}(\star)\ \mathsf{in}\ \mathrm{inc}_{n,k}\big(y_{\sigma(1)},\ldots,y_{\sigma(n)}\big).$$

By repeated use of the commutativity-of-independent-let axiom (together with associativity/pairing laws to justify regrouping), we may reorder the bindings of independent samples without changing denotation. In particular, the term that binds the variables in the order $y_1, \dots, y_n$ is equivalent to the term that binds them in the order $y_{\sigma^{-1}(1)}, \dots, y_{\sigma^{-1}(n)}$. Renaming bound variables, we obtain a program equation identifying $t_{n,k}$ with $t^{\sigma}_{n,k}$ up to the evident reindexing of outputs.

To read this statement on the observed distributions, note that σ acts on $\binom{[n]}{k}$ by $e \mapsto \sigma(e)$, and hence acts on incidence functions by $(\sigma \cdot A)(e) = A(\sigma^{-1}(e))$. The preceding program equality implies that $\Psi([\![t_{n,k}]\!]_{\mathcal{C}})$ is invariant under this output relabeling, i.e.

$$\sigma_* p_n \;=\; p_n.$$

Thus (p_n) is exchangeable, and the only nontrivial content is precisely that commutativity of let provides the syntactic witness to invariance under vertex renaming.

Projectivity from weakening (discard). Write $m_n = \binom{n}{k}$. The term $t_{n+1,k}$ returns a boolean tuple indexed by $\binom{[n+1]}{k}$, hence in particular contains the coordinates indexed by $\binom{[n]}{k} \subseteq \binom{[n+1]}{k}$. Let

$$\pi_{n+1\to n} : \mathsf{bool}^{m_{n+1}} \longrightarrow \mathsf{bool}^{m_n}$$

be the (deterministic) projection term dropping all coordinates corresponding to k-subsets containing $n{+}1$. We claim that the composite $\pi_{n+1\to n} \circ t_{n+1,k}$ is equal, in the program theory, to $t_{n,k}$.

To see this, expand $t_{n+1,k}$ as

$$\mathsf{let}\ x_1 = \mathsf{new}(\star)\ \mathsf{in}\ \cdots\ \mathsf{let}\ x_n = \mathsf{new}(\star)\ \mathsf{in}\ \mathsf{let}\ x_{n+1} = \mathsf{new}(\star)\ \mathsf{in}\ \mathrm{inc}_{n+1,k}(x_1, \dots, x_{n+1}).$$

After applying $\pi_{n+1\to n}$, every remaining output coordinate mentions only $x_1, \dots, x_n$; the bound variable x_{n+1} becomes unused. By the weakening (discard) let-law, we may remove an unused probabilistic binding, obtaining precisely the n-vertex program $t_{n,k}$. Interpreting this equality in $\mathcal{C}$ and observing with Ψ, we conclude that the distribution p_{n+1} marginalizes along restriction $\binom{[n]}{k} \subseteq \binom{[n+1]}{k}$ to p_n, i.e. (p_n) is projective.

k-locality (dissociation) from deterministic substitution. Let $A, B \subseteq [n]$ be disjoint. Consider the deterministic restriction maps

$$\rho_A : \mathsf{bool}^{\binom{n}{k}} \to \mathsf{bool}^{\binom{|A|}{k}}, \qquad \rho_B : \mathsf{bool}^{\binom{n}{k}} \to \mathsf{bool}^{\binom{|B|}{k}},$$

defined by projecting to those coordinates indexed by k-subsets entirely contained in A or in B (using fixed enumerations of $\binom{A}{k}$ and $\binom{B}{k}$). The induced pair

$$(\rho_A, \rho_B) : \mathsf{bool}^{\binom{n}{k}} \longrightarrow \mathsf{bool}^{\binom{|A|}{k}} \times \mathsf{bool}^{\binom{|B|}{k}}$$

is again deterministic. We must show that, under p_n, the random variables ρ_A and ρ_B are independent.

At the level of programs, the key observation is that the coordinates in $\rho_A(\mathrm{inc}_{n,k}(\vec{x}))$ depend only on the subtuple $(x_i)_{i\in A}$, and similarly the coordinates in $\rho_B(\mathrm{inc}_{n,k}(\vec{x}))$ depend only on $(x_j)_{j\in B}$. This dependence claim uses determinism of $\mathsf{hyperedge}_k$ in an essential way: since $\mathsf{hyperedge}_k$ is deterministic, substitution behaves as in the ordinary (non-probabilistic) β-law, and we can treat each query $\mathsf{hyperedge}_k(x_{i_1},\ldots,x_{i_k})$ as a pure function of its inputs when rearranging let-bindings.

Formally, we rewrite $t_{n,k}$ (up to permutation of independent let-bindings) into a form where the vertices indexed by A are sampled first, those indexed by B are sampled second, and all remaining vertices are sampled last:

$$t_{n,k} \;\equiv\; \mathsf{let}\ \vec{x}_A = \mathsf{new}^{|A|}(\star)\ \mathsf{in}\ \mathsf{let}\ \vec{x}_B = \mathsf{new}^{|B|}(\star)\ \mathsf{in}\ \mathsf{let}\ \vec{x}_C = \mathsf{new}^{|C|}(\star)\ \mathsf{in}\ \mathrm{inc}_{n,k}(\vec{x}_A,\vec{x}_B,\vec{x}_C),$$

where $C = [n]\setminus(A\cup B)$ and new^r abbreviates an r-fold product of independent calls to new. Applying (ρ_A,ρ_B) to the output and using the fact that ρ_A and ρ_B ignore every coordinate involving a vertex from C (and moreover, by disjointness, ignore every coordinate involving both an A-vertex and a B-vertex), we obtain a program equivalent to

$$\mathsf{let}\ \vec{x}_A = \mathsf{new}^{|A|}(\star)\ \mathsf{in}\ \mathsf{let}\ \vec{x}_B = \mathsf{new}^{|B|}(\star)\ \mathsf{in}\ \Big(\mathrm{inc}_{A,k}(\vec{x}_A),\ \mathrm{inc}_{B,k}(\vec{x}_B)\Big),$$

where $\mathrm{inc}_{A,k}$ and $\mathrm{inc}_{B,k}$ denote the incidence constructions on the smaller vertex sets (transported along fixed bijections $A\cong[|A|]$ and $B\cong[|B|]$). The remaining sampling of $\vec{x}_C$ has disappeared by weakening, since it is unused after restriction.

Now the displayed term is manifestly a sequential composition of two independent samplers followed by pairing of deterministic outputs. By the pairing and associativity let-laws, its denotation in $\mathcal{C}$ is the tensor/product of the two marginal denotations; observing under Ψ yields

$$(\rho_A,\rho_B)_*p_n \;=\; (\rho_A)_*p_n \;\otimes\; (\rho_B)_*p_n.$$

This is precisely dissociation (unconditional independence) for the induced sub-hypergraphs on A and B.

Summary and transition. We have thus extracted, from the core let-equations together with the k-uniformity and determinism axioms for $\mathsf{hyperedge}_k$, the three structural properties required for the dissociated exchangeable representation theorem: exchangeability (by commutative-let), projectivity (by weakening), and dissociation/k-locality (by deterministic substitution together with the product structure implicit in pairing). In the next section we apply Theorem 3.1 to identify the resulting family (p_n) with a hypergraphon-generated model, and we explain precisely what is and is not identifiable from the induced equational theory.

5 From random hypergraphs to hypergraphons

From Section 4 we have, for each Bernoulli-based interpretation of Lang_k, an induced family of distributions

$$p_n \;\in\; \Delta\big(\{0,1\}^{\binom{[n]}{k}}\big) \qquad (n \geq 1),$$

and we have verified that this family is exchangeable, projective, and k-local (dissociated) in the sense of the preceding definitions. These three properties are exactly the hypotheses under which the dissociated exchangeable-array representation theorem applies in the k-uniform, $\{0,1\}$-valued, symmetric setting.

Hypergraphon-generated models. Recall that a k-uniform hypergraphon (in dissociated array form) is a measurable map

$$W : [0,1]^{2^k-2} \longrightarrow [0,1]$$

invariant under the natural S_k-action on coordinates indexed by the nonempty proper subsets of $[k]$. Given such a W, one defines for each n a random incidence function on $\binom{[n]}{k}$ by the standard sampling scheme: sample i.i.d. uniforms

$$U_I \sim \mathrm{Unif}[0,1] \qquad (\emptyset \neq I \subseteq [n],\ |I| < k),$$

and then, for each $e \in \binom{[n]}{k}$, include the hyperedge e with conditional probability

$$W\big((U_J)_{\emptyset\neq J\subsetneq e}\big),$$

independently over distinct e conditional on the entire family (U_I). We write $p_{W,n}$ for the resulting distribution on $\{0,1\}^{\binom{[n]}{k}}$. The conditional independence built into this construction is precisely the probabilistic shadow of dissociation, and the S_k-invariance is what ensures that the resulting model is symmetric as a k-uniform hypergraph model rather than a model of ordered k-tuples.

Application of the representation theorem. Let us write Theorem 3.1 for a fixed choice of dissociated Aldous–Hoover/Kallenberg theorem specialized to $\{0,1\}$-valued, symmetric k-arrays.[1] Applying it to the family $(p_n)_{n\geq 1}$ produced by $t_{n,k}$ yields the promised hypergraphon semantics.

Corollary 5.1 (Hypergraphon representation). *Let $\mathcal{C}$ be a Bernoulli-based distributive Markov category interpreting* Lang_k *with deterministic* $\mathsf{hyperedge}_k$

[1] Any of the standard formulations suffices for our purposes; we only use existence and the usual uniqueness notion up to measure-preserving change of variables.

satisfying irreflexivity and symmetry. Let (p_n) be the induced random k-uniform hypergraph model. Then there exists a k-uniform hypergraphon W such that

$$p_n \;=\; p_{W,n} \qquad \text{for all } n \geq 1.$$

Moreover, W is unique up to the standard measure-preserving equivalence.

In this corollary, the content is not that W exists as a function on a particular parameter space, but that the entire *projective system of finite distributions* is representable by a single measurable kernel with the Aldous–Hoover coordinate structure. The role of dissociation is to eliminate any additional "global" latent random variable beyond those attached to nonempty proper subsets: intuitively, there is no extra randomness coupling disjoint vertex blocks, so the representation may be taken in the "ergodic" (dissociated) form.

Measure-preserving equivalence and what "uniqueness" means. As in the graphon case ($k = 2$), the representing hypergraphon is not unique as a pointwise function. There are at least two inevitable sources of non-identifiability.

First, W is only defined up to modification on a null set, since $p_{W,n}$ depends on W only through integrals against product Lebesgue measure. Thus we freely identify kernels that agree almost everywhere.

Second, the latent coordinate system admits measure-preserving reparametrizations. A basic family of symmetries is obtained as follows: if $\phi : [0,1] \to [0,1]$ is measure-preserving, then by applying ϕ to each coordinate we obtain a transformed kernel

$$W^{\phi} \;:=\; W \circ \phi^{\times(2^k-2)}.$$

Because (U_I) is an i.i.d. family of uniforms, the transformed family $(\phi(U_I))$ has the same joint law, and therefore

$$p_{W^{\phi},n} \;=\; p_{W,n} \qquad \text{for all } n.$$

In the hypergraphon setting there are, in general, more elaborate "structure-preserving" transformations of the product space $[0,1]^{2^k-2}$ compatible with the subset-indexing and the S_k-action; these are discussed in the exchangeability literature. For our purposes, it is technically cleanest to take the induced model itself as the invariant:

Definition 5.2 (Weak equivalence)**.** Two k-uniform hypergraphons W, W' are *weakly equivalent* if

$$p_{W,n} \;=\; p_{W',n} \qquad \text{for all } n \geq 1.$$

The uniqueness clause in Corollary 5.1 is then read as: the program-induced model determines a unique weak equivalence class of hypergraphons, and any two hypergraphons producing the same projective family of finite-dimensional distributions are related by a measure-theoretic isomorphism of the underlying parameter spaces (in the standard sense for dissociated exchangeable arrays). We will therefore speak of *the* hypergraphon semantics of an equational theory, meaning this weak equivalence class.

Explicit finite-dimensional formulas. Although we do not need an explicit integration formula to invoke the representation theorem, it is useful to record what information about W is visible at finite n. Fix n, and write λ for Lebesgue measure. For each $e \in \binom{[n]}{k}$ let

$$W_e\big((u_I)_{\emptyset \neq I \subseteq [n],\, |I|<k}\big) \;:=\; W\big((u_J)_{\emptyset \neq J \subsetneq e}\big),$$

viewing W_e as the pullback of W along the coordinate projection selecting precisely those subset-variables that lie inside e. Then for any labeled hypergraph $H = ([n], E)$ the probability of sampling exactly H under W is

$$p_{W,n}(H) \;=\; \int \left(\prod_{e \in E} W_e(u)\right) \left(\prod_{e \in \binom{[n]}{k} \setminus E} (1 - W_e(u))\right) d\lambda(u), \qquad (1)$$

where u ranges over $[0,1]^{\sum_{r=1}^{k-1} \binom{n}{r}}$ indexed by the nonempty subsets of $[n]$ of size $< k$. In particular, every $p_{W,n}(H)$ is a polynomial expression in W and $1 - W$ integrated against product measure, and hence depends only on the weak equivalence class of W.

Identifiability from the program theory. The program theory—more precisely, the observed denotations under Ψ of closed numeral terms—determines exactly the family $(p_n)_{n \geq 1}$. Consequently, it determines precisely those invariants of W that can be expressed in terms of the finite laws (1). Concretely:

- For each n and each labeled k-uniform hypergraph H on $[n]$, the real number $p_n(H)$ is determined. Equivalently, the theory determines all joint moments of the incidence indicators $\big(A(e)\big)_{e \in \binom{[n]}{k}}$.
- For each fixed finite k-uniform hypergraph F (say with vertex set $[m]$), the probability that the induced random hypergraph on m sampled vertices equals F is determined; by exchangeability this is independent of which m vertices we choose, and by projectivity it is consistent across $n \geq m$.

- Any two hypergraphons W, W' that differ only by a measure-preserving change of variables (in particular, any pullback W^{ϕ} as above) are indistinguishable by the program theory, since they give identical $(p_{W,n})$.

Conversely, nothing in the program theory can select a canonical representative of the weak equivalence class: pointwise features of W (values on a set of positive measure, continuity, etc.) are not invariants of the induced model unless they can be reconstructed from the collection of all finite distributions. This is the same phenomenon familiar from graphons: the observable content is the distribution on finite graphs (or equivalently the family of substructure densities), not a particular coordinatization of the latent $[0, 1]$-space.

What the dissociated form buys us. It is worth isolating the conceptual consequence of k-locality. Without dissociation, the general Aldous–Hoover representation introduces an additional global random variable (often denoted $U_{\emptyset}$) mediating correlations between disjoint vertex blocks. Our locality lemma rules out precisely such long-range coupling, and therefore forces the representing object to be a single hypergraphon W with only the subset-indexed coordinates. In this sense, the let-equations do not merely ensure symmetry and consistency; they also enforce an "ergodic" exchangeable structure compatible with independent generation on disjoint vertex sets.

Transition to universal semantics. We have now moved from a Bernoulli-based equational theory to a canonical probabilistic invariant: a projective, exchangeable, dissociated model (p_n), equivalently a weak equivalence class of k-uniform hypergraphons. To complete the correspondence in the opposite direction, and to obtain a syntax-free description of all such models at once, we next construct a universal distributive Markov category $\mathcal{U}_k$ for the interface and identify its numeral-fragment Markov functors into $\mathsf{FinStoch}$ with precisely these random hypergraph models.

5.1 A universal semantics for the hypergraph interface

We now give a syntax-free construction which plays, for the interface $(\mathsf{vertex}, \mathsf{new}, \mathsf{hyperedge}_k)$, the same role that the usual "classifying category" plays for ordinary algebraic theories. The point is to isolate a single distributive Markov category $\mathcal{U}_k$ in which the interface is interpreted *universally*, so that every concrete Bernoulli-based interpretation factors through $\mathcal{U}_k$ by a unique distributive Markov functor. In the next section we will restrict attention to the numeral fragment $(\mathcal{U}_k)_{\mathbb{N}}$ and characterize its $\mathsf{FinStoch}$-models as precisely the exchangeable/projective/k-local random hypergraph models.

The category of finite k-uniform hypergraphs. We write Hyp_k for the category whose objects are finite k-uniform hypergraphs $H = (V(H), E(H))$

and whose morphisms are structure-preserving and structure-reflecting maps.

Definition 5.3 (Hyp_k)**.** A morphism $f : H \to H'$ in Hyp_k is a function $f : V(H) \to V(H')$ such that for every k-element subset $e \subseteq V(H)$ one has

$$e \in E(H) \quad \Longleftrightarrow \quad f(e) \in E(H'),$$

where $f(e) := \{f(v) \mid v \in e\}$ (so the right-hand side is understood only when $|f(e)| = k$).

The preservation direction ensures that hyperedges are sent to hyperedges, while reflection ensures that no new hyperedge is created by post-composition. In particular, for labeled hypergraphs on vertex sets $[n]$ this implies that relabelings are exactly the isomorphisms in Hyp_k; injections $[m] \hookrightarrow [n]$ which identify $[m]$ with a subset of $[n]$ correspond precisely to induced-subhypergraph embeddings; and more generally, Hyp_k is a convenient deterministic skeleton in which "remembering the incidence tensor" is functorial.

We will systematically use the opposite category $\mathsf{Hyp}_k^{\mathrm{op}}$, since in semantics we want maps *out of* a larger vertex context to represent restriction along an inclusion. Concretely, an inclusion of vertex sets $\iota : [m] \hookrightarrow [n]$ induces, for each hypergraph H on $[n]$, a restriction hypergraph $H \restriction_\iota$ on $[m]$; this operation is contravariant and therefore naturally lives in $\mathsf{Hyp}_k^{\mathrm{op}}$.

Finite-coproduct completion. The deterministic core of our universal semantics must interpret sum types (in particular $\mathsf{bool} = \mathbf{1} + \mathbf{1}$) and must be distributive. We therefore pass from $\mathsf{Hyp}_k^{\mathrm{op}}$ to its finite-coproduct completion.

Definition 5.4 ($\mathsf{Fam}(-)$)**.** For a category $\mathcal{D}$, let $\mathsf{Fam}(\mathcal{D})$ denote the category of finite families in $\mathcal{D}$. An object is a finite indexed family $(D_i)_{i\in I}$ with I finite. A morphism

$$(D_i)_{i\in I} \longrightarrow (E_j)_{j\in J}$$

is given by a function $\varphi : I \to J$ together with morphisms $D_i \to E_{\varphi(i)}$ in $\mathcal{D}$, one for each $i \in I$. Coproducts are given by concatenation of families, with injections induced by inclusions of index sets.

The category $\mathsf{Fam}(\mathcal{D})$ is extensive, hence distributive, and it contains $\mathcal{D}$ as the full subcategory of singleton families. In particular, $\mathsf{Fam}(\mathsf{Hyp}_k^{\mathrm{op}})$ provides a deterministic setting in which we can form finite case distinctions over finitely many hypergraph-shaped summands. We regard this as the deterministic "world of finite hypergraphs" in which the interface predicate $\mathsf{hyperedge}_k$ will be interpreted.

A canonical family of n-vertex contexts. For each $n \geq 0$ we write $\mathsf{Hyp}_k([n])$ for the (finite) set of all k-uniform hypergraphs with vertex set $[n]$. We introduce in $\mathsf{Fam}(\mathsf{Hyp}_k^{\mathrm{op}})$ the object

$$\mathbf{V}_n \;:=\; (H)_{H \in \mathsf{Hyp}_k([n])},$$

the family consisting of one summand for each labeled k-uniform hypergraph on $[n]$. The intended reading is that $\mathbf{V}_n$ is the "universal n-vertex context": a point of $\mathbf{V}_n$ chooses a concrete hypergraph structure on n labeled vertices. Because $\mathsf{Hyp}_k([n])$ is finite, $\mathbf{V}_n$ is a well-formed object of the finite-coproduct completion.

Given an injection $\iota : [m] \hookrightarrow [n]$, restriction of hypergraphs along ι defines a function

$$\iota^* : \mathsf{Hyp}_k([n]) \longrightarrow \mathsf{Hyp}_k([m]), \qquad H \longmapsto H\restriction_\iota,$$

and for each $H \in \mathsf{Hyp}_k([n])$ the inclusion map $\iota : [m] \to [n]$ itself determines a morphism in Hyp_k from $H\restriction_\iota$ to H (as an induced-subhypergraph embedding), hence a morphism $H \to H\restriction_\iota$ in $\mathsf{Hyp}_k^{\mathrm{op}}$. Assembling these data yields a canonical morphism

$$\mathbf{V}_n \longrightarrow \mathbf{V}_m$$

in $\mathsf{Fam}(\mathsf{Hyp}_k^{\mathrm{op}})$ which implements restriction of an n-vertex hypergraph to the ι-selected m-vertex subhypergraph. Similarly, bijections $\sigma : [n] \to [n]$ act by relabeling and yield endomorphisms of $\mathbf{V}_n$ encoding exchangeability at the deterministic level.

Adjoining new by a monoidal indeterminate. The preceding category is deterministic; it contains no morphism that behaves like a probabilistic sampler. We now freely adjoin such a sampler as a generator. Concretely, we use the monoidal-indeterminate construction from the general theory of distributive Markov categories: given a distributive category $\mathcal{D}$ and an object $X \in \mathcal{D}$, there is a distributive Markov category $\mathcal{D}[\nu]$ equipped with a morphism $\nu : \mathbf{1} \to X$ which is universal among such choices. We apply this with $\mathcal{D} := \mathsf{Fam}(\mathsf{Hyp}_k^{\mathrm{op}})$ and with the distinguished object

$$\mathsf{vertex} \;:=\; \mathbf{V}_1.$$

We write

$$\mathcal{U}_k \;:=\; \mathsf{Fam}(\mathsf{Hyp}_k^{\mathrm{op}})[\nu], \qquad \text{with } \nu : \mathbf{1} \to \mathsf{vertex},$$

and we interpret the interface constant $\mathsf{new} : \mathbf{1} \to \mathsf{vertex}$ as this indeterminate ν.

The universal property we use is the following: for every distributive Markov category $\mathcal{C}$, to give a distributive Markov functor $F : \mathcal{U}_k \to \mathcal{C}$ is equivalently to give (i) a distributive functor $F_0 : \mathsf{Fam}(\mathsf{Hyp}_k^{\mathrm{op}}) \to \mathcal{C}$ interpreting the deterministic hypergraph structure, together with (ii) a morphism

$\mathbf{1} \to F_0(\mathsf{vertex})$ in $\mathcal{C}$ interpreting new. In particular, $\mathcal{U}_k$ is initial among distributive Markov categories equipped with an interpretation of the deterministic hypergraph fragment and a chosen sampler for vertices.

The universal deterministic predicate $\mathsf{hyperedge}_k$. It remains to define, inside $\mathsf{Fam}(\mathsf{Hyp}_k^{\mathrm{op}})$ (hence also inside $\mathcal{U}_k$), a deterministic morphism interpreting

$$\mathsf{hyperedge}_k : (\mathsf{vertex})^k \to \mathsf{bool}.$$

We take $\mathsf{bool} := \mathbf{1} + \mathbf{1}$ in $\mathcal{U}_k$, with coproduct injections denoted $\mathsf{true} : \mathbf{1} \to \mathsf{bool}$ and $\mathsf{false} : \mathbf{1} \to \mathsf{bool}$.

The key observation is that on a fixed set of k labeled vertices there is exactly one potential k-hyperedge, namely $[k]$ itself. Hence $\mathsf{Hyp}_k([k])$ has precisely two elements: the empty k-vertex hypergraph $H_k^{\perp}$ with $E = \emptyset$, and the full k-vertex hypergraph $H_k^{\top}$ with $E = \{[k]\}$. Therefore the corresponding family object

$$\mathbf{V}_k \;=\; (H_k^{\perp},\, H_k^{\top})$$

is canonically a binary coproduct in $\mathsf{Fam}(\mathsf{Hyp}_k^{\mathrm{op}})$. We define $\mathsf{hyperedge}_k$ to be the evident "tag" map which distinguishes these two summands:

$$\mathsf{hyperedge}_k : \mathbf{V}_k \longrightarrow \mathsf{bool}, \qquad \mathsf{hyperedge}_k \circ \iota_{\perp} = \mathsf{false}, \quad \mathsf{hyperedge}_k \circ \iota_{\top} = \mathsf{true},$$

where $\iota_{\perp} : H_k^{\perp} \to \mathbf{V}_k$ and $\iota_{\top} : H_k^{\top} \to \mathbf{V}_k$ are the coproduct injections. By construction this morphism is deterministic (it lies in the base category before adjoining ν).

To evaluate the edge predicate on an arbitrary n-vertex context and a chosen k-subset $e \in \binom{[n]}{k}$, we use the restriction morphism induced by the inclusion $\iota_e : [k] \hookrightarrow [n]$ whose image is e (in increasing order). The composite

$$\mathbf{V}_n \xrightarrow{\iota_e^*} \mathbf{V}_k \xrightarrow{\mathsf{hyperedge}_k} \mathsf{bool}$$

is then the deterministic query returning the incidence value of e in the chosen n-vertex hypergraph. Symmetry under permutations of the k arguments is immediate from the fact that $\mathbf{V}_k$ carries the evident S_k-action by relabeling, and $\mathsf{hyperedge}_k$ is invariant under this action since it depends only on whether the unique k-set is present. Irreflexivity is enforced by stipulating that any attempt to test a k-tuple with repeated vertices factors through the k-vertex discrete context in which no k-edge can be present; equivalently, such a test is identified with the constant map false in $\mathcal{U}_k$.

Summary. The category $\mathcal{U}_k = \mathsf{Fam}(\mathsf{Hyp}_k^{\mathrm{op}})[\nu]$ is thus equipped with distinguished interpretations of vertex, new, and $\mathsf{hyperedge}_k$ satisfying the k-uniformity axioms by construction, and it is universal with this property. In particular, any concrete Bernoulli-based equational theory interpreting the interface determines (and is determined on the deterministic fragment by) a distributive Markov functor out of $\mathcal{U}_k$. In the next section we use this universality to classify the distributive Markov functors $(\mathcal{U}_k)_{\mathbb{N}} \to \mathsf{FinStoch}$ and hence to reconstruct, via a Bernoulli-based quotient, an equational theory realizing any prescribed k-uniform hypergraphon model.

5.2 Hypergraphons as Bernoulli-based quotients

We now restrict the universal category $\mathcal{U}_k$ to its numeral (finite) part and explain how stochastic semantics on this fragment are exactly the same data as exchangeable, projective, and k-local random k-uniform hypergraph models. We then use the resulting functorial presentation to build, for an *arbitrary* hypergraphon W, a Bernoulli-based equational theory whose induced distributions coincide with $(p_{W,n})_{n\geq 1}$.

The numeral fragment and the induced distributions. Write $(\mathcal{U}_k)_{\mathbb{N}}$ for the full subcategory of $\mathcal{U}_k$ on the objects vertex^n ($n \in \mathbb{N}$), regarded as the canonical n-vertex contexts.[2] By Lemma L5 we have a natural identification

$$(\mathcal{U}_k)(\mathbf{1}, \mathsf{vertex}^n) \;\cong\; \mathsf{Hyp}_k([n]),$$

so that a global element of vertex^n is precisely a labeled k-uniform hypergraph on vertex set $[n]$.

Let $F : (\mathcal{U}_k)_{\mathbb{N}} \to \mathsf{FinStoch}$ be a distributive Markov functor. We define

$$p_n \;:=\; F(\nu^{\otimes n}) \;\in\; \mathsf{FinStoch}\big(1, F(\mathsf{vertex}^n)\big),$$

and we view p_n as a probability distribution on $\mathsf{Hyp}_k([n])$ by transporting along the canonical bijection between the underlying finite set of $F(\mathsf{vertex}^n)$ and $\mathsf{Hyp}_k([n])$ induced by the above hom-set identification. Intuitively, $\nu^{\otimes n}$ is the universal program fragment which samples n fresh vertices; applying F produces a distribution on the space of n-vertex hypergraph structures.

Deterministic structure: permutations and restrictions. Two classes of deterministic morphisms in $(\mathcal{U}_k)_{\mathbb{N}}$ play a distinguished role.

First, every bijection $\sigma : [n] \to [n]$ induces (by relabeling in Hyp_k and contravariance) a morphism

$$\rho_\sigma : \mathsf{vertex}^n \longrightarrow \mathsf{vertex}^n$$

encoding relabeling. Second, every injection $\iota : [m] \hookrightarrow [n]$ induces a restriction map

$$r_\iota : \mathsf{vertex}^n \longrightarrow \mathsf{vertex}^m,$$

[2] Equivalently, one may use the isomorphic family objects $\mathbf{V}_n$ introduced in the deterministic core; the choice is inessential for the present discussion.

obtained by passing to the induced subhypergraph on the image of ι. These maps satisfy the evident functorial identities

$$\rho_{\sigma\tau} = \rho_\tau \circ \rho_\sigma, \qquad r_{\iota\circ\kappa} = r_\kappa \circ r_\iota, \qquad r_\iota \circ \rho_\sigma = r_{\sigma^{-1}\circ\iota},$$

and they generate, in a precise sense, the deterministic shape of the n-vertex contexts.

Stochastic structure: locality via monoidality. The Markov structure in $\mathcal{U}_k$ supplies tensorial composition of independent samplers. On numeral contexts this is reflected by canonical comparison maps that split an $(m+n)$-vertex context into two disjoint blocks (first m vertices and last n vertices), together with restriction morphisms projecting to each block. Under a distributive Markov functor F, the monoidality constraints ensure that sampling on disjoint blocks becomes product (independent) sampling in $\mathsf{FinStoch}$. This is the categorical origin of dissociation/k-locality: cross-block dependence is carried precisely by those deterministic maps in $\mathcal{U}_k$ which *remember* edges meeting both blocks, and monoidality forces independence once we forget such cross-terms by restricting to induced subhypergraphs on each block separately.

Theorem 5.5 (Functor classification on numerals). *Distributive Markov functors*

$$F : (\mathcal{U}_k)_{\mathbb{N}} \longrightarrow \mathsf{FinStoch}$$

are in bijection with sequences $(p_n)_{n\geq 1}$ of probability distributions on $\mathsf{Hyp}_k([n])$ that are exchangeable, projective, and k-local.

Proof sketch. Given F, define $p_n := F(\nu^{\otimes n})$ as above. Exchangeability follows from naturality with respect to ρ_σ: since ρ_σ is deterministic, we have

$$F(\rho_\sigma) \circ p_n \;=\; F(\rho_\sigma \circ \nu^{\otimes n}) \;=\; F(\nu^{\otimes n}) \;=\; p_n,$$

which says exactly that p_n is invariant under relabeling by σ.

Projectivity follows similarly from naturality with respect to restriction maps. If $\iota : [n] \hookrightarrow [n+1]$ is the standard inclusion, then

$$F(r_\iota) \circ p_{n+1} \;=\; F(r_\iota) \circ F(\nu^{\otimes(n+1)}) \;=\; F(r_\iota \circ \nu^{\otimes(n+1)}).$$

In $\mathcal{U}_k$, the composite $r_\iota \circ \nu^{\otimes(n+1)}$ agrees with $\nu^{\otimes n}$ (discard the last sampled vertex and then restrict), hence $F(r_\iota) \circ p_{n+1} = p_n$.

For k-locality, let $A, B \subseteq [n]$ be disjoint. Consider the deterministic map $\mathsf{vertex}^n \to \mathsf{vertex}^{|A|} \times \mathsf{vertex}^{|B|}$ obtained by restricting to A and to B and pairing the results. Functoriality yields that the pushforward of p_n along this map is the joint law of the induced subhypergraphs on A and B. In $\mathcal{U}_k$ this paired restriction factors through a tensor decomposition corresponding to the disjointness of A and B, and applying the (strong) monoidality of F identifies the resulting joint law with a product measure, establishing independence.

Conversely, given (p_n) exchangeable, projective, and k-local, we define F on objects by $F(\mathsf{vertex}^n) := |\mathsf{Hyp}_k([n])|$ (as an object of $\mathsf{FinStoch}$) and on the generating deterministic maps ρ_σ and r_ι by the corresponding relabeling and restriction functions on hypergraphs, viewed as deterministic stochastic matrices. On the states $\nu^{\otimes n} : \mathbf{1} \to \mathsf{vertex}^n$ we set $F(\nu^{\otimes n}) := p_n$. Exchangeability and projectivity ensure that these assignments respect the relations among the ρ_σ and r_ι, while k-locality is exactly what is needed to make the tensorial structure on disjoint blocks compatible with the monoidal structure in $\mathsf{FinStoch}$. Extending by distributivity and the Markov axioms then determines F uniquely. □

From hypergraphons to functors. Fix a k-uniform hypergraphon W in the dissociated array form. Its sampling scheme produces, for each n, a distribution $p_{W,n}$ on $\mathsf{Hyp}_k([n])$ which is exchangeable and projective by construction and dissociated (hence k-local) because, conditional on the latent i.i.d. family $(U_J)_{0<|J|<k}$, hyperedges are independent and are measurable in the U_J supported on the corresponding vertex subsets. Therefore Theorem 5.5 yields a canonical distributive Markov functor

$$F_W : (\mathcal{U}_k)_{\mathbb{N}} \longrightarrow \mathsf{FinStoch} \qquad \text{with} \qquad F_W(\nu^{\otimes n}) = p_{W,n}.$$

Contextual-equivalence quotients and Bernoulli bases. To obtain a *full* Bernoulli-based equational theory (not merely a model on numeral contexts) we quotient $\mathcal{U}_k$ by the observational congruence induced by F_W. Concretely, we define an equivalence relation $\sim_W$ on parallel morphisms of $\mathcal{U}_k$ by declaring $f \sim_W g$ whenever, for every numeral context vertex^n and every deterministic "observation" map o out of the codomain into a numeral object, the induced stochastic maps agree after applying F_W on numerals. Equivalently, $f \sim_W g$ holds when f and g are indistinguishable by any closed numeral experiment in the sense of the Bernoulli base.

The quotient category

$$\mathcal{U}_k/F_W$$

is defined by keeping the same objects as $\mathcal{U}_k$ and quotienting each hom-set by $\sim_W$. By construction, the inclusion of the numeral fragment descends to a faithful distributive Markov functor

$$\Psi_W : (\mathcal{U}_k/F_W)_{\mathbb{N}} \hookrightarrow \mathsf{FinStoch}$$

which agrees with F_W on numeral morphisms. The faithfulness of Ψ_W is the categorical manifestation of the fact that we have quotiented *exactly* by observational equality on numerals and no further.

Corollary 5.6 (Realizing hypergraphons by Bernoulli-based theories)**.** *For every k-uniform hypergraphon W, there exists a Bernoulli-based distributive Markov category interpreting Lang_k such that the distributions induced by the incidence-tensor terms $t_{n,k}$ are exactly $(p_{W,n})_{n\geq 1}$.*

Proof sketch. Interpret the interface in $\mathcal{U}_k/F_W$ by the images of $(\mathsf{vertex}, \nu, \mathsf{hyperedge}_k)$ under the quotient functor $\mathcal{U}_k \to \mathcal{U}_k/F_W$. The axioms for k-uniformity hold already in $\mathcal{U}_k$ and are preserved by quotienting. On numerals, the Bernoulli base Ψ_W yields exactly the stochastic semantics prescribed by F_W, hence running $t_{n,k}$ and observing via Ψ_W produces $p_{W,n}$ by definition. □

Combining Corollary 5.6 with the representation direction obtained earlier (every Bernoulli-based theory induces an exchangeable, projective, k-local model and hence a hypergraphon up to equivalence), we obtain the expected completeness statement: the space of Bernoulli-based equational theories for the hypergraph interface, when restricted to numeral observations, is exhausted by dissociated hypergraphon models and nothing else. In the next section we make this correspondence concrete in examples, beginning with the case $k = 2$ and then indicating genuinely higher-order phenomena for $k \geq 3$.

6 Examples and variants

We record a number of basic instances and straightforward extensions of the interface. The purpose is twofold: first, to verify that for $k = 2$ we recover the usual graphon correspondence; second, to indicate which features of the construction are genuinely k-ary (already for $k = 3$), and which are artifacts of the particular choice of a boolean, undirected interface.

6.1 $k = 2$ and the usual graphon picture

When $k = 2$ the term constant $\mathsf{hyperedge}_2 : \mathsf{vertex} \times \mathsf{vertex} \to \mathsf{bool}$ is constrained by irreflexivity and symmetry. The incidence tensor output by $t_{n,2} : \mathbf{1} \to \mathsf{bool}^{\binom{n}{2}}$ is equivalently an adjacency matrix with a forced zero diagonal and symmetric entries. In this case our hypergraphon definition specializes to a measurable map

$$W : [0,1]^{2^2-2} = [0,1]^2 \longrightarrow [0,1],$$

invariant under the action of S_2 swapping the two coordinates; this is precisely a (symmetric) graphon in the standard sense. The sampling scheme reads as follows: for each $i \in [n]$ sample $U_{\{i\}} \sim \mathrm{Unif}[0,1]$ i.i.d., and then for each unordered pair $\{i,j\} \in \binom{[n]}{2}$ include the edge with probability $W(U_{\{i\}}, U_{\{j\}})$, independently over pairs conditional on the vertex latents.

Under this identification, exchangeability becomes invariance under relabeling of vertices, projectivity becomes consistency under restriction to induced subgraphs, and 2-locality coincides with dissociation of exchangeable graph models (independence for induced subgraphs on disjoint vertex sets). Thus the general results above contain the familiar statement that Bernoulli-based equational theories for the undirected graph interface, when observed on numeral outputs, are classified (up to measure-preserving equivalence) by graphons.

6.2 $k = 3$: dependence on lower-dimensional faces

For $k = 3$ the canonical hypergraphon has domain

$$[0,1]^{2^3-2} = [0,1]^6,$$

with coordinates indexed by the nonempty proper subsets of $[3]$, namely $\{1\}, \{2\}, \{3\}, \{1,2\}, \{1,3\}, \{2,3\}$. In the sampling scheme on $[n]$, we therefore draw i.i.d. latents $U_{\{i\}}$ for vertices and $U_{\{i,j\}}$ for unordered pairs, and we include a triple $\{i, j, \ell\}$ with probability

$$W\big(U_{\{i\}}, U_{\{j\}}, U_{\{\ell\}}, U_{\{i,j\}}, U_{\{i,\ell\}}, U_{\{j,\ell\}}\big),$$

independently over triples conditional on all these latents. The S_3-invariance condition forces W to be symmetric under simultaneous permutation of the six inputs according to the induced action on the indexing subsets.

This already exhibits a phenomenon absent at $k = 2$: the probability of a 3-edge may depend not only on per-vertex parameters but also on per-pair parameters. Such dependence is not cosmetic; it expresses the fact that, in general dissociated exchangeable 3-uniform hypergraph models, correlations between hyperedges sharing two vertices can be mediated by a latent variable attached to that shared *pair*. Concretely, consider a symmetric measurable W of the form

$$W\big(u_1, u_2, u_3, u_{12}, u_{13}, u_{23}\big) \;:=\; \mathbf{1}\{u_{12} \le 1/2\} \cdot \mathbf{1}\{u_{13} \le 1/2\} \cdot \mathbf{1}\{u_{23} \le 1/2\},$$

interpreted as a $\{0,1\}$-valued hypergraphon. Then a triple $\{i, j, \ell\}$ is present if and only if all three pair-latents $U_{\{i,j\}}, U_{\{i,\ell\}}, U_{\{j,\ell\}}$ fall in $[0, 1/2]$. In particular, for fixed $i \neq j$ the random variables $\mathsf{hyperedge}_3(i, j, \ell)$ as ℓ varies share the common factor $\mathbf{1}\{U_{\{i,j\}} \le 1/2\}$, yielding a nontrivial correlation pattern which cannot be represented by a model depending only on vertex latents $(U_{\{i\}})$ (equivalently, by a naive kernel $[0,1]^3 \to [0,1]$). Said differently, for $k \geq 3$ the correct limit object is not vertex-parametric; it must allow dependence on all proper faces of a k-simplex, as mandated by the dissociated Aldous–Hoover/Kallenberg theorem.

From the programming perspective, this distinction corresponds to the fact that the interface does not constrain how $\mathsf{hyperedge}_3$ may depend on the history of sampling, beyond determinism and symmetry; therefore a Bernoulli-based interpretation may encode additional shared structure at the level of pairs, while still satisfying 3-locality in the sense of independence for disjoint vertex sets.

6.3 Higher uniformities

For general $k \geq 2$, the number of latent coordinates per potential hyperedge is $2^k - 2$, indexed by the nonempty proper subsets of $[k]$. The combinatorics grows quickly, but the conceptual content remains stable: the latent family contains a variable U_J for each lower-dimensional face J of the k-simplex, and the hyperedge probability is a symmetric measurable function of these face variables. The k-locality/dissociation property is exactly what removes any additional global randomness beyond this face-indexed family (and the independent coin flips used to realize Bernoulli outcomes), and this is why the resulting sampling scheme factors over disjoint vertex blocks once we restrict to induced subhypergraphs.

It is sometimes useful to isolate sub-classes of hypergraphons by imposing further invariances or factorization properties on W. For instance, restricting to those W that depend only on singleton coordinates yields the "vertex-kernel" subclass $W : [0,1]^k \to [0,1]$, which is strictly less expressive for $k \geq 3$ but may serve as a convenient parametric model. Our completeness statement does not privilege such subclasses: the universal category and quotient construction realize the full dissociated class.

6.4 Dropping symmetry: directed and ordered interfaces

The symmetry axioms implement undirectedness by forcing $\mathsf{hyperedge}_k$ to be invariant under permutations of its k arguments, and irreflexivity enforces simplicity. If we remove symmetry, we obtain an interface for *ordered* (or directed) k-ary relations:

$\mathsf{hyperedge}_k : \mathsf{vertex}^k \to \mathsf{bool}$ deterministic, but not assumed symmetric.

The induced distributions are then exchangeable as *jointly exchangeable* k-arrays, i.e. invariant under the simultaneous relabeling action of S_n on each coordinate, but not necessarily invariant under permutations of the k slots. The representation theorem correspondingly yields a measurable

$$W : [0,1]^{2^k - 2} \to [0,1]$$

with no S_k-invariance requirement. Thus, at the level of the universal semantics, the same functorial classification applies after replacing Hyp_k by the appropriate category of finite ordered k-ary relational structures. Irreflexivity may likewise be weakened or removed, leading to looped structures and the corresponding modification of the deterministic core.

6.5 Colored hyperedges and non-boolean relations

The boolean codomain is inessential for the classification method; what matters is that observations on numerals are finite. Fix a finite set of colors $\mathbf{m} = \{1, \ldots, m\}$ (as a numeral type) and replace the interface constant by a deterministic color assignment

$$\mathsf{hyperedge}_k^{(m)} : \mathsf{vertex}^k \to \mathbf{m},$$

still subject to symmetry/irreflexivity as desired. A closed term $t_{n,k}$ then produces a tensor in $\mathbf{m}^{\binom{n}{k}}$, i.e. an m-colored k-uniform hypergraph. In the hypergraphon representation this corresponds to a measurable map into the simplex,

$$W : [0,1]^{2^k-2} \longrightarrow \Delta_m,$$

so that, conditional on the face latents, each hyperedge receives a color drawn from the categorical distribution $W(\cdots)$. Equivalently one may represent W as m functions $W_c : [0,1]^{2^k-2} \to [0,1]$ with $\sum_c W_c = 1$.

From the categorical viewpoint, nothing essential changes: $\mathsf{FinStoch}$ already supports finite-valued observations, and the quotient-by-contextual-equivalence construction proceeds verbatim. The only modification is in the deterministic core used to define $\mathcal{U}_k$, where Hyp_k is replaced by the category of finite colored k-uniform hypergraphs (with color-preserving and -reflecting maps if one wishes to keep the deterministic semantics conservative).

6.6 Multi-sorted array interfaces: bipartite and higher arities

A particularly useful variant replaces the single vertex sort by several independent sorts. For example, introduce two type constants $\mathsf{row}, \mathsf{col}$ with samplers

$$\mathsf{newRow} : \mathbf{1} \to \mathsf{row}, \qquad \mathsf{newCol} : \mathbf{1} \to \mathsf{col},$$

and a deterministic predicate $\mathsf{entry} : \mathsf{row} \times \mathsf{col} \to \mathsf{bool}$. The numeral contexts are then $\mathsf{row}^m \times \mathsf{col}^n$, and the induced distributions describe random bipartite graphs (or $\{0,1\}$-matrices) with separate exchangeability under $S_m \times S_n$, projectivity under restriction in either coordinate, and a suitable locality condition expressing independence for disjoint row-blocks and disjoint column-blocks.

On the analytic side, this recovers the usual bipartite graphon picture: sample i.i.d. uniforms $(U_i)_{i\in[m]}$ for rows and $(V_j)_{j\in[n]}$ for columns, and include (i,j) with probability $W(U_i, V_j)$ for a measurable $W : [0,1]^2 \to [0,1]$ (or, if one prefers to state it in Aldous–Hoover form, as a function of the singleton latents together with the independent per-entry coin flips). More generally, higher-sorted variants yield the standard limit objects for separately exchangeable tensors, and can be treated by the same "universal category + functors on numerals + quotient" pipeline after replacing Hyp_k by the appropriate category of finite multi-sorted relational structures.

6.7 Random-free semantics and $\{0,1\}$-valued limit objects

Finally, it is worth isolating the degenerate situation in which the ambient Markov structure carries no genuine randomness beyond deterministic computation (for instance, in a purely measurable/predicate semantics where new is interpreted deterministically). In such models the induced distributions p_n are Dirac measures, hence the associated hypergraphon W may be chosen $\{0,1\}$-valued (up to null sets). This mirrors the graph case: a deterministic semantics cannot produce "gray" edge densities, and the quotient construction collapses to a deterministic equational theory. The Bernoulli-based setting is therefore not merely technical; it is exactly what allows the interface to express nontrivial probabilistic mixtures while retaining an extensional, observation-based account of program equality on numerals.

7 Discussion and future directions

Our results give a complete account, on numeral observations, of Bernoulli-based equational theories for the interface $(\mathsf{new}, \mathsf{hyperedge}_k)$ under the structural assumptions of exchangeability, projectivity, and k-locality/dissociation. The corresponding analytic object is the (dense) k-uniform hypergraphon $W : [0,1]^{2^k-2} \to [0,1]$, unique up to measure-preserving equivalence, and the categorical object is the universal distributive Markov category $\mathcal{U}_k$ together with its Bernoulli-based quotients. In this section we indicate several directions in which the present framework can be extended, and where new technical input is likely required.

7.1 Sparse regimes

The hypergraphon representation used above is intrinsically a dense limit theory: for fixed n the output of $t_{n,k}$ is a full incidence tensor in $\mathsf{bool}^{\binom{n}{k}}$, and in the associated models the expected number of hyperedges typically scales as $\Theta(n^k)$. In sparse regimes one instead expects $\Theta(n^\alpha)$ hyperedges for some $\alpha < k$, or even bounded average degree. Analytically, the appropriate limit objects are not bounded kernels on $[0,1]^{2^k-2}$ but rather variants of integrable kernels or random measures (for $k = 2$ one encounters L^p-graphons and graphex processes), and the correct notion of sampling consistency is often not "take the induced subhypergraph on the first n vertices" but rather a thinning or restriction operation compatible with a Poissonized construction.

From the semantic side, two obstacles appear immediately. First, sparsity is naturally expressed by allowing the edge probability to scale with n (or by allowing an unbounded latent intensity), whereas our current interface fixes a single $\mathsf{hyperedge}_k : (\mathsf{vertex})^k \to \mathsf{bool}$ used uniformly for all n. Second, common sparse constructions produce random counts and are better expressed by point processes or random measures; this suggests moving beyond the numeral fragment of FinStoch to a setting supporting countable coproducts or standard Borel spaces, together with an observation mechanism compatible with those. A plausible approach is to replace projectivity with a more flexible consistency axiom (e.g. sampling by independent thinning of vertices, or by restricting a random measure to $[0, t]$), and to rebuild the universal category so that its deterministic core captures finite *partial* hypergraphs or finite configurations of hyperedges rather than full incidence tensors. Establishing a precise "universal semantics $\leftrightarrow$ sparse limit objects" correspondence would then require importing an appropriate sparse representation theorem for dissociated exchangeable k-arrays or exchangeable random measures in k dimensions.

7.2 Approximate and metric observation

We have assumed a faithful Bernoulli base $\Psi : \mathcal{C}_{\mathbb{N}} \hookrightarrow \mathsf{FinStoch}$, so that equality of closed numeral programs is witnessed by literal equality of finite distributions. This is appropriate when the equational theory is intended to capture exact observational equivalence. In applications, however, one often observes programs only approximately (finite sampling error, numerical rounding, or deliberate relaxation), and on the analytic side hypergraphons are usually considered modulo a metric such as an L^1-distance or a cut-type distance. It is therefore natural to ask for a quantitative variant in which programs are compared by an *observational pseudometric*.

One route is to replace the discrete equality induced by Ψ with a family of distances d_n on $\mathsf{FinStoch}(1, \mathbf{2}^{\binom{n}{k}})$, for instance total variation distance, and to define a contextual distance on closed terms by

$$d(t, t') := \sup_{n \geq 1} d_n\big(\Psi(\llbracket t \rrbracket), \Psi(\llbracket t' \rrbracket)\big),$$

or by restricting to a designated class of contexts. Categorically, this suggests enriching the numeral fragment of $\mathcal{C}$ over extended pseudometric spaces and requiring the Markov structure and distributive structure to be nonexpansive. Analytically, one would like to relate the induced program distance to a metric on hypergraphons; for $k = 2$ the relevant comparison is between finite subgraph distributions and the cut distance. For general k there are higher-order cut norms and associated metrics. A satisfactory statement would identify the metric completion of programs modulo distance 0 with hypergraphons modulo the corresponding metric equivalence, together with explicit continuity bounds for the passage from W to $(p_{W,n})_n$. We emphasize that such bounds are nontrivial: they mix combinatorial blowups (coming from $\binom{n}{k}$) with analytic control of convergence.

7.3 Computability and effective reconstruction

Our completeness direction is existential and uses a quotient-by-contextual-equivalence construction together with a representation theorem that is itself nonconstructive in general (due to measurability choices and quotienting by measure-preserving transformations). If one is interested in algorithmic extraction of limit objects, or in learnability from finite observations, an effective refinement is required.

A basic question is the following. Suppose we are given, for each n, a computable description of the finite distribution p_n (for example as rational probabilities in $\mathsf{FinStoch}$) satisfying the exchangeability/projectivity/k-locality constraints. Under what conditions can we compute, to any desired precision, a representative of the associated hypergraphon W (say as a step function on a dyadic partition of $[0,1]^{2^k-2}$)? Even ignoring measure-preserving nonuniqueness, such a procedure amounts to an effective inverse to the Aldous–Hoover/Kallenberg correspondence in the dissociated case. One expects to approximate W by finite models via regularity-type lemmas, but the quantitative bounds in hypergraph regularity are large and often ineffective in strong forms. It is therefore of interest to identify subclasses of programs (or of equational theories) giving rise to more tractable W (e.g. those definable by finite mixtures of simple kernels, or those arising from finite latent-variable models). On the categorical side, a related question is whether the quotient $\mathcal{U}_k/F_W$ can be presented effectively from finite data about F_W on numerals, and whether equivalence checking for a restricted fragment of programs is decidable or semidecidable.

7.4 Interacting with other effects and richer base types

The present interface is deliberately austere: the only source of randomness is new, and the only observable outputs are finite. A common extension in probabilistic programming is to add real-valued primitives, continuous distributions, and conditioning. Each of these interacts nontrivially with the structural assumptions underpinning dissociation.

If one adds a base type $\mathbb{R}$ and allows sampling of real random variables, then the appropriate observation category is no longer $\mathsf{FinStoch}$ but a category of Markov kernels on standard Borel spaces (or an alternative such as quasi-Borel spaces). One can still ask for a "Bernoulli base" on numerals, but it is no longer faithful on the full language; moreover, approximate observation becomes unavoidable. If one adds an operation $\mathsf{sample} : \mathbf{1} \to \mathbb{R}$ and allows $\mathsf{hyperedge}_k$ to depend deterministically on sampled reals, then the induced hypergraph model may acquire a *global* latent variable, violating dissociation unless additional axioms enforce that such global randomness is not shared across disjoint vertex blocks. Thus, to maintain a clean hypergraphon correspondence, one must either restrict the additional effects so that they are generated locally from vertices (for example, per-vertex reals sampled inside the scope of each new), or else generalize the target analytic object to include global mixing (leading to mixtures of dissociated models).

Conditioning introduces further complications because it typically breaks the affine/weakening structure and can destroy exchangeability unless handled with care. A precise account would likely require moving from Markov categories to a setting supporting disintegration and Bayesian inversion, and stating explicitly which equational principles remain sound. We regard this as compatible with the present approach but not formalized here: the universal-category method isolates the deterministic core and the free generators, and it should extend provided the additional structure admits a similar universal characterization and an adequate observational semantics.

7.5 Prospects for mechanization

Several points in the development are routine for a human reader but involve substantial bookkeeping, notably the index management in k-locality and the naturality/monoidality arguments in the universal semantics. These are prime candidates for mechanization. On the syntactic side, one may implement a normal form for nested let-expressions in a proof assistant and use rewriting tactics for the probabilistic let-laws (associativity, commutativity of independent lets, and weakening) together with deterministic substitution. The goal would be to make arguments such as "reorder the n vertex samplings" and "factor the computation over disjoint blocks" fully formal and reusable.

On the categorical side, mechanization requires a concrete encoding of Hyp_k, its opposite, and the finite-coproduct completion $\mathsf{Fam}(-)$, together with explicit combinatorial lemmas about $\binom{[n]}{k}$ under injections and disjoint unions. Lemma L5 (hom-set identification) is especially well-suited to this: it reduces to a finite combinatorial bijection that can be proved by explicit constructions. The more delicate step is the correspondence "monoidality $\Rightarrow$ k-locality", which hinges on understanding how disjoint vertex blocks induce a product decomposition of incidence tensors and how this interacts with the Markov tensor product. A successful mechanization here would not merely increase confidence; it would provide a library of reusable lemmas for other relational interfaces (directed, colored, multi-sorted) and would clarify which parts of the proof rely essentially on distributivity and which rely on specific properties of bool.

In summary, the dense hypergraphon correspondence should be viewed as a baseline classification theorem. Extending it to sparse regimes, quantitative observation, effective semantics, and richer effect combinations appears feasible, but each direction requires importing nontrivial additional structure—either analytic (new representation theorems and metrics) or categorical (new universal constructions and enriched notions of observation). The universal-category viewpoint is intended to keep these extensions modular: one modifies the interface and the observation principle, and the resulting semantic invariants should follow by the same pattern of "universal object + functors on a designated fragment + quotient."