THE UNIVERSITY OF CHICAGO

APPLICATIONS OF CONTINUOUS COMBINATORICS TO QUASIRANDOMNESS AND EXTREMAL COMBINATORICS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF COMPUTER SCIENCE
AND
DEPARTMENT OF MATHEMATICS

BY
LEONARDO NAGAMI COREGLIANO

CHICAGO, ILLINOIS
JUNE 2021

Table of Contents

LIST OF FIGURES . . . v

LIST OF TABLES . . . vi

ACKNOWLEDGMENTS . . . vii

ABSTRACT . . . viii

1 INTRODUCTION . . . 1

2 A REVIEW OF CONTINUOUS COMBINATORICS . . . 6
- 2.1 Model theory . . . 6
- 2.2 Densities and convergence . . . 11
- 2.3 Flag algebra: the syntax . . . 14
- 2.4 Theons: the semantics . . . 16
- 2.5 Limit object operators for open interpretations . . . 20
- 2.6 Uniqueness . . . 22

3 ADVANCES IN CONTINUOUS COMBINATORICS . . . 25
- 3.1 Rank and independence . . . 25
- 3.2 Theon lifting, couplings and amalgamations . . . 28
- 3.3 L_1-topology . . . 35
- 3.4 Rank function in density topology and in L_1-topology . . . 38
- 3.5 Low rank theories . . . 46
- 3.6 Strengthening theon lifting . . . 55
- 3.7 Amalgamations over more general diagrams . . . 66
- 3.8 Concluding remarks and open problems . . . 80

4 NATURAL QUASIRANDOMNESS . . . 84
- 4.1 Preliminaries . . . 93
  - 4.1.1 Quasirandomness properties . . . 93
  - 4.1.2 Useful theories and objects . . . 96
- 4.2 Main results on natural quasirandomness . . . 98
  - 4.2.1 Comparison to ad hoc quasirandomness theories . . . 102
- 4.3 Basic properties and the first equivalence . . . 105
- 4.4 Unique inducibility . . . 110
- 4.5 Unique coupleability . . . 124
- 4.6 Separations . . . 140
- 4.7 Top level quasirandomness . . . 151
- 4.8 Compatibility . . . 171
- 4.9 Concluding remarks and open problems . . . 175

5 ABSTRACT CHROMATIC NUMBER 179
5.1 Preliminaries 180
5.1.1 The general Turán density and the abstract chromatic number 180
5.1.2 Partite Ramsey numbers 182
5.1.3 Non-induced setting 187
5.2 Main results on abstract chromatic number 188
5.3 Abstract Turán densities from abstract chromatic number 191
5.4 Partite Ramsey numbers 192
5.5 Ramsey-based formula for the abstract chromatic number 197
5.6 The non-induced case 202
5.7 Applications to concrete theories 205
5.8 Concluding remarks and open problems 211

REFERENCES 213

List of Figures

3.1 Pictorial view of constructions of Definition 3.7.7. 73
3.2 Example of construction of shape $\widehat{\mathbf{S}}$ from $\mathbf{S}$, identity morphisms are omitted. . . 74
3.3 Commutative diagram of morphisms and amalgamations constructed in reduction to the case when $(D(*), (D(g_A))_{A\in\mathrm{Obj}(\mathbf{S})})$ is the colimit of D_1. 75
3.4 Commutative diagram of morphisms and amalgamations constructed in the case when A_0 is an isolated vertex of F. 77
3.5 Commutative diagram of morphisms and amalgamations constructed in the case when there is an oriented edge f in $\vec{F}$ from some A_1 to A_0. 79

4.1 Implications between quasirandomness properties. 89
4.2 Hasse diagram of quasirandomness hypergraph properties in arity k. 92

5.1 Pictorial view of a Q-uniform model 186

List of Tables

4.1 Proof locations for theorems of Section 4.2. 103

ACKNOWLEDGMENTS

I am grateful

to Prof. Alexander Razborov

for advising me very wisely; and

for the immense patience;

to Prof. Yoshiharu Kohayakawa

for introducing me to the field of asymptotic combinatorics;

to Irroko Nagami

for all the love and support;

to Valdemir Aparecido Coregliano

for all the love and support.

ABSTRACT

The theory of limits of dense combinatorial objects studies the asymptotic behavior of densities of small templates in an increasing sequence of combinatorial objects. The inaugural limit theory of graphons captures limits of dense graph sequences in a semantic/geometric limit object that can be thought of as a measurable fractional version of an adjacency matrix over $[0,1]$. Since the theory of graphons is specifically tailored to graphs, to study limits of other combinatorial objects, limit theories have been developed in a case-by-case basis. On the other hand, the theory of flag algebras explored a syntactic/algebraic approach to the subject, producing limit objects for general combinatorial objects (more specifically, for models of any universal first-order theory on a finite relational language). While the minimalist nature of the syntactic approach generates an elegant and clean theory, it has the drawback of losing the geometric intuition of the underlying objects. To address this issue, in a joint work with A. Razborov, we have developed the theory of theons, a semantic/geometric limit that works in the same general setting of universal theories. In this dissertation we review the theory of theons and apply these tools of limit theory to two different settings.

Our first application of limit theory, which uses both flag algebras and theons, is to quasirandomness. The existing theory of quasirandomness provides a plethora of quasirandomness properties with their rich implications and separations for several different combinatorial objects such as graphs, hypergraphs, permutations, tournaments, etc. However, such study of quasirandomness in the literature, much like in the early limit theory, has been made in case-by-case fashion for each type of combinatorial object. We develop a more general and systematic study of quasirandomness in the same setting of universal theories. Our main motivation is to study "natural" quasirandomness properties in the sense that they are preserved under local combinatorial constructions, which are captured by open interpretations. Our properties mainly revolve around the notion of couplings of limit objects, which are alignments of limit objects in the combined theory, and uniquely coupleable limit objects, which are limit objects such that every coupling is equivalent to the random coupling. We prove

several implications, separations and characterizations of our quasirandomness properties and we show the best possible separation between our properties and the quasirandomness properties of the literature.

Our second application of limit theory is a generalization of the celebrated Erdős–Stone–Simonovits Theorem and its generalization by Alon–Shikhelman that characterize the asymptotic behavior of the maximum density $\pi^t_{\mathcal{F}}$ of the t-clique K_t in a graph without non-induced copies of graphs in a family $\mathcal{F}$ in terms of the chromatic numbers of the graphs in $\mathcal{F}$. We show that these theorems extend to the general setting of any local combinatorial construction encoded by an open interpretation $I\colon T_{\mathrm{Graph}} \rightsquigarrow T$ in the sense that we can characterize the maximum density π^t_I of a t-clique K_t obtained from a limit graph interpreted from a limit object of T in terms of an abstract chromatic number $\chi(I)$. This in particular covers the case where the copies of graphs in $\mathcal{F}$ are instead assumed to be induced, and the case where we have graphs with extra structure (e.g., a linear order, a cyclic order, a coloring) and we want to maximize the density of t-cliques (with any structure) while forbidding some induced or non-induced family $\mathcal{F}$ of graphs with structure. We also show that if T is finitely axiomatizable (for the example of graphs with extra structure, this includes the case when the family $\mathcal{F}$ is finite), then $\chi(I)$ (and hence also π^t_I) is computable from the list of axioms of T and a description of I.

Part of this dissertation is based on a joint work with Alexander A. Razborov.

CHAPTER 1

INTRODUCTION

The theory of limits of discrete combinatorial objects has been thriving over the last couple of decades. The main thrust of the theory is that some properties of extremely large combinatorial objects can be encoded in a continuous limit that is susceptible to analytical and continuous tools. For this reason limit theory is also sometimes associated with the name continuous combinatorics. One of the first limit theories was that of graphons [54] (see also [53]), which encoded all properties captured by graph homomorphism densities by graphons, i.e., symmetric measurable functions $W\colon [0,1]^2 \to [0,1]$. The field was drastically expanded with the theory of flag algebras [59], where not only was it shown that by considering induced subgraph densities one greatly reduces the redundancy of the description allowing more concrete results to be proven, but also the theory was extended to capture "general" combinatorial objects (more specifically, for models of any universal first-order theory on a finite relational language). This uniform treatment of combinatorial objects of different nature also provided operators in the limit world corresponding to usual local combinatorial constructions such as graph of inversions of a permutation, 3-hypergraph of triangles of a graph, graph induced by the common neighborhood of two vertices (these are formally captured by open interpretations between different universal theories).

Other than the different levels of generality, the theories of graphons and flag algebras differ in a more fundamental level: the former is a geometric/semantic limit as a graphon can be thought of as a fractional graph over $[0,1]$ and the latter is an algebraic/syntactic limit as the flag algebra homomorphisms that encode limits are essentially just lists of sampling densities satisfying some polynomial relations. This means that proofs in the former theory often have a geometric intuition while proofs in the latter theory are often comprised of algebraic manipulations with almost no intuition at all. To address this issue, in a joint work with A. Razborov, we developed the theory of theons [24], a geometric/semantic limit theory in the same general setting of universal theories used by flag algebras. Other geometric/semantic limit theories had also been developed in an ad hoc manner for several particular objects such as digraphs [31], hypergraphs [34], permutations [44], interval graphs [30], etc. and for general universal theories in [4, 5]. However, the theory of theons also provides geometric/semantic limit world operators that capture local combinatorial constructions between different universal theories. In this dissertation, we use tools of continuous combinatorics to study two different applications: natural quasirandomness properties and the abstract chromatic number.

One of the main motivations of the theory of graphons [54] was the theory of graph quasirandomness introduced by Thomason [65] and Chung–Graham–Wilson [17]. The main discovery of the latter theory is that several properties that hold asymptotically almost surely for the sequence of Erdős–Rényi random graphs $(\boldsymbol{G_{n,p}})_{n\in\mathbb{N}}$ can be re-phrased as properties of a deterministic graph sequence $(G_n)_{n\in\mathbb{N}}$. Since then, the theory of quasirandomness has expanded not only within graph theory [16, 61, 62, 60, 71, 45, 20] but also towards studying quasirandomness for other combinatorial objects such as tournaments [15, 46, 25], permutations [22, 23, 49, 10] and hypergraphs [14, 18, 48, 47, 28, 51, 52, 66, 1], etc.

Just as in the case of early limit theory, the theory of quasirandomness has been studied so far in a case-by-case manner, with very few attempts at an intrinsic definition of quasirandomness. As the first application of continuous combinatorics in this dissertation, we initiate a more systematic study of quasirandom properties that can be reasonably identified as "intrinsic" in the sense that they can be formulated for arbitrary models of a universal theory (continuing the theme of generality of the theories of flag algebras [59] and theons [24]) and "natural" in the sense that they are preserved by local combinatorial constructions.

The main motivation behind this natural quasirandomness theory is that several of the quasirandomness properties of the literature can be stated in terms of properties of couplings between two limit objects (i.e., a limit object in the combined theory that "projects" to the previous two objects under the appropriate structure-erasing "projection"). As an example, the equivalence between properties P_1 and P_4 of [17] can be stated in graphon language as: (P_1) a graphon W is p-quasirandom (i.e., the non-induced labeled density

$t(G,W) \stackrel{\text{def}}{=} \int_{[0,1]^V} \prod_{ij\in E} W(x_i,x_j)\ dx$ of any subgraph $G=(V,E)$ with m edges is p^m) if and only if (P_4) in any red/blue coloring of its vertices, red edges have density $c^2 p$, where c is the density of red vertices and p is the edge density of W. Thus, graphon quasirandomness is also equivalent to having all labeled densities of red/blue colored graphs G in any red/blue coloring of W being $c^r(1-c)^b q$, where r and b are the number of red and blue vertices of G, respectively and q is the labeled density of the underlying graph of G in W. This example is paradigmatic of the notion of unique coupleability: it says that a graphon is quasirandom if and only if there is only one "coupling" of it with any given red/blue coloring, namely, any coupling has the same sampling densities as the random coupling. In this format, this unique coupleability property is perfectly generalizable to arbitrary theories and using the theory of theons we can show surprising theorems such as: if a limit object is uniquely coupleable with any red/blue coloring, then it is also uniquely coupleable with any linear order, with any permutation, and in fact with any rank 1 limit object.

Two of the most famous theorems in extremal graph theory are Turán's Theorem [67] characterizing the maximum number of edges in a graph without ℓ-cliques K_ℓ and Ramsey's Theorem [58] that says that for every ℓ, a large enough k-uniform hypergraph must either contain an ℓ-clique $K_\ell^{(k)}$ or an ℓ-independent set $\overline{K}_\ell^{(k)}$. The celebrated Erdős–Stone–Simonovits Theorem [38, 37] generalizes Turán's Theorem by characterizing the maximum asymptotic edge density when we instead forbid a family $\mathcal{F}$ of non-induced subgraphs in terms of the smallest chromatic number $\chi(\mathcal{F})$ of a graph in $\mathcal{F}$. In another direction, Erdős [35] generalized Turán's Theorem by characterizing the maximum number of t-cliques K_t in a graph without ℓ-cliques K_ℓ ($t < \ell$) and Alon–Shikhelman [2] provided an analogue of the Erdős–Stone–Simonovits Theorem that characterizes the maximum asymptotic density of K_t in a graph without any non-induced copies of graphs in a family $\mathcal{F}$ also in terms of $\chi(\mathcal{F})$ (Theorem 5.0.1).

A relatively new type of generalization of the Turán and Erdős–Stone–Simonovits theorems is to study maximization of the asymptotic edge density in graphs with extra structure while forbidding non-induced copies of some family $\mathcal{F}$. This has been done for ordered graphs [57], cyclically ordered graphs [9] and edge-ordered graphs [41] and in all these cases a theorem similar to the Erdős–Stone–Simonovits Theorem is proved in terms of a suitable generalization of the chromatic number (see also [64] for a survey). However, all these cases were done in an ad hoc fashion.

A uniform and general treatment of this problem was first done in [24, Examples 25 and 31]: in the general case, we want to maximize asymptotic edge density in a hereditary family of graphs with some extra structure. Note that even when restricted to the usual case of graphs without extra structure, this is already a generalization of the Erdős–Stone–Simonovits as the forbidden subgraphs are *induced*. This general setting is formally captured by using open interpretations $I\colon T_{\mathrm{Graph}} \rightsquigarrow T$ that provides a combinatorial construction that produces a graph $I(M)$ from a model of a universal theory T; the problem then consists of maximizing the asymptotic edge density of $I(M)$ over all possible choices of M as the size of M goes to infinity for a given fixed I. For example, the aforementioned setting of (cyclically) ordered graphs are captured using the construction I that simply "forgets" the (cyclic) order.

In [24, Example 31], it was shown that in this general setting a result analogous to the Erdős–Stone–Simonovits Theorem still holds for an appropriately defined *abstract chromatic number* $\chi(I)$. However, the formula for $\chi(I)$ presented in [24, Equation (16)] is considerably abstract and it was left open if $\chi(I)$ was (algorithmically) computable even when T is assumed to be finitely axiomatizable. As our second application of continuous combinatorics, we further extend this result giving an analogue of the Alon–Shikhelman Theorem in the general setting of an open interpretation $I\colon T_{\mathrm{Graph}} \rightsquigarrow T$ by characterizing the maximum asymptotic density of K_t in terms of the abstract chromatic number $\chi(I)$ and we provide an alternative, more concrete formula for the abstract chromatic number $\chi(I)$. Such formula allows us to deduce that when T is *finitely* axiomatizable, then $\chi(I)$ is (algorithmically) computable from a list of the axioms of T and a description of I. Our alternative formula is based on a partite version of Ramsey's Theorem for universal theories that informally says that given $\ell, m \in \mathbb{N}$, there exists $n \in \mathbb{N}$ such that for every model M and every partition of M into ℓ parts all of size at least n must have a "uniform" submodel on the same partition with all parts of size m (this version of Ramsey's Theorem for disjoint unions of theories of hypergraphs follows from [42, Section 5] and the non-partite version, when $\ell = 1$, for general theories follows from the general Ramsey theory for systems of [55]). By using these different formulas for $\chi(I)$, we can retrieve the results of [57, 9, 41] on ordered graphs, cyclically ordered graphs and edge-ordered graphs, respectively from the general theory.

The dissertation is organized as follows. In Chapter 2, we review the concepts and results of continuous combinatorics that are used on our diverse applications. In Chapter 3, we present new results of continuous combinatorics that were developed while studying natural quasirandomness properties and that are crucial to some of its proofs. Chapter 4 contains the application of continuous combinatorics to natural quasirandomness. In Chapter 5, we present the results related to the abstract chromatic number. Except for Chapter 2, the last section of each chapter contains some concluding remarks and open problems related to the chapter's topics.

CHAPTER 2

A REVIEW OF CONTINUOUS COMBINATORICS

Throughout this dissertation, we let $\mathbb{N} \stackrel{\text{def}}{=} \{0, 1, 2, \ldots\}$ denote the set of non-negative integers and let $\mathbb{N}_+ \stackrel{\text{def}}{=} \mathbb{N} \setminus \{0\}$. For $n, \ell \in \mathbb{N}$, we let $[n] \stackrel{\text{def}}{=} \{1, \ldots, n\}$ and $(n)_\ell \stackrel{\text{def}}{=} n(n-1)\cdots(n-\ell+1)$ and for a set V, we let $2^V \stackrel{\text{def}}{=} \{A \subseteq V\}$ be the collection of all its subsets, we let $\binom{V}{\ell} \stackrel{\text{def}}{=} \{A \subseteq V \mid |A| = \ell\}$ be the set of all subsets of V of size ℓ and we let $\binom{V}{>\ell} \stackrel{\text{def}}{=} \{A \subseteq V \mid |A| > \ell\}$. For $V \subseteq \mathbb{N}$ and $A \subseteq V$, let $\iota_{A,V} \colon [|A|] \rightarrowtail V$ be the injection that enumerates A in increasing order; we will abuse notation and omit V from the notation (as ι_A) when it is clear from context. We further let $r(V) \stackrel{\text{def}}{=} \bigcup_{\ell \in \mathbb{N}_+} \binom{V}{\ell}$ be the set of all finite non-empty subsets of V and we let $r(V, \ell) \stackrel{\text{def}}{=} \bigcup_{t \in [\ell]} \binom{V}{t}$. The usage of the arrow $\rightarrowtail$ for a function will always presume that the function is injective. We further let $(V)_\ell$ be the set of all injective functions of the form $\alpha \colon [\ell] \rightarrowtail V$ and we let S_V be the group of all bijections $\alpha \colon V \rightarrowtail V$. We will frequently abuse notation by identifying $[n]$ with n in notation similar to the above, e.g., we will use $r(n, \ell)$ as a shorthand for $r([n], \ell)$. Random variables will always be typed in $\boldsymbol{math\ bold\ face}$. For two random variables with values in the same σ-algebra, $\boldsymbol{X} \sim \boldsymbol{Y}$ will mean that $\boldsymbol{X}$ and $\boldsymbol{Y}$ are equidistributed.

2.1 Model theory

As we noted in the introduction, our preferred way of seeing a "general" combinatorial object is based on notions of first-order logic and model theory. In this section, we review some of the concepts of these theories needed for our application (we refer the reader to [11, 7] for a more comprehensive introduction to the topic).

A (first-order) *relational language*[1] is a set $\mathcal{L}$ of *predicate symbols*. Each predicate symbol $P \in \mathcal{L}$ comes along with a positive integer $k(P) \in \mathbb{N}_+$ called its *arity* and designates the

1. Languages are sometimes also called *signatures* or *vocabularies*, and in general (when non-relational) may contain also constant and/or function symbols.

number of variables that P depends on. All of our languages will be assumed to finite first-order relational languages. Given our restrictions on the language $\mathcal{L}$ (no constant or function symbols), *atomic formulas* may only have the form $P(x_{i_1}, \ldots, x_{i_{k(P)}})$ or $x_{i_1} = x_{i_2}$ (we do allow equality), and *open formulas*[2] are made from atomic formulas using standard logical connectives $\neg, \vee, \wedge, \rightarrow, \leftrightarrow$, etc., but not quantifiers. A *universal formula* is a formula of the form $\forall x_1 \cdots \forall x_n F(x_1, \ldots, x_n)$, where F is an open formula.

A *universal (first-order) theory* T in a relational language $\mathcal{L}$ is a set of universal formulas called *axioms*; universal quantifiers in front of the axioms are usually omitted. A *structure* K in a relational language $\mathcal{L}$ consists of a *vertex set*[3] $V(K)$ and a mapping that assigns to every $P \in \mathcal{L}$ a $k(P)$-ary relation $R_P(K) \subseteq V(K)^{k(P)}$; the *size* of K is denoted by $|K| \stackrel{\text{def}}{=} |V(K)|$. A structure K is a *model* of a theory T in the language $\mathcal{L}$, denoted by $K \vDash T$, if all axioms of T are universally true on M (see any textbook in mathematical logic for a formal definition).

As usual, an *embedding* of a structure K_1 in $\mathcal{L}$ in a structure K_2 in $\mathcal{L}$ is an injective function $f\colon V(K_1) \rightarrowtail V(K_2)$ that respects the relations of $\mathcal{L}$, that is, we have

$$\alpha \in R_P(K_1) \iff f \circ \alpha \in R_P(K_2)$$

for every $P \in \mathcal{L}$ and every $\alpha \in V(K_1)^{k(P)}$. A *positive embedding* of K_1 in K_2 is an injective function $f\colon V(K_1) \rightarrowtail V(K_2)$ that is only required to preserve relations but not non-relations, that is, we have

$$\alpha \in R_P(K_1) \implies f \circ \alpha \in R_P(K_2).$$

An *isomorphism* between K_1 and K_2 is a bijective embedding of K_1 in K_2 and when one such isomorphism exists, we say K_1 and K_2 are *isomorphic* (denoted $K_1 \cong K_2$). For a set V, we let $\mathcal{K}_V[T]$ be the set of all models K of T with $V(K) = V$ and for $n \in \mathbb{N}$, we let

2. Sometimes also called quantifier-free formulas.

3. Sometimes also called *universe* or *domain of discourse*.

$\mathcal{M}_n[T] \stackrel{\text{def}}{=} \mathcal{K}_n[T]/\cong$ be the set of all models of T of size n up to isomorphism (we will typically think of elements of $\mathcal{M}_n[T]$ as models that are representatives of their isomorphism classes). We also let $\mathcal{M}[T] \stackrel{\text{def}}{=} \bigcup_{n\in\mathbb{N}} \mathcal{M}_n[T]$ be the set of all finite models of T up to isomorphism. We can think of $\mathcal{K}_V[T]$ as models of T labeled by V. For $n = |V|$ and $K \in \mathcal{K}_V[T]$, we denote by $[K] \in \mathcal{M}_n[T]$ the isomorphism type of K.

Given $n \in \mathbb{N}$ and a model $K \in \mathcal{K}_n[T]$, the *open diagram* of K is the open formula $D_{\text{open}}(K)(x_1,\ldots,x_n)$ given by

$$\left(\bigwedge_{1\leq i<j\leq n} x_i \neq x_j\right) \wedge \left(\bigwedge_{\substack{P\in\mathcal{L}\\ \alpha\in R_P(K)}} P(x_{\alpha_1},\ldots,x_{\alpha_n})\right) \wedge \left(\bigwedge_{\substack{P\in\mathcal{L}\\ \alpha\in V(K)^{k(P)}\setminus R_P(K)}} \neg P(x_{\alpha_1},\ldots,x_{\alpha_n})\right).$$

Under this definition, it follows that $\alpha\colon V(K) \to V(K')$ is an embedding of K in K' if and only if $K' \vDash D_{\text{open}}(K)(\alpha_1,\ldots,\alpha_n)$, that is, if and only if $(\alpha_1,\ldots,\alpha_n)$ satisfies $D_{\text{open}}(K)$ in K'.

Given a structure K in $\mathcal{L}$ and a set $V \subseteq V(K)$, the *substructure of K induced by V*, denoted $K|_V$, is defined by $V(K|_V) \stackrel{\text{def}}{=} V$ and $R_P(K|_V) \stackrel{\text{def}}{=} R_P(K) \cap V^{k(P)}$ for every $P \in \mathcal{L}$. A property that characterizes universal theories is that they are precisely the first-order theories whose class of models is closed under taking induced substructures (and thus we use the name *(induced) submodel*); this property is key for the sampling definition of densities of continuous combinatorics. We say that T *proves* or *entails* a formula F, denoted by $T \vdash F$, if it does so in first-order logic. Using the completeness theorem [11, Theorem 1.3.21] and induced submodel property above, it follows that if T is a universal theory and F is an open formula, then $T \vdash \forall\vec{x}, F(\vec{x})$ is true if and only if the formula $\forall\vec{x}, F(\vec{x})$ is true in every *finite* model of T. A universal theory T is called *degenerate* if $\mathcal{M}_n[T] = \varnothing$ for some $n \in \mathbb{N}$; by the submodel property above, this is equivalent to requiring that T does not have infinite model.

Given two relational languages $\mathcal{L}_1$ and $\mathcal{L}_2$, a translation from $\mathcal{L}_1$ to $\mathcal{L}_2$ is a map I that maps each $P \in \mathcal{L}_1$ to an open formula $I(P)$ in $\mathcal{L}_2$. The map I is extended to open formulas by declaring that it commutes with logical connectives. An *open interpretation* from a universal theory T_1 in $\mathcal{L}_1$ to a universal theory T_2 in $\mathcal{L}_2$ is a translation from $\mathcal{L}_1$ to $\mathcal{L}_2$ such that for each axiom $\forall \vec{x}, F(\vec{x})$ of T_1, we have $T_2 \vdash \forall \vec{x}, I(F)(\vec{x})$; we denote such open interpretations as $I\colon T_1 \rightsquigarrow T_2$. Translations I from $\mathcal{L}_1$ to $\mathcal{L}_2$ give a natural way of constructing a structure $I(K)$ in $\mathcal{L}_1$ from a structure K in $\mathcal{L}_2$ by letting $V(I(K)) \stackrel{\text{def}}{=} V(K)$ and for every $P \in \mathcal{L}_1$, letting

$$R_P(I(K)) \stackrel{\text{def}}{=} \{(v_1, \ldots, v_{k(P)}) \in V(K)^{k(P)} \mid K \vDash I(P)(v_1, \ldots, v_{k(P)})\}$$

be the set of all tuples $(v_1, \ldots, v_{k(P)})$ of vertices of K that satisfy $I(P)(v_1, \ldots, v_{k(P)})$ in K. Using the completeness theorem [11, Theorem 1.3.21] of first-order logic, it follows that for universal theories T_1, T_2 in $\mathcal{L}_1, \mathcal{L}_2$, respectively and a translation I from $\mathcal{L}_1$ to $\mathcal{L}_2$, I is an open interpretation from T_1 to T_2 if and only if for every finite model K of T_2, $I(K)$ is a model of T_1 (i.e., $I(\mathcal{M}[T_2]) \subseteq I(\mathcal{M}[T_1])$). We denote the identity interpretation of a theory T by $\mathrm{id}_T\colon T \rightsquigarrow T$.

Given two universal theories T_1, T_2 in relational languages $\mathcal{L}_1, \mathcal{L}_2$, respectively, their *disjoint union* is the theory $T_1 \cup T_2$ on the disjoint union $\mathcal{L}_1 \mathbin{\dot\cup} \mathcal{L}_2$ of the languages whose axioms are those of T_1 (about symbols in $\mathcal{L}_1$) and of T_2 (about symbols of in $\mathcal{L}_2$). The two most important types of open interpretations are the *structure-erasing* interpretations, which are open interpretations of the form $I\colon T_1 \rightsquigarrow T_1 \cup T_2$ that act identically on the language of T_1 and *axiom-adding* interpretations, which are open interpretations of the form $I\colon T_1 \rightsquigarrow T_2$ when T_2 is obtained from T_1 by adding axioms and I acts identically on the language of T_1. Given two open interpretations $I\colon T_1 \rightsquigarrow T_3$ and $J\colon T_2 \rightsquigarrow T_4$, we denote by $I \cup J\colon T_1 \cup T_2 \rightsquigarrow T_3 \cup T_4$ the amalgamation interpretation that acts as I on T_1 and acts as

J on T_2.

Two open interpretations $I, J\colon T_1 \rightsquigarrow T_2$ are *equivalent* if for every $P \in \mathcal{L}_1$, we have $T_2 \vdash \forall \vec{x}, (I(P)(\vec{x}) \leftrightarrow J(P)(\vec{x}))$ (or, equivalently, if they define the same maps $\mathcal{K}_V[T_2] \to \mathcal{K}_V[T_1]$ for each finite set V). The category whose objects are universal theories on finite relational languages and whose morphisms are open interpretations up to equivalence is denoted $\mathbf{Int}$. Under this definition, every open interpretation $I\colon T_1 \rightsquigarrow T_2$ can be decomposed as $I = J \circ A \circ S$, where $S\colon T_1 \rightsquigarrow T_1 \cup T_2$ is structure-erasing, $A\colon T_1 \cup T_2 \rightsquigarrow T$ is axiom-adding for the theory T obtained from $T_1 \cup T_2$ by adding the axioms

$$\forall \vec{x}, (P(\vec{x}) \leftrightarrow I(P)(\vec{x}))$$

for every P in the language of T_1 and $J\colon T \rightsquigarrow T_2$ is the isomorphism of $\mathbf{Int}$ that acts identically in the language of T_2 and acts as I in the language of T_1 (the inverse of J acts identically on the predicate symbols of T_2, see [24, Remark 2]).

A universal theory T in a relational language $\mathcal{L}$ is *canonical* if for every $P \in \mathcal{L}$, the theory T entails

$$\left(\bigvee_{1 \leq i < j \leq k(P)} x_i = x_j\right) \to \neg P(x_1, \ldots, x_{k(P)}). \tag{2.1}$$

Since every universal theory is isomorphic in $\mathbf{Int}$ to a canonical theory (see [24, Theorem 2.3]), we will assume that all of our theories are canonical.

The *pure canonical theory* in $\mathcal{L}$, denoted $T_{\mathcal{L}}$ is the theory whose axioms are precisely (2.1) for each $P \in \mathcal{L}$. A structure in $\mathcal{L}$ is canonical if it satisfies (2.1) for every $P \in \mathcal{L}$ (equivalently, if it is a model of $T_{\mathcal{L}}$). Since we will only be working with canonical theories, all of our structures will also be assumed to be canonical.

Other important examples of canonical theories include the theory of k-hypergraphs $T_{k\text{-Hypergraph}}$, whose language contains a single predicate symbol E of arity $k(E) \stackrel{\text{def}}{=} k$ and whose axioms are (2.1) for $P = E$ and

$$\forall \vec{x}, (E(x_1, \ldots, x_k) \to E(x_{\sigma(1)}, \ldots, x_{\sigma(k)})) \qquad (\sigma \in S_k); \tag{2.2}$$

the theory of (simple) graphs $T_{\text{Graph}} \stackrel{\text{def}}{=} T_{2\text{-Hypergraph}}$; the theory of (strict) linear orders T_{LinOrder}, whose language contains a single binary predicate symbol $\prec$ and whose axioms are

$$\forall x, \neg(x \prec x);$$
$$\forall \vec{x}, (x_1 \neq x_2 \to (x_1 \prec x_2 \lor x_2 \prec x_1));$$
$$\forall \vec{x}, (x_1 \prec x_2 \land x_2 \prec x_3 \to x_1 \prec x_3);$$

and the theory of c-colorings $T_{c\text{-Coloring}}$, whose language contains c unary predicate symbols $\chi_1, \ldots, \chi_c$ and whose axioms are

$$\forall x, \neg\chi_i(x) \lor \neg\chi_j(x) \qquad (1 \leq i < j \leq c);$$
$$\forall x, \bigvee_{i \in [c]} \chi_i(x).$$

Note that $T_{2\text{-Coloring}}$ and $T_{1\text{-Hypergraph}}$ are isomorphic in **Int**.

2.2 Densities and convergence

In this section, we justify the name "limit theory" by defining the notion of convergence studied by it. This notion of convergence is based on sampling densities of submodels of a large model.

Given models M and N of the same canonical theory T with $|M| \leq |N|$, the *(unlabeled* *induced) density* of M in N is

$$p(M,N) \stackrel{\text{def}}{=} \frac{\left|\left\{V \in \binom{V(N)}{|M|} \;\middle|\; N|_V \cong M\right\}\right|}{\binom{|N|}{|M|}},$$

that is, it is the normalized number of submodels of N that are isomorphic to M. The *labeled (induced) density* of M in N is

$$t_{\text{ind}}(M,N) \stackrel{\text{def}}{=} \frac{|\{\alpha\colon V(M) \rightarrowtail V(N) \mid \alpha \text{ embedding of } M \text{ in } N\}|}{(|N|)_{|M|}},$$

that is, it is the normalized number of embeddings of M in N. The *labeled non-induced density* of M in N is

$$t_{\text{inj}}(M,N) \stackrel{\text{def}}{=} \frac{|\{\alpha\colon V(M) \rightarrowtail V(N) \mid \alpha \text{ positive embedding of } M \text{ in } N\}|}{(|N|)_{|M|}},$$

that is, it is the normalized number of positive embeddings of M in N.

It is easy to see that these quantities are related via the formulas

$$t_{\text{ind}}(M,N) = \frac{|\text{Aut}(M)|}{|M|!} \cdot p(M,N), \tag{2.3}$$

$$t_{\text{inj}}(M,N) = \sum_{\substack{K \in \mathcal{K}_{V(M)}[T] \\ \forall P \in \mathcal{L}, R_P(M) \subseteq R_P(K)}} t_{\text{ind}}(K,N), \tag{2.4}$$

where $\text{Aut}(M)$ denotes the *automorphism* group of M (i.e., the set of all isomorphisms between M and itself).

A sequence $(N_n)_{n\in\mathbb{N}}$ of models of a canonical theory T is *convergent* if $|N_n| < |N_{n+1}|$ for every $n \in \mathbb{N}$ and for every finite model M of T, the sequence of densities $(p(M,N_n))_{n\in\mathbb{N}}$ is convergent (this is equivalent to requiring that $(t_{\text{ind}}(M,N_n))_{n\in\mathbb{N}}$ is convergent for every $M \in \mathcal{M}[T]$ and also equivalent to requiring that $(t_{\text{inj}}(M,N_n))_{n\in\mathbb{N}}$ is convergent for every $M \in \mathcal{M}[T]$, see [24, Proposition 2.5]). Since $\mathcal{M}[T]$ is countable, this notion of convergence is (pre-)compact: if $(N_n)_{n\in\mathbb{N}}$ is a sequence of models of T with $|N_n| < |N_{n+1}|$ for every $n \in \mathbb{N}$, then a simple diagonalization argument shows that $(N_n)_{n\in\mathbb{N}}$ has a convergent subsequence.

As we mentioned in the introduction, we will use the theories of flag algebras [59] and theons [24] that provide limit objects to these convergent sequences from which the limit densities $\lim_{n\to\infty} p(M, N_n)$ can be extracted. The goal of the next two sections is to provide the definitions necessary for the following theorem on the equivalence of these different encodings of convergent sequences.

Theorem 2.2.1 ([54, 59], [24, Theorem 6.3], see also [24, Sct. 7])**.** *Let Ω be an atomless standard probability space and consider the following objects for a canonical theory T.*

i. A convergent sequence $(N_n)_{n\in\mathbb{N}}$ of models of T.

ii. A positive homomorphism $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ from the flag algebra $\mathcal{A}[T]$.

iii. A (weak or strong) T-on $\mathcal{N}$ over Ω.

iv. A local exchangeable array $\boldsymbol{K}$ supported on models of T.

The objects above are cryptomorphic in the sense that given an instance of one of them, one can "explicitly" construct instances of the others that satisfy the following for every $m \in \mathbb{N}$ and every $K \in \mathcal{K}_m[T]$:

$$\lim_{n\to\infty} p(K, N_n) = \phi(K) = \phi_{\mathcal{N}}(K) = \mathbb{P}[\boldsymbol{K}|_{[m]} \cong K];$$
$$\lim_{n\to\infty} t_{\mathrm{ind}}(K, N_n) = \phi(\langle K\rangle) = t_{\mathrm{ind}}(K, \mathcal{N}) = \mathbb{P}[\boldsymbol{K}|_{[m]} = K].$$

Theorems as the above are known in limit theory as two-sided[4] existence theorems.

4. One-sided would be if only constructions in one of the directions were provided.

2.3 Flag algebra: the syntax

In one sentence, the theory of flag algebras can be summarized as the study of relations that the coordinates of $\phi \in [0,1]^{\mathcal{M}[T]}$ must satisfy if $\phi(M) \stackrel{\text{def}}{=} \lim_{n\to\infty} p(M, N_n)$ for some convergent sequence $(N_n)_{n\in\mathbb{N}}$ for its own sake, without any explicit references to the actual limit object. In this section we present the fraction of the theory of flag algebras used in our applications[5]; we refer the interested reader to [59] for a more thorough treatment.

Given models M_1, M_2, N of T with $|M_1| + |M_2| \leq |N|$, the *joint density* of M_1 and M_2 in N is defined as

$$p(M_1, M_2; N) \stackrel{\text{def}}{=} \frac{\left|\left\{(V_1, V_2) \in \binom{V(N)}{|M_1|}\times\binom{V(N)}{|M_2|} \;\middle|\; V_1 \cap V_2 = \varnothing \land N|_{V_1} \cong M_1 \land N|_{V_2} \cong M_2\right\}\right|}{\binom{|N|}{|M_1|}\binom{|N|-|M_1|}{|M_2|}},$$

that is, it is the probability that picking disjoint subsets V_1 and V_2 of $V(N)$ of sizes $|M_1|$ and $|M_2|$, respectively, yields submodels of N isomorphic to M_1 and M_2, respectively.

Let $\mathcal{A}[T]$ be the quotient of the space $\mathbb{R}\mathcal{M}[T]$ of all formal $\mathbb{R}$-linear combinations of finite models of T by the linear subspace generated by elements of the form

$$M - \sum_{M'\in\mathcal{M}_\ell[T]} p(M, M')M'$$

for $M \in \mathcal{M}[T]$ and $\ell \geq |M|$.

Lemma 2.3.1 ([59, Lemma 2.4])**.** *The bilinear mapping $\mathcal{A}[T] \times \mathcal{A}[T] \to \mathcal{A}[T]$ defined by*

$$M_1 \cdot M_2 \stackrel{\text{def}}{=} \sum_{N\in\mathcal{M}_n[T]} p(M_1, M_2; N)N,$$

for $M_1, M_2 \in \mathcal{M}[T]$ and $n \geq |M_1| + |M_2|$ is well-defined and turns $\mathcal{A}[T]$ into an $\mathbb{R}$-algebra

5. The main omissions are the fact that we work only with the 0 type of flag algebras (i.e., there is no partial labeling of models), we skip all material related to homomorphism extensions (this essentially corresponds to random partial labelings) and differential methods.

whose multiplicative identity is (the equivalence class of) $1 \stackrel{\text{def}}{=} \sum_{M \in \mathcal{M}_n[T]} M$ *for any* $n \in \mathbb{N}$.

Furthermore, the theory T *is non-degenerate if and only if the algebra* $\mathcal{A}[T]$ *is not the zero algebra*[6].

The algebra obtained from the lemma above is called the *flag algebra* of the theory T.

Recall that $\operatorname{Hom}(\mathcal{A}[T], \mathbb{R})$ denotes the set of all $\mathbb{R}$-algebra homomorphisms from $\mathcal{A}[T]$ to $\mathbb{R}$ (i.e., maps that preserve the operations and the multiplicative identity). A *positive homomorphism* is a homomorphism $\phi \in \operatorname{Hom}(\mathcal{A}[T], \mathbb{R})$ such that $\phi(M) \geq 0$ for every $M \in \mathcal{M}[T]$ (since $1 = \sum_{M \in \mathcal{M}_n[T]} M$, this in particular implies $\phi(M) \in [0, 1]$); we denote the set of positive homomorphisms by $\operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$. The intuition behind the definition of positive homomorphisms is that the value $\phi(M)$ is the limiting value $\lim_{n \to \infty} p(M, N_n)$ of the density M in some convergent sequence $(N_n)_{n \in \mathbb{N}}$. We typically think of $\operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ as a subset of $[0, 1]^{\mathcal{M}[T]}$; this allows us to equip $\operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ with the *density topology*, which is the topology induced from the product topology of $[0, 1]^{\mathcal{M}[T]}$. With this and (2.3) in mind, for $M \in \mathcal{M}[T]$, we let

$$\langle M \rangle \stackrel{\text{def}}{=} \frac{|\operatorname{Aut}(M)|}{|M|!} \cdot M$$

denote the element of $\mathcal{A}[T]$ that encodes the labeled (induced) density of M.

The main theorem of flag algebras is the cryptomorphism between items (i) and (ii) of Theorem 2.2.1, that says that $\operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ precisely captures the set of limits of convergent sequences[7].

---

6. In [59, Lemma 2.4], it is shown only that if T is non-degenerate, then $\mathcal{A}[T]$ is not the zero algebra, but the converse is obvious as if $\mathcal{M}_n[T]$ is empty, then $1 = \sum_{n \in \mathcal{M}_n[T]} M = 0$.

7. Again, in [59, Theorem 3.3], only the non-degenerate case is considered, but the degenerate case holds trivially since then there are no convergent sequences and the algebra $\mathcal{A}[T]$ is the zero algebra thus $\operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is empty.

2.4 Theons: the semantics

In this section, we present the geometric/semantic limit theory of theons. In the same way that the definition of model of a theory progresses through associating relations to predicate symbols to form structures then requiring the structure to satisfy axioms, the definition of theons can be seen as associating peons to predicate symbols to form Euclidean structures then requiring Euclidean structures to satisfy axioms.

Given an atomless standard probability space $\Omega = (X, \mathcal{A}, \mu)$, a set V and $\ell \in \mathbb{N}$, let $\mathcal{E}_V(\Omega) \stackrel{\text{def}}{=} X^{r(V)}$ and $\mathcal{E}_{V,\ell}(\Omega) \stackrel{\text{def}}{=} X^{r(V,\ell)}$, we equip these sets with the completion[8] of the product measure of $|r(V)|$ and $|r(V,\ell)|$ copies of the measure μ, respectively. By abuse of notation, we will also denote this completion measure by μ. We will further abuse the notation and denote simply by $[0,1]^t$ the space $([0,1]^t, \mathcal{B}_t, \lambda)$, where $\mathcal{B}_t$ is the Borel σ-algebra and λ is the (t-dimensional) Lebesgue measure. When Ω is the space $[0,1]$, we will omit it from the notation. Furthermore, for atomless standard probability spaces Ω and Ω', we let $\Omega \times \Omega'$ be their product and we will abuse notation by identifying the spaces $\mathcal{E}_V(\Omega \times \Omega')$ and $\mathcal{E}_V(\Omega) \times \mathcal{E}_V(\Omega')$ via the correspondence $\mathcal{E}_V(\Omega \times \Omega') \ni x \leftrightarrow (y,z) \in \mathcal{E}_V(\Omega) \times \mathcal{E}_V(\Omega')$ given by $y_A \stackrel{\text{def}}{=} (x_A)_1$ and $z_A \stackrel{\text{def}}{=} (x_A)_2$ $(A \in r(V))$.

The *diagonal* sets are defined as $\mathcal{D}_V(\Omega) \stackrel{\text{def}}{=} \{x \in \mathcal{E}_V(\Omega) \mid \exists v, w \in V, (v \neq w \land x_{\{v\}} = x_{\{w\}})\}$ and $\mathcal{D}_{V,\ell}(\Omega) \stackrel{\text{def}}{=} \{x \in \mathcal{E}_{V,\ell}(\Omega) \mid \exists v, w \in V, (v \neq w \land x_{\{v\}} = x_{\{w\}})\}$, i.e., they are the sets of points that have some repetition in coordinates indexed by singletons. An injective function $\alpha\colon [k] \rightarrowtail V$ defines natural projections $X^{\binom{V}{\ell}} \to X^{\binom{[k]}{\ell}}$ given by the formula $\alpha^*(x)_A \stackrel{\text{def}}{=} x_{\alpha(A)}$ $(A \in \binom{[k]}{\ell})$. By abuse, we also use the same notation α^* for the projections $\mathcal{E}_V(\Omega) \to \mathcal{E}_k(\Omega)$ and $\mathcal{E}_{V,\ell}(\Omega) \to \mathcal{E}_{k,\ell}(\Omega)$ defined by the same formula (but with A ranging either in $r(k)$ or $r(k,\ell)$).

8. In [24, Sct. 7] we carefully considered the difference between equipping these sets with the product σ-algebra or with its completion, cf. the discussion in [53, page 218]. It was needed to differentiate between weak theons (satisfying the axioms a.e.) and strong ones (satisfying them everywhere off-diagonal) and how constructively can one go from the former type to the latter via removal lemmas. As we prefer to avoid dwelling into these issues here, we make the simplifying assumption of completeness once and for all.

Fix an atomless standard probability space $\Omega = (X, \mathcal{A}, \mu)$ and a finite relational language $\mathcal{L}$. For a predicate symbol $P \in \mathcal{L}$, a *P-on* over Ω is a measurable subset of $\mathcal{E}_{k(P)}(\Omega)$. An *Euclidean structure* in $\mathcal{L}$ over Ω is a function $\mathcal{N}$ that maps each predicate symbol $P \in \mathcal{L}$ to a P-on $\mathcal{N}_P$ over Ω. Given an Euclidean structure over Ω, a finite set V and a structure $K \in \mathcal{K}_V[T_{\mathcal{L}}]$, define the following measurable subsets of $\mathcal{E}_{V(K)}(\Omega)$:

$$T_{\text{inj}}(K, \mathcal{N}) \stackrel{\text{def}}{=} \bigcap_{P \in \mathcal{L}} \bigcap_{\alpha \in R_P(K)} (\alpha^*)^{-1}(\mathcal{N}_P);$$

$$T_{\text{ind}}(K, \mathcal{N}) \stackrel{\text{def}}{=} T_{\text{inj}}(K, \mathcal{N}) \cap \bigcap_{P \in \mathcal{L}} \bigcap_{\alpha \in (V(K))_{k(P)} \setminus R_P(K)} (\alpha^*)^{-1}(\mathcal{E}_{k(P)} \setminus \mathcal{N}_P).$$

If we interpret elements of $\mathcal{E}_k(\Omega)$ as "limit k-tuples", then $T_{\text{inj}}(K, \mathcal{N})$ is the set of all "limit $|V|$-tuples" that are positive embeddings of K in $\mathcal{N}$ and the set $T_{\text{ind}}(K, \mathcal{N})$ is the set of all "limit $|V|$-tuples" that are embeddings of K in $\mathcal{N}$. This and (2.3) motivate the following density definitions:

$$t_{\text{inj}}(K, \mathcal{N}) \stackrel{\text{def}}{=} \mu(T_{\text{inj}}(K, \mathcal{N}));$$

$$t_{\text{ind}}(K, \mathcal{N}) \stackrel{\text{def}}{=} \mu(T_{\text{ind}}(K, \mathcal{N}));$$

$$\phi_{\mathcal{N}}(K) \stackrel{\text{def}}{=} \frac{|V(K)|!}{|\text{Aut}(K)|} t_{\text{ind}}(K, \mathcal{N}).$$

For a canonical theory T in $\mathcal{L}$, a *(weak) T-on* over Ω is an Euclidean structure $\mathcal{N}$ in $\mathcal{L}$ over Ω such that for every $M \in \mathcal{M}[T_{\mathcal{L}}] \setminus \mathcal{M}[T]$, we have $\phi_{\mathcal{N}}(M) = 0$ (equivalently, $t_{\text{ind}}(M, \mathcal{N}) = 0$), that is, every structure that is not a model of T has zero density in $\mathcal{N}$. A *strong T-on* over Ω is an Euclidean structure $\mathcal{N}$ in $\mathcal{L}$ over Ω such that for every $M \in \mathcal{M}[T_{\mathcal{L}}] \setminus \mathcal{M}[T]$, we have $T_{\text{ind}}(M, \mathcal{N}) \subseteq \mathcal{D}_{V(M)}(\Omega)$, that is, $\mathcal{N}$ does not have any copy of non-models of T, except in the diagonal. We use the names peons and theons for P-ons and T-ons when referring to these objects for generic P or T.

In this language, a k-hypergraphon of [34] is simply a strong $T_{k\text{-Hypergraph}}$-on and there is a (not one-to-one) correspondence between graphons W of [54] and T_{Graph}-ons $\mathcal{N}$ that preserves densities given by

$$\begin{aligned} W &\mapsto \{x \in \mathcal{E}_2 \mid x_{\{1,2\}} < W(x_{\{1\}}, x_{\{2\}})\} \\ W_{\mathcal{N}} &\mapsfrom \mathcal{N}, \end{aligned}$$

where

$$W_{\mathcal{N}}(x_{\{1\}}, x_{\{2\}}) \stackrel{\text{def}}{=} \lambda(\{x_{\{1,2\}} \mid (x_{\{1\}}, x_{\{2\}}, x_{\{1,2\}}) \in \mathcal{N}\}). \tag{2.5}$$

Note that the definition of theons, weak or strong, are ensuring that Euclidean structures satisfy the axioms of the theory in an indirect way, by requiring that its "submodels" are all models of the theory. To state the equivalent formulation as Euclidean structures that satisfy the axioms of the theory we need a couple more definitions.

For an open formula $F(x_1, \ldots, x_n)$ and an equivalence relation $\approx$ on $[n]$ with m equivalence classes, we let $F_{\approx}(y_1, \ldots, y_m) \stackrel{\text{def}}{=} F(y_{\nu_1}, \ldots, y_{\nu_n})$, where ν_i is the equivalence class of $i \in [n]$. A canonical theory T is called *substitutionally closed* if for every axiom $\forall x F(x_1, \ldots, x_n)$ and every equivalence relation $\approx$ on $[n]$, T proves $\forall \vec{y}, F_{\approx}(\vec{y})$ using only propositional rules and possibly renaming variables in its axioms (but substitutions of the same variable for two different variables are disallowed). It is important to note that this is a technical property of the axiomatization of T that can easily be obtained by adding all theorems of T to its axioms (and this preserves the class of models of T).

For an Euclidean structure $\mathcal{N}$ in $\mathcal{L}$ over Ω and an open formula $F(x_1, \ldots, x_n)$, the *truth set* of F in $\mathcal{N}$ is the set $T(F, \mathcal{N}) \subseteq \mathcal{E}_n(\Omega)$ defined inductively as follows.

i. If F is $P(x_{i_1}, \ldots, x_{i_k})$ and $i_1, \ldots, i_k$ are not pairwise distinct, or F is $x_i = x_j$ with $i \neq j$, then $T(F, \mathcal{N}) \stackrel{\text{def}}{=} \varnothing$.

ii. $T(x_i = x_i, \mathcal{N}) \stackrel{\text{def}}{=} \mathcal{E}_n(\Omega)$.

iii. If F is $P(x_{i_1}, \ldots, x_{i_k})$ and $i_1, \ldots, i_k$ are pairwise distinct, then $T(F, \mathcal{N}) \stackrel{\text{def}}{=} (i^*)^{-1}(\mathcal{N}_P)$, where i is viewed as a function $i\colon [k] \rightarrowtail [n]$.

iv. $T(-, \mathcal{N})$ commutes with logical connectives (e.g., we have $T(F_1 \vee F_2, \mathcal{N}) \stackrel{\text{def}}{=} T(F_1, \mathcal{N}) \cup T(F_2, \mathcal{N})$).

Theorem 2.4.1 ([24, Theorem 3.7], see also [24, Sct. 7])**.** *Let $\Omega = (X, \mathcal{A}, \mu)$ be an atomless standard probability space, let T be a substitutionally closed canonical theory in a language $\mathcal{L}$ and let $\mathcal{N}$ be an Euclidean structure in $\mathcal{L}$ over Ω. Then $\mathcal{N}$ is a weak [strong] T-on if and only if for every axiom $\forall \vec{x}, F(x_1, \ldots, x_n)$ of T, we have $\mu(T(F, \mathcal{N})) = 1$ [$T(F, \mathcal{N}) \supseteq \mathcal{E}_n(\Omega) \setminus \mathcal{D}_n(\Omega)$, respectively].*

Naturally, the main theorem of the theory of theons is the addition of weak and strong theons to the list of objects of Theorem 2.2.1 that are cryptomorphic to convergent sequences. The particular cryptomorphism between strong and weak theons is given by the Induced Euclidean Removal Lemma [24, Theorem 3.3] that says that any weak theon can be turned into a strong theon by changing only a zero measure set of its peons. In Chapter 3, we will prove a stronger version of this theorem, the Euclidean Robustness Lemma, Theorem 3.6.6 (see also Lemma 3.6.5).

The other cryptomorphism is actually proved by adding another intermediate object to the list of cryptomorphisms: local exchangeable arrays. This connection was first explored for the case of (di)graphs in [31].

Note first that there is a natural (left) action of $S_{\mathbb{N}_+}$ on $\mathcal{K}_{\mathbb{N}_+}[T_{\mathcal{L}}]$ given by $R_P(\sigma \cdot K) \stackrel{\text{def}}{=} \{\sigma \circ \alpha \mid \alpha \in R_P(K)\}$ for every $\sigma \in S_{\mathbb{N}_+}$, every $K \in \mathcal{K}_{\mathbb{N}_+}[T_{\mathcal{L}}]$ and every $P \in \mathcal{L}$. An *exchangeable array* in $\mathcal{L}$ is a random variable $\boldsymbol{K}$ with values in $\mathcal{K}_{\mathbb{N}_+}[T_{\mathcal{L}}]$ whose distribution is invariant under the action of $S_{\mathbb{N}_+}$. The exchangeable array $\boldsymbol{K}$ is called *local* if the marginals $\boldsymbol{K}|_U$ and $\boldsymbol{K}|_V$ are independent whenever $U, V \subseteq \mathbb{N}_+$ are disjoint.

One of the (easy) directions of the cryptomorphism of Theorem 2.2.1 will be of particular importance to us, namely, how to construct a local exchangeable array $\boldsymbol{K}$ from a given T-on

$\mathcal{N}$ over $\Omega = (X, \mathcal{A}, \mu)$. Intuitively, the only thing we have to do is to independently sample countably many points from our theon. Formally, let $\boldsymbol{\theta} = (\boldsymbol{\theta}_A)_{A \in r(\mathbb{N}_+)}$ be picked in $\mathcal{E}_{\mathbb{N}_+}(\Omega)$ according to μ. The *exchangeable array* $\boldsymbol{K}$ *corresponding to* $\mathcal{N}$ *with respect to* $\boldsymbol{\theta}$ is defined by

$$V(\boldsymbol{K}) \stackrel{\text{def}}{=} \mathbb{N}_+, \qquad R_P(\boldsymbol{K}) \stackrel{\text{def}}{=} \{\alpha \in (\mathbb{N}_+)_{k(P)} \mid \alpha^*(\boldsymbol{\theta}) \in \mathcal{N}_P\}. \tag{2.6}$$

and we have $\phi_{\mathcal{N}}(M) = \mathbb{P}[\boldsymbol{K}|_{[|M|]} \cong M]$ for every $M \in \mathcal{M}[T]$.

2.5 Limit object operators for open interpretations

As we have seen in Section 2.1, open interpretations between canonical theories $I \colon T_1 \rightsquigarrow T_2$ give rise to maps $\mathcal{K}_V[T_2] \to \mathcal{K}_V[T_1]$ that correspond to local combinatorial constructions. It is easy to see that the notion of convergence of Section 2.2 is preserved under these constructions, namely, if $(N_n)_{n\in\mathbb{N}}$ is a convergent sequence of models of T_2, then $(I(N_n))_{n\in\mathbb{N}}$ is a convergent sequence of models of T_1. The goal of this section is to provide the definitions of the corresponding operators for flag algebra homomorphisms, theons and exchangeable arrays.

For exchangeable arrays, the answer is trivial: if $\boldsymbol{K}$ corresponds to $(N_n)_{n\in\mathbb{N}}$ under Theorem 2.2.1, then $I(\boldsymbol{K})$ corresponds to $(I(N_n))_{n\in\mathbb{N}}$ under the same theorem. For theons, the corresponding operator is obtained via truth sets: for a T_2-on $\mathcal{N}$ corresponding to $(N_n)_{n\in\mathbb{N}}$, it is straightforward to check that the T_1-on $I(\mathcal{N})$ defined by $I(\mathcal{N})_P \stackrel{\text{def}}{=} T(I(P), \mathcal{N})$ corresponds to $(I(N_n))_{n\in\mathbb{N}}$.

For the case of flag algebra homomorphisms, we will recall a more general version that will also be needed for different purposes. Given canonical theories T_1, T_2 in $\mathcal{L}_1, \mathcal{L}_2$, respectively, a *conditional open interpretation* from T_1 to T_2 is a pair (U, I) (denoted $(U, I) \colon T_1 \rightsquigarrow T_2$), where I is a translation from $\mathcal{L}_1$ to $\mathcal{L}_2$ and U is an open formula in $\mathcal{L}_2$ with one free variable such that for each axiom $\forall x_1 \cdots \forall x_n, F(x_1, \ldots, x_n)$ of T_1, we have $T_2 \vdash \forall x_1 \cdots \forall x_n, (U(x_1) \land \cdots \land U(x_n) \to I(F)(x_1, \ldots, x_n))$. Clearly, a conditional open interpretation when $U(x)$ is a tautology, say $x = x$, is simply an open interpretation[9]. A *U-model* of T_2 is a model M that satisfies $\forall x, U(x)$ and we let $\mathcal{M}_n^U[T_2] \subseteq \mathcal{M}_n[T_2]$ be the set of all U-models of T_2 of size n up to isomorphism.

Theorem 2.5.1 ([59, Theorem 2.6])**.** *Let $(U, I)\colon T_1 \rightsquigarrow T_2$ be a conditional open interpretation, let*

$$u \stackrel{\text{def}}{=} \sum_{M \in \mathcal{M}_1^U[T_2]} M$$

and let $\mathcal{A}_u[T_2]$ be the localization of the algebra $\mathcal{A}[T_2]$ with respect to the multiplicative system $\{u^n \mid n \in \mathbb{N}\}$. Then the map $\pi^{(U,I)}\colon \mathcal{A}[T_1] \to \mathcal{A}_u[T_2]$ defined by

$$\pi^{(U,I)}(M) \stackrel{\text{def}}{=} \frac{\sum\{M' \in \mathcal{M}_{|M|}^U[T_2] \mid I(M') \cong M\}}{u^{|M|}}$$

is well-defined and is an algebra homomorphism[10].

As a corollary of this theorem, note that if $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ is such that $\phi(u) > 0$ (which in particular implies u is not nilpotent, thus $\mathcal{A}_u[T_2]$ is not the zero algebra), then it naturally extends to a homomorphism from $\mathcal{A}_u[T_2]$ to $\mathbb{R}$ as $\phi(f/u^n) \stackrel{\text{def}}{=} \phi(f)/\phi(u)^n$, thus $\phi \circ \pi^{(U,I)} \in \operatorname{Hom}^+(\mathcal{A}[T_1], \mathbb{R})$ (the non-negativity of $(\phi \circ \pi^{(U,I)})(M)$ for $M \in \mathcal{M}[T_1]$ is obvious).

When U is a tautology (i.e., when $I\colon T_1 \rightsquigarrow T_2$ is an open interpretation), then $u = 1$ so $\mathcal{A}_u[T_2] = \mathcal{A}[T_2]$. In this case, we denote $\pi^{(U,I)}$ simply by π^I and we abbreviate $\phi^I \stackrel{\text{def}}{=} \phi \circ \pi^I$. This is precisely the flag-algebraic operator that respects Theorem 2.2.1: if ϕ corresponds to $(N_n)_{n \in \mathbb{N}}$ under this theorem, then ϕ^I corresponds to $(I(N_n))_{n \in \mathbb{N}}$ under the same theorem.

9. In [59], conditional open interpretations are simply called open interpretations and U is instantiated to a tautology when the non-conditional version is used, but since we rarely use this more general form, we elected add the adjective "conditional" to this more general version.

10. In [59, Theorem 2.6], there is the extra hypothesis that u is not a zero divisor, but this hypothesis is only used to ensure that $\mathcal{A}_u[T_2]$ is not the zero algebra, which although not necessary for this theorem, is necessary for most of its applications.

In the case of conditional open interpretations, the mapping $\phi \mapsto \phi \circ \pi^{(U,I)}$ informally corresponds to applying the local combinatorial construction only to the "submodel" of ϕ induced by "vertices" that satisfy U and renormalizing the densities (this is precisely the role of the localization).

2.6 Uniqueness

After existence theorems, the other type of theorem that is of utmost importance in limit theory are uniqueness theorems. More specifically, these describe necessary and sufficient conditions for two limit objects to represent the same limit. In the case of flag algebra homomorphisms, the answer is trivial: if $\phi_1, \phi_2 \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ satisfy $\lim_{n\to\infty} p(M, N_n) = \phi_i(M)$ ($M \in \mathcal{M}[T]$, $i \in [2]$), then clearly $\phi_1 = \phi_2$. For local exchangeable arrays, the answer is also easy: if $\boldsymbol{K_1}, \boldsymbol{K_2}$ satisfy $\mathbb{P}[\boldsymbol{K_i}|_{[m]} \cong M] = \lim_{n\to\infty} p(M, N_n)$ ($m \in \mathbb{N}$, $M \in \mathcal{M}_m[T]$, $i \in [2]$), then clearly $\boldsymbol{K_1}$ and $\boldsymbol{K_2}$ have the same distribution ($\boldsymbol{K_1} \sim \boldsymbol{K_2}$).

For theons, however, the uniqueness theorem is much more complicated and technical as it needs to handle examples such as the fact that for $p \in [0, 1]$, the $T_{3\text{-Hypergraph}}$-ons $\mathcal{N}^p$ and $\mathcal{H}^p$ defined by

$$\mathcal{N}^p_E \stackrel{\text{def}}{=} \{x \in \mathcal{E}_3 \mid x_{[3]} < p\},$$
$$\mathcal{H}^p_E \stackrel{\text{def}}{=} \left\{x \in \mathcal{E}_3 \;\middle|\; \left(\sum_{A\in r(3)} x_A\right) \bmod 1 < p\right\}$$

represent the same limit, namely, the quasirandom 3-hypergraphon of density p.

To formally state the theon uniqueness theorem, we will first need some more definitions.

For an atomless standard probability space $\Omega = (X, \mathcal{A}, \mu)$ and a set V, there is a natural (right) action of S_V on $\mathcal{E}_V(\Omega)$ given by $(x \cdot \sigma)_A \stackrel{\text{def}}{=} x_{\sigma(A)}$ ($x \in \mathcal{E}_V(\Omega)$, $\sigma \in S_V$, $A \in r(V)$). If we are further given another atomless standard probability space $\Omega' = (X', \mathcal{A}', \mu')$ and a function $f\colon \mathcal{E}_k(\Omega) \to \Omega'$ ($k \in \mathbb{N}$), we say that f is *symmetric* if f is S_k-invariant and we say that f is *measure-preserving on highest order argument* (h.o.a.) if f is measurable and for every $x \in \mathcal{E}_{k,k-1}(\Omega)$, the restriction $f(x,-)\colon \Omega \to \Omega'$ (where we identify $X^{\binom{[k]}{k}}$ with X) is measure-preserving. Given a family $f = (f_1, \ldots, f_k)$ of symmetric functions with $f_d\colon \mathcal{E}_d(\Omega) \to \Omega'$, we define a family of functions $\widehat{f} = (\widehat{f}_1, \ldots, \widehat{f}_k)$ with $\widehat{f}_d\colon \mathcal{E}_d(\Omega) \to \mathcal{E}_d(\Omega')$ given by

$$\widehat{f}_d(x)_A \stackrel{\text{def}}{=} f_{|A|}(\iota_A^*(x))$$

(recall that $\iota_A\colon [|A|] \rightarrowtail [d]$ enumerates A in increasing order). Note that $\widehat{f}_d$ is S_d-equivariant.

Theorem 2.6.1 ([24, Theorems 3.9 and 3.11, Proposition 7.7])**.** *Let Ω and Ω' be atomless standard probability spaces, let $k \in \mathbb{N}_+$, let T be a canonical theory in a language $\mathcal{L}$ with $k(P) \leq k$ for every $P \in \mathcal{L}$ and let $\mathcal{N}$ and $\mathcal{N}'$ be T-ons over Ω and Ω', respectively. The following are equivalent.*

i. *We have $\phi_{\mathcal{N}} = \phi_{\mathcal{N}'}$.*

ii. *There exist families $f = (f_1, \ldots, f_k)$ and $g = (g_1, \ldots, g_k)$ of symmetric functions measure-preserving on h.o.a. ($f_d\colon \mathcal{E}_d \to \Omega$ and $g_d\colon \mathcal{E}_d \to \Omega'$) such that*

$$\widehat{f}_{k(P)}(x) \in \mathcal{N}_P \iff \widehat{g}_{k(P)}(x) \in \mathcal{N}'_P$$

for every $P \in \mathcal{L}$ and a.e. $x \in \mathcal{E}_{k(P)}$.

iii. *There exists a family $h = (h_1, \ldots, h_k)$ of symmetric functions measure-preserving on h.o.a. ($h_d\colon \mathcal{E}_d(\Omega' \times \Omega') \to \Omega$) such that*

$$\widehat{h}_{k(P)}(x, y) \in \mathcal{N}_P \iff x \in \mathcal{N}'_P$$

for every $P \in \mathcal{L}$ and a.e. $(x, y) \in \mathcal{E}_{k(P)}(\Omega') \times \mathcal{E}_{k(P)}(\Omega')$.

Note that in item (iii) we are using the aforementioned identification between $\mathcal{E}_{k(P)}(\Omega' \times \Omega')$ and $\mathcal{E}_{k(P)}(\Omega') \times \mathcal{E}_{k(P)}(\Omega')$.

CHAPTER 3

ADVANCES IN CONTINUOUS COMBINATORICS

In this chapter, we present some continuous combinatorics results that were obtained while studying natural quasirandomness. The results of Sections 3.1, 3.2 and 3.3 will be sufficient for all of our needs in Chapter 4 regarding natural quasirandomness. Sections 3.4, 3.5, 3.6 and 3.7 are devoted to a further exploration of the notions and results of the previous sections; these can be safely skipped if the reader is only interested in the applications of Chapters 4 and 5.

3.1 Rank and independence

We start by introducing the notion of rank of a limit object, which can be seen as a notion of complexity of it, and the dual notion of independence. These notions will play a key role in Chapter 4.

Definition 3.1.1 (rank and independence)**.** The *rank* of a peon $\mathcal{N} \subseteq \mathcal{E}_k(\Omega)$ over $\Omega = (X, \mathcal{A}, \mu)$, denoted $\operatorname{rk}(\mathcal{N})$, is the minimum $r \in \mathbb{N}$ such that $\mathcal{N}$ can be written as $\mathcal{N} = \mathcal{H} \times X^{\binom{[k]}{>r}}$ for some $\mathcal{H} \subseteq \mathcal{E}_{k,r}(\Omega)$. The *rank* of an Euclidean structure $\mathcal{N}$ is the maximum rank $\operatorname{rk}(\mathcal{N})$ of its peons.

Dually, for $\ell \in \mathbb{N}$, a peon $\mathcal{N} \subseteq \mathcal{E}_k(\Omega)$ is called *ℓ-independent* if it can be written as $\mathcal{N} = \mathcal{E}_{k,\ell}(\Omega) \times \mathcal{H}$ for some $\mathcal{H} \subseteq X^{\binom{[k]}{>\ell}}$ and an Euclidean structure is called *ℓ-independent* if all of its peons are ℓ-independent.

Given $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$, the *rank* of ϕ, denoted $\operatorname{rk}(\phi)$, is the minimum rank of a T-on $\mathcal{N}$ such that $\phi_{\mathcal{N}} = \phi$. Dually, we say $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is *ℓ-independent* if there exists an ℓ-independent T-on $\mathcal{N}$ such that $\phi_{\mathcal{N}} = \phi$. We will refer to this property as `Independence`$[\ell]$.

It is important to note that these definitions require only the *existence of some* geometric realization of the limit object with the required properties. As an example, the T_{Graph}-ons

$$\begin{aligned} \mathcal{G} &\stackrel{\text{def}}{=} \left\{ (x_{\{1\}}, x_{\{2\}}, x_{\{1,2\}}) \;\middle|\; x_{\{1,2\}} \leq p \right\}; \\ \mathcal{G}' &\stackrel{\text{def}}{=} \left\{ (x_{\{1\}}, x_{\{2\}}, x_{\{1,2\}}) \;\middle|\; (x_{\{1\}} + x_{\{2\}} + x_{\{1,2\}}) \bmod 1 \leq p \right\}; \end{aligned} \tag{3.1}$$

both represent the quasirandom limit of graphs of density p, but the second one is far from being 1-independent. The next proposition says that for rank, the situation is precisely the opposite: any representation of a low rank limit object must be of low rank except for a zero-measure change.

Proposition 3.1.2. *For every peon $\mathcal{N} \subseteq \mathcal{E}_k(\Omega)$ there exists another peon $\mathcal{H} \subseteq \mathcal{E}_k(\Omega)$ such that $\operatorname{rk}(\mathcal{H}) = \operatorname{rk}(\phi_{\mathcal{N}})$ and $\mathcal{H} = \mathcal{N}$ a.e. Moreover, if $\mathcal{N}$ is ℓ-independent for some $\ell \leq k$, then $\mathcal{H}$ can be taken to also be ℓ-independent.*

Proof. Let μ be the measure of Ω and X be its underlying space, let $r \stackrel{\text{def}}{=} \operatorname{rk}(\phi_{\mathcal{N}})$ and define the function $W\colon \mathcal{E}_{k,r}(\Omega) \to [0,1]$ by

$$W(x) \stackrel{\text{def}}{=} \mu(\{y \in X^{\binom{[k]}{>r}} \mid (x,y) \in \mathcal{N}\}), \tag{3.2}$$

defining it arbitrarily when this set is not measurable. Fubini's Theorem ensures that this function is measurable so we define

$$\mathcal{H} \stackrel{\text{def}}{=} W^{-1}(1) \times X^{\binom{[k]}{>r}}.$$

Clearly $\operatorname{rk}(\mathcal{H}) \leq r$. Hence, to prove that $\mathcal{H} = \mathcal{N}$ a.e., it is sufficient to show that W is 0-1 valued a.e.

Since $\operatorname{rk}(\phi_{\mathcal{N}}) = r$, we know that there exists a peon $\mathcal{G}$ over some space $\Omega' = (X', \mathcal{A}', \mu')$ such that $\phi_{\mathcal{G}} = \phi_{\mathcal{N}}$ and $\operatorname{rk}(\mathcal{G}) = r$. By theon uniqueness, Theorem 2.6.1, there exist sequences $f = (f_d)_{d=1}^k, g = (g_d)_{d=1}^k$ of symmetric measure preserving on h.o.a. functions ($f_d\colon \mathcal{E}_d \to \Omega$

and $g_d\colon \mathcal{E}_d \to \Omega')$ such that

$$\widehat{f}_k(z) \in \mathcal{N} \iff \widehat{g}_k(z) \in \mathcal{G} \tag{3.3}$$

for almost every $z \in \mathcal{E}_k$. From the structure of the function $\widehat{f}_k$, we can decompose it as

$$\widehat{f}_k(x, y) = (F_1(x), F_2(x, y)),$$

for every $(x, y) \in \mathcal{E}_{k,r} \times [0,1]^{\binom{[k]}{>r}}$, where $F_1\colon \mathcal{E}_{k,r} \to \mathcal{E}_{k,r}(\Omega)$ and $F_2\colon \mathcal{E}_k \to X^{\binom{[k]}{>r}}$ are given by

$$F_1(x)_A \stackrel{\text{def}}{=} f_{|A|}(\iota_A^*(x)), \qquad\qquad F_2(x, y)_A \stackrel{\text{def}}{=} f_{|A|}(\iota_A^*(x, y)).$$

We perform a similar decomposition of $\widehat{g}_k$ in terms of functions $G_1\colon \mathcal{E}_{k,r} \to \mathcal{E}_{k,r}(\Omega')$ and $G_2\colon \mathcal{E}_k \to (X')^{\binom{[k]}{>r}}$.

Since the functions f_d are measure preserving on h.o.a., it follows that F_1 is measure preserving and for every $x \in \mathcal{E}_{k,r}$ the restriction $F_2(x, -)\colon [0,1]^{\binom{[k]}{>r}} \to X^{\binom{[k]}{>r}}$ is measure preserving. Hence Fubini's Theorem applied to (3.3) implies

$$W(F_1(x)) = \lambda(\{y \in [0,1]^{\binom{[k]}{>r}} \mid (G_1(x), G_2(x, y)) \in \mathcal{G}\})$$

for almost every $x \in \mathcal{E}_{k,r}$. But since $\mathrm{rk}(\mathcal{G}) = r$, the measure above can only be 0 or 1 (as $G_2(x, y)$ contains only coordinates with $|A| > r$). Since F_1 is measure preserving, this implies that $W(z) \in \{0, 1\}$ for almost every $z \in \mathcal{E}_{k,r}(\Omega)$ and thus $\mathcal{H} = \mathcal{N}$ a.e.

We have already shown that $\mathrm{rk}(\mathcal{H}) \leq r$ and since $\mathcal{H} = \mathcal{N}$ a.e. implies $\phi_{\mathcal{H}} = \phi_{\mathcal{N}}$, the other inequality must also hold.

The last statement is obvious from the construction. ∎

Remark 1. Note that the proof of Proposition 3.1.2 above in particular implies that a theon

$\mathcal{N}$ over $\Omega = (X, \mathcal{A}, \mu)$ satisfies $\operatorname{rk}(\phi_{\mathcal{N}}) \leq k$ if and only if for each peon $\mathcal{N}_P$, the function

$$W_P(x) \stackrel{\text{def}}{=} \mu(\{y \in X^{\binom{[k]}{>r}} \mid (x, y) \in \mathcal{N}_P\}) \tag{3.4}$$

(defined arbitrarily when the set is not measurable) is 0-1 valued a.e.

3.2 Theon lifting, couplings and amalgamations

We now proceed to an application of theon uniqueness, Theorem 2.6.1, that allows us to lift theons through interpretations. As we have seen in Section 2.5, given an open interpretation $I\colon T_1 \rightsquigarrow T_2$ and a T_2-on $\mathcal{H}$, the T_1-on $I(\mathcal{H})$ represents the limit object constructed from $\phi_{\mathcal{H}}$ via I, i.e., we have $\phi_{I(\mathcal{H})} = \phi_{\mathcal{H}}^I$. However, the next example adapted from [24, Example 45] shows that given a T_1-on $\mathcal{N}$ and $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ such that $\phi^I = \phi_{\mathcal{N}}$, it is *not* true that there exists a T_2-on $\mathcal{H}$ such that both $I(\mathcal{H}) = \mathcal{N}$ a.e. and $\phi_{\mathcal{H}} = \phi$.

Example 1. Consider the $(T_{2\text{-Coloring}} \cup T_{\text{LinOrder}})$-on $\mathcal{G}$ over $[0,1]^2$ given by

$$\begin{aligned}
\mathcal{G}_{\prec} &\stackrel{\text{def}}{=} \{(x, y) \in \mathcal{E}_2 \times \mathcal{E}_2 \mid x_{\{1\}} < x_{\{2\}}\};\\
\mathcal{G}_{\chi 1} &\stackrel{\text{def}}{=} \{(x, y) \in \mathcal{E}_1 \times \mathcal{E}_1 \mid y_{\{1\}} < 1/2\};\\
\mathcal{G}_{\chi 2} &\stackrel{\text{def}}{=} \{(x, y) \in \mathcal{E}_1 \times \mathcal{E}_1 \mid y_{\{1\}} \geq 1/2\};
\end{aligned}$$

and let $I\colon T_{\text{LinOrder}} \rightsquigarrow T_{2\text{-Coloring}} \cup T_{\text{LinOrder}}$ be the structure-erasing interpretation. It is clear that $\phi_{\mathcal{H}}^I$ is the unique element of $\operatorname{Hom}^+(\mathcal{A}[T_{\text{LinOrder}}], \mathbb{R})$, which is also represented by the T_{LinOrder}-on $\mathcal{N}$ over $[0,1]$ given by

$$\mathcal{N}_{\prec} \stackrel{\text{def}}{=} \{x \in \mathcal{E}_2 \mid x_{\{1\}} < x_{\{2\}}\}.$$

However, there does not exist a $(T_{2\text{-Coloring}} \cup T_{\text{LinOrder}})$-on $\mathcal{H}$ such that both $\phi_{\mathcal{H}} = \phi_{\mathcal{G}}$ and $I(\mathcal{H})_{\prec} = \mathcal{N}_{\prec}$ a.e. Indeed, if such $\mathcal{H}$ existed then $\mathcal{H}_{\chi 1}$ would have to be a measurable subset of $[0,1]$ with Lebesgue density $1/2$ in every interval, contradicting the Lebesgue Density Theorem (see [8, I-5.6(ii)] and [56, Theorem 3.21]).

The next proposition says in essence that Example 1 is the worst that can happen: at the cost of adding an extra dummy variable, we can find an $\mathcal{H}$ such that $I(\mathcal{H})_P = \mathcal{N}_P \times \mathcal{E}_{k(P)}$ a.e. and $\phi_{\mathcal{H}} = \phi$.

Proposition 3.2.1. *Let $I\colon T_1 \rightsquigarrow T_2$, let $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$, and let $\mathcal{N}$ be a T_1-on over Ω such that $\phi^I = \phi_{\mathcal{N}}$. Then there exists a T_2-on $\mathcal{H}$ over $\Omega \times \Omega$ such that $\phi_{\mathcal{H}} = \phi$ and $I(\mathcal{H})_P = \mathcal{N}_P \times \mathcal{E}_{k(P)}(\Omega)$ a.e., for every predicate symbol P in the language of T_1.*

Furthermore, if $T_2 = T_1 \cup T'$ for some T' and I is the structure-erasing interpretation, then $\mathcal{H}$ can be taken to satisfy $I(\mathcal{H})_P \stackrel{\text{def}}{=} \mathcal{H}_P = \mathcal{N}_P \times \mathcal{E}_{k(P)}(\Omega)$ everywhere for every predicate symbol P in the language of T_1.

Proof. For $i \in [2]$, let $\mathcal{L}_i$ be the language of T_i and let $k_i \stackrel{\text{def}}{=} \max_{P \in \mathcal{L}_i} k(P)$. Let $\mathcal{G}$ be a T_2-on over Ω such that $\phi_{\mathcal{G}} = \phi$. Since $\phi_{I(\mathcal{G})} = \phi^I = \phi_{\mathcal{N}}$, by theon uniqueness, Theorem 2.6.1, there exists a sequence $h = (h_d)_{d=1}^{k_1}$ of symmetric measure preserving on h.o.a. functions $(h_d\colon \mathcal{E}_d(\Omega) \times \mathcal{E}_d(\Omega) \to \Omega)$ such that

$$\widehat{h}_{k(P)}(x, \widehat{x}) \in I(\mathcal{G})_P \iff x \in \mathcal{N}_P, \tag{3.5}$$

for every $P \in \mathcal{L}_1$ and almost every $(x, \widehat{x}) \in \mathcal{E}_{k(P)}(\Omega) \times \mathcal{E}_{k(P)}(\Omega)$. Extend the family h by defining $h_d\colon \mathcal{E}_d(\Omega) \times \mathcal{E}_d(\Omega) \to \Omega$ for $k_1 < d \leq \max\{k_1, k_2\}$ as $h_d(x, \widehat{x}) \stackrel{\text{def}}{=} x_{[d]}$, and note that h_d is symmetric and measure preserving on h.o.a.

Define then the T_2-on $\mathcal{H}$ over $\Omega \times \Omega$ by

$$\mathcal{H}_Q \stackrel{\text{def}}{=} \widehat{h}_{k(Q)}^{-1}(\mathcal{G}_Q) \tag{3.6}$$

for every $Q \in \mathcal{L}_2$. By (the easy direction of) theon uniqueness, Theorem 2.6.1, it follows that

$\phi_{\mathcal{H}} = \phi_{\mathcal{G}} = \phi$. On the other hand, the definition of $\mathcal{H}$ ensures that

$$(x, \widehat{x}) \in I(\mathcal{H})_P \iff \widehat{h}_{k(P)}(x, \widehat{x}) \in I(\mathcal{G})_P$$

for every $P \in \mathcal{L}_1$ and every $(x, \widehat{x}) \in \mathcal{E}_{k(P)}(\Omega) \times \mathcal{E}_{k(P)}(\Omega)$, which together with (3.5) implies $I(\mathcal{H})_P = \mathcal{N}_P \times \mathcal{E}_{k(P)}(\Omega)$ a.e.

For the case when $T_2 = T_1 \cup T'$ for some T' and I is the structure-erasing interpretation, we define $\mathcal{H}$ instead by using (3.6) only for $Q \in \mathcal{L}_2 \setminus \mathcal{L}_1$ and use $\mathcal{H}_P \stackrel{\text{def}}{=} \mathcal{N}_P \times \mathcal{E}_{k(P)}(\Omega)$ for every $P \in \mathcal{L}_1$ (as required). By (3.5) this is only a zero-measure change from the previous definition so we still have $\phi_{\mathcal{H}} = \phi$. ■

This proposition is fundamental to prove an amalgamation property of limits. Recall from [24, Sct. 2.2] that the category **Int** has pushouts (otherwise known as amalgamated sums, fibred coproducts, etc.). More concretely, for open interpretations $I_1 \colon T \rightsquigarrow T_1$ and $I_2 \colon T \rightsquigarrow T_2$, a pushout of (I_1, I_2) is given by the theory T' obtained from $T_1 \cup T_2$ by adding the axioms

$$\forall \vec{x}, (I_1(P)(\vec{x}) \leftrightarrow I_2(P)(\vec{x})) \tag{3.7}$$

for every P in the language of T and the open interpretations $J_i \colon T_i \rightsquigarrow T'$ ($i \in [2]$) that act identically on the language of T_i so that

$$\begin{array}{ccc} T & \xrightarrow{I_1} & T_1 \\ {\scriptstyle I_2}\downarrow & & \downarrow{\scriptstyle J_1} \\ T_2 & \xrightarrow[J_2]{} & T' \end{array} \tag{3.8}$$

is commutative and has the standard universality property (see Proposition 3.7.1 for the more general case of finite colimits).

One case of the pushout above that will be of particular importance in Chapter 4 is when the theory T is trivial, in which case $T' = T_1 \cup T_2$ and J_1 and J_2 are structure-erasing.

We are now interested in amalgamating limit objects along (3.8), that is, given $\phi_1 \in \operatorname{Hom}^+(\mathcal{A}[T_1],\mathbb{R})$ and $\phi_2 \in \operatorname{Hom}^+(\mathcal{A}[T_2],\mathbb{R})$ such that $\phi_1^{I_1} = \phi_2^{I_2}$, can we construct $\xi \in \operatorname{Hom}^+(\mathcal{A}[T'],\mathbb{R})$ such that $\xi^{J_1} = \phi_1$ and $\xi^{J_2} = \phi_2$? In the case when T is trivial, this question reduces to the question of whether there exists a coupling of ϕ_1, ϕ_2 defined below.

Definition 3.2.2 (couplings)**.** Given canonical theories $T_1,\ldots,T_t$ and $\phi_i \in \operatorname{Hom}^+(\mathcal{A}[T_i],\mathbb{R})$ ($i \in [t]$), a *coupling* of $\phi_1,\ldots,\phi_t$ is a positive homomorphism $\xi \in \operatorname{Hom}^+(\mathcal{A}[\bigcup_{i\in[t]} T_i],\mathbb{R})$ such that $\xi^{I_i} = \phi_i$ for every $i \in [t]$, where $I_i\colon T_i \rightsquigarrow \bigcup_{j\in[t]} T_j$ is the structure-erasing interpretation.

It is easy to construct a coupling of $\phi_1,\ldots,\phi_t$ by simply aligning any geometric representations of them. Namely, if $\mathcal{N}^i$ ($i \in [t]$) is a T_i-on such that $\phi_i = \phi_{\mathcal{N}^i}$, then the $(\bigcup_{i\in[t]} T_i)$-on $\mathcal{H}$ defined by letting $\mathcal{H}_P \stackrel{\text{def}}{=} \mathcal{N}^i_P$ whenever P is in the language of T_i gives a coupling $\xi \stackrel{\text{def}}{=} \phi_{\mathcal{H}}$ of $\phi_1,\ldots,\phi_t$. However, note that in this construction, ξ might depend on the particular choice of $\mathcal{N}^1,\ldots,\mathcal{N}^t$ (this potential dependence will be further explored in Chapter 4). The more natural notion of an independent coupling defined below is given functorially, that is, it depends only on $\phi_1,\ldots,\phi_t$.

Definition 3.2.3 (independent coupling, syntactic version)**.** For every $i \in [t]$, let $\phi_i \in \operatorname{Hom}^+(\mathcal{A}[T_i],\mathbb{R})$. The *independent coupling* $\phi_1\otimes\cdots\otimes\phi_t \in \operatorname{Hom}^+(\mathcal{A}[\bigcup_{i\in[t]} T_i],\mathbb{R})$ of $\phi_1,\ldots,\phi_t$ is defined by

$$(\phi_1 \otimes \cdots \otimes \phi_t)(\langle M\rangle) \stackrel{\text{def}}{=} \prod_{i\in[t]} \phi_i(\langle I_i(M)\rangle), \tag{3.9}$$

for every $M \in \mathcal{M}[\bigcup_{i\in[t]} T_i]$, where $I_i\colon T_i \rightsquigarrow \bigcup_{j\in[t]} T_j$ is the structure-erasing interpretation.

One can check by calculations that (3.9) indeed satisfies all flag-algebraic constraints, but it is much simpler to give a theon representation of the independent coupling.

Definition 3.2.4 (independent coupling, semantic version)**.** For $i \in [t]$, let $\mathcal{N}^i$ be a T_i-on over Ω_i. The *independent coupling* of $\mathcal{N}^1,\ldots,\mathcal{N}^t$ is the $(\bigcup_{i\in[t]} T_i)$-on $\mathcal{N}^1 \otimes \cdots \otimes \mathcal{N}^t$ over

$\prod_{i\in[t]} \Omega_i$ defined by

$$(\mathcal{N}^1 \otimes \cdots \otimes \mathcal{N}^t)_P \stackrel{\text{def}}{=} \left\{ x \in \prod_{j\in[t]} \mathcal{E}_{k(P)}(\Omega_j) \;\middle|\; \pi_i(x) \in \mathcal{N}^i_P \right\},$$

whenever P is in the language of T_i and where π_i denotes the natural projection on the i-th coordinate.

It is easy to see that if $\mathcal{N}^i$ is a T_i-on over Ω_i such that $\phi_{\mathcal{N}^i} = \phi_i$ ($i \in [t]$), then $(\phi_1 \otimes \cdots \otimes \phi_t) = \phi_{\mathcal{N}^1\otimes\cdots\otimes\mathcal{N}^t}$. In particular, this implies that $\phi_1 \otimes \cdots \otimes \phi_t \in \operatorname{Hom}^+(\mathcal{A}[\bigcup_{i\in[t]} T_i], \mathbb{R})$.

The following theorem says that we can also amalgamate limit objects along general pushouts. Let us warn that unless the theory T is trivial (in which case a "canonical" amalgamation is provided by the independent coupling), we are not aware of any natural, functorial construction here.

Theorem 3.2.5. *Let*

$$\begin{array}{ccc} T & \xrightarrow{I_1} & T_1 \\ {\scriptstyle I_2}\downarrow & & \downarrow{\scriptstyle J_1} \\ T_2 & \xrightarrow[J_2]{} & T' \end{array}$$

be a pushout of **Int** *and let* $\phi_1 \in \operatorname{Hom}^+(\mathcal{A}[T_1], \mathbb{R})$ *and* $\phi_2 \in \operatorname{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ *be such that* $\phi_1^{I_1} = \phi_2^{I_2}$. *Then there exists* $\psi \in \operatorname{Hom}^+(\mathcal{A}[T'], \mathbb{R})$ *such that* $\psi^{J_1} = \phi_1$ *and* $\psi^{J_2} = \phi_2$.

Proof. First we claim that it is enough to show the case when T' is obtained from $T_1 \cup T_2$ by adding the axioms (3.7). Indeed, if ψ is constructed for such particular case, then we can get our desired element of $\operatorname{Hom}^+(\mathcal{A}[T'], \mathbb{R})$ for a general pushout T' as ψ^I for the universal isomorphism I between the pushout theories.

Let us prove then the particular case. Let $\mathcal{L}$, $\mathcal{L}_1$ and $\mathcal{L}_2$ be the languages of T, T_1 and T_2, respectively. For $i \in [2]$, let $\mathcal{N}^i$ be a T_i-on (over $[0,1]$) such that $\phi_i = \phi_{\mathcal{N}^i}$. Since $\phi_{I_1(\mathcal{N}^1)} = \phi_1^{I_1} = \phi_2^{I_2} = \phi_{I_2(\mathcal{N}^2)}$, by Proposition 3.2.1, there exists a T_1-on $\mathcal{H}^1$ over $[0,1]^2$ such that $I_1(\mathcal{H}^1)_P = I_2(\mathcal{N}^2)_P \times \mathcal{E}_{k(P)}$ λ-a.e. for every $P \in \mathcal{L}$.

Define then the Euclidean structure $\mathcal{H}$ on $\mathcal{L}_1 \mathbin{\dot\cup} \mathcal{L}_2$ over $[0,1]^2$ by

$$\mathcal{H}_P \stackrel{\text{def}}{=} \begin{cases} \mathcal{H}^1_P, & \text{if } P \in \mathcal{L}_1; \\ \mathcal{N}^2_P \times \mathcal{E}_{k(P)}, & \text{if } P \in \mathcal{L}_2. \end{cases}$$

Let us show that $\mathcal{H}$ is a (weak) T'-on. To show this, by Theorem 2.4.1 (by reaxiomatizing T, T_1, T_2 to be substitutionally closed, T' also becomes substitutionally closed) it is enough to show that $T(I_1(P), \mathcal{H}) = T(I_2(P), \mathcal{H})$ λ-a.e. for every $P \in \mathcal{L}$. But this follows from

$$T(I_1(P), \mathcal{H}) = T(I_1(P), \mathcal{H}^1) = I_1(\mathcal{H}^1)_P;$$
$$T(I_2(P), \mathcal{H}) = T(I_2(P), \mathcal{N}^2) \times \mathcal{E}_{k(P)} = I_2(\mathcal{N}^2) \times \mathcal{E}_{k(P)}.$$

Finally, since we trivially have $J_1(\mathcal{H}) = \mathcal{H}^1$ and $J_2(\mathcal{H})_P = \mathcal{N}^2_P \times \mathcal{E}_{k(P)}$ for every $P \in \mathcal{L}_2$, it follows that $\psi \stackrel{\text{def}}{=} \phi_{\mathcal{H}}$ satisfies $\psi^{J_1} = \phi_1$ and $\psi^{J_2} = \phi_2$. ■

The amalgamation property of Theorem 3.2.5 above in particular implies that couplings can be "lifted" through interpretations.

Proposition 3.2.6 (Coupling lifting)**.** *Let* $I\colon T_1 \rightsquigarrow T_2$ *be an open interpretation, let* T *be a canonical theory and let* $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ *and* $\phi_2 \in \operatorname{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$*. If* ξ *is a coupling of* ϕ_2^I *and* ϕ*, then there exists a coupling* $\widehat{\xi}$ *of* ϕ_2 *and* ϕ *such that* $\xi = \widehat{\xi}^{I \cup \mathrm{id}_T}$*.*

Proof. This follows from Theorem 3.2.5 and the fact that

$$\begin{array}{ccc} T_1 & \xrightarrow{\;I\;} & T_2 \\ \downarrow & & \downarrow \\ T_1 \cup T & \xrightarrow[I \cup \mathrm{id}_T]{} & T_2 \cup T \end{array}$$

is a pushout in **Int**, where the vertical arrows are the structure-erasing interpretations. ■

Definition 3.2.7 (unique coupleability)**.** We say that $\phi_1, \ldots, \phi_t$ are *uniquely coupleable* if the independent coupling is their only coupling.

We will see in Chapter 4 that the easiest case of unique coupleability is between $\phi_1 \in \texttt{Independence}[\ell]$ and ϕ_2 with $\mathrm{rk}(\phi_2) \leq \ell$ (see Definition 3.1.1). The notion of unique coupleability is fundamental to Chapter 4, but in this chapter we will concentrate on a more abstract study of this property (the reader might want to skip momentarily to the beginning of Chapter 4 for some intuition and motivating examples).

The next lemma says that unique coupleability satisfies a "chain rule" analogous to the chain rule for mutual independence of random variables.

Lemma 3.2.8. *Let $\phi_i \in \mathrm{Hom}^+(\mathcal{A}[T_i], \mathbb{R})$ for $i \in [t]$ and suppose that for every $i \in [t-1]$, ϕ_{i+1} is uniquely coupleable with $\phi_1 \otimes \cdots \otimes \phi_i$. Then $\phi_1, \ldots, \phi_t$ are uniquely coupleable.*

Proof. The proof is by induction in t. The result for $t = 1$ is trivial. For $t \geq 2$, let $\xi \in \mathrm{Hom}^+(\mathcal{A}[\bigcup_{i=1}^t T_i], \mathbb{R})$ be a coupling of $\phi_1, \ldots, \phi_t$ and let $I\colon \bigcup_{i=1}^{t-1} T_i \rightsquigarrow \bigcup_{i=1}^t T_i$ be the structure-erasing interpretation. Since ξ^I is a coupling of $\phi_1, \ldots, \phi_{t-1}$, by inductive hypothesis we must have $\xi^I = \phi_1 \otimes \cdots \otimes \phi_{t-1}$ so ξ is also a coupling of $\phi_1 \otimes \cdots \otimes \phi_{t-1}$ and ϕ_t, hence we must have $\xi = \phi_1 \otimes \cdots \otimes \phi_t$. ∎

Proposition 3.2.6 and Lemma 3.2.8 allow us to show that unique coupleability is preserved under open interpretations.

Proposition 3.2.9. *For $i \in [n]$, let $\phi_i \in \mathrm{Hom}^+(\mathcal{A}[T_i], \mathbb{R})$ and $I_i\colon T'_i \rightsquigarrow T_i$ be an open interpretation. If $\phi_1, \ldots, \phi_n$ are uniquely coupleable, then $\phi_1^{I_1}, \ldots, \phi_n^{I_n}$ are uniquely coupleable.*

Proof. The proof is by induction in n. The case $n = 1$ is trivial.

Consider now the case when $n = 2$, $T_1 = T'_1$ and $I_1 = \mathrm{id}_{T_1}$. In this case, for a coupling ξ of $\phi_1 \in \mathrm{Hom}^+(\mathcal{A}[T_1], \mathbb{R})$ and $\phi_2^{I_2} \in \mathrm{Hom}^+(\mathcal{A}[T'_2], \mathbb{R})$, Proposition 3.2.6 gives us a coupling $\widehat{\xi}$ of ϕ_1 with ϕ_2 such that $\xi = \widehat{\xi}^{\mathrm{id}_{T_1} \cup I}$. Since ϕ_1, ϕ_2 are uniquely coupleable, we must have $\widehat{\xi} = \phi_1 \otimes \phi_2$, from which we get $\xi = \widehat{\xi}^{\mathrm{id}_{T_1} \cup I} = \phi_1 \otimes \phi_2^I$, hence $\phi_1, \phi_2^{I_2}$ are uniquely coupleable.

For the case $n = 2$ but I_1 arbitrary, note that unique coupleability of ϕ_1, ϕ_2 implies that $\phi_1, \phi_2^{I_2}$ are uniquely coupleable by the case above, which in turn implies that $\phi_1^{I_1}, \phi_2^{I_2}$ are uniquely coupleable by the symmetric of the case above.

Finally, for the general case $n \geq 3$ by Lemma 3.2.8, it is enough to show that for every $t \in [n-1]$, $\phi_{t+1}^{I_{t+1}}$ is uniquely coupleable with $\phi_1^{I_1} \otimes \cdots \otimes \phi_t^{I_t}$.

First, we claim that $\phi_1, \ldots, \phi_{t+1}$ are uniquely coupleable. Indeed, any coupling ξ of $\phi_1, \ldots, \phi_{t+1}$ can be lifted to a coupling $\widehat{\xi}$ of $\phi_1, \ldots, \phi_n$ using Proposition 3.2.6 and since $\widehat{\xi}$ must be $\phi_1 \otimes \cdots \otimes \phi_n$, it follows that $\xi = \phi_1 \otimes \cdots \otimes \phi_{t+1}$. Note that this in particular implies that ϕ_{t+1} is uniquely coupleable with $\phi_1 \otimes \cdots \otimes \phi_t$ (as any coupling of these must also be a coupling of $\phi_1, \ldots, \phi_{t+1}$).

Now let $I\colon \bigcup_{i\in[t]} T_i' \rightsquigarrow \bigcup_{i\in[t]} T_i$ act as I_i in T_i, then since ϕ_{t+1} is uniquely coupleable with $\phi_1 \otimes \cdots \otimes \phi_t$, the case $n = 2$ implies that $\phi_{t+1}^{I_{t+1}}$ is uniquely coupleable with $(\phi_1 \otimes \cdots \otimes \phi_t)^I = \phi_1^{I_1} \otimes \cdots \otimes \phi_t^{I_t}$. ■

3.3 L_1-topology

Recall that the set of limit objects $\mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ comes equipped with the density topology, that is, the topology induced from the inclusion $\mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R}) \subseteq [0,1]^{\mathcal{M}[T]}$. Let us now introduce the L_1-topology that is a direct analogue of the L_1-topology on graphons [53, Sct. 8.2.5 and Sct. 8.3].

Definition 3.3.1. If T is a theory in a language $\mathcal{L}$ and $\phi_1, \phi_2 \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$, then the L_1-*distance* between ϕ_1 and ϕ_2 is defined as

$$\delta_1(\phi_1, \phi_2) \stackrel{\text{def}}{=} \min_{\mathcal{N}^1, \mathcal{N}^2} \sum_{P \in \mathcal{L}} \mu(\mathcal{N}_P^1 \triangle \mathcal{N}_P^2), \tag{3.10}$$

where the minimum is taken over T-ons $\mathcal{N}^1$ and $\mathcal{N}^2$ over the same space such that $\phi_1 = \phi_{\mathcal{N}^1}$ and $\phi_2 = \phi_{\mathcal{N}^2}$.

It is not immediately clear from this definition that the minimum in (3.10) is actually attained, nor is it clear why δ_1 is a metric.

The first issue is easy to address by giving an alternative purely algebraic definition.

Namely, for any $P \in \mathcal{L}$, introduce the element $d_P \in \mathcal{A}[T \cup T]$ as

$$d_P \stackrel{\text{def}}{=} \sum_{\substack{K \in \mathcal{K}_{k(P)}[T \cup T] \\ \mathrm{id}_{k(P)} \in R_{P_1}(K) \triangle R_{P_2}(K)}} \langle K \rangle,$$

where P_1 and P_2 are the two copies of P in $\mathcal{L} \dot\cup \mathcal{L}$, and let

$$d_T \stackrel{\text{def}}{=} \sum_{P \in \mathcal{L}} d_P. \tag{3.11}$$

This element measures the distance in a coupling of ϕ_1, ϕ_2 so we have

$$\delta_1(\phi_1, \phi_2) = \inf_{\xi} \xi(d_T), \tag{3.12}$$

where ξ runs over all couplings of ϕ_1 and ϕ_2. Their set is determined in $\mathrm{Hom}^+(\mathcal{A}[T \cup T], \mathbb{R})$ by countably many linear equations and hence is compact in the density topology. Therefore the minimum in (3.12) and (3.10) is actually achieved (as $\xi \mapsto \xi(d_T)$ is continuous in the density topology).

The second issue is trickier, and the proof is similar to the analogous proof that δ_1 is a metric in the case of graphons. Fortunately, we already did most of the necessary (and notationally heavy) work in the proof of Proposition 3.2.1.

Lemma 3.3.2. *The L_1-distance δ_1 is a metric on $\mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ and generates a finer topology than the density topology.*

Proof. Let us first check the triangle inequality. Let ξ be a coupling of ϕ_1 and ϕ_2 and ζ be a coupling of ϕ_2 and ϕ_3 attaining the L_1-distances in (3.12). Let also $J_i \colon T \rightsquigarrow T \cup T$ be the structure-erasing interpretation corresponding to coordinate i and $I_{ij} \colon T \cup T \rightsquigarrow T \cup T \cup T$ be the structure-erasing interpretation corresponding to coordinates i and j. Since ξ is a coupling of ϕ_1 and $\phi_2 = \zeta^{J_1}$, Proposition 3.2.6 gives us a coupling $\widehat{\xi}$ of ϕ_1 and ζ such that

$\widehat{\xi}^{\mathrm{id}_T \cup J_1} = \xi$. Since $\mathrm{id}_T \cup J_1 = I_{12}$, we get that $\widehat{\xi}$ is a coupling of ϕ_1, ϕ_2 and ϕ_3 such that $\widehat{\xi}^{I_{12}} = \xi$ and $\widehat{\xi}^{I_{23}} = \zeta$. But $\widehat{\xi}^{I_{13}}$ is a coupling of ϕ_1 and ϕ_3 and for each $P \in \mathcal{L}$ we have

$$\widehat{\xi}^{I_{13}}(d_P) \leq \widehat{\xi}^{I_{12}}(d_P) + \widehat{\xi}^{I_{23}}(d_P),$$

hence by (3.12) we get $\delta_1(\phi_1, \phi_3) \leq \delta_1(\phi_1, \phi_2) + \delta_1(\phi_2, \phi_3)$.

Finally, note that by (3.10) we have

$$|\phi_1(\langle M \rangle) - \phi_2(\langle M \rangle)| \leq \delta_1(\phi_1, \phi_2) \sum_{P \in \mathcal{L}} (|M|)_{k(P)},$$

for every $M \in \mathcal{M}[T]$. This implies both $\delta_1(\phi_1, \phi_2) = 0 \implies \phi_1 = \phi_2$ and that the L_1-topology is finer than the density topology. ■

Since the density topology is Hausdorff (as it is metrizable), compact and coarser than the L_1-topology, it follows that the L_1-topology is compact if and only if it is equal to the density topology[1]. Such equality of topologies typically does not happen as we will see in Section 3.4.

Remark 2. Given a T-on $\mathcal{N}$ over some space $\Omega = (X, \mathcal{A}, \mu)$ and some $\psi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$, the L_1-distance between $\phi_{\mathcal{N}}$ and ψ can be alternatively computed only optimizing over T-ons corresponding to ψ by the following formula

$$\delta_1(\phi_{\mathcal{N}}, \psi) = \min_{\mathcal{H}} \sum_{P \in \mathcal{L}} \mu((\mathcal{N}_P \times \mathcal{E}_{k(P)}(\Omega)) \triangle \mathcal{H}_P),$$

where the minimum is taken over all T-ons $\mathcal{H}$ over $\Omega \times \Omega$ such that $\phi_{\mathcal{H}} = \psi$. To see this, we form a coupling ξ of ϕ and ψ attaining the minimum in (3.12) and use Proposition 3.2.1 to produce a $(T \cup T)$-on $\mathcal{G}$ over $\Omega \times \Omega$ such that $\phi_{\mathcal{G}} = \xi$ and $\mathcal{G}_{P_1} = \mathcal{N}_P \times \mathcal{E}_{k(P)}$ for every $P \in \mathcal{L}$,

1. The non-trivial direction follows by noting that the identity map with L_1-topology in the domain and density topology in the codomain is continuous, so if the L_1-topology is compact, the same map must also be closed (as the density topology is Hausdorff), thus a homeomorphism.

where P_1 is the first copy of P in $\mathcal{L} \mathbin{\dot\cup} \mathcal{L}$. If $\mathcal{H} \stackrel{\text{def}}{=} I(\mathcal{G})$ for the structure-erasing interpretation $I\colon T \rightsquigarrow T \cup T$ that keeps the second copy, then

$$\sum_{P \in \mathcal{L}} \mu((\mathcal{N}_P \times \mathcal{E}_{k(P)}(\Omega)) \mathbin{\triangle} \mathcal{H}_P), = \sum_{P} \xi(d_P) = \xi(d_T) = \delta_1(\phi_{\mathcal{N}}, \psi).$$

3.4 Rank function in density topology and in L_1-topology

As we have seen in Lemma 3.3.2, the L_1-topology is finer than the density topology. In this section, we illustrate some of the differences between these topologies with respect to the rank function: while the sets $\{\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R}) \mid \operatorname{rk}(\phi) \leq r\}$ are closed in the L_1-topology (i.e., the rank function is lower semi-continuous in the L_1-topology), in pure canonical theories $T = T_{\mathcal{L}}$, these sets are dense in $\operatorname{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R})$ in the density topology as long as $r \geq 1$.

Proposition 3.4.1. *The rank function in* $\operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ *is lower semi-continuous in the* L_1*-topology.*

Proof. Let $\mathcal{L}$ be the language of T, let $(\phi_n)_{n \in \mathbb{N}}$ be a sequence in $\operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ converging to $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ in L_1-topology. To show lower semi-continuity of the rank function, we need to show that if $\operatorname{rk}(\phi_n) \leq r$ for every $n \in \mathbb{N}$, then $\operatorname{rk}(\phi) \leq r$.

For each $n \in \mathbb{N}$, let ξ_n be a coupling of ϕ_n and ϕ such that $\delta_1(\phi_n, \phi) = \xi_n(d_T)$ (see (3.11)). Let also $I_i\colon T \rightsquigarrow T \cup T$ be the structure-erasing interpretation that keeps the i-th copy so that $\xi_n^{I_1} = \phi_n$ and $\xi_n^{I_2} = \phi$.

For a fixed T-on $\mathcal{N}$ (over $[0,1]$) such that $\phi_{\mathcal{N}} = \phi$, by Proposition 3.2.1, for each $n \in \mathbb{N}$, there exists a $(T \cup T)$-on $\mathcal{H}^n$ over $[0,1]^2$ such that $\phi_{\mathcal{H}^n} = \xi_n$ and $I_2(\mathcal{H}^n)_P = \mathcal{N}_P \times \mathcal{E}_{k(P)}$ for every predicate symbol $P \in \mathcal{L}$. By Proposition 3.1.2, we may change the peons of $\mathcal{H}^n$ in a zero-measure set so that $\operatorname{rk}(I_1(\mathcal{H}^n)) = \operatorname{rk}(\phi_n) \leq r$.

For each $P \in \mathcal{L}$ and $n \in \mathbb{N}$, let $W_P^n, U_P\colon \mathcal{E}_{k(P),r}([0,1]^2) \to [0,1]$ be the functions defined by

$$W^n_P(x) \stackrel{\text{def}}{=} \lambda(\{y \in ([0,1]^2)^{\binom{[k]}{>r}} \mid (x,y) \in I_1(\mathcal{H}^n)_P\}); \tag{3.13}$$

$$U_P(x) \stackrel{\text{def}}{=} \lambda(\{y \in ([0,1]^2)^{\binom{[k]}{>r}} \mid (x,y) \in \mathcal{N}_P \times \mathcal{E}_{k(P)}\}); \tag{3.14}$$

and defined arbitrarily when the respective sets are not measurable. Since $\operatorname{rk}(I_1(\mathcal{H}^n)) \leq r$, we know that W^n_P is 0-1 valued for every $P \in \mathcal{L}$. To show that $\operatorname{rk}(\phi) \leq r$, we need to show that U_P is 0-1 valued a.e. for every $P \in \mathcal{L}$ (see Remark 1).

But note that if d_1 is the usual L_1-distance of functions, then we have

$$\sum_{P \in \mathcal{L}} d_1(W^n_P, U_P) \leq \sum_{P \in \mathcal{L}} \lambda(I_1(\mathcal{H}^n)_P \mathbin{\triangle} I_2(\mathcal{H}^n)_P) = \xi_n(d_T) = \delta_1(\phi_n, \phi),$$

so W^n_P converges to U_P in the usual L_1-distance of functions and thus U_P must be 0-1 valued a.e., so $\operatorname{rk}(\phi) \leq r$. ∎

The next proposition implies that for a pure canonical theory $T_{\mathcal{L}}$, the rank function is not lower semi-continuous in $\operatorname{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R})$ in the density topology as long as its image has values greater than 1 (otherwise, the rank function is lower semi-continuous for trivial reasons: $\{\phi \in \operatorname{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R}) \mid \operatorname{rk}(\phi) = 0\}$ is closed).

Proposition 3.4.2. *For a pure canonical theory $T_{\mathcal{L}}$ and $r \geq 1$, the set $\{\phi \in \operatorname{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R}) \mid rk(\phi) \leq r\}$ is dense in $\operatorname{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R})$ in the density topology.*

Proof. It is enough to show the case $r = 1$. Let $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_{\mathcal{L}}], \mathbb{R})$ and let $(N_n)_{n \in \mathbb{N}}$ be a convergent sequence of models of $T_{\mathcal{L}}$ converging to ϕ. Without loss of generality, suppose $V(N_n) = [m_n]$ for some $m_n \in \mathbb{N}$ and define the $T_{\mathcal{L}}$-ons $\mathcal{N}^n$ by

$$\mathcal{N}^n_P \stackrel{\text{def}}{=} \{x \in \mathcal{E}_{k(P)} \mid (\lceil m_n \cdot x_{\{1\}} \rceil, \ldots, \lceil m_n \cdot x_{\{k(P)\}} \rceil) \in R_P(N_n)\}$$

and let $\phi_n \stackrel{\text{def}}{=} \phi_{\mathcal{N}^n}$. Clearly $\operatorname{rk}(\phi_n) \leq \operatorname{rk}(\mathcal{N}^n) \leq 1$.

Note that for every $M \in \mathcal{M}[T]$, we have

$$|t_{\text{ind}}(M, \mathcal{N}^n) - t_{\text{ind}}(M, N_n)| \leq O\left(\frac{1}{m_n}\right),$$

where the hidden constant depends on M. Thus, since $(N_n)_{n\in\mathbb{N}}$ converges to ϕ, it follows that ϕ_n converges to ϕ in the density topology. ■

The remainder of this section is devoted to the following generalization of Proposition 3.4.2, which can be seen as a version relative to $\texttt{Independence}[\ell]$.

Proposition 3.4.3. *For a pure canonical theory $T_\mathcal{L}$ and $\ell, r \in \mathbb{N}$ such that $\ell < r$, the set $\{\phi \in \operatorname{Hom}^+(\mathcal{A}[T_\mathcal{L}], \mathbb{R}) \mid rk(\phi) \leq r \wedge \phi \in \texttt{Independence}[\ell]\}$ is dense in $\operatorname{Hom}^+(\mathcal{A}[T_\mathcal{L}], \mathbb{R}) \cap \texttt{Independence}[\ell]$ in the density topology.*

Note that the condition $\ell < r$ is required as if ϕ satisfies $\texttt{Independence}[\ell]$ for some $\ell \geq \operatorname{rk}(\phi)$, then ϕ must be a trivial limit (i.e., $\operatorname{rk}(\phi) = 0$, or equivalently, all its peons have measure either 0 or 1).

If we take a step back on the proof of Proposition 3.4.2 and recall that one way of producing a convergent sequence $(N_n)_{n\in\mathbb{N}}$ converging to ϕ is to consider a theon $\mathcal{N}$ representing ϕ and sampling points from it, i.e., producing the exchangeable array $\boldsymbol{K}$ corresponding to $\mathcal{N}$, we see can see the theons $\mathcal{N}^n$ as "rank 1 blow-ups" of the marginal $\boldsymbol{K}|_{[n]}$. The next definition generalizes this concept to higher ranks: the idea is to preserve all rank less than k information and randomize the information of rank at least k by an approximation that takes place at rank exactly k. In particular, the rank 1 approximation corresponds precisely to the "rank 1 blow-ups" of the marginal $\boldsymbol{K}|_{[n]}$ described above.

Definition 3.4.4 (Rank k approximation)**.** Let $\mathcal{N}$ be an Euclidean structure in a language $\mathcal{L}$ (over $[0,1]$) and let $k, n \in \mathbb{N}_+$. The *rank k approximation of $\mathcal{N}$ at step n* is the random the Euclidean structure $\boldsymbol{\mathcal{N}^{k,n}}$ defined as follows.

For $\ell \in \mathbb{N}_+$, let F_k^ℓ be the set of all functions $\alpha\colon \binom{[\ell]}{k} \to \mathbb{N}_+$ and consider the natural (right) action of S_ℓ on F_k^ℓ given by $\alpha \cdot \sigma \stackrel{\text{def}}{=} \alpha \circ \sigma$ ($\alpha \in F_k^\ell, \sigma \in S_\ell$). For $\alpha \in F_k^\ell$, let $\widehat{\alpha}$ denote the orbit of α under the action of S_ℓ. Let also $\widehat{F}_k^\ell \stackrel{\text{def}}{=} \{\widehat{\alpha} \mid \alpha \in F_k^\ell\}$ be the set of all orbits of F_k^ℓ and let $\widehat{F}_k \stackrel{\text{def}}{=} \bigcup_{\ell\in\mathbb{N}_+} \widehat{F}_k^\ell$.

Given $x \in [0,1]^{\binom{[\ell]}{k}}$, let $\alpha_x^n \colon \binom{[\ell]}{k} \to \mathbb{N}_+$ be defined by

$$\alpha_x^n(A) \stackrel{\text{def}}{=} \max\{\lceil n \cdot x_A \rceil, 1\} \qquad \left(A \in \binom{[\ell]}{k}\right).$$

Recall that for $A \in r(\ell)$, we denote by $\iota_{A,\ell} \colon [|A|] \to [\ell]$ the function that enumerates A in increasing order and that it induces the natural projection $\iota_{A,\ell}^* \colon [0,1]^{\binom{[\ell]}{k}} \to [0,1]^{\binom{[|A|]}{k}}$ given by $\iota_{A,\ell}^*(x)_B \stackrel{\text{def}}{=} x_{\iota_{A,\ell}(B)}$ $(B \in \binom{[|A|]}{\ell})$.

For every $\widehat{\alpha} \in \widehat{F}_k$, pick $\boldsymbol{y}_{\widehat{\alpha}}$ independently and at random according to λ (so $\boldsymbol{y}$ is distributed according to the product measure $\lambda^{\widehat{F}_k}$) and for every $(x,x') \in \mathcal{E}_{\ell,k-1} \times [0,1]^{\binom{[\ell]}{k}}$, define the random point $w_\ell^{k,n}((x,x'),\boldsymbol{y})$ of $\mathcal{E}_\ell$ by

$$(w_\ell^{k,n}((x,x'),\boldsymbol{y}))_A \stackrel{\text{def}}{=} \begin{cases} x_A, & \text{if } |A| < k; \\ \boldsymbol{y}_{\widehat{\alpha}^n_{\iota^*_{A,\ell}(x')}}, & \text{if } |A| \geq k. \end{cases} \qquad (A \in r(\ell)).$$

Finally, we let

$$\boldsymbol{\mathcal{N}}_P^{\boldsymbol{k,n}} \stackrel{\text{def}}{=} \{(x,x',x'') \in \mathcal{E}_{k(P),k-1} \times [0,1]^{\binom{[k(P)]}{k}} \times [0,1]^{\binom{[k(P)]}{>k}} \mid w_{k(P)}^{k,n}((x,x'),\boldsymbol{y}) \in \mathcal{N}_P\} \tag{3.15}$$

for every $P \in \mathcal{L}$.

Remark 3. Since the formula in (3.15) does not depend on x'', which accounts for all coordinates indexed by sets of size larger than k, it follows that $\operatorname{rk}(\boldsymbol{\mathcal{N}}^{\boldsymbol{k,n}}) \leq k$. Furthermore, since the formula (3.15) depends on x only via $\mathcal{N}$, it follows that if $\mathcal{N}$ is ℓ-independent for some $\ell \in \mathbb{N}$, then so is $\boldsymbol{\mathcal{N}}^{\boldsymbol{k,n}}$.

We will show that $\phi_{\boldsymbol{\mathcal{N}}^{\boldsymbol{k,n}}}$ converges to $\phi_{\mathcal{N}}$ with probability 1, which together with Remark 3 will give Proposition 3.4.3. Let us first recall a basic fact of probability theory.

Lemma 3.4.5. *If A and B are events in a probability space and $\mathbb{P}[B] > 0$, then $|\mathbb{P}[A \mid B] - \mathbb{P}[A]| \leq 1 - \mathbb{P}[B]$.*

Proof. If $\mathbb{P}[B] = 1$, the result is trivial, otherwise, we have

$$\mathbb{P}[A] - \mathbb{P}[A \mid B] = (\mathbb{P}[A \mid B^c] - \mathbb{P}[A \mid B])(1 - \mathbb{P}[B]),$$

where B^c is the complement of B. Taking absolute values and noting that $|\mathbb{P}[A \mid B^c] - \mathbb{P}[A \mid B]| \leq 1$ yields the result. ■

As a first step, the next lemma says that $\phi_{\boldsymbol{\mathcal{N}}^{\boldsymbol{k},\boldsymbol{n}}}$ converges to $\phi_{\mathcal{N}}$ at least in expected value (in the density topology).

Lemma 3.4.6. *Let $\mathcal{N}$ be an Euclidean structure in a language $\mathcal{L}$ and let $k \in \mathbb{N}_+$. Then*

$$\lim_{n\to\infty} \mathbb{E}[\phi_{\boldsymbol{\mathcal{N}}^{\boldsymbol{k},\boldsymbol{n}}}(M)] = \phi_{\mathcal{N}}(M)$$

for every $M \in \mathcal{M}[T]$.

Proof. It is enough to show that for every $m \in \mathbb{N}$ and every $K \in \mathcal{K}_m[T_{\mathcal{L}}]$, we have

$$\lim_{n\to\infty} \mathbb{E}[t_{\text{ind}}(K, \boldsymbol{\mathcal{N}}^{\boldsymbol{k},\boldsymbol{n}})] = t_{\text{ind}}(K, \mathcal{N}).$$

For every $n \in \mathbb{N}$, let us define the set G_n of "good" points of $[0,1]^{\binom{[m]}{k}}$ at stage n as the set of points $x \in [0,1]^{\binom{[m]}{k}}$ such that α_x^n is injective, that is, let

$$G_n \stackrel{\text{def}}{=} \left\{ x \in [0,1]^{\binom{[m]}{k}} \;\middle|\; \forall A, B \in \binom{[m]}{k}, (A \neq B \to \alpha_x^n(A) \neq \alpha_x^n(B)) \right\}.$$

We define the labeled density of K in $\boldsymbol{\mathcal{N}}^{\boldsymbol{k},\boldsymbol{n}}$ relative to "good" points as

$$t'_{\text{ind}}(K, \boldsymbol{\mathcal{N}}^{\boldsymbol{k},\boldsymbol{n}}) \stackrel{\text{def}}{=} \frac{\lambda(T_{\text{ind}}(K, \boldsymbol{\mathcal{N}}^{\boldsymbol{k},\boldsymbol{n}}) \cap (\mathcal{E}_{m,k-1} \times G_n \times [0,1]^{\binom{[m]}{>k}}))}{\lambda(G_n)}.$$

Note that Lemma 3.4.5 implies

$$|t_{\mathrm{ind}}(K,\boldsymbol{\mathcal{N}^{k,n}}) - t'_{\mathrm{ind}}(K,\boldsymbol{\mathcal{N}^{k,n}})| \leq 1 - \lambda(G_n) = 1 - \frac{(n)_{\binom{m}{k}}}{n^{\binom{n}{k}}} \leq O\left(\frac{1}{n}\right),$$

thus, by defining $t'_{\mathrm{ind}}(K,\mathcal{N})$ analogously, it is sufficient to show that $\mathbb{E}[t'_{\mathrm{ind}}(K,\boldsymbol{\mathcal{N}^{k,n}})] = t'_{\mathrm{ind}}(K,\mathcal{N})$ for every $n \in \mathbb{N}$.

Let us then partition G_n as follows: for every $\alpha\colon \binom{[m]}{k} \to \mathbb{N}_+$ (i.e., $\alpha \in F_k^m$), we let

$$X_\alpha^n \stackrel{\mathrm{def}}{=} \{x \in G_n \mid \alpha_x^n = \alpha\}$$

(in particular, $X_\alpha^n \neq \varnothing$ if and only if α is injective).

Note that for $A \in r(\ell)$ and $B \in \binom{[|A|]}{k}$, we have $\alpha_x^n(\iota_{A,\ell}(B)) = \alpha^n_{\iota^*_{A,\ell}(x)}(B)$. This implies that if $\alpha \in F_k^m$ is injective, then for every $x \in X_\alpha^n$ and every $A \in r(m)$, the function $\alpha^n_{\iota^*_{A,m}(x)}$ is also injective and we further have

$$\forall \beta \in ([m])_\ell, \forall A, B \in r(\ell), (A = B \leftrightarrow \widehat{\alpha}^n_{\iota^*_{A,\ell}(\beta^*(x))} = \widehat{\alpha}^n_{\iota^*_{B,\ell}(\beta^*(x))}). \tag{3.16}$$

Pick $\boldsymbol{x}$, $\boldsymbol{x'}$ and $\boldsymbol{x''}$ independently uniformly at random in $\mathcal{E}_{m,k-1}$, $[0,1]^{\binom{[m]}{k}}$ and $[0,1]^{\binom{[m]}{>k}}$, respectively, and independently from $\boldsymbol{y}$ of the definition of $\boldsymbol{\mathcal{N}^{k,n}}$ and note that

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{y}}[t'_{\mathrm{ind}}(K,\boldsymbol{\mathcal{N}^{k,n}})] \\
&\quad = \mathbb{E}_{\boldsymbol{y}}\left[\mathbb{P}_{\boldsymbol{x},\boldsymbol{x'},\boldsymbol{x''}}\left[\bigwedge_{P\in\mathcal{L}}\bigwedge_{\beta\in([m])_{k(P)}} (\beta^*(\boldsymbol{x},\boldsymbol{x'},\boldsymbol{x''}) \in \boldsymbol{\mathcal{N}^{k,n}_P} \leftrightarrow \beta \in R_P(K)) \,\middle|\, \boldsymbol{x'} \in G_n\right]\right] \\
&\quad = \mathbb{P}_{\boldsymbol{x},\boldsymbol{x'},\boldsymbol{y}}\left[\bigwedge_{P\in\mathcal{L}}\bigwedge_{\beta\in([m])_{k(P)}} (w^{k,n}_{k(P)}(\beta^*(\boldsymbol{x},\boldsymbol{x'}),\boldsymbol{y}) \in \mathcal{N}_P \leftrightarrow \beta \in R_P(K)) \,\middle|\, \boldsymbol{x'} \in G_n\right].
\end{aligned}$$

Let us now analyze the restriction of the event above to the event $\boldsymbol{x'} \in X_\alpha^n$ for some injective $\alpha \in F_k^m$. Note that (3.16) implies that the coordinate of $\boldsymbol{y}$ indexed by $\widehat{\alpha}^n_{\iota^*_{A,\ell}(\beta^*(\boldsymbol{x'}))}$ used for a particular coordinate $w^{k,n}_{k(P)}(\beta^*(\boldsymbol{x}, \boldsymbol{x'}), \boldsymbol{y})_A$ for some $A \in \binom{[k(P)]}{>k-1}$, some $\beta \in ([m])_{k(P)}$ and some $P \in \mathcal{L}$ depends only on $\beta(A) \subseteq [m]$ and is distinct from other coordinates with a different value of $\beta(A)$. Since the coordinates of $\boldsymbol{y}$ are i.i.d. uniform in $[0, 1]$ and independent of $(\boldsymbol{x}, \boldsymbol{x'})$, it follows that if $\boldsymbol{z}$ is picked uniformly in $[0, 1]^{\binom{[m]}{>k}}$, then

$$(w^{k,n}_{k(P)}(\beta^*(\boldsymbol{x}, \boldsymbol{x'}), \boldsymbol{y}) \mid P \in \mathcal{L}, \beta \in ([m])_{k(P)})$$

has the same conditional distribution as $(\beta^*(\boldsymbol{x}, \boldsymbol{x'}, \boldsymbol{z}) \mid P \in \mathcal{L}, \beta \in ([m])_{k(P)})$ when given $\boldsymbol{x'}$ and the event $\boldsymbol{x'} \in X^n_\alpha$. Thus, by conditioning, we get

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{y}}[t'_{\text{ind}}(K, \boldsymbol{\mathcal{N}^{k,n}})] \\
&\quad = \frac{1}{\lambda(G_n)} \sum_{\substack{\alpha \in F^m_k \\ \lambda(X^n_\alpha) > 0}} \mathbb{P}_{\boldsymbol{x'}}[\boldsymbol{x'} \in X^n_\alpha] \cdot \mathbb{P}_{\boldsymbol{x},\boldsymbol{x'},\boldsymbol{z}}[(\boldsymbol{x}, \boldsymbol{x'}, \boldsymbol{z}) \in T_{\text{ind}}(K, \mathcal{N}) \mid \boldsymbol{x'} \in X^n_\alpha] \\
&\quad = t'_{\text{ind}}(K, \mathcal{N}),
\end{aligned}$$

as desired. ■

We will now use Azuma's Inequality (see e.g. [3, Theorem 7.2.1]) for martingale concentration to show that $\boldsymbol{\phi_{\mathcal{N}^{k,n}}}$ asymptotically concentrates on its expected value.

Lemma 3.4.7. *Let $\mathcal{N}$ be an Euclidean structure in a language $\mathcal{L}$ and let $k \in \mathbb{N}_+$. Then $\phi_{\boldsymbol{\mathcal{N}^{k,n}}}$ converges to $\phi_{\mathcal{N}}$ in density topology with probability* 1.

Proof. By the union bound, it is enough to show that for every $m \in \mathbb{N}$ and every $K \in \mathcal{K}_m[T_{\mathcal{L}}]$, we have

$$\mathbb{P}\left[\lim_{n\to\infty} t_{\text{ind}}(K, \boldsymbol{\mathcal{N}^{k,n}}) = t_{\text{ind}}(K, \mathcal{N})\right] = 1. \tag{3.17}$$

For every $t \in \{0, \ldots, n\}$, let $\mathcal{B}_t$ be the σ-algebra generated by $(\boldsymbol{y}_{\widehat{\alpha}} \mid \widehat{\alpha} \in \widehat{F}_k \wedge \operatorname{im}(\alpha) \subseteq [t])$

(note that $\mathcal{B}_0$ is the trivial σ-algebra and $\boldsymbol{\mathcal{N}}^{\boldsymbol{k},\boldsymbol{n}}$ is $\mathcal{B}_n$-measurable). For every $t \in [n]$, let also

$$C_t \stackrel{\text{def}}{=} \left\{ (x, x', x'') \in \mathcal{E}_{m,k-1} \times [0,1]^{\binom{[m]}{k}} \times [0,1]^{\binom{[m]}{>k}} \;\middle|\; t \in \operatorname{im}(\alpha_{x'}) \right\}$$

and note that

$$\lambda(C_t) = 1 - \left(1 - \frac{1}{n}\right)^{\binom{m}{k}} \leq O_{m,k}\left(\frac{1}{n}\right).$$

For every $t \in \{0, \ldots, n\}$, define the random variable

$$\boldsymbol{X}_t \stackrel{\text{def}}{=} \mathbb{E}[t_{\text{ind}}(K, \boldsymbol{\mathcal{N}}^{\boldsymbol{k},\boldsymbol{n}}) \mid \mathcal{B}_t]$$

so that $(\boldsymbol{X}_t)_{t=0}^n$ forms a (Doob) martingale such that

$$\boldsymbol{X}_n = t_{\text{ind}}(K, \boldsymbol{\mathcal{N}}^{\boldsymbol{k},\boldsymbol{n}}); \qquad \boldsymbol{X}_0 = \mathbb{E}[t_{\text{ind}}(K, \boldsymbol{\mathcal{N}}^{\boldsymbol{k},\boldsymbol{n}})]. \tag{3.18}$$

Note also that for every $t \in [n]$, we have $|\boldsymbol{X}_t - \boldsymbol{X}_{t-1}| \leq \lambda(C_t) \leq O_{m,k}(n^{-1})$, so by Azuma's Inequality (see e.g. [3, Theorem 7.2.1]) and (3.18), we get

$$\begin{aligned}
\mathbb{P}[|t_{\text{ind}}(K, \boldsymbol{\mathcal{N}}^{\boldsymbol{k},\boldsymbol{n}}) - \mathbb{E}[t_{\text{ind}}(K, \boldsymbol{\mathcal{N}}^{\boldsymbol{k},\boldsymbol{n}})]| > \varepsilon] &\leq 2\exp\left(-\frac{\varepsilon^2}{n \cdot O_{m,k}(n^{-1})^2}\right) \\
&= 2\exp(-\varepsilon^2 \cdot \Omega_{m,k}(n))
\end{aligned}$$

for every $\varepsilon > 0$. Thus (3.17) follows Lemma 3.4.6 and a standard Borel–Cantelli argument. ■

We can finally derive Proposition 3.4.3.

Proof of Proposition 3.4.3. Follows immediately from Remark 3 and Lemma 3.4.7. ■

3.5 Low rank theories

In this section we explore how the axioms of a theory T can force all of its limit objects to have low rank; this is captured by the following definition.

Definition 3.5.1. For a theory T, the *rank* of T the maximum rank $\mathrm{rk}(\phi)$ of some $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ (if T is degenerate, we declare $\mathrm{rk}(T) = -\infty$).

The classic example is that of T_{LinOrder}, whose axioms force its unique limit object, represented by the T_{LinOrder}-on $\mathcal{N} \stackrel{\mathrm{def}}{=} \{x \in \mathcal{E}_2 \mid x_{\{1\}} < x_{\{2\}}\}$, to have rank 1 (even though the arity is 2), so $\mathrm{rk}(T_{\mathrm{LinOrder}}) = 1$. The objective of this section is to study examples of theories T obtained from $T_{k\text{-Hypergraph}}$ by adding axioms that reduce $\mathrm{rk}(T)$ to some fixed $r \leq k$. We start with a some examples of low rank theories obtained by using the notion of interpreted theories defined below.

Definition 3.5.2. Let $I\colon T_1 \rightsquigarrow T_2$ be an open interpretation and let T_2' be a universal theory obtained from T_2 by adding axioms. The *interpreted theory* $I(T_2')$ is the universal theory of all models M of T_1 such that there exists some model N of T_2' with $M \cong I(N)$. Formally, we let $I(T_2')$ be the theory whose axioms are

$$\forall x_1 \cdots \forall x_\ell, \left(\bigwedge_{1\leq i<j\leq \ell} x_i \neq x_j\right) \to \left(\bigvee_{K\in\mathcal{K}_\ell[T_2']} D_{\mathrm{open}}(I(K))(x_1, \ldots, x_\ell)\right),$$

for every $\ell \in \mathbb{N}_+$.

Remark 4. Note that every $\phi \in \mathrm{Hom}^+(\mathcal{A}[I(T_2')], \mathbb{R})$ is of the form $\phi = \psi^I$ for some $\psi \in \mathrm{Hom}^+(\mathcal{A}[T_2'], \mathbb{R}) \subseteq \mathrm{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$. This follows since if a sequence $(N_n)_{n\in\mathbb{N}}$ of models of T_1 converges to ϕ, then there must exist models $\widehat{N}_n$ of T_2' such that $I(\widehat{N}_n) = N_n$, then any convergent subsequence of $(\widehat{N}_n)_{n\in\mathbb{N}}$ converges to a $\psi \in \mathrm{Hom}^+(\mathcal{A}[T_2'], \mathbb{R}) \subseteq \mathrm{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ with $\phi = \psi^I$.

Proposition 3.5.3. *Let $k \in \mathbb{N}_+$ and let $r \in \{-\infty, 0, 1, \ldots, k\}$. Then there exists a theory T obtained from $T_{k\text{-Hypergraph}}$ by adding axioms such that $\mathrm{rk}(T) = r$.*

Proof. If $r = -\infty$, form a degenerate T by adding a contradiction, say, $\forall x, x \neq x$. If $r = 0$, add the axiom $\forall \vec{x}, \neg E(\vec{x})$ so that T consists of the theory of empty graphs, which clearly has rank 0.

Suppose then that $r \in [k]$ and consider the open interpretation $I\colon T_{k\text{-Hypergraph}} \rightsquigarrow T_{r\text{-Hypergraph}}$ that declares k-edges to be k-cliques, that is, it is given by

$$I(E)(x_1, \ldots, x_k) \stackrel{\text{def}}{=} \bigwedge_{1 \leq i_1 < \cdots < i_r \leq k} E(x_{i_1}, \ldots, x_{i_r}).$$

Let $\psi_{r,1/2} \in \operatorname{Hom}^+(\mathcal{A}[T_{r\text{-Hypergraph}}], \mathbb{R})$ be the quasirandom r-hypergraphon of density 1/2, that is, it is represented by

$$\mathcal{N} \stackrel{\text{def}}{=} \left\{ x \in \mathcal{E}_r \mid x_{[r]} < \frac{1}{2} \right\}.$$

Let also[2] $T \stackrel{\text{def}}{=} I(T_{r\text{-Hypergraph}})$. Since $\psi_{r,1/2}^I$ is represented by

$$I(\mathcal{N}) \stackrel{\text{def}}{=} \left\{ x \in \mathcal{E}_k \;\middle|\; \forall A \in \binom{[k]}{r}, x_A < \frac{1}{2} \right\},$$

it follows that $\operatorname{rk}(T) \geq \operatorname{rk}(\psi_{r,1/2}^I) = r$ by Remark 1.

On the other hand, since any $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ must be of the form $\phi = \psi^I$ for some $\psi \in \operatorname{Hom}^+(\mathcal{A}[T_{r\text{-Hypergraph}}], \mathbb{R})$ (see Remark 4), we get $\operatorname{rk}(\phi) \leq \operatorname{rk}(\psi) \leq r$. Therefore $\operatorname{rk}(T) = r$. ■

Recall that a point $x \in \mathcal{E}_n([0,1]^d)$ is called a *Lebesgue density point* of a measurable set $A \subseteq \mathcal{E}_n([0,1]^d)$ (relative to $\mathcal{E}_n([0,1]^d)$) if

$$\lim_{\varepsilon \to 0} \frac{\lambda(B(x, \varepsilon) \cap A)}{\lambda(B(x, \varepsilon) \cap \mathcal{E}_n([0,1]^d))} = 1,$$

2. Even though the case $r = k$ is trivial, this construction still works as then $I = \operatorname{id}_{T_{k\text{-Hypergraph}}}$ and $T = T_{k\text{-Hypergraph}}$.

where $B(x, \varepsilon)$ denotes the ℓ_∞-ball[3] of radius ε around x. We denote the set of Lebesgue density points of A by $D(A)$. It is easy to see that $D(D(A)) = D(A)$ and the Lebesgue Density Theorem (see [8, I-5.6(ii)] and [56, Theorem 3.21]) implies that for every measurable set A we have $\lambda(D(A) \triangle A) = 0$.

The next proposition gives a sufficient condition for a theory obtained from $T_{k\text{-Hypergraph}}$ to have low rank.

Proposition 3.5.4. *Let $r, k \in \mathbb{N}$ with $r < k$, let $V_1, \ldots, V_k$ be pairwise disjoint finite non-empty sets and let $F, \overline{F} \subseteq \prod_{i\in[k]} V_i$ be disjoint sets satisfying the following consistency condition: if $\alpha_1, \alpha_2 \in F \cup \overline{F}$ have strictly more than r coordinates in common, then $\alpha_1 \in F \iff \alpha_2 \in F$.*

Let T be a theory obtained from $T_{k\text{-Hypergraph}}$ by adding axioms and suppose that T entails

$$\forall \vec{x}, \neg \left(\bigwedge_{i \neq j} x_i \neq x_j \wedge \bigwedge_{(i_1,\ldots,i_k)\in F} E(x_{i_1}, \ldots, x_{i_k}) \wedge \bigwedge_{(i_1,\ldots,i_k)\in \overline{F}} \neg E(x_{i_1}, \ldots, x_{i_k}) \right), \tag{3.19}$$

where the variables are indexed by $\bigcup_{i\in[k]} V_i$. Then rk$(T) \leq r$.

Proof. Suppose toward a contradiction that there exists $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ with $\mathrm{rk}(\phi) \geq r + 1$, let $\mathcal{N}$ be a T-on such that $\phi = \phi_{\mathcal{N}}$ and let $G(x_1, \ldots, x_n)$ be the open formula

$$\bigwedge_{i \neq j} x_i \neq x_j \wedge \bigwedge_{(i_1,\ldots,i_k)\in F} E(x_{i_1}, \ldots, x_{i_k}) \wedge \bigwedge_{(i_1,\ldots,i_k)\in \overline{F}} \neg E(x_{i_1}, \ldots, x_{i_k})$$

Our objective is to show that $\lambda(T(G, \mathcal{N})) > 0$.

By Remark 1, the function $W(x) \stackrel{\text{def}}{=} \lambda(\{y \in [0,1]^{\binom{[k]}{>r}} \mid (x, y) \in \mathcal{N}\})$ is *not* 0-1 valued a.e. This means that there exists $\delta > 0$ such that the set $W^{-1}((\delta, 1-\delta))$ has positive Lebesgue

3. In fact, one can use other norms to define Lebesgue density points and get an a.e. equivalent definition, but for us it will be slightly more convenient to use the ℓ_∞-norm.

measure, which implies that the set

$$\mathcal{G} \stackrel{\text{def}}{=} \{(x,y,z) \in \mathcal{E}_{k,r} \times [0,1]^{\binom{[k]}{>r}} \times [0,1]^{\binom{[k]}{>r}} \mid (x,y) \in \mathcal{N} \land (x,z) \in \mathcal{E}_k \setminus \mathcal{N}\}$$

has measure at least

$$\lambda(W^{-1}((\delta, 1-\delta)))\delta^2 > 0.$$

Let $(x^*, y^*, z^*) \in D(\mathcal{G})$ be a Lebesgue density point of $\mathcal{G}$ and let $\varepsilon > 0$ be small enough so that $B((x^*, y^*, z^*), \varepsilon) \subseteq \mathcal{E}_{k,r} \times [0,1]^{\binom{[k]}{>r}} \times [0,1]^{\binom{[k]}{>r}}$ and

$$\frac{\lambda(B((x^*, y^*, z^*), \varepsilon) \cap \mathcal{G})}{\lambda(B((x^*, y^*, z^*), \varepsilon)} > 1 - \frac{1}{|F \cup \overline{F}|}. \tag{3.20}$$

Let $V \stackrel{\text{def}}{=} \bigcup_{i \in [k]} V_i$, let $\chi\colon V \to [k]$ be the unique function such that $v \in V_{\chi(v)}$ for every $v \in V$. Let us call a set $A \in r(V)$ *χ-transversal* if $\chi|_A$ is injective and let $\mathcal{T} \subseteq r(V)$ be the set of all χ-transversal sets. Consider the set

$$\mathcal{U} \stackrel{\text{def}}{=} \{x \in \mathcal{E}_{V,r} \mid \forall A \in r(V,r) \cap \mathcal{T}, x_A \in (x^*_{\chi(A)} - \varepsilon, x^*_{\chi(A)} + \varepsilon)\}.$$

Let us now define a function $g\colon \binom{V}{>r} \cap \mathcal{T} \to \{F, \overline{F}, *\}$ by letting $g(A)$ be equal to

i. F, if there exists $\alpha \in F$ such that $A \subseteq \operatorname{im}(\alpha)$.

ii. $\overline{F}$, if there exists $\alpha \in \overline{F}$ such that $A \subseteq \operatorname{im}(\alpha)$.

iii. $*$ if for every $\alpha \in F \cup \overline{F}$, we have $A \not\subseteq \operatorname{im}(\alpha)$.

Note now that our consistency condition implies that for each $A \in \binom{V}{>r} \cap \mathcal{T}$, *exactly one* of the above holds, so g is well-defined. We now let

$$\mathcal{R} \overset{\text{def}}{=} \Big\{ y \in [0,1]^{\binom{V}{>r}} \;\Big|\; \forall A \in \binom{V}{>r} \cap \mathcal{T}, \big(g(A) = F \to y_A \in (y^*_{\chi(A)} - \varepsilon, y^*_{\chi(A)} + \varepsilon)\big) \wedge \big(g(A) = \overline{F} \to y_A \in (z^*_{\chi(A)} - \varepsilon, z^*_{\chi(A)} + \varepsilon)\big) \Big\}.$$

Finally, we will show that $\lambda(T(G, \mathcal{N})) > 0$ by showing that a point $(\boldsymbol{x}, \boldsymbol{y})$ picked uniformly at random in $\mathcal{U} \times \mathcal{R}$ has positive probability of belonging to $T(G, \mathcal{N})$.

Note first that the union bound gives

$$\begin{aligned} \mathbb{P}[(\boldsymbol{x}, \boldsymbol{y}) \in T(G, \mathcal{N})] &= \mathbb{P}[(\forall \alpha \in F, \alpha^*(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{N}) \wedge (\forall \alpha \in \overline{F}, \alpha^*(\boldsymbol{x}, \boldsymbol{y}) \notin \mathcal{N})] \\ &= 1 - \sum_{\alpha \in F} (1 - \mathbb{P}[\alpha^*(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{N})]) - \sum_{\alpha \in \overline{F}} (1 - \mathbb{P}[\alpha^*(\boldsymbol{x}, \boldsymbol{y}) \notin \mathcal{N}]), \end{aligned}$$

where we interpret α in the above as an element of $(V)_k$.

Now, for $\alpha \in F$, note that $\alpha^*(\boldsymbol{x}, \boldsymbol{y})$ has uniform distribution in $B((x^*, y^*), \varepsilon)$, so from (3.20), we get

$$\mathbb{P}[\alpha^*(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{N}] \geq \frac{\lambda(B((x^*, y^*, z^*), \varepsilon) \cap \mathcal{G})}{\lambda(B((x^*, y^*, z^*), \varepsilon))} > 1 - \frac{1}{|F \cup \overline{F}|}.$$

On the other hand, for $\alpha \in \overline{F}$, $\alpha^*(\boldsymbol{x}, \boldsymbol{y})$ has uniform distribution in $B((x^*, z^*), \varepsilon)$, so from (3.20), we get

$$\mathbb{P}[\alpha^*(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{N}] \geq \frac{\lambda(B((x^*, y^*, z^*), \varepsilon) \cap \mathcal{G})}{\lambda(B((x^*, y^*, z^*), \varepsilon))} > 1 - \frac{1}{|F \cup \overline{F}|}.$$

Putting these together, we get $\mathbb{P}[(\boldsymbol{x}, \boldsymbol{y}) \in T(G, \mathcal{N})] > 0$ as desired. ■

Remark 5. Of course, an easy way to ensure that the consistency condition of Proposition 3.5.4 holds is to require the stronger condition that any distinct $\alpha_1, \alpha_2 \in F \cup \overline{F}$ have at most r coordinates in common. However, there are very natural low rank theories that are not captured by the analogue of Proposition 3.5.4 with this stronger condition for $r \leq k-2$.

An example for $r \leq k-2$ is the theory T of ℓ-linear k-hypergraphs, that is, k-hypergraphs in which any two distinct hyperedges intersect in at most ℓ points. If $\ell \leq k-2$, all such hypergraphs are sparse, so the only $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is the empty k-hypergraphon, which has rank 0, hence $\mathrm{rk}(T) = 0$. However, if $V_1, \ldots, V_k, F, \overline{F}$ satisfy the stronger condition for some $r \leq \ell$, then the k-hypergraph H defined by

$$V(H) \stackrel{\text{def}}{=} \bigcup_{i \in [k]} V_i;$$

$$R_E(H) \stackrel{\text{def}}{=} \{\alpha \in (V(H))_k \mid \exists \beta \in F, \mathrm{im}(\alpha) = \mathrm{im}(\beta)\}$$

is ℓ-linear and violates (3.19), so $\mathrm{rk}(T) = 0$ is not captured by a version of Proposition 3.5.4 with the stronger condition.

However, $\mathrm{rk}(T) = 0$ is captured by the normal version of Proposition 3.5.4 by taking V_i with a single element for $i \leq \ell + 1$, taking V_i with two elements for $i > \ell + 1$, letting F have exactly two k-tuples with exactly the first $\ell + 1$ coordinates in common and letting $\overline{F}$ be empty. It is clear that T entails (3.19) for this choice of $(V_1, \ldots, V_k, F, \overline{F})$ and since $\overline{F} = \varnothing$, the consistency condition is satisfied trivially for any $r \in \mathbb{N}$, so Proposition 3.5.4 is able to conclude that $\mathrm{rk}(T) = 0$.

We finish this section by showing a converse to Proposition 3.5.4 when $r = k-1$ (note that the consistency condition is omitted because it is trivially satisfied when $r = k-1$).

Theorem 3.5.5. *For pairwise disjoint finite non-empty sets $V_1, \ldots, V_k$ and disjoint sets $F, \overline{F} \subseteq \prod_{i \in [k]} V_i$, let $G_{V_1, \ldots, V_k, F, \overline{F}}(\vec{x})$ be the formula*

$$\bigwedge_{i \neq j} x_i \neq x_j \wedge \bigwedge_{(i_1, \ldots, i_k) \in F} E(x_{i_1}, \ldots, x_{i_k}) \wedge \bigwedge_{(i_1, \ldots, i_k) \in \overline{F}} \neg E(x_{i_1}, \ldots, x_{i_k})$$

The following are equivalent for a theory T obtained from $T_{k\text{-Hypergraph}}$ by adding axioms.

i. $rk(T) \leq k - 1$.

ii. *There exist pairwise disjoint finite non-empty sets* $V_1, \ldots, V_k$ *and disjoint sets* $F, \overline{F} \subseteq \prod_{i\in[k]} V_i$ *such that* T *entails* $\forall \vec{x}, \neg G_{V_1,\ldots,V_k,F,\overline{F}}(\vec{x})$.

iii. *There exists* $d \in \mathbb{N}_+$ *such that* T *entails* $\forall \vec{x}, \neg G_{V_1,\ldots,V_k,F,\overline{F}}(\vec{x})$ *for*

$$\begin{aligned}
V_i &\stackrel{\text{def}}{=} [d] \times \{i\} \qquad (i \in [k-1]);\\
V_k &\stackrel{\text{def}}{=} 2^{\prod_{i\in[k-1]} V_i};\\
F &\stackrel{\text{def}}{=} \left\{ (v_1, \ldots, v_{k-1}, A) \in \prod_{i\in[k]} V_i \;\middle|\; (v_1, \ldots, v_{k-1}) \in A \right\};\\
\overline{F} &\stackrel{\text{def}}{=} \left(\prod_{i\in[k]} V_k \right) \setminus F.
\end{aligned} \tag{3.21}$$

Proof. Implication (iii) $\implies$ (i) follows from Proposition 3.5.4.

For implication (ii) $\implies$ (iii), by considering the contra-positive, note that if there exists a model M of T in which $G_{V_1,\ldots,V_k,F,\overline{F}}(\vec{a})$ holds for some choice of $\vec{a}$ and $(V_1, \ldots, V_k, F, \overline{F})$ as in (3.21) for some $d \in \mathbb{N}_+$ then for any choice of $(V_1', \ldots, V_k', F', \overline{F}')$ as in item (ii) with $\max\{|V_i'| \mid i \in [k]\} \leq d/2$, we can find a solution $\vec{b}$ of $G_{V_1',\ldots,V_k',F',\overline{F}'}$ in M using the tuple $\vec{a}$.

Namely, for each $v \in V_i'$ and $i \in [k-1]$, we assign to b_v a distinct a_{w_v} with $w_v \in V_i$ (this is possible as $d \geq |V_i'|$). Then we assign to each b_v with $v \in V_k$ a distinct a_{A_v} with $A_v \in V_k$ such that

$$\begin{aligned}
\forall (v_1, \ldots, v_{k-1}) \in \prod_{i\in[k-1]} V_i', &\big((v_1, \ldots, v_{k-1}, v) \in F \to (w_{v_1}, \ldots, w_{v_{k-1}}) \in A_v\big)\\
&\wedge \big((v_1, \ldots, v_{k-1}, v) \in \overline{F} \to (w_{v_1}, \ldots, w_{v_{k-1}}) \notin A_v\big).
\end{aligned}$$

The choice of the A_v can be made distinct since $d - |V_i'| \geq d/2 \geq |V_k'|$ for every $i \in [k-1]$.

Finally, let us show implication (i) $\implies$ (ii).

Suppose not, that is, suppose that for each choice of $C = (n, V_1, \ldots, V_k, F, \overline{F})$ as in item (ii), there exists a model $K_C \in \mathcal{K}_n[T]$ such that $F \subseteq R_E(K_C)$ and $\overline{F} \cap R_E(K_C) = \varnothing$.

We can rephrase the property above as follows. Let T' be the theory obtained from $T_{k\text{-Hypergraph}} \cup T_{k\text{-Coloring}}$ by forbidding any non-rainbow edges, that is, we add the axiom

$$\forall x_1 \cdots \forall x_k, E(x_1, \ldots, x_k) \to \bigwedge_{1 \leq i < j \leq k} \bigwedge_{t \in [k]} \neg(\chi_t(x_i) \land \chi_t(x_j)).$$

Consider the interpretation $I\colon T' \rightsquigarrow T \cup T_{k\text{-Coloring}}$ that acts identically on the coloring and removes non-rainbow edges, that is, it is given by

$$I(\chi_i)(x) \stackrel{\text{def}}{=} \chi_i(x) \qquad (i \in [k]);$$
$$I(E)(x_1, \ldots, x_k) \stackrel{\text{def}}{=} E(x_1, \ldots, x_k) \land \bigwedge_{1 \leq i < j \leq k} \bigwedge_{t \in [k]} \neg(\chi_t(x_i) \land \chi_t(x_j)).$$

Then the property above can be restated as follows: for every $M \in \mathcal{M}[T']$, there exists $N \in \mathcal{M}[T \cup T_{k\text{-Coloring}}]$ such that $I(N) = M$, that is, we have $I(T \cup T_{c\text{-Coloring}}) \vdash T'$ (see Definition 3.5.2). Indeed, assuming without loss of generality that $V(M) = [n]$ and that each $R_{\chi_i}(M)$ is non-empty (otherwise, we can use a larger model M' with $R_\chi(M')$ non-empty and with M as an induced submodel), one such N is obtained as K_C for the choice $C = (n, R_{\chi_1}(M), \ldots, R_{\chi_k}(M), F, \overline{F})$, where

$$F \stackrel{\text{def}}{=} R_E(M) \cap \prod_{i \in [k]} R_{\chi_i}(M)$$
$$\overline{F} \stackrel{\text{def}}{=} \left(\prod_{i \in [k]} R_{\chi_i}(M)\right) \setminus R_E(M).$$

Now consider the T'-on $\mathcal{N}$ given by

$$\mathcal{N}_{\chi_i} \stackrel{\text{def}}{=} \left\{x \in \mathcal{E}_1 \;\middle|\; \frac{i-1}{k} \leq x < \frac{i}{k}\right\} \qquad (i \in [k]);$$
$$\mathcal{N}_E \stackrel{\text{def}}{=} \left\{x \in \mathcal{E}_k \;\middle|\; x_{[k]} < \frac{1}{2} \land \forall i, j, t \in [k], (i \neq j \to x_{\{i\}} \notin \mathcal{N}_{\chi_t} \lor x_{\{i\}} \notin \mathcal{N}_{\chi_t})\right\}$$

and let $\phi = \phi_{\mathcal{N}} \in \operatorname{Hom}^+(\mathcal{A}[T'], \mathbb{R})$. Our property implies that there exists $\psi \in \operatorname{Hom}^+(\mathcal{A}[T \cup T_{k\text{-Coloring}}], \mathbb{R})$ such that $\phi = \psi^I$ (cf. Remark 4) and by Proposition 3.2.1, there exists a $(T \cup T_{k\text{-Coloring}})$-on $\mathcal{H}$ over $[0,1]^2$ such that $\phi_{\mathcal{H}} = \psi$, $I(\mathcal{H})_E = \mathcal{N}_E \times \mathcal{E}_k$ a.e. and $\mathcal{H}_{\chi_i} = \mathcal{N}_{\chi_i} \times \mathcal{E}_1$ a.e. for every $i \in [k]$.

Note now that for a.e. $(x, y) \in \mathcal{E}_{k,k-1} \times \mathcal{E}_{k,k-1}$, if $(x_{\{i\}}, y_{\{i\}}) \in \mathcal{H}_{\chi_i}$ for every $i \in [t]$, then

$$\lambda(\{(\widehat{x}, \widehat{y}) \in [0,1]^2 \mid (x, \widehat{x}, y, \widehat{y}) \in \mathcal{H}_E\}) = \lambda(\{\widehat{x} \in [0,1] \mid (x, \widehat{x}) \in \mathcal{N}_E\}) = \frac{1}{2}.$$

Since for the structure-erasing interpretation $J\colon T \rightsquigarrow T \cup T_{c\text{-Coloring}}$, the homomorphism $\psi^J \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is represented by $J(\mathcal{H})$ and since $(x, y) \in J(\mathcal{H}) \iff (x, y) \in \mathcal{H}_E$ whenever $(x_{\{i\}}, y_{\{i\}}) \in \mathcal{H}_{\chi_i}$ for every $i \in [t]$, by Remark 1, we have $\operatorname{rk}(\psi^J) \geq k$ so $\operatorname{rk}(T) \geq k$. ■

Remark 6. Note that in the case $k = 2$, item (iii) with parameter d is equivalent to saying that the Vapnik–Chervonenkis dimension (VC dimension, see [69, 70, 68]) of neighborhoods of vertices in models of T is at most $d - 1$, that is, if M is a model of T and

$$\mathcal{F} \stackrel{\text{def}}{=} \{N_M(v) \mid v \in V(M)\};$$
$$N_M(v) \stackrel{\text{def}}{=} \{w \in V(M) \mid (v, w) \in R_E(M)\};$$

then for every $V \in \binom{V(M)}{d}$, we have

$$|\{V \cap F \mid F \in \mathcal{F}\}| < 2^d,$$

i.e., $\mathcal{F}$ does not *shatter* V.

In turn, this property in model theory is known as saying that the formula $E(x, y)$ satisfies d-NIP in T, which stands for "not d-independence property", but it is not directly related to our `Independence` property of Definition 3.1.1.

When $k > 2$, item (iii) is equivalent to the VC dimension of neighborhoods of vertices being bounded by some $d' \in \mathbb{N}$ (there is a loss in the parameter), that is, there exists $d' \in \mathbb{N}$ such that if M is a model of T and

$$\mathcal{F} \stackrel{\text{def}}{=} \{N_M(v) \mid v \in V(M)\};$$
$$N_M(v) \stackrel{\text{def}}{=} \left\{\{v_1, \ldots, v_{k-1}\} \in \binom{V(M)}{k-1} \;\middle|\; (v, v_1, \ldots, v_{k-1}) \in R_E(M)\right\};$$

then for every $V \subseteq \binom{V(M)}{k-1}$ with $|V| = d'$, we have

$$|\{V \cap N \mid N \in \mathcal{F}\}| < 2^{d'},$$

i.e., $\mathcal{F}$ does not *shatter* V. This is equivalent to saying that $E(x, \vec{y})$ satisfies d'-NIP in T.

3.6 Strengthening theon lifting

Recall that Proposition 3.2.1 said that if $\mathcal{N}$ is a T_1-on and $\phi_{\mathcal{N}} = \phi^I$ for some $I\colon T_1 \rightsquigarrow T_2$, then we can find a T_2-on $\mathcal{H}$ such that $I(\mathcal{H})_P = \mathcal{N}_P \times \mathcal{E}_{k(P)}$ a.e. and $\phi_{\mathcal{H}} = \phi$. It is natural to ask if we can strengthen this proposition to require that $I(\mathcal{H})_P = \mathcal{N}_P \times \mathcal{E}_{k(P)}$ *everywhere* (except for the diagonal) even when I is not necessarily a structure-erasing interpretation. The example below illustrates the main obstacle to such a generalization.

Example 2. Consider the interpretation $I\colon T_{3\text{-Hypergraph}} \rightsquigarrow T_{\text{Graph}}$ of triangles of a graph given by

$$I(E)(x_1, x_2, x_3) = \bigwedge_{1 \le i < j \le 3} (E(x_i, x_j) \vee E(x_j, x_i)).$$

(We write it in a slightly different way so that one cannot violate the symmetry axiom of T_{Graph} to avoid the problem illustrated here.)

Let $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_{\text{Graph}}], \mathbb{R})$ and take any $T_{3\text{-Hypergraph}}$-on $\mathcal{N}$ with $\phi^I = \phi_{\mathcal{N}}$ but such that $T_{\text{ind}}(K_4^-, \mathcal{N}) \not\subseteq \mathcal{D}_4$, where K_4^- is the 3-hypergraph with 4 vertices and 3 edges; this can be done by adding a zero-measure amount of off-diagonal copies of K_4^- to $I(\mathcal{N}')$ for a T_{Graph}-on $\mathcal{N}'$ representing ϕ. Then no T_{Graph}-on $\mathcal{H}$ satisfies an off-diagonal everywhere version of Proposition 3.2.1 simply because $I(\mathcal{H})$ can never contain any off-diagonal copies of K_4^-. This remains the case even if we assume $\mathcal{N}$ to be a strong $T_{3\text{-Hypergraph}}$-on.

The obstacle illustrated by Example 2 is that there is a "hidden axiom" coming from the fact that no 3-hypergraph obtained as $I(G)$ for some graph G can have a copy of K_4^-. To present to natural hypotheses that could be added we need one more definition.

Definition 3.6.1. For $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$, let $\operatorname{Th}(\phi)$ be the theory obtained from T by adding the axiom $\forall \vec{x}, \neg D_{\text{open}}(M)(\vec{x})$ for every $M \in \mathcal{M}[T]$ such that $\phi(M) = 0$, i.e., it is the theory whose models are precisely the ones that have positive density in ϕ.

To surpass the "hidden axiom" obstacle, a natural extra hypothesis would be to require $\mathcal{N}$ to be a strong $I(T_2)$-on. Alternatively, a more intrinsic condition on $\mathcal{N}$ would be to require it to be as strong as it can be without reference to I, that is, we could require it to be a strong $\operatorname{Th}(\phi_{\mathcal{N}})$-on. The next lemma says that this intrinsic hypothesis implies the natural hypothesis.

Lemma 3.6.2. *Let $I\colon T_1 \rightsquigarrow T_2$ be an open interpretation, let $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ and let $\mathcal{H}$ be a strong $\operatorname{Th}(\phi^I)$-on over Ω. Then $\mathcal{H}$ is a strong $I(\operatorname{Th}(\phi))$-on.*

Proof. We have to show that for every canonical structure M in the language of T_1 that is *not* a model of $I(\mathrm{Th}(\phi))$, we have $T_{\mathrm{ind}}(M,\mathcal{H}) \subseteq \mathcal{D}_{V(M)}(\Omega)$. We will show this by the contrapositive.

Let M be a canonical structure in the language of T_1 such that $T_{\mathrm{ind}}(M,\mathcal{H}) \not\subseteq \mathcal{D}_{V(M)}(\Omega)$. Since $\mathcal{H}$ is a strong $\mathrm{Th}(\phi^I)$-on, we must have $\phi^I(M) > 0$. But this in particular implies that there exists $N \in \mathcal{M}_{V(M)}[T_2]$ such that $I(N) \cong M$ and $\phi(N) > 0$, which in turn implies that $N \in \mathcal{M}[\mathrm{Th}(\phi)]$, so $M \in \mathcal{M}[I(\mathrm{Th}(\phi))]$. ■

Unfortunately, to actually show a generalization of Proposition 3.2.1, we will need the much stronger and less natural conditions on our theons that we define below.

Definition 3.6.3. Fix $d \in \mathbb{N}_+$. A theon $\mathcal{N}$ over $[0,1]^d$ is called *sound* if for every open formula $F(x_1,\ldots,x_n)$ the set $T(F,\mathcal{N})$ contains all of its off-diagonal Lebesgue density points (relative to $\mathcal{E}_n$), that is, we have $D(T(F,\mathcal{N})) \setminus \mathcal{D}_n([0,1]^d) \subseteq T(F,\mathcal{N})$.

Let $\varepsilon > 0$ and let $Z = (z^1,\ldots,z^t)$ be a finite sequence of points with $z^i \in \mathcal{E}_{n_i}([0,1]^d) \setminus \mathcal{D}_{n_i}([0,1]^d)$ for each $i \in [t]$. We define the *random ε-perturbation* $\boldsymbol{Z^\varepsilon} = (\boldsymbol{z^{1,\varepsilon}},\ldots,\boldsymbol{z^{t,\varepsilon}})$ of Z as follows. We let $C_Z \subseteq [0,1]^d$ be the set of all coordinates of the points in Z, that is, we let

$$C_Z = \{z^i_A \mid i \in [t], A \in r(n_i)\}$$

and for $x \in C_Z$ and $X \subseteq C_Z$, we introduce independent random variables $\boldsymbol{\xi^\varepsilon_Z}(x,X)$ that are uniformly distributed (according to λ^d) in $B(x,\varepsilon) \cap [0,1]^d$ (where $B(x,\varepsilon) \subseteq \mathbb{R}^d$ is the ℓ^∞-ball of radius ε centered at x). We then define the random variable $\boldsymbol{z^{i,\varepsilon}}$ in $\mathcal{E}_{n_i}([0,1]^d)$ by

$$\boldsymbol{z^{i,\varepsilon}_A} = \boldsymbol{\xi^\varepsilon_Z}(z^i_A, \{z^i_{\{j\}} \mid j \in A\}).$$

A theon $\mathcal{N}$ over $[0,1]^d$ is called *robust* if for every finite sequence $((F_i,z^i))_{i=1}^t$ where $F_i(x_1,\ldots,x_{n_i})$ is an open formula with n_i variables and $z^i \in \mathcal{E}_{n_i}([0,1]^d) \setminus \mathcal{D}_{n_i}([0,1]^d)$ (for the same $n_i \in \mathbb{N}_+$), we have

$$\limsup_{\varepsilon \to 0} \mathbb{P}[\forall i \in [t], \boldsymbol{z^{i,\varepsilon}} \in T(F_i, \mathcal{N}) \leftrightarrow z^{i,\varepsilon} \in T(F_i, \mathcal{N})] > 0.$$

The intuition of theon soundness is that for each open formula $F(x_1, \ldots, x_n)$, if "almost all" small perturbations of a point $z \in \mathcal{E}_n$ satisfy F, then z is required to satisfy F as well. Robustness takes this one step further by saying that small consistent perturbations of several points $z^1, \ldots, z^n$ should match the behavior of the points $z^1, \ldots, z^n$ with probability not going to 0.

Before proving properties on theon soundness and robustness, we need some basic properties of random ε-perturbations.

Lemma 3.6.4. *The following properties hold for random ε-perturbations of a sequence $Z = (z^1, \ldots, z^t)$.*

i. If $\alpha^*(z^i) = \beta^*(z^j)$ for some $\alpha\colon [k] \rightarrowtail [n_i]$ and some $\beta\colon [k] \rightarrowtail [n_j]$, then $\alpha^*(\boldsymbol{z^{i,\varepsilon}}) = \beta^*(\boldsymbol{z^{i,\varepsilon}})$.

ii. For $i \in [t]$ and $\alpha\colon [k] \rightarrowtail [n_i]$, the point $\alpha^*(\boldsymbol{z^{i,\varepsilon}})$ is uniformly distributed in $B(\alpha^*(z^i), \varepsilon) \cap \mathcal{E}_k([0,1]^d)$.

iii. If $Y = (y^1, \ldots, y^\ell)$ is such that $y^j = z^{i_j}$ for every $j \in [\ell]$, then $\boldsymbol{Y^\varepsilon}$ has the same distribution as $(\boldsymbol{z^{i_1,\varepsilon}}, \ldots, \boldsymbol{z^{i_\ell,\varepsilon}})$.

Proof. For item (i), note that for $A \in r(k)$, we have

$$\begin{aligned}\alpha^*(\boldsymbol{z^{i,\varepsilon}})_A = \boldsymbol{z^{i,\varepsilon}_{\alpha(A)}} = \boldsymbol{\xi^\varepsilon_Z}(z^i_{\alpha(A)}, \{z^i_{\{j\}} \mid t \in \alpha(A)\}) = \boldsymbol{\xi^\varepsilon_Z}(z^i_{\alpha(A)}, \{z^i_{\{\alpha(t)\}} \mid t \in A\}) \\ = \boldsymbol{\xi^\varepsilon_Z}(z^j_{\beta(A)}, \{z^j_{\{\beta(t)\}} \mid t \in A\}) = \boldsymbol{z^{j,\varepsilon}_{\beta(A)}} = \beta^*(\boldsymbol{z^{j,\varepsilon}})_A,\end{aligned}$$

where the fourth equality follows since $z^i_{\alpha(B)} = z^j_{\beta(B)}$ for every $B \in r(k)$ as $\alpha^*(z^i) = \beta^*(z^j)$.

For item (ii), note that if $A \in r(k)$, then $\alpha^*(\boldsymbol{z^{i,\varepsilon}})_A = \boldsymbol{\xi^{\varepsilon}_{Z}}(z^i_{\alpha(A)}, \{z^i_{\{\alpha(t)\}} \mid t \in A\})$, which is picked uniformly at random in $B(z^i_{\alpha(A)}, \varepsilon) \cap [0,1]^d$, thus it is sufficient to show that the coordinates of $\alpha^*(\boldsymbol{z^{i,\varepsilon}})$ are mutually independent. But indeed, each coordinate of $\alpha^*(\boldsymbol{z^{i,\varepsilon}})$ uses a different $\boldsymbol{\xi^{\varepsilon}_{Z}}$ random variable as the second parameter is always different: if $A, B \in r(k)$ are distinct, then $\{z^i_{\alpha(t)} \mid t \in A\} \neq \{z^i_{\beta(t)} \mid t \in B\}$ as $z^i \notin \mathcal{D}_{n_i}([0,1]^d)$.

The last item (iii) follows easily from construction. ■

The next lemma shows that theon robustness is stronger than theon soundness, which in turn is stronger than (the maximum possible) theon strength.

Lemma 3.6.5. *The following hold for a theon $\mathcal{N}$ over $[0,1]^d$.*

i. If $\mathcal{N}$ is robust, then $\mathcal{N}$ is sound.

ii. If $\mathcal{N}$ is sound, then $\mathcal{N}$ is a strong $\mathrm{Th}(\phi_{\mathcal{N}})$*-on.*

Proof. Suppose $\mathcal{N}$ is robust, let $F(x_1, \ldots, x_n)$ be an open formula and let $z \in D(T(F, \mathcal{N})) \setminus \mathcal{D}_n([0,1]^d)$. Since $\mathcal{N}$ is robust, for the sequence $Z = (z^1)$ with $z^1 = z$, we have

$$\limsup_{\varepsilon \to 0} \mathbb{P}[\boldsymbol{z^{1,\varepsilon}} \in T(F, \mathcal{N}) \leftrightarrow z \in T(F, \mathcal{N})] > 0,$$

but $z \in D(T(F, \mathcal{N}))$ implies $\lim_{\varepsilon \to 0} \mathbb{P}[\boldsymbol{z^{1,\varepsilon}} \in T(F, \mathcal{N})] = 1$, so we must have $z \in T(F, \mathcal{N})$.

Suppose now that $\mathcal{N}$ is sound and let M be a canonical structure on $[n]$ such that $\phi_{\mathcal{N}}(M) = 0$. Then $\lambda(T(\neg D_{\mathrm{open}}(M), \mathcal{N})) = 1$, so $D(T(\neg D_{\mathrm{open}}(M), \mathcal{N})) = \mathcal{E}_n([0,1]^d)$, which by soundness of $\mathcal{N}$ implies $T(\neg D_{\mathrm{open}}(M), \mathcal{N}) \supseteq \mathcal{E}_n([0,1]^d) \setminus \mathcal{D}_n([0,1]^d)$, hence $T_{\mathrm{ind}}(M, \mathcal{N}) \subseteq \mathcal{D}_n([0,1]^d)$. ■

Before we proceed to the generalization of Proposition 3.2.1, let us show that theons can be made robust by changing only a zero-measure set. The proof is based on the Induced Euclidean Removal Lemma [24, Theorem 3.3]. In fact, it is worth noting that the Induced Euclidean Removal Lemma actually produced theons that satisfied the soundness condition at least when F is a literal.

Theorem 3.6.6 (Euclidean Robustness Lemma)**.** *Let T be a canonical theory over a language $\mathcal{L}$. If $\mathcal{N}$ is a weak T-on over $[0,1]^d$, then there exists a robust T-on $\mathcal{N}'$ over $[0,1]^d$ such that*

$$\lambda(\mathcal{N}_P \mathbin{\triangle} \mathcal{N}'_P) = 0$$

for every predicate symbol $P \in \mathcal{L}$.

Proof. We prove the case $d = 1$, the general case $d \in \mathbb{N}_+$ is completely analogous. For $P \in \mathcal{L}$, let us call a point $y \in \mathcal{E}_{k(P)} \setminus \mathcal{D}_{k(P)}$ *bad for P* if $y \notin D(\mathcal{N}_P) \cup D(\mathcal{E}_{k(P)} \setminus \mathcal{N}_P)$ (i.e., if y is neither a density point of $\mathcal{N}_P$ nor of its complement) and let $\mathcal{B}_P \subseteq \mathcal{E}_{k(P)} \setminus \mathcal{D}_{k(P)}$ be the set of all points that are bad for P.

We introduce an uncountable set of propositional variables $v = (v_{P,y} \mid P \in \mathcal{L}, y \in \mathcal{B}_P)$ and define the Euclidean structure $\mathcal{N}^v$ by

$$\mathcal{N}^v_P \stackrel{\text{def}}{=} D(\mathcal{N}_P) \cup \{y \in \mathcal{B}_P \mid v_{P,y} = 1\} \qquad (P \in \mathcal{L}).$$

It is clear that for any assignment u of the variables v, we have $\lambda(\mathcal{N}_P \mathbin{\triangle} \mathcal{N}^u_P) = 0$, which in particular implies that $\mathcal{N}^u$ is a T-on. This property also extends: for an open formula $F(x_1, \ldots, x_n)$, we have $\lambda(T(F, \mathcal{N}) \mathbin{\triangle} T(F, \mathcal{N}^u)) = 0$.

Our objective is to find an assignment u of the variables v so that $\mathcal{N}^u$ is robust. For this, we introduce uncountably many constraints on these variables. For each finite sequence $\mathcal{F} = (F_i, z^i)_{i=1}^t$, where $F_i(x_1, \ldots, x_{n_i})$ is an open formula with n_i variables and $z^i \in \mathcal{E}_{n_i} \setminus \mathcal{D}_{n_i}$, we introduce a constraint $\mathcal{R}(\mathcal{F})$ on the variables $v_{P,y}$ encoding

$$\limsup_{\varepsilon \to 0} \mathbb{P}[\forall i \in [t], \boldsymbol{z^{i,\varepsilon}} \in T(F_i, \mathcal{N}) \leftrightarrow z^i \in T(F_i, \mathcal{N}^v)] > 0. \tag{3.22}$$

Note that we replaced $\mathcal{N}^v$ with $\mathcal{N}$ in the first term since $\mathbb{P}[\boldsymbol{z^{i,\varepsilon}} \in T(F_i, \mathcal{N}) \leftrightarrow \boldsymbol{z^{i,\varepsilon}} \in T(F_i, \mathcal{N}^u)] = 1$ as $\lambda(T(F_i, \mathcal{N}) \mathbin{\triangle} T(F_i, \mathcal{N}^u)) = 0$ for any assignment u of the variables v.

By the compactness theorem for propositional logic [11, Corollary 1.2.12], to show that this system is satisfiable, it is enough to show that any *finite* subsystem $\{\mathcal{R}(\mathcal{F}_1), \ldots, \mathcal{R}(\mathcal{F}_\ell)\}$ is satisfiable. On the other hand, by Lemma 3.6.4(iii), by taking the concatenation $\mathcal{F}$ of the sequences $\mathcal{F}_1, \ldots, \mathcal{F}_\ell$, the constraint $\mathcal{R}(\mathcal{F})$ implies all constraints $\mathcal{R}(\mathcal{F}_1), \ldots, \mathcal{R}(\mathcal{F}_\ell)$, so it is enough to show that $\mathcal{R}(\mathcal{F})$ is satisfiable for any *single* finite sequence $\mathcal{F} = (F_i, z^i)_{i=1}^t$.

For $P \in \mathcal{L}$, let Y_P be the set of all $y \in \mathcal{E}_{k(P)}$ such that the variable $v_{P,y}$ appears in the constraint $\mathcal{R}(\mathcal{F})$. This means that there exists $i \in [t]$ and $\alpha\colon [k(P)] \rightarrowtail [n_i]$ such that $y = \alpha^*(z^i)$ (in particular, we have $Y_P \cap \mathcal{D}_{k(P)} = \varnothing$). For $P \in \mathcal{L}$ and $y \in Y_P$, from the random ε-perturbation $\boldsymbol{Z^\varepsilon} = (\boldsymbol{z^{1,\varepsilon}}, \ldots, \boldsymbol{z^{t,\varepsilon}})$ of $Z = (z^1, \ldots, z^t)$, define the random variable $\boldsymbol{y^\varepsilon} = \alpha^*(\boldsymbol{z^{i,\varepsilon}})$, where z^i and $\alpha\colon [k(P)] \rightarrowtail [n_i]$ are such that $y = \alpha^*(z^i)$. From Lemma 3.6.4(i), this definition does not depend on the choice of α and z^i.

We define a random partial assignment $\boldsymbol{u^\varepsilon}$ of the variables v that assigns values only to the finitely many propositional variables $v' \stackrel{\text{def}}{=} (v_{P,y} \mid P \in \mathcal{L}, y \in Y_P)$ by letting $\boldsymbol{u^\varepsilon_{P,y}} = 1$ if and only if $\boldsymbol{y^\varepsilon} \in \mathcal{N}_P$. As there are only finitely many partial assignments assigning values only to v', there exists a partial assignment u' of v assigning only values to v' such that

$$\limsup_{\varepsilon \to 0} \mathbb{P}[\boldsymbol{u^\varepsilon} = u'] > 0. \tag{3.23}$$

From definition, the partial assignment u' determines all points of $\mathcal{N}^v$ that appear in the constraint $\mathcal{R}(\mathcal{F})$. Our objective is to show that any complete assignment u that extends u' satisfies $\mathcal{R}(\mathcal{F})$.

For $i \in [t]$, $P \in \mathcal{L}$ and $\alpha\colon [k(P)] \rightarrowtail [n_i]$, let $E^\varepsilon(i, P, \alpha)$ be the event $\alpha^*(\boldsymbol{z^{i,\varepsilon}}) \in \mathcal{N}_P \leftrightarrow \alpha^*(z^i) \in \mathcal{N}_P^u$. By our definition of $\boldsymbol{y^\varepsilon}$, the conjunction

$$\bigwedge\{E^\varepsilon(i, P, \alpha) \mid i \in [t], P \in \mathcal{L}, \alpha\colon [k(P)] \rightarrowtail [n_i], \alpha^*(z^i) \in \mathcal{B}_P\}$$

is implied by the event $\boldsymbol{u}^{\varepsilon} = u'$ in (3.23). Thus

$$\limsup_{\varepsilon\to 0} \mathbb{P}\left[\bigwedge\{E^{\varepsilon}(i,P,\alpha) \mid i \in [t], P \in \mathcal{L}, \alpha\colon [k(P)] \rightarrowtail [n_i], \alpha^*(z^i) \in \mathcal{B}_P\}\right] > 0. \tag{3.24}$$

On the other hand, if $\alpha^*(z^i) \notin \mathcal{B}_P$, then $\lim_{\varepsilon\to 0} \mathbb{P}[E^{\varepsilon}(i,P,\alpha)] = 1$ by the definition of Lebesgue density point and Lemma 3.6.4(ii). Putting this together with (3.24) gives

$$\limsup_{\varepsilon\to 0} \mathbb{P}\left[\bigwedge\{E^{\varepsilon}(i,P,\alpha) \mid i \in [t], P \in \mathcal{L}, \alpha\colon [k(P)] \rightarrowtail [n_i]\}\right] > 0. \tag{3.25}$$

Finally, since the event $\forall i \in [t], \boldsymbol{z^{i,\varepsilon}} \in T(F_i,\mathcal{N}) \leftrightarrow z^i \in T(F_i,\mathcal{N}^u)$ from (3.22) is implied by the conjunction in (3.25), we get that any complete assignment u extending u' satisfies the constraint $\mathcal{R}(\mathcal{F})$. ■

We can finally prove a generalization of Proposition 3.2.1 that works everywhere except for a weak version of the diagonal. We note that the idea and structure of the proof is quite similar to that of Theorem 3.6.6.

Theorem 3.6.7. *Let $I\colon T_1 \rightsquigarrow T_2$ be an open interpretation, let $\phi \in \mathrm{Hom}^+(\mathcal{A}[T_2],\mathbb{R})$ and let $\mathcal{N}$ be a robust T_1-on over $[0,1]^d$ such that $\phi_{\mathcal{N}} = \phi^I$. Then there exists a robust T_2-on $\mathcal{H}$ over $[0,1]^{2d}$ such that $\phi = \phi_{\mathcal{H}}$ and*

$$I(\mathcal{H})_P \mathbin{\triangle} (\mathcal{N}_P \times \mathcal{E}_{k(P)}([0,1]^d)) \subseteq \mathcal{D}_{k(P)}([0,1]^d) \times \mathcal{E}_{k(P)}([0,1]^d) \tag{3.26}$$

for every P in the language of T_1.

Proof. We prove the case $d = 1$, the general case $d \in \mathbb{N}_+$ is completely analogous. Let $\mathcal{L}_i$ be the language of T_i.

By Proposition 3.2.1, there exists a (weak) T_2-on $\mathcal{H}$ over $[0,1]^2$ such that $\phi = \phi_{\mathcal{H}}$ and $I(\mathcal{H})_P = \mathcal{N}_P \times \mathcal{E}_{k(P)}$ a.e. for every $P \in \mathcal{L}_1$.

For $Q \in \mathcal{L}_2$, let us call a point $y \in \mathcal{E}_{k(Q)}([0,1]^2) \setminus \mathcal{D}_{k(Q)}([0,1]^2)$ *bad for Q* if $y \notin$

$D(\mathcal{H}_Q) \cup D(\mathcal{E}_{k(Q)}([0,1]^2) \setminus \mathcal{H}_Q)$ and let $\mathcal{B}_Q \subseteq \mathcal{E}_{k(Q)}([0,1]^2) \setminus \mathcal{D}_{k(Q)}([0,1]^2)$ be the set of all points that are bad for Q.

Again, we introduce an uncountable set of propositional variables $v = (v_{Q,y} \mid Q \in \mathcal{L}_2, y \in \mathcal{B}_Q)$ and define the Euclidean structure $\mathcal{H}^v$ in $\mathcal{L}_2$ over $[0,1]^2$ by

$$\mathcal{H}^v_Q \stackrel{\text{def}}{=} D(\mathcal{H}_Q) \cup \{y \in \mathcal{B}_Q \mid v_{Q,y} = 1\} \qquad (Q \in \mathcal{L}_2), \tag{3.27}$$

and for any assignment u of the variables v and for an open formula $F(x_1, \ldots, x_n)$, we have $\lambda(T(F, \mathcal{H}^u) \triangle T(F, \mathcal{H})) = 0$, so $\mathcal{H}^u$ is a T_2-on with $\phi_{\mathcal{H}^u} = \phi_{\mathcal{H}} = \phi$.

Our objective is to find an assignment u of the variables v so that $\mathcal{H}^u$ is both robust and satisfies (3.26). For this, we introduce the following (uncountably many) constraints on these variables.

i. For each finite sequence $\mathcal{F} = (F_i, z^i)_{i=1}^t$, where $F_i(x_1, \ldots, x_{n_i})$ is an open formula on $\mathcal{L}_2$ with n_i variables and $z^i \in \mathcal{E}_{n_i}([0,1]^2) \setminus \mathcal{D}_{n_i}([0,1]^2)$, we introduce a constraint $\mathcal{R}(\mathcal{F})$ on the variables $v_{Q,y}$ encoding

$$\limsup_{\varepsilon \to 0} \mathbb{P}[\forall i \in [t], \boldsymbol{z^{i,\varepsilon}} \in T(F_i, \mathcal{H}) \leftrightarrow z^i \in T(F_i, \mathcal{H}^v)] > 0. \tag{3.28}$$

ii. For each $P \in \mathcal{L}_1$ and each $w \in (\mathcal{E}_{k(P)} \setminus \mathcal{D}_{k(P)}) \times \mathcal{E}_{k(P)}$, we introduce a constraint $\mathcal{I}(P, w)$ on the variables $v_{Q,y}$ encoding

$$w \in I(\mathcal{H}^v)_P \iff w \in \mathcal{N}_P \times \mathcal{E}_{k(P)}.$$

Again, we replaced $\mathcal{H}^v$ with $\mathcal{H}$ in the first term of (3.28) since $\mathbb{P}[\boldsymbol{z^{i,\varepsilon}} \in T(F_i, \mathcal{H}) \leftrightarrow \boldsymbol{z^{i,\varepsilon}} \in T(F_i, \mathcal{H}^u)] = 1$ as $\lambda(T(F_i, \mathcal{H}) \triangle T(F_i, \mathcal{H}^u)) = 0$ for any assignment u of the variables v.

By the compactness theorem for propositional logic [11, Corollary 1.2.12], to show that this system of restrictions is satisfiable, it is enough to show that any *finite* subsystem $\{\mathcal{R}(\mathcal{F}_1), \ldots, \mathcal{R}(\mathcal{F}_\ell), \mathcal{I}(P_1, w^1), \ldots, \mathcal{I}(P_m, w^m)\}$ is satisfiable. Again, by Lemma 3.6.4(iii), for the concatenation $\mathcal{F}$ of the sequences $\mathcal{F}_1, \ldots, \mathcal{F}_\ell$, the constraint $\mathcal{R}(\mathcal{F})$ implies all constraints $\mathcal{R}(\mathcal{F}_1), \ldots, \mathcal{R}(\mathcal{F}_\ell)$, so it is enough to show that any finite subsystem of the form $\{\mathcal{R}(\mathcal{F}), \mathcal{I}(P_1, w^1), \ldots, \mathcal{I}(P_m, w^m)\}$ is satisfiable, where $\mathcal{F} = (F_i, z^i)_{i=1}^t$. In fact, we can augment the sequence $\mathcal{F}$ so that for every $j \in [m]$ there exists $i_j \in [t]$ such that $F_{i_j}(x_1, \ldots, x_{k(P_j)})$ is $I(P_j)(x_1, \ldots, x_{k(P_j)})$ and $z^{i_j} = w^j$.

Consider the random ε-perturbation $\boldsymbol{Z^\varepsilon} = (\boldsymbol{z^{1,\varepsilon}}, \ldots, \boldsymbol{z^{t,\varepsilon}})$ of $Z = (z^1, \ldots, z^t)$ and for convenience, let us denote $\boldsymbol{w^{j,\varepsilon}} = \boldsymbol{z^{i_j,\varepsilon}}$ for $j \in [m]$.

For $Q \in \mathcal{L}_2$, let Y_Q be the set of all $y \in \mathcal{E}_{k(Q)}([0,1]^2)$ such that the variable $v_{Q,y}$ appears either in the constraint $\mathcal{R}(\mathcal{F})$ or in some constraint $\mathcal{I}(P_j, w^j)$. This means that there exists $i \in [t]$ and $\alpha\colon [k(Q)] \rightarrowtail [n_i]$ such that $y = \alpha^*(z^i)$, so we can define $\boldsymbol{y^\varepsilon} = \alpha^*(\boldsymbol{z^{i,\varepsilon}})$ and Lemma 3.6.4(i) implies that this does not depend on the choice of α and z^i.

We now define a random partial assignment $\boldsymbol{u^\varepsilon}$ of the variables v assigning values only to the finitely many variables $v' \stackrel{\text{def}}{=} (v_{Q,y} \mid Q \in \mathcal{L}_2, y \in Y_Q)$ by letting $\boldsymbol{u^\varepsilon_{Q,y}} = 1$ if and only if $\boldsymbol{y^\varepsilon} \in \mathcal{H}_Q$. Let $\mathcal{U}$ be the set of partial assignments u' of the variables v assigning values only to $(v_{Q,y} \mid Q \in \mathcal{L}_2, y \in Y_Q)$ and such that

$$\limsup_{r \to 0} \mathbb{P}[\boldsymbol{u^\varepsilon} = u'] > 0. \tag{3.29}$$

As there are only finitely many partial assignments assigning values only to v', we know that $\mathcal{U}$ is non-empty. Again, our definition ensures that the partial assignments $u' \in \mathcal{U}$ determine all points of $\mathcal{N}^v$ that appear in the constraints $\mathcal{R}(\mathcal{F}), \mathcal{I}(P_1, w^1), \ldots, \mathcal{I}(P_m, w^m)$ and our objective is to show that there exists $u' \in \mathcal{U}$ such that any complete assignment u extending u' satisfies these constraints.

Fix a partial assignment $u' \in \mathcal{U}$. For $i \in [t]$, $Q \in \mathcal{L}_2$ and $\alpha\colon [k(Q)] \rightarrowtail [n_i]$, let $E^\varepsilon_{u'}(i, Q, \alpha)$ be the event $\alpha^*(\boldsymbol{z^{i,\varepsilon}}) \in \mathcal{H}_Q \leftrightarrow \alpha^*(z^i) \in \mathcal{H}^u_Q$. By an argument analogous to that of

Theorem 3.6.6, we have

$$\limsup_{\varepsilon\to 0} \mathbb{P}\left[\bigwedge\{E^{\varepsilon}_{u'}(i,Q,\alpha) \mid i\in[t], Q\in\mathcal{L}_2, \alpha\colon [k(P)]\rightarrowtail[n_i]\}\right] > 0. \tag{3.30}$$

Since the event $\forall i \in [t], \boldsymbol{z^{i,\varepsilon}} \in T(F_i,\mathcal{H}) \leftrightarrow z^i \in T(F_i,\mathcal{H}^u)$ from (3.28) is implied by the conjunction in (3.30), it follows that for every $u' \in \mathcal{U}$, every complete assignment u extending u' satisfies the constraint $\mathcal{R}(\mathcal{F})$.

To satisfy the constraints $\mathcal{I}(P_1,w^1),\ldots,\mathcal{I}(P_m,w^m)$, we will have to choose $u' \in \mathcal{U}$ more carefully. For $j \in [m]$, let S^{ε}_j be the event $\boldsymbol{w^{j,\varepsilon}} \in \mathcal{N}_{P_j} \times \mathcal{E}_{k(P_j)} \leftrightarrow w^j \in \mathcal{N}_{P_j} \times \mathcal{E}_{k(P_j)}$. Recalling that $w^j \in (\mathcal{E}_{k(P_j)} \setminus \mathcal{D}_{k(P_j)}) \times \mathcal{E}_{k(P_j)}$, note that robustness of $\mathcal{N}$ implies that

$$\limsup_{\varepsilon\to 0} \mathbb{P}\left[\bigwedge_{j\in[m]} S^{\varepsilon}_j\right] > 0.$$

Using again the fact that there are only finitely many assignments, it follows that there exists $u' \in \mathcal{U}$ such that

$$\limsup_{\varepsilon\to 0} \mathbb{P}\left[\boldsymbol{u^{\varepsilon}} = u' \wedge \bigwedge_{j\in[m]} S^{\varepsilon}_j\right] > 0.$$

But note that within the event in the above, we can deduce the following chain of a.e. equivalences

$$\begin{aligned} w^j \in \mathcal{N}_{P_j} \times \mathcal{E}_{k(P_j)} &\iff \boldsymbol{w^{j,\varepsilon}} \in \mathcal{N}_{P_j} \times \mathcal{E}_{k(P_j)} \iff \boldsymbol{w^{j,\varepsilon}} \in I(\mathcal{H})_{P_j} \\ &\iff \boldsymbol{w^{j,\varepsilon}} \in T(I(P_j),\mathcal{H}) \iff \boldsymbol{z^{i_j,\varepsilon}} \in T(F_{i_j},\mathcal{H}) \end{aligned}$$

where the second equivalence holds a.e. within the event since $\mathcal{N}_{P_j} \times \mathcal{E}_{k(P_j)} = I(\mathcal{H})_{P_j}$ a.e.

On the other hand, we also have

$$z^{i_j} \in T(F_{i_j}, \mathcal{H}^u) \iff w^j \in T(I(P_j), \mathcal{H}^u) \iff w^j \in I(\mathcal{H}^u)_{P_j},$$

so the fact that any complete assignment u extending u' satisfies $\mathcal{R}(\mathcal{F})$ implies

$$w^j \in \mathcal{N}_{P_j} \times \mathcal{E}_{k(P_j)} \iff w^j \in I(\mathcal{H}^u)_{P_j},$$

that is, u also satisfies $\mathcal{I}(P_j, w^j)$. ■

3.7 Amalgamations over more general diagrams

In this section we study a generalization of Theorem 3.2.5 and Proposition 3.2.9 to diagrams with more complicated shapes. Our first order of business is to show that **Int** is closed under finite colimits.

Proposition 3.7.1. *Let $D\colon \mathbf{S} \to \textbf{Int}$ be a finite diagram and let T be the theory obtained from $\bigcup_{A\in\mathrm{Obj}(J)} D(A)$ by adding the axioms*

$$P(x_1, \ldots, x_{k(P)}) \leftrightarrow D(f)(P)(x_1, \ldots, x_{k(P)}) \tag{3.31}$$

for every $\mathbf{S}$-morphism $f\colon A_1 \to A_2$ and every predicate symbol P in the language of $D(A_1)$. For each $A \in \mathrm{Obj}(\mathbf{S})$, let also $I_A\colon D(A) \rightsquigarrow T$ be the interpretation that acts identically on the language of $D(A)$.

Then $(T, (I_A)_{A\in\mathrm{Obj}(\mathbf{S})})$ is a colimit of D.

Proof. The fact that $(T, (I_A)_{A\in\mathrm{Obj}(\mathbf{S})})$ is a cone from D follows directly from the axioms (3.31).

For every $A \in \mathrm{Obj}(\mathbf{S})$, let $T_A \stackrel{\text{def}}{=} D(A)$ and let $\mathcal{L}_A$ be the language of T_A so that the language of T is $\mathcal{L} \stackrel{\text{def}}{=} \dot{\bigcup}_{A\in\mathrm{Obj}(\mathbf{S})} \mathcal{L}_A$. To show universality, for a cone $(U, (J_A)_{A\in\mathrm{Obj}(\mathbf{S})})$

from D, define a translation I from the language of T to the language of U by

$$I(P)(x_1, \ldots, x_{k(P)}) \stackrel{\text{def}}{=} J_A(P)(x_1, \ldots, x_{k(P)}),$$

for every $A \in \mathrm{Obj}(\mathbf{S})$ and every $P \in \mathcal{L}_A$. From this definition, it is trivial that for every $A \in \mathrm{Obj}(\mathbf{S})$ we have $I \circ I_A = J_A$.

It remains to show that $I\colon T \rightsquigarrow U$ is the unique element of $\mathrm{Hom}_{\mathbf{INT}}(T, U)$ with this property. Suppose $I'\colon T \rightsquigarrow U$ also satisfies $I' \circ I_A = J_A$ for every $A \in \mathrm{Obj}(\mathbf{S})$ and let us show that $I = I'$ in **Int** (recall that in this category we factor by interpretation equivalence). We need to show that for every $P \in \mathcal{L}$, we have

$$U \vdash \forall \vec{x}, (I(P)(\vec{x}) \leftrightarrow I'(P)(\vec{x})). \tag{3.32}$$

But indeed, if $P \in \mathcal{L}_A$ for some $A \in \mathrm{Obj}(\mathbf{S})$, then the above follows since $I \circ I_A = J_A = I' \circ I_A$ in **Int** and I_A acts identically on P. ■

Just as couplings, amalgamations over pushouts and unique coupleability, we can define the analogous notions over a general finite diagram.

Definition 3.7.2. Let $(T, (I_A)_{A \in \mathrm{Obj}(\mathbf{S})})$ be a colimit of a finite diagram $D\colon \mathbf{S} \to \mathbf{INT}$ and for each $A \in \mathrm{Obj}(\mathbf{S})$ let $\phi_A \in \mathrm{Hom}^+(\mathcal{A}[D(A)], \mathbb{R})$.

We say that $(\phi_A)_{A \in \mathrm{Obj}(\mathbf{S})}$ *respects* D if for every $\mathbf{S}$-morphism $f\colon A_1 \to A_2$, we have $\phi_{A_2}^{D(f)} = \phi_{A_1}$.

An *amalgamation* of $(\phi_A)_{A \in \mathrm{Obj}(\mathbf{S})}$ over D (with respect to $(T, (I_A)_{A \in \mathrm{Obj}(\mathbf{S})})$) is an element $\xi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ such that $\xi^{I_A} = \phi_A$ for every $A \in \mathrm{Obj}(\mathbf{S})$.

We say that the family $(\phi_A)_{A \in \mathrm{Obj}(\mathbf{S})}$ is *uniquely amalgamatable* over D if for each colimit $(T, (I_A)_{A \in \mathrm{Obj}(\mathbf{S})})$ of D, there exists a unique amalgamation of $(\phi_A)_{A \in \mathrm{Obj}(\mathbf{S})}$ over D with respect to $(T, (I_A)_{A \in \mathrm{Obj}(\mathbf{S})})$.

Remark 7. Note that the universal isomorphisms allows us to translate between amalgamations over D with respect to different colimits, thus existence (resp., uniqueness) of amalgamations over some colimit is equivalent to existence (resp., uniqueness) of amalgamations over every fixed colimit. Thus, we will typically omit the colimit when we talk about existence/uniqueness of amalgamations.

It is obvious that for an amalgamation of $(\phi_A)_{A\in\mathrm{Obj}(\mathbf{S})}$ over D to exist, $(\phi_A)_{A\in\mathrm{Obj}(\mathbf{S})}$ must respect D, but the following examples show that this condition is not sufficient.

Example 3. Consider the finite shape $\mathbf{S} \stackrel{\text{def}}{=} 0 \rightrightarrows 1$ consisting of two parallel arrows $f, g\colon 0 \to 1$ (plus identity morphisms $\mathrm{id}_0, \mathrm{id}_1$) and consider the diagram $D\colon \mathbf{S} \to \textbf{I}\textsc{nt}$ given by letting $D(0) \stackrel{\text{def}}{=} D(1) \stackrel{\text{def}}{=} T_{2\text{-Coloring}}$, letting $D(f) \stackrel{\text{def}}{=} \mathrm{id}_{T_{2\text{-Coloring}}}$ and letting $D(g)$ be the open interpretation that swaps the colors, that is, it is given by

$$D(g)(\chi_i)(x) \stackrel{\text{def}}{=} \chi_{3-i}(x) \qquad (i \in [2]). \tag{3.33}$$

Consider the limit $\psi_{1/2,1/2} \in \mathrm{Hom}^+(\mathcal{A}[T_{2\text{-Coloring}}], \mathbb{R})$ that assigns density 1/2 to each color, that is, it is represented by the $T_{2\text{-Coloring}}$-on

$$\mathcal{N}_{\chi_1} \stackrel{\text{def}}{=} \left\{x \in \mathcal{E}_1 \;\middle|\; x_{\{1\}} < \frac{1}{2}\right\}; \qquad \mathcal{N}_{\chi_2} \stackrel{\text{def}}{=} \left\{x \in \mathcal{E}_1 \;\middle|\; x_{\{1\}} \geq \frac{1}{2}\right\}. \tag{3.34}$$

Then letting $\phi_0 \stackrel{\text{def}}{=} \phi_1 \stackrel{\text{def}}{=} \psi_{1/2,1/2}$, it follows that (ϕ_0, ϕ_1) respects D.

However, letting (T, I_0, I_1) be as in Proposition 3.7.1 and denoting by χ_1, χ_2 the predicate symbols of T corresponding to $D(0)$ and denoting by χ'_1, χ'_2 the ones corresponding to $D(1)$, note that (3.31) applied to f and g respectively imply

$$\chi_i(x) \iff \chi'_i(x) \qquad (i \in [2]);$$
$$\chi_i(x) \iff \chi'_{3-i}(x) \qquad (i \in [2]);$$

which contradicts the axioms of T inherited from $T_{2\text{-Coloring}} \cup T_{2\text{-Coloring}}$, so T is degenerate and thus no amalgamation ξ of (ϕ_0, ϕ_1) exists.

To avoid the problem illustrated by the example above, one could ask for the diagram to be commutative, that is, ask for $\mathbf{S}$ to be a finite poset category (i.e., a category $\mathbf{S}$ in which for all $A_1, A_2 \in \mathrm{Obj}(\mathbf{S})$, there is at most one $\mathbf{S}$-morphism of the form $A_1 \to A_2$ and whose isomorphisms are all identities), but the example below shows that this extra hypothesis is still not sufficient to ensure existence of an amalgamation.

Example 4. Consider the poset category $\mathbf{S}$ with shape

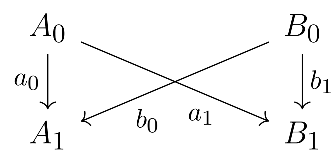

and let $D\colon \mathbf{S} \to \textsc{Int}$ be the diagram given by letting $D(A_i) \stackrel{\text{def}}{=} D(B_i) \stackrel{\text{def}}{=} T_{2\text{-Coloring}}$ ($i \in [2]$), letting $a_0 \stackrel{\text{def}}{=} a_1 \stackrel{\text{def}}{=} b_0 \stackrel{\text{def}}{=} \mathrm{id}_{T_{2\text{-Coloring}}}$ and letting b_1 be the open interpretation that swaps that colors (see (3.33)).

Again, for the coloring $\psi_{1/2,1/2} \in \mathrm{Hom}^+(\mathcal{A}[T_{2\text{-Coloring}}], \mathbb{R})$ that assigns density 1/2 to each color (see (3.34)), $(\psi_{1/2,1/2}, \psi_{1/2,1/2}, \psi_{1/2,1/2}, \psi_{1/2,1/2})$ respects D but it has no amalgamation over D since, just as in Example 3, the colimit of D is a degenerate theory.

Let us note that the fact that the colimit theory is degenerate is only for proof convenience: one can take any of Examples 3 and 4 and form a diagram D' by replacing $T_{2\text{-Coloring}}$ with the pure canonical theory $T_{\{\chi_1,\chi_2\}}$ and the colimit T' of D' will have several unary predicate symbols that must always agree, hence T' is isomorphic to $T_{2\text{-Coloring}}$, which is not degenerate. However, the families of limits still cannot be amalgamated over D' as any such amalgamation will necessarily also be an amalgamation over D. We will return to this in Section 3.8.

It is natural to ask then what shapes $\mathbf{S}$ of the diagram D ensure that any family of limits respecting D can be amalgamated.

Definition 3.7.3. Let $\mathbf{S}$ be a finite category. We say that $\mathbf{S}$ *amalgamates theons* if for every diagram $D\colon \mathbf{S} \to \textsc{Int}$ and every $(\phi_A)_{A\in\mathrm{Obj}(\mathbf{S})}$ that respects D, there exists an amalgamation of $(\phi_A)_{A\in\mathrm{Obj}(\mathbf{S})}$ over D.

The next theorem says that finite forest-like categories (defined below) amalgamate theons.

Definition 3.7.4. Let F be a forest (i.e., an acyclic graph) and let $\vec{F}$ be an orientation of the edges of F. The category $C(\vec{F})$ is the small poset category whose objects are $V(\vec{F})$, whose morphisms are directed paths of $\vec{F}$ with identity morphisms given by length 0 paths and with composition given by path concatenation.

A category is *forest-like* if it is of the form $C(\vec{F})$ for some orientation $\vec{F}$ of some forest F. Equivalently, a small category is forest-like if it is a poset category such that the corresponding Hasse diagram does not have any (undirected) cycles.

Theorem 3.7.5. *If* $\mathbf{S}$ *is a finite forest-like category, then* $\mathbf{S}$ *amalgamates theons.*

Proof. Let $(T, (I_A)_{A\in\mathrm{Obj}(\mathbf{S})})$ be a colimit of a finite diagram $D\colon \mathbf{S} \to \mathbf{Int}$, let $(\phi_A)_{A\in\mathrm{Obj}(\mathbf{S})}$ respect D and let us show that there exists an amalgamation of $(\phi_A)_{A\in\mathrm{Obj}(\mathbf{S})}$ over D with respect to $(T, (I_A)_{A\in\mathrm{Obj}(\mathbf{S})})$.

By Remark 7, it is enough to show the case when $(T, (I_A)_{A\in\mathrm{Obj}(\mathbf{S})})$ is as in Proposition 3.7.1.

Let $\vec{F}$ be an orientation of a forest F such that $\mathbf{S} = C(\vec{F})$ and let us show the result by induction in $|V(F)|$. If $V(F) = \varnothing$, then T is the trivial theory and its unique limit is an amalgamation of the empty family over D, so suppose $V(F)$ is non-empty.

Since F is a forest, there exists a vertex A_0 of degree at most 1. Let $\vec{F}' \stackrel{\text{def}}{=} \vec{F} - A_0$ and let $D'\colon \mathbf{S}' \to \mathbf{Int}$ be the restriction of D to $\mathbf{S}' \stackrel{\text{def}}{=} C(\vec{F}')$. Let also $(T', (I'_A)_{A\in\mathrm{Obj}(\mathbf{S}')})$ be the colimit of D' given by Proposition 3.7.1 and let $I\colon T' \rightsquigarrow T$ be the universal $\mathbf{Int}$-morphism from T' (it acts identically on the language of T') so that $I_A = I \circ I'_A$. By inductive hypothesis, let ψ be an amalgamation of $(\phi_A)_{A\in\mathrm{Obj}(\mathbf{S}')}$ over D' with respect to $(T', (I'_A)_{A\in\mathrm{Obj}(\mathbf{S}')})$.

Suppose first that A_0 is an isolated vertex of F. Since the only $\mathbf{S}$-morphism that is not a $\mathbf{S}'$-morphism is id_{A_0}, it follows that $T = T' \cup T_{A_0}$. This means that if ξ is a coupling of ϕ_{A_0} and ψ, then ξ is an amalgamation of $(\phi_A)_{A\in\mathrm{Obj}(\mathbf{S}')}$ over D with respect to $(T, (I_A)_{A\in\mathrm{Obj}(\mathbf{S})})$.

Consider now the case when there is an oriented edge f in $\vec{F}$ from A_0 to some $A_1 \in \mathrm{Obj}(\mathbf{S}')$.

In this case, since all paths from A_0 to some vertex of $\vec{F}'$ must go through f, it follows that T is the theory obtained from $T' \cup T_{A_0}$ by adding the axioms

$$P(x_1, \ldots, x_{k(P)}) \leftrightarrow D(f)(P)(x_1, \ldots, x_{k(P)})$$

for every P in the language of $D(A_0)$ (the axioms corresponding to longer paths can all be entailed from these and the axioms of T'). This implies that I is an isomorphism (I^{-1} acts identically on the language of T' and acts as $D(f)$ on the language of T_{A_0}). We claim that $\xi \stackrel{\text{def}}{=} \psi^{I^{-1}}$ is an amalgamation of $(\phi_A)_{A \in \mathrm{Obj}(\mathbf{S})}$ over D with respect to $(T, (I_A)_{A \in \mathrm{Obj}(\mathbf{S})})$. Note that for every $A \in \mathrm{Obj}(\mathbf{S}')$, we have $\xi^{I_A} = \phi_A$, since $I_A = I \circ I'_A$. On the other hand, we also have $\xi^{I_{A_0}} = \xi^{I_{A_1} \circ D(f)} = \phi_{A_1}^{D(f)} = \phi_{A_0}$.

Finally, for the case when there is an oriented edge f in $\vec{F}$ from some $A_1 \in \mathrm{Obj}(\mathbf{S}')$ to A_0, note that the diagram

$$\begin{array}{ccc} D(A_1) & \xrightarrow{D(f)} & D(A_0) \\ {\scriptstyle I'_{A_1}}\downarrow & & \downarrow{\scriptstyle I_{A_0}} \\ T' & \xrightarrow{\quad I \quad} & T \end{array}$$

is a pushout in $\textbf{I}\textsc{nt}$, so by Theorem 3.2.5 there exists $\xi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ such that $\xi^{I_{A_0}} = \phi_{A_0}$ and $\xi^I = \psi$. Since the latter implies $\xi^{I_A} = \psi^{I'_A} = \phi_A$ for every $A \in \mathrm{Obj}(\mathbf{S}')$, it follows that ξ is an amalgamation of $(\phi_A)_{A \in \mathrm{Obj}(\mathbf{S})}$ over D with respect to $(T, (I_A)_{A \in \mathrm{Obj}(\mathbf{S})})$. ■

The remainder of this section is devoted to showing the following generalization of Proposition 3.2.9 that says that unique amalgamation is preserved under natural transformations as long as the shape is forest-like.

Theorem 3.7.6. *Let $\tau\colon D_1 \to D_2$ be a natural transformation between finite diagrams $D_1, D_2\colon \mathbf{S} \to \textbf{I}\textsc{nt}$ such that $\mathbf{S}$ is forest-like and let $(\phi_A)_{A \in \mathrm{Obj}(\mathbf{S})}$ be uniquely amalgamatable over D_2. Then $(\phi_A^{\tau_A})_{A \in \mathrm{Obj}(\mathbf{S})}$ is uniquely amalgamatable over D_1.*

Just as in Proposition 3.2.9: the theorem above will follow from the fact that amalgamations lift through natural transformations. This involves considering amalgamations over a diagram that encodes both diagrams D_1, D_2, the colimit of D_1 and the natural transformation $\tau\colon D_1 \to D_2$.

Definition 3.7.7. Let $\mathbf{S}$ be a poset category and let $\preceq_{\mathbf{S}}$ be the underlying partial order on $\mathrm{Obj}(\mathbf{S})$ given by

$$A_1 \preceq_{\mathbf{S}} A_2 \iff \mathrm{Hom}_{\mathbf{S}}(A_0, A_1) \neq \varnothing.$$

The category $\widehat{\mathbf{S}}$ (see Figures 3.1a and 3.2) is the poset category obtained from $\mathbf{S}$ by forming the product poset with $([2], \leq)$ and adding an element $*$ that is greater than every element in the first copy of $\mathbf{S}$; more formally, we let $\mathrm{Obj}(\mathbf{S}) \stackrel{\text{def}}{=} (\mathrm{Obj}(\mathbf{S}) \times [2]) \mathbin{\dot\cup} \{*\}$ and let the underlying partial order be given by

$$(A_1, i_1) \preceq_{\widehat{\mathbf{S}}} (A_2, i_2) \iff A_1 \preceq_{\mathbf{S}} A_2 \wedge i_1 \leq i_2 \qquad (A_1, A_2 \in \mathrm{Obj}(\mathbf{S}), i_1, i_2 \in [2]);$$
$$(A, 1) \preceq_{\widehat{\mathbf{S}}} * \qquad (A \in \mathrm{Obj}(\mathbf{S}));$$
$$* \preceq_{\widehat{\mathbf{S}}} *.$$

The $\widehat{\mathbf{S}}$-morphisms can be subclassified according to their relation to $*$ into the following three types.

i. Each $\mathbf{S}$-morphism $f\colon A_1 \to A_2$ gives rise to three $\widehat{\mathbf{S}}$-morphisms $f_1\colon (A_1, 1) \to (A_2, 1)$, $f_2\colon (A_1, 2) \to (A_2, 2)$ and $f_{12}\colon (A_1, 1) \to (A_2, 2)$ (note that $\mathrm{id}_{(A,i)} = (\mathrm{id}_A)_i$).

ii. For each $A \in \mathrm{Obj}(\mathbf{S})$, we have the $\widehat{\mathbf{S}}$-morphisms $g_A\colon (A, 1) \to *$.

iii. The final $\widehat{\mathbf{S}}$-morphism is the identity id_*.

Given a natural transformation $\tau\colon D_1 \to D_2$ between finite diagrams $D_1, D_2\colon \mathbf{S} \to \textbf{Int}$ and a colimit $C = (T, (I_A)_{A \in \mathrm{Obj}(\mathbf{S})})$ of D_1, we let $D_{\tau,C}\colon \widehat{\mathbf{S}} \to \textbf{Int}$ (see Figure 3.13.1b) be the natural commutative diagram that contains all morphisms in D_1, D_2 and C; more formally, it is given by

$$
\begin{aligned}
D_{\tau,C}((A,i)) &\stackrel{\text{def}}{=} D_i(A) \qquad (A \in \operatorname{Obj}(\mathbf{S}), i \in [2]);\\
D_{\tau,C}(*) &\stackrel{\text{def}}{=} T;\\
D_{\tau,C}(f_i) &\stackrel{\text{def}}{=} D_i(f) \qquad (f \in \operatorname{Hom}(\mathbf{S}), i \in [2]);\\
D_{\tau,C}(f_{12}) &\stackrel{\text{def}}{=} \tau_{A_2} \circ D_1(f) \qquad (f \in \operatorname{Hom}_{\mathbf{S}}(A_1, A_2), A_1, A_2 \in \operatorname{Obj}(\mathbf{S}));\\
D_{\tau,C}(g_A) &\stackrel{\text{def}}{=} I_A \qquad (A \in \operatorname{Obj}(\mathbf{S}));\\
D_{\tau,C}(\mathrm{id}_*) &\stackrel{\text{def}}{=} \mathrm{id}_T\,.
\end{aligned}
$$

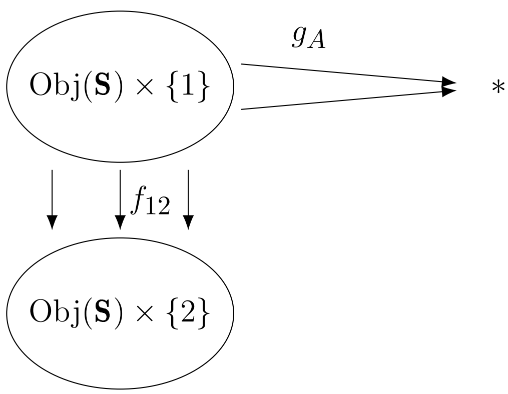

(a) Shape $\widehat{S}$.

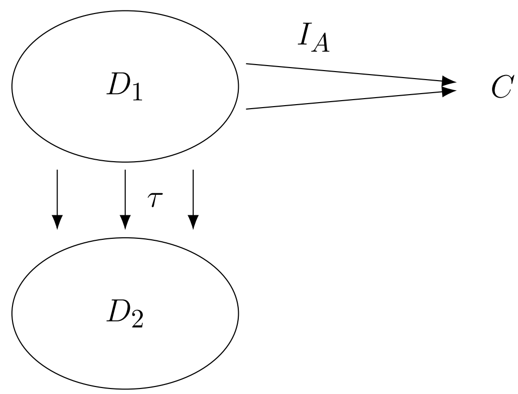

(b) Image of diagram $D_{\tau,C}$.

Figure 3.1: Pictorial view of constructions of Definition 3.7.7.

Note that if $\mathbf{S}$ has at least one non-identity morphism, then $\widehat{\mathbf{S}}$ is not forest-like (even if $\mathbf{S}$ is forest-like), so we cannot use Theorem 3.7.5 to ensure that $\widehat{\mathbf{S}}$ amalgamates theons and must instead prove this ad hoc.

Theorem 3.7.8. *If $\mathbf{S}$ is a finite forest-like category, then $\widehat{\mathbf{S}}$ amalgamates theons.*

Proof. Throughout this proof, let us assume that all colimits are as in Proposition 3.7.1 (see Remark 7), we will also use the notation of Definition 3.7.7 for the objects and morphisms of $\widehat{\mathbf{S}}$.

Let $\vec{F}$ be an orientation of a forest F such that $\mathbf{S} = C(\vec{F})$ and let us show the result by

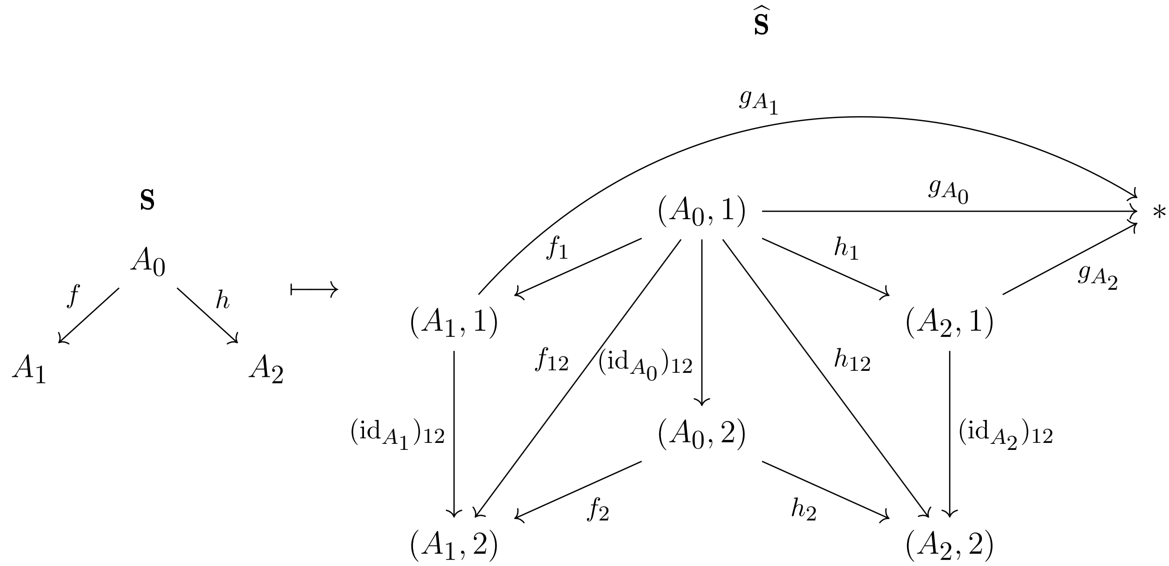

Figure 3.2: Example of construction of shape $\widehat{\mathbf{S}}$ from $\mathbf{S}$, identity morphisms are omitted.

induction in $|V(F)|$. If $V(F) = \varnothing$, then $\widehat{S}$ has a single object, namely $*$, thus it amalgamates theons (e.g., by Theorem 3.7.5), so suppose $V(F)$ is non-empty.

Let $(T, (I_A)_{A \in \mathrm{Obj}(\mathbf{S})})$ be the colimit of some diagram $D\colon \widehat{\mathbf{S}} \to \textbf{Int}$ of shape $\widehat{\mathbf{S}}$, let $(\phi_A)_{A \in \mathrm{Obj}(\widehat{\mathbf{S}})}$ respect D and let us show that there exists an amalgamation of $(\phi_A)_{A \in \mathrm{Obj}(\widehat{\mathbf{S}})}$ over D.

First, we claim that it is enough to show the case when $(D(*), (D(g_A))_{A \in \mathrm{Obj}(\mathbf{S})})$ is the colimit of the restriction $D_1 \stackrel{\text{def}}{=} D|_{\mathrm{Obj}(\mathbf{S}) \times \{1\}}$ of D to the first copy $\mathrm{Obj}(\mathbf{S}) \times \{1\}$ of $\mathbf{S}$ in $\widehat{\mathbf{S}}$.

To show the claim, we will construct several **Int**-morphisms and amalgamations (which are pictorially represented in Figure 3.3) by diagram chasing using Theorem 3.2.5.

Let $I\colon C \to D(*)$ be the universal **Int**-morphism from the colimit $(C, (J_A)_{A \in \mathrm{Obj}(\mathbf{S}) \times \{1\}})$ of D_1 such that $D(g_A) = I \circ J_A$ for every $A \in \mathrm{Obj}(\mathbf{S}) \times \{1\}$, and let $D'\colon \widehat{\mathbf{S}} \to \textbf{Int}$ coincide

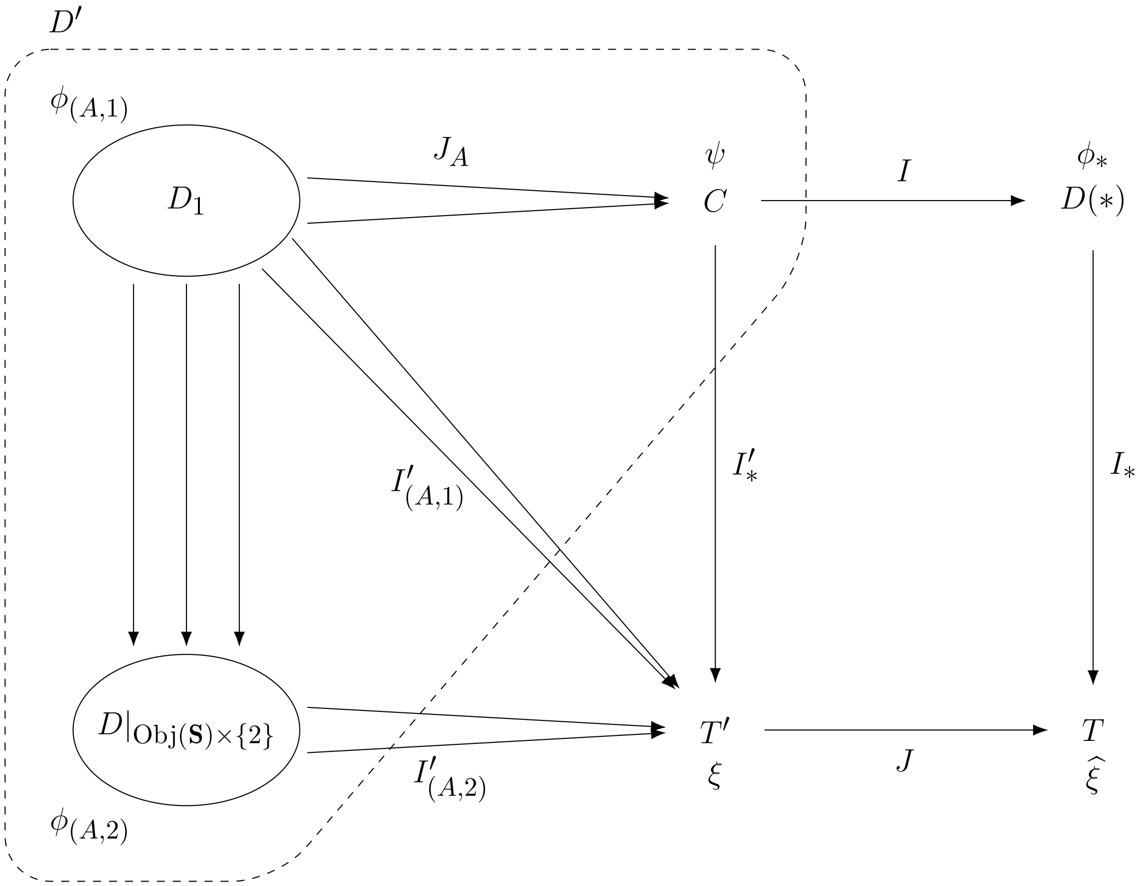

Figure 3.3: Commutative diagram of morphisms and amalgamations constructed in reduction to the case when $(D(*),(D(g_A))_{A\in\mathrm{Obj}(\mathbf{S})})$ is the colimit of D_1; some compositions are omitted. C is the colimit of D_1, T' is the colimit of D' and T is the colimit of D. The square on the right is a pushout.

with D in $\widehat{\mathbf{S}}-*$ and

$$D'(*)\stackrel{\text{def}}{=}C;$$
$$D'(g_A)\stackrel{\text{def}}{=}J_A \qquad (A\in\mathrm{Obj}(\mathbf{S}));$$
$$D'(\mathrm{id}_*)\stackrel{\text{def}}{=}\mathrm{id}_C\,.$$

Let now $\psi\stackrel{\text{def}}{=}\phi_*^I\in\mathrm{Hom}^+(\mathcal{A}[C],\mathbb{R})$ and note that ψ is an amalgamation of the family $(\phi_A)_{A\in\mathrm{Obj}(\mathbf{S})\times\{1\}}$ over D_1. This in particular implies that $((\phi_A)_{A\in\mathrm{Obj}(\widehat{\mathbf{S}})\setminus\{*\}},\psi)$ respects D', so our hypothesis gives an amalgamation $\xi\in\mathrm{Hom}^+(\mathcal{A}[T'],\mathbb{R})$ of this family over D' with respect to its colimit $(T', (I'_A)_{A\in\mathrm{Obj}(\widehat{\mathbf{S}})})$. Note now that the diagram

$$\begin{array}{ccc} C & \xrightarrow{I} & D(*) \\ {\scriptstyle I'_*}\downarrow & & \downarrow{\scriptstyle I_*} \\ T' & \xrightarrow[J]{} & T \end{array}$$

is commutative, where J is the universal **Int**-morphism from the colimit T'.

Since $\phi_*^I = \psi = \xi^{I'_*}$, Theorem 3.2.5 gives us $\widehat{\xi} \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ such that $\widehat{\xi}^J = \xi$ and $\widehat{\xi}^{I_*} = \phi_*$. Note that for $A \in \mathrm{Obj}(\mathbf{S})$, we have

$$\begin{aligned} \widehat{\xi}^{I_{(A,1)}} &= \widehat{\xi}^{I_* \circ D(g_A)} = \phi_*^{D(g_A)} = \phi_{(A,1)}; \\ \widehat{\xi}^{I_{(A,2)}} &= \widehat{\xi}^{J \circ I'_{(A,2)}} = \xi^{I'_{(A,2)}} = \phi_{(A,2)}. \end{aligned}$$

Thus $\widehat{\xi}$ is an amalgamation of $(\phi_A)_{A\in\mathrm{Obj}(\widehat{\mathbf{S}})}$ over D and the claim is proved.

Let us now show the case when $(D(*), (D(g_A))_{A\in\mathrm{Obj}(\mathbf{S})})$ is the colimit of D_1. Again, this is shown by diagram chasing (see Figures 3.4 and 3.5).

Since F is a forest, there exists a vertex A_0 of degree at most 1. Let $\vec{F}' \stackrel{\text{def}}{=} \vec{F} - *$, let $\mathbf{S}' \stackrel{\text{def}}{=} C(\vec{F}')$ and let $D' \colon \widehat{\mathbf{S}}' \to \mathbf{Int}$ be the diagram that coincides with D in $\mathrm{Obj}(\mathbf{S}') \times [2]$ and maps $(*, (g_A)_{A\in\mathrm{Obj}(\mathbf{S})})$ to the colimit $(C, (J_A)_{A\in\mathrm{Obj}(\mathbf{S}')})$ of the restriction of $D'_1 \stackrel{\text{def}}{=} D|_{\mathrm{Obj}(\mathbf{S}')\times\{1\}}$ of D to $\mathrm{Obj}(\mathbf{S}') \times \{1\}$. Let also $J \colon C \rightsquigarrow D(*)$ be the universal **Int**-morphism from C so that $D(g_A) = J \circ J_A$ for every $A \in \mathrm{Obj}(\mathbf{S}')$.

Let further $(T', (I'_A)_{A\in\mathrm{Obj}(\mathbf{S}')})$ be the colimit of D' and let $I \colon T' \rightsquigarrow T$ be the universal **Int**-morphism from T' so that $I_A = I \circ I'_A$. Let $\psi \stackrel{\text{def}}{=} \phi_*^J \in \mathrm{Hom}^+(\mathcal{A}[C], \mathbb{R})$ and note that ψ is an amalgamation of $(\phi_A)_{A\in\mathrm{Obj}(\mathbf{S})\times\{1\}}$ over D'_1, which in particular implies that $((\phi_A)_{A\in\mathrm{Obj}(\widehat{\mathbf{S}}')\setminus\{*\}}, \psi)$ respects D', so by inductive hypothesis, there exists an amalgamation $\xi \in \mathrm{Hom}^+(\mathcal{A}[T'], \mathbb{R})$ of this family over D'.

Consider the case when A_0 is an isolated vertex of F (see Figure 3.4). We claim that in this case T is isomorphic to $T' \cup D((A_0, 2))$. Indeed, the theory T is obtained from

$T' \cup D((A_0, 1)) \cup D((A_0, 2))$ by adding the axioms

$$P(x_1, \ldots, x_{k(P)}) \leftrightarrow D((\mathrm{id}_{A_0})_{12})(P)(x_1, \ldots, x_{k(P)})$$

for every P in the language of $D((A_0, 1))$. This means that the open interpretation $I' \colon T' \cup D((A_0, 2)) \rightsquigarrow T$ that acts identically is an isomorphism (its inverse $(I')^{-1}$ acts identically on the language of $T' \cup D((A_0, 2))$ and acts as $D((\mathrm{id}_{A_0})_{12})$ on the language of $D((A_0, 1))$).

Let then $\xi_2 \in \mathrm{Hom}^+(\mathcal{A}[T' \cup D((A_0, 2))], \mathbb{R})$ be any coupling of ξ and $\phi_{(A_0,2)}$ and note that $\widehat{\xi} \stackrel{\text{def}}{=} \xi_2^{(I')^{-1}} \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is an amalgamation of $(\phi_A)_{A \in \mathrm{Obj}(\widehat{\mathbf{S}})}$ over D.

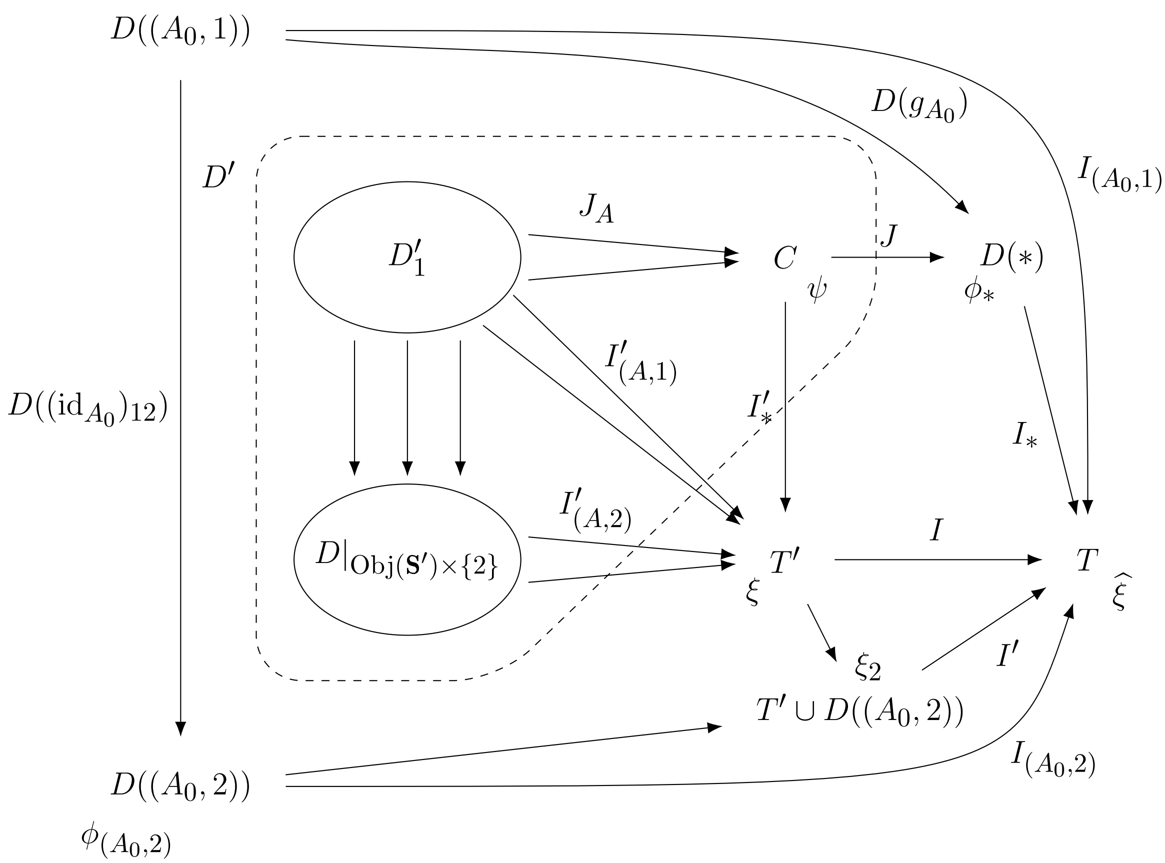

Figure 3.4: Commutative diagram of morphisms and amalgamations constructed in the case when A_0 is an isolated vertex of F; some compositions are omitted. C is the colimit of D'_1, T' is the colimit of D' and T is the colimit of D.

Consider now the case when there is an oriented edge f in $\vec{F}$ from A_0 to some $A_1 \in \mathrm{Obj}(\mathbf{S}')$.

In this case, since all paths from A_0 to some vertex of $\vec{F}'$ must go through f, it follows that T is the theory obtained from $T' \cup D((A_0,1)) \cup D((A_0,2))$ by adding the axioms

$$P(x_1,\dots,x_{k(P)}) \leftrightarrow D(f_2)(P)(x_1,\dots,x_{k(P)});$$
$$Q(x_1,\dots,x_{k(Q)}) \leftrightarrow D((\mathrm{id}_{A_0})_{12})(Q)(x_1,\dots,x_{k(Q)})$$

for every P in the language of $D((A_0,2))$ and every Q in the language of $D((A_0,1))$ (the axioms corresponding to other paths are entailed from these and the axioms of T'). This implies that I is an isomorphism (I^{-1} acts identically on the language of T', acts as $D(f_2)$ on the language of $D((A_0,2))$ and acts as $D(f_2 \circ (\mathrm{id}_{A_0})_{12})$ on the language of $D((A_0,1))$). Then $\widehat{\xi} \stackrel{\text{def}}{=} \xi^{I^{-1}}$ is clearly an amalgamation of $(\phi_A)_{A\in\mathrm{Obj}(\widehat{\mathbf{S}})}$ over D.

Finally, for the case when there is an oriented edge f in $\vec{F}$ from some $A_1 \in \mathrm{Obj}(\mathbf{S}')$ to A_0 (see Figure 3.5), we form the pushout

$$\begin{array}{ccc} D((A_1,2)) & \xrightarrow{I'_{(A_1,2)}} & T' \\ {\scriptstyle D(f_2)}\big\downarrow & & \big\downarrow{\scriptstyle I'_2} \\ D((A_0,2)) & \xrightarrow[I_2]{} & T_2 \end{array} \tag{3.35}$$

Since $\phi_{(A_0,2)}^{D(f_2)} = \phi_{(A_1,2)} = \xi^{I'_{(A_1,2)}}$, by Theorem 3.2.5, there exists $\xi_2 \in \mathrm{Hom}^+(\mathcal{A}[T_2],\mathbb{R})$ such that $\xi_2^{I_2} = \phi_{(A_0,2)}$ and $\xi_2^{I'_2} = \xi$.

We claim that T_2 is isomorphic to T. Indeed, the open interpretation $J_2\colon T_2 \rightsquigarrow T$ that acts identically has as inverse the interpretation $J_2^{-1}\colon T \rightsquigarrow T_2$ that acts identically on the language of $T' \cup D((A_0,2))$ and acts as $D((\mathrm{id}_{A_0})_{12})$ on the language of $D((A_0,1))$. Thus $\widehat{\xi} = \xi_2^{J_2^{-1}}$ is an amalgamation of $(\phi_A)_{A\in\mathrm{Obj}(\widehat{\mathbf{S}})}$ over D.

■

We can now show that amalgamations over forest-like diagrams lift through natural transformations.

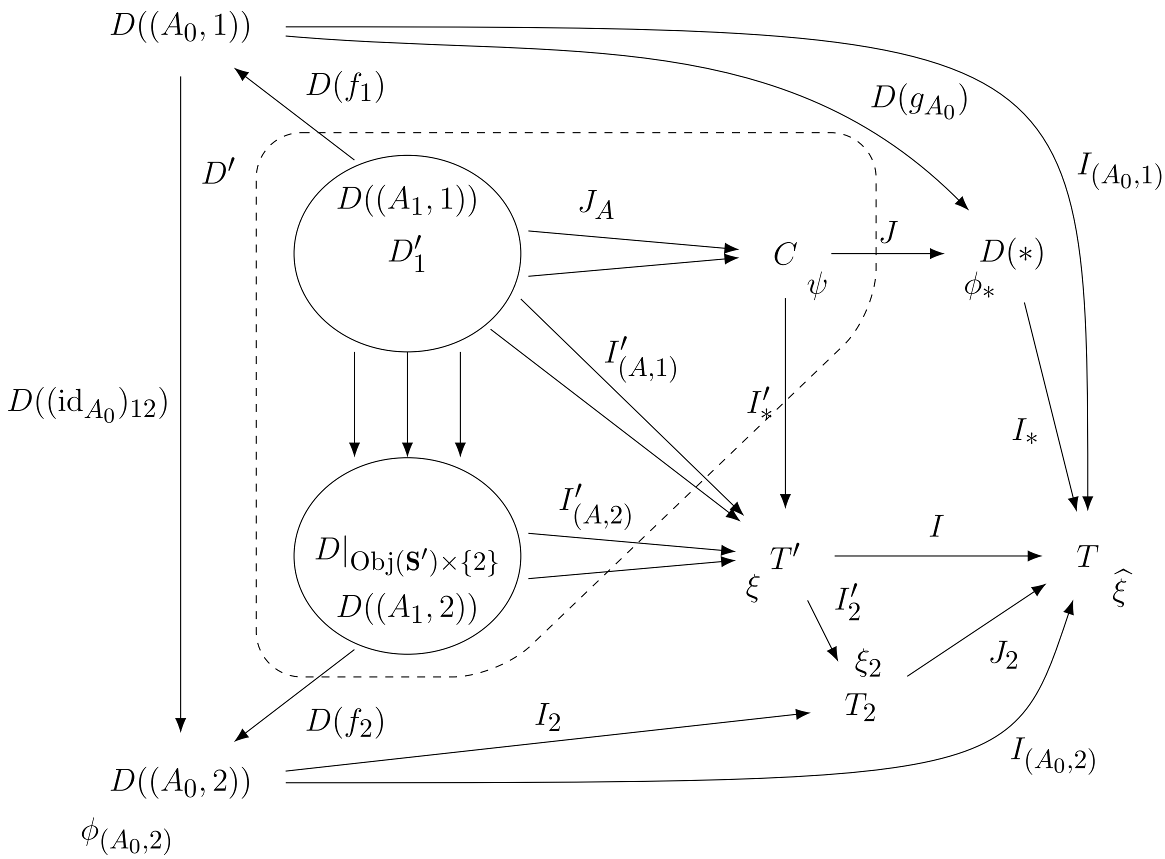

Figure 3.5: Commutative diagram of morphisms and amalgamations constructed in the case when there is an oriented edge f in $\vec{F}$ from some A_1 to A_0; some compositions are omitted. C is the colimit of D'_1, T' is the colimit of D', T is the colimit of D and T_2 is the pushout of $I'_{(A_1,2)}$ and $D(f_2)$.

Proposition 3.7.9. *Let $\tau\colon D_1 \to D_2$ be a natural transformation between finite diagrams $D_1, D_2\colon \mathbf{S} \to \textbf{Int}$ such that $\mathbf{S}$ is forest-like and let $(\phi_A)_{A\in\mathrm{Obj}(\mathbf{S})}$ respect D_2. Let also ξ be an amalgamation of $(\phi_A^{\tau_A})_{A\in\mathrm{Obj}(\mathbf{S})}$ over D_1 with respect to a colimit $(C_1, (I^1_A)_{A\in\mathrm{Obj}(\mathbf{S})})$, let $(C_2, (I^2_A)_{A\in\mathrm{Obj}(\mathbf{S})})$ be a colimit of D_2 and let $I\colon C_1 \to C_2$ be the universal* **Int**-*morphism such that $I^2_A \circ \tau_A = I \circ I^1_A$ for every $A \in \mathrm{Obj}(\mathbf{S})$.*

Then there exists an amalgamation $\widehat{\xi}$ of $(\phi_A)_{A\in\mathrm{Obj}(\mathbf{S})}$ over D_2 with respect to the colimit $(C_2, (I^2_A)_{A\in\mathrm{Obj}(\mathbf{S})})$ such that $\widehat{\xi}^I = \xi$.

Proof. Note that the colimit C_2 of D_2 is isomorphic to the colimit of D_{τ,C_1} and that the family of limit objects $((\phi_A^{\tau_A})_{A\in\mathrm{Obj}(\mathbf{S})}, (\phi_A)_{A\in\mathrm{Obj}(\mathbf{S})}, \xi)$ (indexed in order by $\mathrm{Obj}(\mathbf{S}) \times$

$\{1\}, \mathrm{Obj}(\mathbf{S}) \times \{2\}, \{*\})$ respects D_{τ,C_1}, so by Theorem 3.7.8, there exists an amalgamation $\widehat{\xi} \in \mathrm{Hom}^+(\mathcal{A}[C_2], \mathbb{R})$ of this family over D_{τ,C_1} and such $\widehat{\xi}$ satisfies the required properties. ■

Just as Proposition 3.2.6 is used to show Proposition 3.2.9, we can use Proposition 3.7.9 to show Theorem 3.7.6.

Proof of Theorem 3.7.6. It is clear that $(\phi_A^{\tau_A})_{A \in \mathrm{Obj}(\mathbf{S})}$ respects D_1, so by Theorem 3.7.5, there exists at least one amalgamation of this family over D_1.

On the other hand, if ξ is one such amalgamation, Proposition 3.7.9 gives us an amalgamation $\widehat{\xi}$ of $(\phi_A)_{A \in \mathrm{Obj}(\mathbf{S})}$ over D_2 such that $\widehat{\xi}^I = \xi$ for the universal **Int**-morphism I from the colimit of D_1 to the colimit of D_2 that factors through τ. Since this latter family is uniquely amalgamatable over D_2, it follows that every amalgamation ξ of the former family over D_1 is of the form $\xi = \widehat{\xi}^I$. ■

3.8 Concluding remarks and open problems

In this section we have introduced the basic concepts of continuous combinatorics of rank, independence, couplings, amalgamations and the L_1-topology. We have seen that rank behaves very differently in the L_1-topology and in the density topology and we have seen that the axioms of a theory can force it to have rank much lower than its arity. Finally, we have seen that theons, couplings and amalgamations can be lifted through open interpretations, which is a fundamental property to the study of unique coupleability and unique amalgamatability.

Section 3.4 was completely devoted to the study of semi-continuity of rank with respect to the L_1-topology and the failure of its continuity in the density topology, in other words, the property $\mathrm{rk}(-) \leq \ell$ is closed in L_1-topology but not necessarily closed in density topology. However, we do not know if the dual set of `Independence`$[\ell]$ is closed in either topology. We will return to this question and its importance in Section 4.9.

In Proposition 3.5.4, we provided a sufficient condition for a theory T obtained from $T_{k\text{-Hypergraph}}$ to satisfy $\mathrm{rk}(T) \leq r$ for any fixed $r < k$. In the particular case of $r = k-1$, we showed in Theorem 3.5.5 that this condition is also necessary, and is equivalent to requiring that the VC dimension of neighborhoods of vertices are bounded by some constant $d = d(T)$ (see Remark 6). In a recent work, Chernikov and Towsner [13] completely characterized theories of k-hypergraphs of rank at most r as the theories that have finite VC_r dimension (see the aforementioned paper for the definition) in the language of regularity lemmas and graded probability spaces.

In Section 3.6, we have seen how to strengthen Proposition 3.2.1 on lifting theons through open interpretations to a version that holds everywhere except for a version of the diagonal (Theorem 3.6.7). For such lifting to be possible, we required the theon to be robust. Even though Theorem 3.6.6 says that we can get theon robustness by changing only a zero-measure set, the concept of theon robustness (or even the weaker notion of theon soundness) is not very "natural" as it only makes sense for theons over $[0,1]^d$ and it is not preserved under measure-automorphisms of $[0,1]^d$. Of course, one could simply close this property under measure-isomorphisms between any two spaces, but that would make the definition of robustness even more technical and arguably less "natural".

On the other hand, the discussion in the beginning of Section 3.6 suggests that the main obstacle of "hidden axioms" to Theorem 3.6.7 is already surpassed with the weaker condition of $\mathrm{Th}(\phi_{\mathcal{N}})$-strength (or the even weaker condition of $I(\mathrm{Th}(\phi))$-strength), which brings us to the question of whether it is possible to replace the robustness condition by any of these two conditions (of course, one should also drop the robustness result of the constructed T_2-on $\mathcal{H}$ as well). Even replacing the robustness condition by the much simpler notion of soundness in Theorem 3.6.7 would already be interesting.

In Section 3.7, we have seen a generalization of the notions of couplings and amalgamations to general diagrams in the category **Int**. We have seen that not every diagram shape amalgamates theons, but at least finite forest-like shapes have this property. However, it is easy to see that there are many shapes that amalgamate theons that are not forest-like. For example, any shape **S** containing a terminal object A_1 trivially amalgamates theons as for every diagram $D\colon \mathbf{S} \to \textbf{Int}$, $D(A_1)$ is a colimit of D, thus if $(\phi_A)_{A \in \mathrm{Obj}(\mathbf{S})}$ respects D, then ϕ_{A_1} is an amalgamation of this family. An interesting problem would be characterizing exactly which shapes amalgamate theons.

We have also seen in Theorem 3.7.6 that for a finite forest-like shape $\mathbf{S}$, the fact that $\widehat{\mathbf{S}}$ amalgamates theons implies that unique amalgamatability of a diagram of shape $\mathbf{S}$ is preserved under natural transformations. Unfortunately, the ad hoc proof in Theorem 3.7.8 that $\widehat{\mathbf{S}}$ amalgamates theons heavily used the fact that $\mathbf{S}$ is forest-like, so we would like to ask if $\widehat{\mathbf{S}}$ amalgamates theons whenever $\mathbf{S}$ amalgamates theons. This would immediately imply a generalization of Theorem 3.7.6 to every $\mathbf{S}$ that amalgamates theons.

Finally, let us point out that several of the results of Section 3.7 can be proven in the more general setting of category theory. More specifically, suppose we are given a category $\mathbf{X}$ and a (covariant) functor $F\colon \mathbf{X} \to \textbf{Set}$ to the category $\textbf{Set}$ of sets. Then we can say that for a diagram $D\colon \mathbf{S} \to \mathbf{X}$, a limit[4] $(L, (f_A)_{A \in \mathrm{Obj}(\mathbf{S})})$ of D and a family $(x_A)_{A \in \mathrm{Obj}(\mathbf{S})}$ with $x_A \in F(D(A))$ for every $A \in \mathrm{Obj}(\mathbf{S})$, an *amalgamation* of $(x_A)_{A \in \mathrm{Obj}(\mathbf{S})}$ over D with respect to $(L, (f_A)_{A \in \mathrm{Obj}(\mathbf{S})})$ is a $y \in F(D(L))$ such that $F(D(f_A))(y) = x_A$ for every $A \in \mathrm{Obj}(\mathbf{S})$ and the other notions are defined analogously. We can also say that a shape $\mathbf{S}$ *amalgamates objects* of F when every family that respects a diagram $D\colon \mathbf{S} \to \mathbf{X}$ has an amalgamation over it.

Under these definitions, it is easy to see that if $\mathbf{X}$ contains pullbacks and a terminal object (in particular, it contains limits of all finite forest-like shapes) and the shape of pullbacks amalgamates objects of F (in particular, finite discrete shapes amalgamate objects of F), then Theorems 3.7.5, 3.7.6 and 3.7.8 and Proposition 3.7.9 hold for $F\colon \mathbf{X} \to \textbf{Set}$ after appropriate dualization: forest-like is a self-dual notion, $\widehat{\mathbf{S}}$ is replaced by $\widehat{\mathbf{S}}^{\mathrm{op}}$, amalgamations of $(F(\tau_A)(x_A))_{A \in \mathrm{Obj}(\mathbf{S})}$ over $D_2\colon \mathbf{S} \to \mathbf{X}$ can be lifted to amalgamations of $(x_A)_{A \in \mathrm{Obj}(\mathbf{S})}$ over $D_1\colon \mathbf{S} \to \mathbf{X}$ whenever $\tau\colon D_1 \to D_2$ is a natural transformation and $\mathbf{S}$ is finite forest-like,

4. It is more natural to state the dual version of what we used for $\textbf{Int}$ since the functor $T \mapsto \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ and $I \mapsto -^I$ is contravariant; alternatively, it can be viewed as a covariant functor $F\colon \textbf{Int}^{\mathrm{op}} \to \textbf{Set}$ from the dual category $\textbf{Int}^{\mathrm{op}}$ to $\textbf{Set}$.

and in the same setup unique amalgamatability of the latter family over D_1 implies unique amalgamatability of the former family over D_2. This means that all questions about unique amalgamatability and related concepts can also be studied in the more general setting of category theory.

CHAPTER 4

NATURAL QUASIRANDOMNESS

In this chapter, we present an attempt at a more systematic study of quasirandom properties that can be reasonably identified as "intrinsic" (for reasons that will become clear very shortly, we will also use in this context the word "natural"). As we mentioned in Chapter 1, so far the theory of quasirandomness has mostly been studied in a case-by-case manner, with very few attempts at an intrinsic definition of quasirandomness. One of the equivalent properties in the seminal paper [17] (P_3) was of spectral nature, namely it requested the second largest eigenvalue of G_n to be $o(|G_n|)$. This spectral theme was further continued for (linear) quasirandom hypergraphs in [50, 52].

Even though most other quasirandomness properties in the literature are stated in terms of counting, it is still possible to extract from them something intrinsic. For example, the property P_4 in [17] (see also [61, Theorem 2.4]) implies that quasirandom limits W are the only graphons with the following unique inducibility property: if $(G_n)_{n\in\mathbb{N}}$ converges to W then the sequence of induced graphs $(G_n|_{U_n})_{n\in\mathbb{N}}$ also converges to W as long as $|U_n| \geq \Omega(|G_n|)$. As another example, using graphon language [54], we can extract a trivial intrinsic characterization of quasirandom limits in terms of an independence property: a graphon $W\colon [0,1]^2 \to [0,1]$ is quasirandom if and only if W a.e. does not depend on its variables, that is, it is a.e. constant.

Let us now explain what we mean by "intrinsic" or "natural" quasirandomness. Our explanation will be deliberately informal and open-ended; instead of trying to give a rigorous definition, we present a set of tests that in our view have to be passed and then describe some concrete properties we will be studying in this chapter that pass these tests.

First and foremost, in line with the generality of the theories of flag algebras [59] and theons [24] in continuous combinatorics, we require qualifying properties to be formulated in a uniform way for arbitrary universal theories in a finite relational language.

The next two requirements are somewhat derivative of the first.

We require that the property should not refer to densities of concrete models and their explicit values (thus, this is more about the *formulation* of the property than the class of objects defined by it.) The reason is that any such definition is necessarily somewhat arbitrary. For example, there is no such thing as "edge densities" in the theories of tournaments and permutations so their ad hoc analogues had to be found when defining quasirandom objects in those contexts. Of the quasirandom graph properties mentioned above, the description as a constant graphon definitely satisfies this criterion, and so does the inducibility property (the tweak of P_4 in [17]). Spectral properties also pass the test but unfortunately they fail (given our current state of knowledge) the previous universality test.

The next requirement is that we want the property to be preserved under open interpretations, and this is where the word "natural" (like in "natural transformations"; recall that open interpretations form a category **Int**) comes in. In plain words, everything that can be syntactically defined in a quasirandom object must display proportionally strong quasirandom properties. Again, in an implicit form this requirement was exploited in the previous literature both in positive and negative manner. For example, the proofs of the implications $P_{10} \Longrightarrow P_{11} \Longrightarrow P_1(s)$ in the seminal paper on quasirandom tournaments [15] can be viewed as divided into two parts. First one proves that the quasirandom graph is uniquely coupleable (see Definition 3.2.7) with the linear order, then the tournament obtained from the resulting quasirandom ordered graph via the "arc-orientation" interpretation must be quasirandom. As for "negative" use, let us note that most separations in the hierarchy of quasirandom hypergraphs [1, 51, 66] can be viewed as coming from the fact that these properties are *not* preserved under open interpretations between the theories of hypergraphs of possibly different arity. We will elaborate on this in Section 4.6 (see Theorem 4.2.16).

Our final requirement is more "traditional", and it is well-rooted in the previous literature. Namely, we require that the property should be satisfied asymptotically almost surely for some "natural" random model of some "natural" theory T. Examples of "natural" random models include, of course, the Erdős–Rényi model and its generalization to hypergraphs, the random tournament, the random permutation, etc.

This list of requirements may appear to be rather restrictive, so let us describe quasirandom properties we are studying; they are essentially far-reaching generalizations of what we already discussed above. Several more remarks are in place before we begin.

1. We have deliberately decided against attempting to state our properties in the language of finite combinatorial objects and their asymptotic behavior – it is probably possible but the result might be rather ugly and disappointing. Instead, we use the language of graphons [54], hypergraphons [34] and theons [24] for the geometric view of our objects and that of flag algebras [59] for a concise algebraic description. The advantages of using the continuous setting are illustrated by the fact that such proofs are often more elegant and less technical than their finite world counterparts [45, 49, 66]. This view is more instructive, too: for example, by looking back through the lenses of graphons, we can extract an elegant graphon proof of quasirandomness of property $P_2(4)$ of [17] based on the Lebesgue Density Theorem from a paper as early as [33, Theorem 3.10].

   However, for the benefit of more combinatorially-oriented reader we try to inject as much of "finite intuition" as possible in appropriate places.

2. Our properties are not equivalent with those previously studied in the literature even for hypergraphs (see Figure 4.2). Hence the reader interested only in this case can safely assume that our base theory is $T_{k\text{-Hypergraph}}$ for some $k \geq 3$, and the objects are just hypergraphons. But let us mention that more complicated objects like colorings, orderings, couplings, etc. will pop up in the statements and the proofs anyway.

3. Finally, the description below is loose and sweeps under the rug some important technicalities. Proper definitions are deferred to Section 4.1.1.

`Independence`[ℓ]**.** As we have seen in (3.1), the quasirandom graphon of density p can be represented by a 1-independent T_{Graph}-on (see Definition 3.1.1). More generally, the quasirandom k-hypergraphon of density p can be represented by the $(k-1)$-independent $T_{k\text{-Hypergraph}}$-on

$$\mathcal{G} \stackrel{\text{def}}{=} \{x \in \mathcal{E}_k \mid x_{[k]} < p\}.$$

This is the strongest in the hierarchy of our quasirandomness properties, and it relatively easily implies all the others, with the same value of the parameter ℓ.

`UCouple`$[\ell]$ (Unique ℓ-coupleability). Our next property is based on the notion of unique coupleability of Definition 3.2.7 and the notion of rank of Definition 3.1.1. We say that ϕ is *uniquely ℓ-coupleable* if it is uniquely coupleable with all objects ψ such that $\mathrm{rk}(\psi) \leq \ell$. Intuitively, this means that ϕ "looks random" from the perspective of any low rank limit objects ψ, as they cannot detect any pattern in ϕ via couplings.

`UInduce`$[\ell]$ (Unique ℓ-inducibility). One equivalent way to view the induced subgraph $G|_V$ is this: we first color the vertices into two colors, say, green (corresponding to V) and red. Then instead of removing red vertices, we remove all edges adjacent to at least one red vertex. In this form, it has a perfect generalization in higher dimensions. Namely, we consider couplings ξ of a limit object $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ with an ℓ-hypergraph limit $\psi \in \mathrm{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$ (note that $\mathrm{rk}(\psi) \leq \ell$). The unique coupleability requires that $\xi(M) = (\phi \otimes \psi)(M)$ for any model M of the combined theory $T \cup T_{\ell\text{-Hypergraph}}$, where $\phi \otimes \psi$ is the independent coupling of ϕ and ψ (see Definitions 3.2.3 and 3.2.4). Unique inducibility by ψ relaxes this property by requiring that $\xi(M) = (\phi \otimes \psi)(M)$ holds only for those M that are based on a clique in the hypergraph theory. The limit object $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is *uniquely ℓ-inducible* if it is uniquely inducible by any ℓ-hypergraph limit $\psi \in \mathrm{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$.

From the loose formulation of the properties above, one can already see that the first two "naturality" requirements are satisfied: the formulations are made for arbitrary theories and do not refer to densities of concrete models and their explicit values. As for the third "naturality" requirement (Theorem 4.2.3), we will see that for `Independence`$[\ell]$ it trivially follows from the general theory, for `UCouple`$[\ell]$ it trivially follows from Proposition 3.2.9 and for `UInduce`$[\ell]$ it will follow from coupling lifting (Proposition 3.2.6).

As we mentioned before, the quasirandom k-hypergraph satisfies `Independence`$[k-1]$. The situation for asymmetric combinatorial objects is more diverse. For example, the quasirandom tournament satisfies `UCouple`$[1]$ but not `Independence`$[1]$ and this example can be generalized to higher values of ℓ. One interesting example for unique inducibility is the linear order as it satisfies `UInduce`$[\ell]$ for every ℓ without being a trivial object.

All our properties are anti-monotone in ℓ in the sense that for any of the above, we have the implications $P[\ell] \implies P[\ell-1]$ (see Theorem 4.2.1) and as for relations between the properties (Theorem 4.2.2), we show that `Independence`$[\ell]$ implies `UCouple`$[\ell]$ and that `UCouple`$[\ell]$ implies[1] `UInduce`$[\ell]$ (see Figure 4.1).

In terms of separations, we show that no upward implication holds, that is, none of the studied quasirandomness properties with parameter ℓ can imply the same, or for that matter any other, property with parameter $\ell+1$ (Theorem 4.2.5). As for separations between different families of properties, we show that `UCouple`$[\ell]$ does not imply `Independence`$[\ell]$ (Theorem 4.2.6) and `UInduce`$[\ell]$ does not imply even `UCouple`$[1]$ (Theorem 4.2.7). At an initial stage, we left open the relation between `UCouple`$[\ell]$ and `Independence`$[\ell-1]$, but after personal communication with Henry Towsner, we obtained an argument for `UCouple`$[\ell] \implies$ `Independence`$[\ell-1]$, which will appear in a future joint work. All these separations are relatively easy when we are working with arbitrary theories, but we show that they still hold even if we restrict ourselves to the theory of k-hypergraphs, for $k \geq \ell+2$ (Theorems 4.2.8 and 4.2.9).

Next, we provide the following alternate characterizations (summarized in Theorems 4.2.10 and 4.2.11) of these classes.

1. This implication is obvious from the definition.

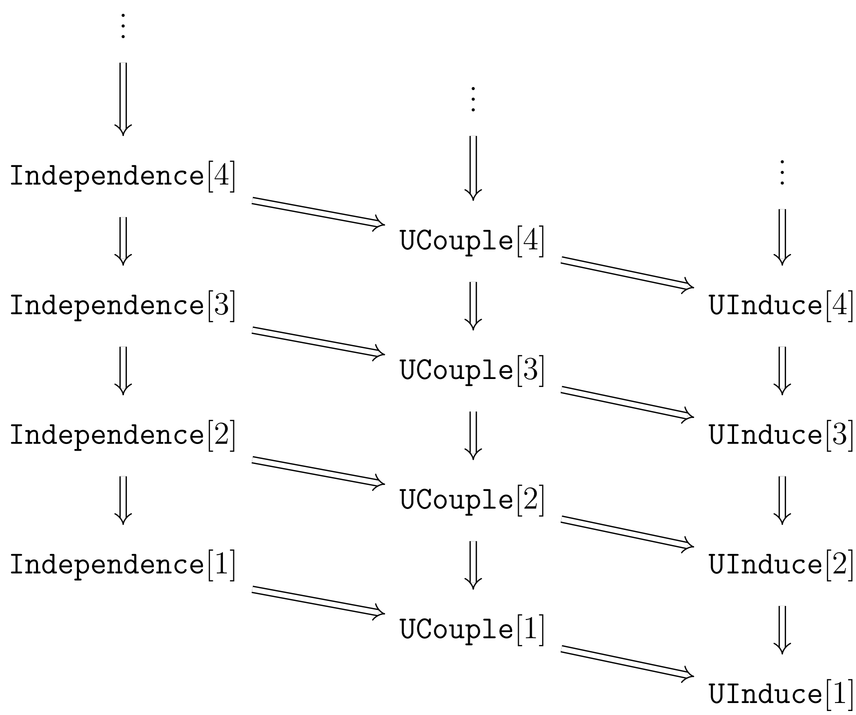

Figure 4.1: Implications between quasirandomness properties. The implications $\texttt{UCouple}[\ell] \implies \texttt{Independence}[\ell-1]$ will appear in a future work.

Weak ℓ-independence. Recall that every limit object $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ can be alternatively represented by a local exchangeable array $\boldsymbol{K}$ defined from a collection of independent random variables $(\boldsymbol{\theta}_A)_A$ indexed by finite non-empty subsets of $\mathbb{N}_+$ (see Theorem 2.2.1 and the end of Section 2.4). We say that ϕ is *weakly ℓ-independent* if $\boldsymbol{K}$ is independent from $(\boldsymbol{\theta}_A \mid |A| \leq \ell)$ **as a random variable** (full $\texttt{Independence}[\ell]$ requires this to happen "pointwise"). This weak version of independence turns out to be equivalent to $\texttt{UCouple}[\ell]$ (Theorem 4.2.10(iv)).

ℓ-Locality. Recall that the locality property of the exchangeable array $\boldsymbol{K}$ required the marginals $(\boldsymbol{K}|_{V_i} \mid i \in I)$ to be mutually independent whenever the collection of finite sets $(V_i)_{i\in I}$ is pairwise disjoint. The notion of *ℓ-locality* strengthens this property to require mutual independence of $(\boldsymbol{K}|_{V_i} \mid i \in I)$ whenever the collection of finite sets $(V_i)_{i\in I}$ have pairwise intersections of size at most ℓ. It is clear that weak ℓ-independence implies ℓ-locality, but we prove that the converse also holds, hence ℓ-locality is also equivalent to $\texttt{UCouple}[\ell]$ (Theorem 4.2.10(vi)).

Symmetric ℓ-locality. The notion of *symmetric ℓ-locality* relaxes the notion of ℓ-locality by requiring only mutual independence of the events $(\boldsymbol{K}|_{V_i} \cong M_i \mid i \in I)$ for all choices of $(V_i)_{i\in I}$ with pairwise intersections of size at most ℓ and all choices of models M_i, i.e., we only care about the submodels $\boldsymbol{K}|_{V_i}$ up to isomorphism. We show that symmetric ℓ-locality is equivalent to $\texttt{UInduce}[\ell]$ (Theorem 4.2.11(iii)).

The right way to view the definitions of unique coupleability and unique inducibility is that each ψ of rank $\leq \ell$ generates a test for the respective property that ϕ has to pass. It is natural to ask for a smaller and more explicit set of universal tests that guarantees each property. We show (Theorem 4.2.10(ii)) that $\phi \in \texttt{UCouple}[\ell]$ is equivalent to ϕ being uniquely coupleable with a non-degenerate **quasirandom** ℓ'-hypergraphon $\psi_{\ell',p}$ in every dimension $\ell' \leq \ell$. We further prove (Theorem 4.2.10(iii)) that it is also equivalent to ϕ being uniquely coupleable with their independent coupling $\psi_{1,p_1} \otimes \ldots \otimes \psi_{\ell,p_\ell}$; for the reasons explained right after the statement of the theorem, it does not immediately follow from the previous item (ii). In the particular case $\ell = 1$, this means that the fact that ϕ is uniquely coupleable with a single non-trivial vertex-coloring implies it must also be uniquely coupleable with any rank 1 limit object, such as linear orders, permutations, etc.

Our findings for unique inducibility are by far less conclusive but at least we can show that it is sufficient to consider only hypergraphons ψ with any fixed non-trivial edge density $p \in (0, 1)$ (Theorem 4.2.11(ii)).

Of all choices of parameters, arguably the most interesting one is when ℓ is exactly one less than the maximum arity k of a predicate of the language. In the theory of k-hypergraphs the three classes with $\ell = k - 1$ become the same and are satisfied only by the full quasirandom hypergraph, that is, the almost sure limit of the generalization of the Erdős–Rényi model. If we consider general theories of arity at most k, it is not hard to see (Theorem 4.2.12) that $(k-1)$-independent objects are (essentially) quasirandom colored k-hypergraphs. The property $\texttt{UCouple}[k-1]$ in arity at most k corresponds to independent couplings of quasirandom colored k-hypergraphs with generalizations of quasirandom tournaments (Theorem 4.2.13). The case of unique inducibility is (again) considerably more complicated: we only deal with arity at most 2, in which case $\texttt{UInduce}[1]$ corresponds to (essentially) independent couplings of quasirandom colored graphs with a linear order (Theorem 4.2.14).

Finally, let us compare our properties to the known hypergraph quasirandomness properties (Figure 4.2). In [66], Towsner defined k-hypergraph quasirandomness properties $\texttt{Disc}[\mathcal{A}]$ for every antichain $\mathcal{A}$ of non-empty subsets of $[k] \stackrel{\text{def}}{=} \{1, \ldots, k\}$ and showed that $\texttt{Disc}[\binom{[k]}{\ell}]$ and $\texttt{Disc}[\mathcal{A}_\ell]$ are equivalent to $\texttt{CliqueDisc}[\ell]$ and $\texttt{Dev}[\ell]$ of [51], respectively, where $\mathcal{A}_\ell \stackrel{\text{def}}{=} \{A \in \binom{[k]}{k-1} \mid [k-\ell] \subseteq A\}$. It is immediate from definitions that $\texttt{UInduce}[\ell]$ implies $\texttt{CliqueDisc}[\ell]$ (Theorem 4.2.15). In terms of separations between our properties and the ones from the literature, we show the strongest separation possible. The strongest $\texttt{Disc}[\mathcal{A}]$ property that is not equivalent to full quasirandomness is $\texttt{Dev}[k-1]$ and this does not imply even $\texttt{UInduce}[1]$ (Theorem 4.2.16). In the other direction, the weakest $\texttt{Disc}[\mathcal{A}]$ property that is not implied by

CliqueDisc[ℓ] is Disc[{[$\ell+1$]}] and this is not implied by Independence[ℓ] (Theorem 4.2.17).

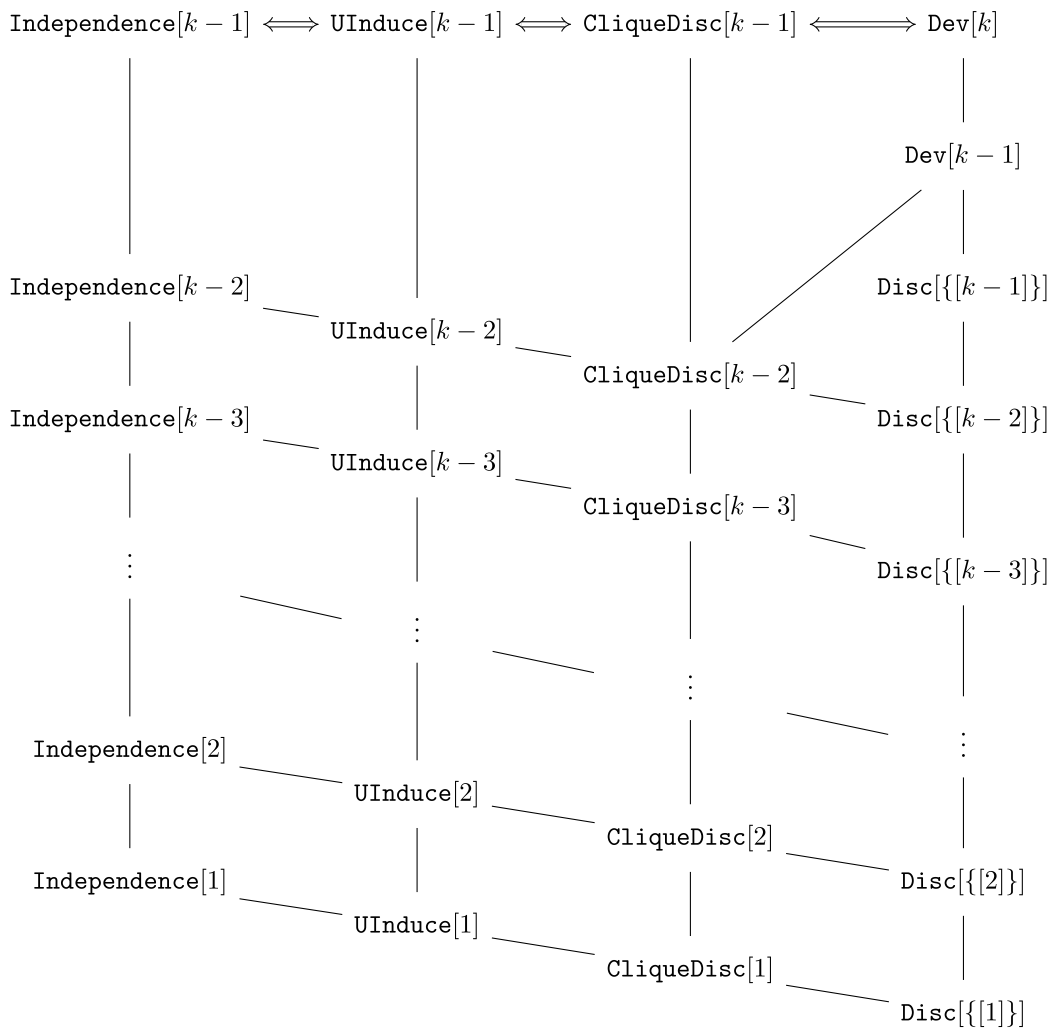

Figure 4.2: Hasse diagram of quasirandomness hypergraph properties in arity k. The top four equivalent properties represent full quasirandomness.

This chapter is organized as follows. In Section 4.1 we formally define our quasirandomness properties and some limit objects needed to state our main theorems. In Section 4.2 we formally state our main results on quasirandomness. In Section 4.3, we prove some basic facts that will be used throughout the chapter. In Section 4.4 we prove the alternative formulations of UInduce, and in Section 4.5 we prove the alternative formulations of UCouple. The proofs are done in this slightly reversed order because they are simpler for the unique inducibility;

besides, some auxiliary statements we need for that part are later re-used for the unique coupleability. In Section 4.6, we show separations between different classes of properties. In Section 4.7, we completely classify the properties `Independence`$[k-1]$ and `UCouple`$[k-1]$ when all arities are at most k and classify `UInduce`$[1]$ when all arities are at most 2. In Section 4.8 we discuss a generalization of the notions of rank and `Independence` and which results can easily be transferred. The chapter is concluded with a few remarks and open problems in Section 4.9.

4.1 Preliminaries

4.1.1 Quasirandomness properties

In this subsection we formalize all notions of quasirandomness presented in the beginning of the chapter.

Definition 4.1.1 (weak independence)**.** For $\ell \in \mathbb{N}$, an Euclidean structure $\mathcal{N}$ on $\mathcal{L}$ over Ω is *weakly ℓ-independent* if the exchangeable array $\boldsymbol{K}$ corresponding to $\mathcal{N}$ with respect to $\boldsymbol{\theta}$ picked in $\mathcal{E}_{\mathbb{N}_+}(\Omega)$ according to μ (see (2.6)) is independent from $(\boldsymbol{\theta}_A \mid A \in r(\mathbb{N}_+, \ell))$ as a random variable.

We say $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is *weakly ℓ-independent* if there exists a weakly ℓ-independent T-on $\mathcal{N}$ such that $\phi_{\mathcal{N}} = \phi$.

Definition 4.1.2 (unique coupleability and inducibility)**.** Recall from Definition 3.2.7 that ϕ_1, ϕ_2 are uniquely coupleable if the independent coupling $\phi \otimes \psi$ is their only coupling. For $\ell \in \mathbb{N}$, we say that $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is *uniquely ℓ-coupleable* if for every theory T' and every $\psi \in \mathrm{Hom}^+(\mathcal{A}[T'], \mathbb{R})$ with $\mathrm{rk}(\psi) \leq \ell$, ϕ and ψ are uniquely coupleable. We will be using the abbreviation `UCouple`$[\ell]$ for this property.

Given $\ell \in \mathbb{N}_+$, $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ and $\psi \in \mathrm{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$, we say that ϕ is *uniquely inducible* by ψ if for any coupling ξ of ϕ and ψ and for every $M \in \mathcal{M}[T \cup T_{\ell\text{-Hypergraph}}]$ such that $I(M)$ is a complete ℓ-hypergraph, we have $\xi(M) = (\phi \otimes \psi)(M)$, where

$I\colon T_{\ell\text{-Hypergraph}} \rightsquigarrow T \cup T_{\ell\text{-Hypergraph}}$ is the structure-erasing interpretation. We say that ϕ is *uniquely ℓ-inducible* if it is uniquely inducible by every $\psi \in \mathrm{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$, and we will be using the abbreviation $\texttt{UInduce}[\ell]$. For completeness, we declare every ϕ to satisfy $\texttt{UInduce}[0]$.

Remark 8. Since $T_{1\text{-Hypergraph}} \cong T_{2\text{-Coloring}}$, for $\ell = 1$ we prefer to work with the following equivalent formulation of $\texttt{UInduce}[1]$ that can be deduced from this isomorphism. $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is uniquely inducible by $\psi \in \mathrm{Hom}^+(\mathcal{A}[T_{2\text{-Coloring}}], \mathbb{R})$ if for any coupling ξ of ϕ and ψ and for every $M \in \mathcal{M}[T \cup T_{2\text{-Coloring}}]$ such that $R_{\chi_1}(M) = V(M)$, we have $\xi(M) = (\phi \otimes \psi)(M)$. Then ϕ is uniquely 1-inducible if it is uniquely inducible by every $\psi \in \mathrm{Hom}^+(\mathcal{A}[T_{2\text{-Coloring}}], \mathbb{R})$.

Also, as we will see below (Theorem 4.2.1), $\texttt{UInduce}[\ell]$ implies $\texttt{UInduce}[\ell']$ for any $\ell' \leq \ell$. Hence, we could have equivalently required in this definition unique inducibility by every $\psi \in \mathrm{Hom}^+(\mathcal{A}[T_{\ell'\text{-Hypergraph}}], \mathbb{R})$ with $\ell' \leq \ell$.

These properties are central to our study of quasirandomness. If P is any of them, we will say interchangeably that ϕ satisfies $P[\ell]$ or that $\phi \in P[\ell]$.

Definition 4.1.3 (locality)**.** Let $\mathcal{N}$ be a T-on over $\Omega = (X, \mathcal{A}, \mu)$ and let $\boldsymbol{K}$ be the exchangeable array corresponding to $\mathcal{N}$ with respect to $\boldsymbol{\theta}$ picked in $\mathcal{E}_{\mathbb{N}_+}(\Omega)$ according to μ (see (2.6)).

We say that $\mathcal{N}$ is *ℓ-local* if for every collection $(V_i)_{i\in I}$ of finite subsets of $\mathbb{N}_+$ with pairwise intersections of size at most ℓ, the marginals $(\boldsymbol{K}|_{V_i} \mid i \in I)$ are mutually independent.

We say that $\mathcal{N}$ is *symmetrically ℓ-local* if for every collection $(V_i)_{i\in I}$ of finite subsets of $\mathbb{N}_+$ with pairwise intersections of size at most ℓ, the random variables $([\boldsymbol{K}|_{V_i}] \mid i \in I)$ (recall that $[K]$ is the isomorphism type of K) are mutually independent.

We say that $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is *ℓ-local* (resp., *symmetrically ℓ-local*) if there exists an ℓ-local (resp., symmetrically ℓ-local) T-on $\mathcal{N}$ such that $\phi = \phi_{\mathcal{N}}$.

Note that both the notions of 0-locality and symmetric 0-locality coincide with the notion of locality for $\boldsymbol{K}$ (see Section 2.4). Besides, it is very easy to give an explicit purely syntactic description of both locality and symmetric locality in the style of Definition 3.2.3; this in particular implies that for an ℓ-local (resp., symmetrically ℓ-local) $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$, every T-on $\mathcal{N}$ with $\phi = \phi_{\mathcal{N}}$ must necessarily be ℓ-local (resp., symmetrically ℓ-local).

Finally, let us state the properties `CliqueDisc`$[\ell]$ and `Disc`$[\mathcal{A}]$ in the limit language.

Definition 4.1.4. Let $K_n^{(t)} \in \mathcal{M}_n[T_{t\text{-Hypergraph}}]$ be the complete t-uniform hypergraph on n vertices and let $\rho_t \stackrel{\text{def}}{=} K_t^{(t)}$. Let $\phi \in \mathrm{Hom}^+(\mathcal{A}[T_{k\text{-Hypergraph}}], \mathbb{R})$ and $\ell \in [k]$.

We say that ϕ satisfies `CliqueDisc`$[\ell]$ ([51]) if for every $\psi \in \mathrm{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$ and every coupling ξ of ϕ and ψ, we have

$$\xi(K_k^{(k,\ell)}) = \phi(\rho_k)\psi(K_k^{(\ell)}),$$

where $K_k^{(k,\ell)} \in \mathcal{M}_k[T_{k\text{-Hypergraph}} \cup T_{\ell\text{-Hypergraph}}]$ is the model obtained by aligning ρ_k and $K_k^{(\ell)}$ (i.e., the model of size k that is a complete hypergraph in both theories).

Given an antichain $\mathcal{A} \subseteq r(k)$, let $\mathcal{L}_{\mathcal{A}}$ be the language containing one predicate symbol P_A of arity $k(P_A) \stackrel{\text{def}}{=} |A|$ for every $A \in \mathcal{A}$. We say that ϕ satisfies `Disc`$[\mathcal{A}]$ ([66, 1]) if for every $\psi \in \mathrm{Hom}^+(\mathcal{A}[T_{\mathcal{L}_{\mathcal{A}}}], \mathbb{R})$ and every coupling ξ of ϕ and ψ, if $\boldsymbol{K}$ is the exchangeable array in $\mathcal{K}_{\mathbb{N}_+}[T_{k\text{-Hypergraph}} \cup T_{\mathcal{L}_{\mathcal{A}}}]$ associated with ξ, then we have

$$\begin{aligned}\mathbb{P}[(1, \ldots, k) \in R_E(\boldsymbol{K}) \wedge \forall A \in \mathcal{A}, \iota_A \in R_{P_A}(\boldsymbol{K})] \\ = \phi(\rho_k) \cdot \mathbb{P}[\forall A \in \mathcal{A}, \iota_A \in R_{P_A}(\boldsymbol{K})],\end{aligned}$$

that is, the events $(1, 2, \ldots, k) \in R_E(\boldsymbol{K})$ and $\forall A \in \mathcal{A}, \iota_A \in R_{P_A}(\boldsymbol{K})$ are independent.

In [66], the definition of `Disc`$[\mathcal{A}]$ further requires symmetry of the predicate symbols P_A, but it was shown in [1] that this condition can be dropped.

4.1.2 Useful theories and objects

In this subsection, we define some theories and limit objects that are necessary to formally state some of our main results. We will denote by ψ_{lin} the (unique) element of $\text{Hom}^+(\mathcal{A}[T_{\text{LinOrder}}], \mathbb{R})$. As for the rest, we start with a very general definition (that nonetheless will be used in full generality in Theorem 4.2.13) and then derive all others as special cases.

For $c \geq 2$, let $\Pi_c \stackrel{\text{def}}{=} \{p = (p_i)_{i=1}^c \in (0,1)^c \mid \sum_{i=1}^c p_i = 1\}$ be the interior of the standard $(c-1)$-dimensional simplex. Also, given $x \in \mathcal{E}_n$, let $\sigma_x \in S_n$ be the unique permutation such that $x_{\{\sigma_x^{-1}(1)\}} < \cdots < x_{\{\sigma_x^{-1}(n)\}}$ when the coordinates $(x_{\{i\}} \mid i \in [n])$ are distinct, and define it arbitrarily otherwise.

Definition 4.1.5 (S_k-action theories)**.** Let $k \in \mathbb{N}_+$, let $\mathcal{L}$ be a language containing only predicate symbols of arity exactly k, let $\Theta\colon S_k \times \mathcal{L} \to \mathcal{L}$ be a (left) action of S_k on $\mathcal{L}$ and write $\sigma \cdot P \stackrel{\text{def}}{=} \Theta(\sigma, P)$. The canonical theory T_Θ is defined as the theory over $\mathcal{L}$ with axioms

$$\left(\bigwedge_{1\leq i<j\leq k} x_i \neq x_j\right) \leftrightarrow \left(\bigvee_{P\in\mathcal{L}} P(x_1,\ldots,x_k)\right); \tag{4.1}$$

$$P(x_{\sigma(1)},\ldots,x_{\sigma(k)}) \leftrightarrow (\sigma\cdot P)(x_1,\ldots,x_k) \qquad (P\in\mathcal{L}, \sigma\in S_k); \tag{4.2}$$

$$\neg P(x_1,\ldots,x_k) \vee \neg P'(x_1,\ldots,x_k) \qquad (P,P'\in\mathcal{L}, P\neq P'). \tag{4.3}$$

Given a $p = (p_P)_{P\in\mathcal{L}} \in [0,1]^{\mathcal{L}}$ with $\sum_{P\in\mathcal{L}} p_P = 1$, the (Θ,p)-*quasirandom homomorphism* is the homomorphism $\psi_{\Theta,p} = \phi_{\mathcal{N}^Z} \in \text{Hom}^+(\mathcal{A}[T_\Theta], \mathbb{R})$, where $\mathcal{N}^Z$ is the T_Θ-on given by[2]

$$\mathcal{N}_P^Z \stackrel{\text{def}}{=} \{x \in \mathcal{E}_k \mid x_{[k]} \in Z_{\sigma_x \cdot P}\} \qquad (P\in\mathcal{L}), \tag{4.4}$$

where $Z = (Z_P)_{P\in\mathcal{L}}$ is a measurable partition of $[0,1]$ with $\lambda(Z_P) = p_P$ $(P\in\mathcal{L})$. When p is

2. We will check that all axioms of T_Θ are satisfied and provide an alternate syntactic description as part of Proposition 4.7.1.

Θ-invariant, we say that $\psi_{\Theta,p}$ is *unbiased*, and in this case $\psi_{\Theta,p}$ corresponds to picking at random for each k-set A, independently of other k-sets, an orbit $O \subseteq \mathcal{L}$ of the action Θ with probability $\sum_{P\in O} p_P$ then uniformly at random choosing an S_k-equivariant assignment of the k-tuples with image A to the elements of O.

Let us now note a few special cases that will play an active role in our study of quasirandomness.

Definition 4.1.6 (c-colored k-hypergraphs)**.** Let $\mathcal{L} = \{E_1, \ldots, E_c\}$ and assume that the action Θ is trivial. In that case we will denote the theory T_Θ by $T_{c,k}$ and call it the *theory of c-colored k-hypergraphs*. The (unbiased) (Θ, p)-quasirandom homomorphism will be called *quasirandom c-colored k-hypergraphon with densities* p and denoted by $\psi_{k,p}$.

Definition 4.1.7 (quasirandom k-hypergraphons)**.** Let us further specify $c = 2$ in the previous definition. Since E_2 is the negation of E_1 and hence can be safely removed, the theory T_Θ is isomorphic to $T_{k\text{-Hypergraph}}$. For $p \in (0, 1)$, the (unbiased) $(\Theta, (p, 1-p))$-quasirandom homomorphism is called the *quasirandom k-hypergraphon of density* p; it will also be denoted by $\psi_{k,p}$.

Definition 4.1.8 (Colorings)**.** Letting instead $k = 1$ in Definition 4.1.5, and keeping the action Θ trivial, we see that T_Θ is naturally isomorphic to the theory $T_{c\text{-Coloring}}$. The (unbiased) quasirandom object will be called *c-coloring with densities* p, $p \in \Pi_c$, and denoted by $\psi_p \in \mathrm{Hom}^+(\mathcal{A}[T_{c\text{-Coloring}}], \mathbb{R})$. For $c = 2$ and $p \in (0, 1)$, $\psi_{(p,1-p)}$ will be often abbreviated to ψ_p (which, in view of Remark 8, is also the same as $\psi_{1,p} \in \mathrm{Hom}^+(\mathcal{A}[T_{1\text{-Hypergraph}}], \mathbb{R})$).

Definition 4.1.9 (k-tournaments)**.** Let now $\mathcal{L} = \{E_1, E_2\}$ and $k \geq 2$, but this time the action Θ is not trivial but instead given by the sign homomorphism $\mathrm{sgn}\colon S_k \to S_2$. Then the only Θ-invariant p is $p_1 = p_2 = 1/2$ and, as in the case of hypergraphons, we can exclude E_2 from the theory. We call it the theory of *k-tournaments* and denote by $T_{k\text{-Tournament}}$; intuitively, this theory corresponds to choosing one of the two possible orientations for every k-set. The (unbiased) quasirandom object $\psi_{\Theta,(1/2,1/2)}$ will then be called the *quasirandom* *k-tournamon* and denoted by ψ_k; thus, $\psi_k \in \mathrm{Hom}^+(\mathcal{A}[T_{k\text{-Tournament}}], \mathbb{R})$, and ψ_2 is the ordinary quasirandom tournamon.

4.2 Main results on natural quasirandomness

In this section we present the main results on quasirandomness. We remark that some of these results follow trivially from definitions and we will point these out as we go along.

Theorem 4.2.1. *The properties `Independence`, `UCouple` and `UInduce` are anti-monotone in the sense that* $P[\ell] \implies P[\ell-1]$.

For `Independence` and `UCouple`, this theorem trivially follows from definitions. Even though it is possible to give an ad hoc proof that `UInduce` is also anti-monotone, this follows trivially from its equivalence with symmetric locality (Theorem 4.2.11 below) and the fact that symmetric locality is trivially anti-monotone.

Theorem 4.2.2. *For any* $\ell \in \mathbb{N}$, `Independence`$[\ell] \implies$ `UCouple`$[\ell] \implies$ `UInduce`$[\ell]$.

The second implication follows trivially from the definitions.

The next theorem concerns preservation of properties under open interpretations.

Theorem 4.2.3 (Naturality)**.** *Let* $I\colon T_1 \rightsquigarrow T_2$ *be an open interpretation and let* $\ell \in \mathbb{N}$. *The following hold for any* $\phi \in \mathrm{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$.

i. *If* $\phi \in$ `Independence`$[\ell]$, *then* $\phi^I \in$ `Independence`$[\ell]$.

ii. *If* $\phi \in$ `UCouple`$[\ell]$, *then* $\phi^I \in$ `UCouple`$[\ell]$.

iii. *If* $\phi \in$ `UInduce`$[\ell]$, *then* $\phi^I \in$ `UInduce`$[\ell]$.

Item (i) follows trivially from the definition of $I(\mathcal{N})$ applied to an ℓ-independent T_2-on $\mathcal{N}$ such that $\phi = \phi_{\mathcal{N}}$ and item (ii) follows trivially from Proposition 3.2.9. Furthermore, applying this theorem to the axiom-adding interpretation $I\colon T_{\mathcal{L}} \rightsquigarrow T$, where $\mathcal{L}$ is the language of T, we see that all our main notions do not depend on non-logical axioms. Nonetheless, using theories and theons (as opposed to arbitrary Euclidean structures) helps to better orient ourselves and put many of the results in the "right" focus.

The next theorem says that both `Independence` and `UCouple` are preserved under independent couplings.

Theorem 4.2.4. *Let $\phi_1 \in \mathrm{Hom}^+(\mathcal{A}[T_1], \mathbb{R})$ and $\phi_2 \in \mathrm{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$. The following hold for $\ell \in \mathbb{N}$.*

i. If $\phi_1, \phi_2 \in$ `Independence`$[\ell]$, then $\phi_1 \otimes \phi_2 \in$ `Independence`$[\ell]$.

ii. If $\phi_1, \phi_2 \in$ `UCouple`$[\ell]$, then $\phi_1 \otimes \phi_2 \in$ `UCouple`$[\ell]$.

Remarkably, this is not true for `UInduce`, and a good example is provided by the quasirandom permuton (see the end of this section).

The next five theorems concern separations between properties, either allowing general theories or restricted to the theory of hypergraphs.

Theorem 4.2.5. `Independence`$[\ell]$ *does not imply* `UInduce`$[\ell + 1]$, *not even when restricted to the theory of k-hypergraphs as long as $k > \ell$.*

In fact, this theorem is a consequence of Theorems 4.2.15 and 4.2.17 below.

The following two theorems are included since the separating objects are quite natural and explicit and the proofs are simpler. But in a sense they will be superseded by Theorems 4.2.8 and 4.2.9.

Theorem 4.2.6. *For every $\ell \in \mathbb{N}_+$, the quasirandom $(\ell + 1)$-tournamon $\psi_{\ell+1}$ satisfies* `UCouple`$[\ell]$ *but does not satisfy* `Independence`$[\ell]$.

Theorem 4.2.7. *The linear order $\psi_{\mathrm{lin}} \in \mathrm{Hom}^+(\mathcal{A}[T_{\mathrm{LinOrder}}], \mathbb{R})$ satisfies* `UInduce`$[\ell]$ *for every $\ell \in \mathbb{N}$ but does not satisfy* `UCouple`$[1]$.

Theorem 4.2.8. *For $\ell \geq 1$, there exists $\phi \in \mathrm{Hom}^+(\mathcal{A}[T_{(\ell+2)\text{-Hypergraph}}], \mathbb{R})$ satisfying* `UCouple`$[\ell]$ *but not satisfying* `Independence`$[\ell]$.

Theorem 4.2.9. *For $\ell \geq 1$ odd, there exists $\phi \in \mathrm{Hom}^+(\mathcal{A}[T_{(\ell+2)\text{-Hypergraph}}], \mathbb{R})$ satisfying* $\texttt{UInduce}[\ell]$ *but not satisfying* $\texttt{UCouple}[1]$.

The next theorem lists several properties that are equivalent to $\texttt{UCouple}[\ell]$. These include both alternative formulations and complete sets of tests for unique coupleability.

Theorem 4.2.10 (Characterization of $\texttt{UCouple}$). *Let $\ell \in \mathbb{N}_+$. The following are equivalent for $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$.*

i. $\phi \in \texttt{UCouple}[\ell]$.

ii. *For every $\ell' \in [\ell]$, there exists $p \in (0,1)$ such that ϕ is uniquely coupleable with the quasirandom ℓ'-hypergraphon $\psi_{\ell',p}$.*

iii. *There exist $p_1, \ldots, p_\ell \in (0,1)$ such that ϕ is uniquely coupleable with the independent coupling $\psi_{1,p_1} \otimes \cdots \otimes \psi_{\ell,p_\ell}$ of the quasirandom ℓ'-hypergraphons $\psi_{\ell',p_{\ell'}}$ for $\ell' \in [\ell]$.*

iv. *ϕ is weakly ℓ-independent.*

v. *Every T-on $\mathcal{N}$ with $\phi_{\mathcal{N}} = \phi$ is weakly ℓ-independent.*

vi. *ϕ is ℓ-local.*

vii. *$\phi \otimes \psi_{\text{lin}}$ satisfies* $\texttt{UInduce}[\ell]$.

Note that since ℓ'-hypergraphons have rank at most ℓ', a posteriori, we can also strengthen items (ii) and (iii) by replacing existential quantifiers on $p, p_1, \ldots, p_\ell$ with universal ones. Also, since the linear order has rank 1, a posteriori, we can strengthen item (vii) to say that *every* coupling of ϕ with the linear order satisfies $\texttt{UInduce}[\ell]$. In the actual proof of the implication (ii) $\implies$ (i) (that, arguably, is our technically most difficult result), we go in the opposite direction and painstakingly "bootstrap" the premise in (ii) to the unique coupleability with increasingly larger families of objects.

Let us also point out that, given Theorem 4.2.4(ii), one might expect that, in general, if each one of $\psi_1, \ldots, \psi_t$ is uniquely coupleable with a given ϕ, then the same should hold for their independent coupling $\psi_1 \otimes \cdots \otimes \psi_t$; this would immediately give (ii) $\implies$ (iii) in Theorem 4.2.10. However, this question has turned out surprisingly difficult in full generality (see Section 4.9 for a discussion).

The next, more modest, theorem provides properties equivalent to $\texttt{UInduce}[\ell]$.

Theorem 4.2.11 (Characterization of $\texttt{UInduce}$)**.** *The following are equivalent for $\ell \in \mathbb{N}_+$ and $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$.*

i. $\phi \in \texttt{UInduce}[\ell]$.

ii. There exists $p \in (0,1)$ such that ϕ is uniquely inducible by every hypergraphon $\psi \in \operatorname{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$ with $\psi(\rho_\ell) = p$.

iii. ϕ is symmetrically ℓ-local.

The next two theorems completely classify $\texttt{Independence}[k-1]$ and $\texttt{UCouple}[k-1]$ when all arities are at most k. These can be thought of as analogues of full quasirandomness for these families of properties.

Theorem 4.2.12. *Let $k \in \mathbb{N}_+$ and suppose that $k(P) \leq k$ for all $P \in \mathcal{L}$. Let T be a theory over $\mathcal{L}$ and $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$. Then $\phi \in \texttt{Independence}[k-1]$ if and only if there exist $c \in \mathbb{N}_+$, $p \in \Pi_c$ and an open interpretation $I\colon T \rightsquigarrow T_{c,k}$ such that $\phi = \psi_{k,p}^I$.*

Theorem 4.2.13. *Let $k \in \mathbb{N}_+$ and suppose that $k(P) \leq k$ for all $P \in \mathcal{L}$. Let T be a theory over $\mathcal{L}$ and $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$. Then $\phi \in \texttt{UCouple}[k-1]$ if and only if there exist a language $\mathcal{L}'$ whose predicate symbols have arity exactly k, an action $\Theta\colon S_k \times \mathcal{L}' \to \mathcal{L}'$, a Θ-invariant $p = (p_P)_{P\in\mathcal{L}'} \in [0,1]^{\mathcal{L}'}$ with $\sum_{P\in\mathcal{L}'} p_P = 1$ and an open interpretation $I\colon T \rightsquigarrow T_\Theta$ such that $\phi = \psi_{\Theta,p}^I$.*

The next, more modest, theorem classifies $\texttt{UInduce}[1]$ when all arities are at most 2.

Theorem 4.2.14. *Let $\mathcal{L}$ be a language such that $k(P) \leq 2$ for every $P \in \mathcal{L}$ and let T be a theory over $\mathcal{L}$. The following are equivalent for $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$.*

i. $\phi \in \texttt{UInduce}[1]$.

ii. There exist a language $\mathcal{L}'$ whose predicate symbols have arity exactly 2, an action $\Theta\colon S_2 \times \mathcal{L}' \to \mathcal{L}'$, a $p = (p_P)_{P \in \mathcal{L}'} \in [0,1]^{\mathcal{L}'}$ with $\sum_{P \in \mathcal{L}'} p_P = 1$ and a translation *$I\colon T_{\mathcal{L}} \rightsquigarrow T_{\mathcal{L}'}$ from $\mathcal{L}$ to $\mathcal{L}'$ such that $\phi^A = \psi_{\Theta,p}^{A' \circ I}$, where $A'\colon T_{\mathcal{L}'} \rightsquigarrow T_\Theta$ and $A\colon T_{\mathcal{L}} \rightsquigarrow T$ are the axiom-adding interpretations.*

iii. There exist $c \in \mathbb{N}_+$, $p \in \Pi_c$ and an open interpretation $I\colon T \rightsquigarrow T_{c,2} \cup T_{\mathrm{LinOrder}}$ such that $\phi = (\psi_{k,p} \otimes \psi_{\mathrm{lin}})^I$.

4.2.1 Comparison to ad hoc quasirandomness theories

Hypergraphs. The theory of hypergraphons has been most inspirational to our work as it also pertains to quasirandomness of "different strength", arranged in hierarchies like ours. In fact, the last three theorems compare our notions with the hierarchies based on various discrepancy properties from the literature.

As we remarked in the beginning of the chapter, the results of [66] imply that $\texttt{Dev}[k-1] = \texttt{Disc}[\mathcal{A}_{k-1}]$ is the strongest discrepancy property below full quasirandomness and $\texttt{Disc}[\{[\ell+1]\}]$ is the weakest discrepancy property above $\texttt{CliqueDisc}[\ell]$. This together with Theorems 4.2.1, 4.2.2 and 4.2.9 and the three theorems below justify the Hasse diagram of Figure 4.2 between the families `Independence` and `UInduce` and the discrepancy properties in the literature.

The following theorem trivially follows from definitions.

Theorem 4.2.15. *For every $k \geq \ell \geq 1$ and every $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_{k\text{-Hypergraph}}], \mathbb{R})$, if $\phi \in \texttt{UInduce}[\ell]$, then $\phi \in \texttt{CliqueDisc}[\ell]$.*

Theorem 4.2.16. *For every $k \in \mathbb{N}_+$, there exists $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_{k\text{-Hypergraph}}], \mathbb{R})$ satisfying $\texttt{Dev}[k-1]$ but not satisfying $\texttt{UInduce}[1]$.*

Theorem 4.2.17. *For every $k > \ell \geq 1$, there exists $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_{k\text{-Hypergraph}}], \mathbb{R})$ satisfying $\texttt{Independence}[\ell]$ but not satisfying $\texttt{Disc}[\{[\ell+1]\}]$.*

Theorem		Proof location
4.2.1		Section 4.4
4.2.2		Section 4.3
4.2.3		Section 4.3
4.2.4		Section 4.3
4.2.5		Section 4.6
4.2.6		Section 4.6
4.2.7		Section 4.6
4.2.8		Section 4.6
4.2.9		Section 4.6
4.2.10	(i)$\Leftrightarrow$(ii)$\Leftrightarrow$(iii)	Lemma 4.5.7
	(i)$\Leftrightarrow$(iv)$\Leftrightarrow$(v)	Lemma 4.3.2
	(iv) $\Longrightarrow$ (vi)	Lemma 4.3.4
	(vi) $\Longrightarrow$ (vii)	Lemma 4.5.8
	(vii) $\Longrightarrow$ (ii)	Lemma 4.5.9
4.2.11	(i)$\Leftrightarrow$(ii)	Lemma 4.4.1
	(iii) $\Longrightarrow$ (i)	Lemma 4.4.3
	(i) $\Longrightarrow$ (iii)	Lemma 4.4.13
4.2.12		Section 4.7
4.2.13		Section 4.7
4.2.14		Section 4.7
4.2.15		Trivial (see Definitions 4.1.2 and 4.1.4)
4.2.16		Section 4.6
4.2.17		Section 4.6

Table 4.1: Proof locations for theorems of Section 4.2.

Table 4.1 contains pointers to where each of the theorems (or their parts) are proved.

Permutations. In our language, the quasirandom permuton [22, 49] is simply $\psi_{\text{lin}} \otimes \psi_{\text{lin}}$ (see [24, Example 6]). It does not satisfy even the weakest of our properties $\texttt{UInduce}[1]$. This can be easily verified by a direct computation, but a more instructive way would be to apply Theorem 4.2.7 and Theorem 4.2.10(i)≡(vii). Since, on the other hand, $\psi_{\text{lin}} \in \texttt{UInduce}[1]$, we see that the analogue of Theorem 4.2.4 is not true for unique inducibility.

These observations suggest an interesting research direction; we will return to it in Section 4.9.

Words. In our language, quasirandom words defined in [43] are simply $\psi_{\text{lin}} \otimes \psi_p$ ($p \in (0,1)$, $\psi_p \in \text{Hom}^+(\mathcal{A}[T_{2\text{-Coloring}}], \mathbb{R})$). This is clearly generalizable to more colors by considering $\psi_p \in \text{Hom}^+(\mathcal{A}[T_{c\text{-Coloring}}], \mathbb{R})$ ($p \in \Pi_c$), corresponding to quasirandom word sequences over the alphabet $[c]$ with given letter frequencies $(p_1, \ldots, p_c)$. In this way, one can immediately recover existence and uniqueness of the limits of arbitrary (not necessarily quasirandom) convergent sequences from the general theory in [24].

In terms of comparisons, since $\psi_p \notin \texttt{UInduce}[1]$, the same is true for the quasirandom "wordeons" $\psi_{\text{lin}} \otimes \psi_p$.

Latin squares. This is a very interesting example since it is the first time we have encountered an ad hoc theory of limit objects that is provably different from what might be extracted from our framework.

Recall (see e.g. [29]) that there are two major forms of representing a Latin square: as a multiplication table of a quasigroup and as an orthogonal array. As it turns out, they lead to different theories.

The limit theory of Latin squares based on the tabular representation was developed in [40], and the corresponding theory of quasirandomness was continued in [21]. In the language of theons, this theory can be handled only after a fashion, in the same vein as limits of functions on finite vector spaces [24, Sct. 7.5], that is by introducing countably many auxiliary predicate symbols. In this way one immediately gets existence and uniqueness, but other than that the result will be somewhat ugly and not particularly instructive.

The orthogonal array representation opens up another possibility. Recall that in this representation a Latin square is simply an n^2-subset of $[n] \times [n] \times [n]$ such that its projection onto every two coordinates is bijective. Uniformly sampling from this set, we will get a model of $T_{\text{LinOrder}} \cup T_{\text{LinOrder}} \cup T_{\text{LinOrder}}$. Hence a "Borromean" (as in "Borromean rings") view of limits of Latin squares would be simply an element of $\text{Hom}^+(\mathcal{A}[T_{\text{LinOrder}} \cup T_{\text{LinOrder}} \cup T_{\text{LinOrder}}], \mathbb{R})$ such that all three permutons obtained from it by erasing one of the orders are quasirandom.

One obvious example is the quasirandom limit of Latin squares $\psi_{\text{lin}} \otimes \psi_{\text{lin}} \otimes \psi_{\text{lin}}$. But there are others. Indeed, in complete analogy with permutons, limits of Latin squares (in our sense) can be uniquely identified with probability distributions on $[0,1]^3$ such that all three 2-dimensional marginals are uniform. Under this identification, $\psi_{\text{lin}} \otimes \psi_{\text{lin}} \otimes \psi_{\text{lin}}$ corresponds to the uniform probability measure on $[0,1]^3$ and a non-quasirandom example is provided, say, by the uniform probability measure supported on the skewed quasi-random graphon $\mathcal{G}'$ of (3.1) with $p = 1/2$.

Finally, since the quasirandom permuton does not satisfy $\texttt{UInduce}[1]$, it follows that no limit of Latin squares satisfies $\texttt{UInduce}[1]$ as well.

4.3 Basic properties and the first equivalence

In this section we present some initial properties about the notions we have defined, we prove the easiest equivalence in Theorem 4.2.10 between items (i), (iv) and (v) and we prove Theorem 4.2.3 on the naturality of our properties. The first proposition says that only trivial objects can have unique coupleability parameter greater or equal to its rank; this stems from the fact that non-trivial objects are not uniquely coupleable with themselves.

Proposition 4.3.1. *Let $\phi \in \text{Hom}^+(\mathcal{A}[T], \mathbb{R})$ and $r \stackrel{\text{def}}{=} rk(\phi)$.*

i. $r = 0$ if and only if $\phi \in \bigcap_{\ell \in \mathbb{N}} \texttt{UCouple}[\ell]$.

ii. If $r > 0$ *then* $\phi \notin \texttt{UCouple}[r]$.

Proof. Note that $r = 0$ if and only if all peons $\mathcal{N}_P$ are trivial (that is, $\mathcal{N}_P = \varnothing$ or $\mathcal{N}_P = \mathcal{E}_{k(P)}$ a.e.), which in turn is equivalent to having $\phi(\langle K\rangle) \in \{0, 1\}$ for every finite set V and every $K \in \mathcal{K}_V[T]$. This implies that there is a unique $K \in \mathcal{K}_V[T]$ with $\phi(\langle K\rangle) = 1$ and this K must further have full automorphism group $\operatorname{Aut}(K) = S_V$.

Let now $\psi \in \operatorname{Hom}^+(\mathcal{A}[T'], \mathbb{R})$ for some theory T', and assume that ξ is a coupling of ϕ and ψ. Fix a $(T \cup T')$-on $\mathcal{N}$ such that $\xi = \phi_{\mathcal{N}}$. Then for every $K \in \mathcal{K}_V[T \cup T']$ with V finite we have $T_{\text{ind}}(K, \mathcal{N}) = T_{\text{ind}}(I(K), I(\mathcal{N})) \cap T_{\text{ind}}(I'(K), I'(\mathcal{N}))$, where $I\colon T \rightsquigarrow T \cup T'$ and $I'\colon T' \rightsquigarrow T \cup T'$ are the structure-erasing interpretations.

If $r = 0$, we get $\xi(\langle K\rangle) = \phi(\langle I(K)\rangle)\psi(\langle I'(K)\rangle)$ (since ϕ is 0-1 valued) so the forward direction of item (i) follows.

The backward direction of item (i) clearly follows from item (ii), so let us prove the latter by contradiction. Suppose that $\phi \in \texttt{UCouple}[r]$ and fix a T-on $\mathcal{N}$ such that $\phi = \phi_{\mathcal{N}}$ and $\operatorname{rk}(\mathcal{N}) = r$. Consider the $(T \cup T)$-on $\mathcal{H} \stackrel{\text{def}}{=} \mathcal{N} \mathbin{\dot\cup} \mathcal{N}$ in which both copies of each predicate symbol P get mapped to $\mathcal{N}_P$, i.e., $\mathcal{H}$ is the coupling of $\mathcal{N}$ with itself. Since $\operatorname{rk}(\mathcal{H}) = \operatorname{rk}(\mathcal{N}) = r$ and $\phi \in \texttt{UCouple}[r]$, we must have $\phi_{\mathcal{H}} = \phi \otimes \phi$.

Fix a finite set V and $K \in \mathcal{K}_V[T]$ and let $K_2 \in \mathcal{K}_V[T \cup T]$ be given by setting $R_P(K_2) \stackrel{\text{def}}{=} R_P(K)$ for both copies of each predicate symbol P. Then we have

$$\phi(\langle K\rangle) = t_{\text{ind}}(K, \mathcal{N}) = t_{\text{ind}}(K_2, \mathcal{H}) = (\phi \otimes \phi)(\langle K_2\rangle) = \phi(\langle K\rangle)^2,$$

so we must have $\phi(\langle K\rangle) \in \{0, 1\}$. Hence $r = 0$, and item (ii) follows. ∎

We will now use Propositions 3.1.2 and 3.2.1 to show the equivalence in Theorem 4.2.10 between items (i), (iv) and (v).

Lemma 4.3.2 (Theorem 4.2.10(i)⇔(iv)⇔(v))**.** *The following are equivalent for* $\ell \in \mathbb{N}$ *and* $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$.

i. $\phi \in \texttt{UCouple}[\ell]$.

ii. ϕ *is weakly* ℓ*-independent.*

iii. Every T*-on* $\mathcal{N}$ *with* $\phi = \phi_{\mathcal{N}}$ *is weakly* ℓ*-independent.*

Proof. (iii) $\implies$ (ii) is trivial.

(ii) $\implies$ (i).

Let $\mathcal{N}$ be a T-on over some space $\Omega = (X, \mathcal{A}, \mu)$ such that the exchangeable array $\boldsymbol{K}$ corresponding to $\mathcal{N}$ with respect to $\boldsymbol{\theta}$ picked in $\mathcal{E}_{\mathbb{N}_+}(\Omega)$ according to μ is independent from $(\boldsymbol{\theta}_A \mid A \in r(\mathbb{N}_+, \ell))$. Let $\psi \in \operatorname{Hom}^+(\mathcal{A}[T'], \mathbb{R})$ for some theory T' be such that $\operatorname{rk}(\psi) \leq \ell$ and let $\xi \in \operatorname{Hom}^+(\mathcal{A}[T \cup T'], \mathbb{R})$ be any coupling of ϕ and ψ. We have to prove that $\xi = \phi \otimes \psi$.

Let also $I\colon T \rightsquigarrow T \cup T'$ and $I'\colon T' \rightsquigarrow T \cup T'$ be the structure-erasing interpretations. By Proposition 3.2.1, there exists a $(T \cup T')$-on $\mathcal{H}$ over $\Omega \times \Omega$ such that $\xi = \phi_{\mathcal{H}}$ and

$$\mathcal{H}_P = \mathcal{N}_P \times \mathcal{E}_{k(P)}(\Omega) \tag{4.5}$$

for every P in the language of T. By possibly changing zero-measure sets of the peons corresponding to T' using Proposition 3.1.2, we may also assume $\operatorname{rk}(I'(\mathcal{H})) = \operatorname{rk}(\psi) \leq \ell$.

Let us pick $\boldsymbol{\eta}$ in $\mathcal{E}_{\mathbb{N}_+}(\Omega)$ according to μ and independently from $\boldsymbol{\theta}$; we view $(\boldsymbol{\theta}, \boldsymbol{\eta})$ as a $\mathcal{E}_{\mathbb{N}_+}(\Omega \times \Omega)$-valued random variable distributed according to $\mu \otimes \mu$. Let $\boldsymbol{L}$ be the exchangeable array corresponding to $\mathcal{H}$ with respect to $(\boldsymbol{\theta}, \boldsymbol{\eta})$. Note that (4.5) implies that $I(\boldsymbol{L}) = \boldsymbol{K}$, which in turn implies that $I(\boldsymbol{L})$ is independent from $((\boldsymbol{\theta}_A \mid A \in r(\mathbb{N}_+, \ell)), \boldsymbol{\eta})$. On the other hand, since $\operatorname{rk}(I'(\mathcal{H})) \leq \ell$, it follows that $I'(\boldsymbol{L})$ is completely determined by $((\boldsymbol{\theta}_A, \boldsymbol{\eta}_A) \mid A \in r(\mathbb{N}_+, \ell))$, so $I(\boldsymbol{L})$ is independent from $I'(\boldsymbol{L})$. This means that for $m \in \mathbb{N}_+$ and $K \in \mathcal{K}_m[T \cup T']$, we have

$$\begin{aligned}\xi(\langle K \rangle) = \mathbb{P}[\boldsymbol{L}|_{[m]} = K] &= \mathbb{P}[I(\boldsymbol{L})|_{[m]} = I(K) \wedge I'(\boldsymbol{L})|_{[m]} = I'(K)] \\ &= \mathbb{P}[I(\boldsymbol{L})|_{[m]} = I(K)] \cdot \mathbb{P}[I'(\boldsymbol{L})|_{[m]} = I'(K)] = \phi(\langle I(K) \rangle) \cdot \psi(\langle I'(K) \rangle),\end{aligned}$$

so $\xi = \phi \otimes \psi$, hence item (i) follows.

Let us prove (i) $\implies$ (iii). Let $\Omega = (X, \mathcal{A}, \mu)$ be an atomless complete probability space and $\mathcal{N}$ be a T-on over Ω with $\phi = \phi_{\mathcal{N}}$. We have to prove that the exchangeable array $\boldsymbol{K}$ corresponding to $\mathcal{N}$ with respect to $\boldsymbol{\theta}$ picked in $\mathcal{E}_{\mathbb{N}_+}(\Omega)$ according to μ is independent from $(\boldsymbol{\theta}_A \mid A \in r(\mathbb{N}_+, \ell))$. For that, it is sufficient to show that for any $m \in \mathbb{N}$, any $K \in \mathcal{K}_m[T]$ and any measurable set $B \subseteq \mathcal{E}_{m,\ell}(\Omega)$, the events $\boldsymbol{K}|_{[m]} = K$ and $(\boldsymbol{\theta}_A \mid A \in r(m, \ell)) \in B$ are independent.

Let Q be a new m-ary predicate symbol and consider the $(T \cup T_{\{Q\}})$-on $\mathcal{H}$ over Ω given by $\mathcal{H}_P \stackrel{\text{def}}{=} \mathcal{N}_P$ for every P in the language of T and $\mathcal{H}_Q \stackrel{\text{def}}{=} B \times X^{\binom{[m]}{>\ell}}$. Let also $I\colon T \rightsquigarrow T \cup T_{\{Q\}}$ and $I'\colon T_{\{Q\}} \rightsquigarrow T \cup T_{\{Q\}}$ be the structure-erasing interpretations so that $\phi_{\mathcal{H}}$ is a coupling of ϕ and $\phi_{\mathcal{H}}^{I'}$. Since $\operatorname{rk}(\phi_{\mathcal{H}}^{I'}) \leq \operatorname{rk}(\mathcal{H}_Q) \leq \ell$ and $\phi \in \texttt{UCouple}[\ell]$, we have $\phi_{\mathcal{H}} = \phi \otimes \phi_{\mathcal{H}}^{I'}$. Finally, let S be the set of all $L \in \mathcal{K}_m[T \cup T_{\{Q\}}]$ such that $I(L) = K$ and $(1, 2, \ldots, m) \in R_Q(L)$. Then we have

$$
\begin{aligned}
\mathbb{P}[\boldsymbol{K}|_{[m]} = K \wedge (\boldsymbol{\theta}_A \mid A \in r(m, \ell)) \in B] &= \sum_{L \in S} \phi_{\mathcal{H}}(\langle L \rangle) \\
&= \phi(\langle K \rangle) \sum_{L \in S} \phi_{\mathcal{H}}^{I'}(\langle I'(L) \rangle) = \mathbb{P}[\boldsymbol{K}|_{[m]} = K] \cdot \mathbb{P}[(\boldsymbol{\theta}_A \mid A \in r(m, \ell)) \in B],
\end{aligned}
$$

which completes the proof. ■

The alternative characterization of `UCouple` via weak independence gives easy proofs of Theorems 4.2.2 and 4.2.4.

Proof of Theorem 4.2.2. $\texttt{Independence}[\ell] \implies \texttt{UCouple}[\ell]$.

Let $\mathcal{N}$ be an ℓ-independent T-on, and let $\boldsymbol{K}$ be the exchangeable array corresponding to $\mathcal{N}$. Then each $R_P(\boldsymbol{K})$ depends only on the coordinates $\boldsymbol{\theta}_A$ with $|A| > \ell$ (see (2.6)) and hence is independent from $(\boldsymbol{\theta}_A \mid A \in r(A, \ell))$. Therefore, $\mathcal{N}$ is weakly ℓ-independent and $\texttt{Independence}[\ell] \implies \texttt{UCouple}[\ell]$ follows from Lemma 4.3.2.

The implication $\texttt{UCouple}[\ell] \implies \texttt{UInduce}[\ell]$ follows trivially from the definitions. ■

Proof of Theorem 4.2.4. For item (i), if $\mathcal{N}^1$ and $\mathcal{N}^2$ are ℓ-independent theons then $\mathcal{N}^1 \otimes \mathcal{N}^2$

is also ℓ-independent, from which the statement follows.

For item (ii), pick arbitrarily theons $\mathcal{N}^1$ and $\mathcal{N}^2$ such that $\phi_i = \phi_{\mathcal{N}^i}$. Let $(\boldsymbol{\theta}^1, \boldsymbol{\theta}^2)$ be uniformly distributed in $\mathcal{E}_{\mathbb{N}_+} \times \mathcal{E}_{\mathbb{N}_+}$, and let $\boldsymbol{K}$ be the exchangeable array corresponding to $\mathcal{N}^1 \otimes \mathcal{N}^2$ with respect to $(\boldsymbol{\theta}^1, \boldsymbol{\theta}^2)$. Note that for $i \in [2]$ and for the structure-erasing interpretation $I_i\colon T_i \rightsquigarrow T_1 \cup T_2$, the exchangeable array corresponding to $\mathcal{N}^i$ with respect $\boldsymbol{\theta}^i$ is $I_i(\boldsymbol{K})$.

By Lemma 4.3.2, it is sufficient to show that if $I_i(\boldsymbol{K})$ is independent from $(\boldsymbol{\theta^i_A} \mid A \in r(\mathbb{N}_+, \ell))$ for $i \in [2]$, then $\boldsymbol{K}$ is independent from $((\boldsymbol{\theta^1_A}, \boldsymbol{\theta^2_A}) \mid A \in r(\mathbb{N}_+, \ell))$. This immediately follows from the following easily verifiable general fact:

Claim 4.3.3. *Let $\boldsymbol{X_1}, \boldsymbol{X_2}, \boldsymbol{Y_1}, \boldsymbol{Y_2}$ be mutually independent random variables, and let $f_1(X_1, Y_1)$, $f_2(X_2, Y_2)$ be functions such that $f_i(\boldsymbol{X_i}, \boldsymbol{Y_i})$ is independent from $\boldsymbol{X_i}$ $(i = 1, 2)$. Then $(f_1(\boldsymbol{X_1}, \boldsymbol{Y_1}), f_2(\boldsymbol{X_2}, \boldsymbol{Y_2}))$ is independent from $(\boldsymbol{X_1}, \boldsymbol{X_2})$.*

In our context, we set $\boldsymbol{X_i} = (\boldsymbol{\theta^i_A} \mid |A| \leq \ell)$, $\boldsymbol{Y_i} = (\boldsymbol{\theta^i_A} \mid |A| > \ell)$ and let f_i compute the array $I_i(K)$ from (X_i, Y_i) (thus $(f_1(X_1, Y_1), f_2(X_2, Y_2))$ computes the array K from (X_1, X_2, Y_1, Y_2)). ■

Let us now show (almost trivial) implication (iv) $\Longrightarrow$ (vi) of Theorem 4.2.10.

Lemma 4.3.4 (Theorem 4.2.10(iv) $\Longrightarrow$ (vi))**.** *Let $\ell \in \mathbb{N}$. If ϕ is weakly ℓ-independent, then ϕ is ℓ-local.*

Proof. Let $\boldsymbol{K}$ be the exchangeable array corresponding to some theon $\mathcal{N}$ with respect to $\boldsymbol{\theta}$ picked in $\mathcal{E}_{\mathbb{N}_+}(\Omega)$ according to μ such that $\phi = \phi_{\mathcal{N}}$ and suppose $\boldsymbol{K}$ is independent from $(\boldsymbol{\theta}_A \mid A \in r(\mathbb{N}_+, \ell))$. Since for $V \in r(\mathbb{N}_+)$ the marginal $\boldsymbol{K}|_V$ depends only on $(\boldsymbol{\theta}_A \mid A \in r(V))$, the marginals $(\boldsymbol{K}|_{V_i} \mid i \in I)$ are mutually independent as long as the sets V_i have pairwise intersections of size at most ℓ. This follows from the following general observation.

Claim 4.3.5. *Let $\boldsymbol{X}, \boldsymbol{Y_1}, \ldots, \boldsymbol{Y_n}$ be mutually independent random variables and $f_i(X, Y_i)$ be functions such that $(f_1(\boldsymbol{X}, \boldsymbol{Y_1}), \ldots, f_n(\boldsymbol{X}, \boldsymbol{Y_n}))$ is independent of $\boldsymbol{X}$. Then the random variables $f_1(\boldsymbol{X}, \boldsymbol{Y_1}), \ldots, f_n(\boldsymbol{X}, \boldsymbol{Y_n})$ are mutually independent.*

In our situation, $\boldsymbol{X} = (\boldsymbol{\theta}_A \mid A \in r(\mathbb{N}_+, \ell))$, $\boldsymbol{Y_i} = (\boldsymbol{\theta}_A \mid A \in r(V_i) \setminus r(\mathbb{N}_+, \ell))$ and f_i computes the marginal $K|_{V_i}$ from $(\theta_A \mid A \in r(V_i))$.

This completes the proof that ϕ is ℓ-local. ■

Finally, from Propositions 3.2.6 and 3.2.9, we can prove Theorem 4.2.3 about naturality of our properties.

Proof of Theorem 4.2.3. Item (i) follows trivially from the fact that if $\mathcal{N}$ is an ℓ-independent T_2-on with $\phi = \phi_{\mathcal{N}}$, then $I(\mathcal{N})$ is an ℓ-independent T_1-on with $\phi_{I(\mathcal{N})} = \phi^I$.

Item (ii) follows trivially from Proposition 3.2.9.

For item (iii), we let $\psi \in \operatorname{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$ and ξ be a coupling of ϕ^I with ψ. Then by Proposition 3.2.6 there exists a coupling $\widehat{\xi}$ of ϕ and ψ such that $\xi = \widehat{\xi}^{I \cup \mathrm{id}_{T_{\ell\text{-Hypergraph}}}}$. For $i \in [2]$, let $I_i \colon T_i \rightsquigarrow T_i \cup T_{\ell\text{-Hypergraph}}$ and $J_i \colon T_{\ell\text{-Hypergraph}} \rightsquigarrow T_i \cup T_{\ell\text{-Hypergraph}}$ be the structure-erasing interpretation and note that if $M \in \mathcal{M}[T_1 \cup T_{\ell\text{-Hypergraph}}]$ is such that $J_1(M) \cong K_{|M|}^{(\ell)}$, then we have

$$\begin{aligned}
\xi(M) &= \widehat{\xi}^{I \cup \mathrm{id}_{T_{\ell\text{-Hypergraph}}}}(M) \\
&= \widehat{\xi}\left(\sum\left\{M' \in \mathcal{M}_{|M|}[T_2 \cup T_{\ell\text{-Hypergraph}}] \;\middle|\; I(I_2(M')) \cong I_1(M) \land J_2(M') \cong K_{|M|}^{(\ell)}\right\}\right) \\
&= \psi(K_{|M|}^{(\ell)}) \cdot \phi\left(\sum\left\{M' \in \mathcal{M}_{|M|}[T_2] \;\middle|\; I(M') \cong I_1(M)\right\}\right) \\
&= \psi(K_{|M|}^{(\ell)}) \cdot \phi^I(I_1(M)) \\
&= (\phi^I \otimes \psi)(M),
\end{aligned}$$

where the third equality follows from the fact that $\phi \in \mathtt{UInduce}[\ell]$. Hence $\phi^I \in \mathtt{UInduce}[\ell]$. ■

4.4 Unique inducibility

In this section we prove Theorem 4.2.11. We start by showing the equivalence between items (i) and (ii). Curiously, the case $\ell = 1$ is the hardest one to prove.

Lemma 4.4.1 (Theorem 4.2.11(i)⇔(ii))**.** *Let $\ell \in \mathbb{N}_+$ and $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$. Then $\phi \in$ `UInduce`$[\ell]$ if and only if there exists $p \in (0,1)$ such that ϕ is uniquely inducible by every $\psi \in \operatorname{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$ with $\psi(\rho_\ell) = p$.*

Proof. The forward implication is obvious.

For $p \in (0,1)$, let us say that ϕ is uniquely p-inducible if it is uniquely inducible by every $\psi \in \operatorname{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$ with $\psi(\rho_\ell) = p$. Then the backward implication amounts to showing that unique p-inducibility implies unique q-inducibility for every $p, q \in (0,1)$ (the cases $q \in \{0,1\}$ are trivial).

Let $I\colon T \rightsquigarrow T \cup T_{\ell\text{-Hypergraph}}$ and $J\colon T_{\ell\text{-Hypergraph}} \rightsquigarrow T \cup T_{\ell\text{-Hypergraph}}$ be the structure-erasing interpretations. Let us assume that ϕ is uniquely p-inducible and let us show that ϕ is uniquely inducible by any $\psi \in \operatorname{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$ with $\psi(\rho_\ell) = q$. Let ξ be a coupling of ϕ and ψ.

Our objective is to prove that for every $m \in \mathbb{N}$ and every $M \in \mathcal{M}_m[T \cup T_{\ell\text{-Hypergraph}}]$ with $J(M) \cong K_m^{(\ell)}$ we have

$$\xi(M) = \phi(I(M))\psi(K_m^{(\ell)}). \tag{4.6}$$

For $m < \ell$ this is trivial (as $\psi(K_m^{(\ell)}) = 1$), so suppose $m \geq \ell$.

Let $I'\colon T_{\ell\text{-Hypergraph}} \rightsquigarrow T_{\ell\text{-Hypergraph}} \cup T_{2\text{-Coloring}}$ be an open interpretation (to be specified later); note that the diagram

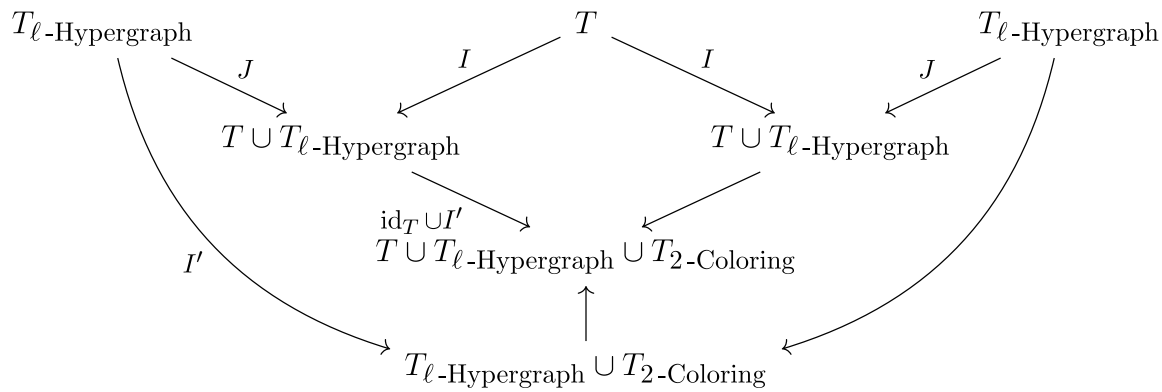

(4.7)

is commutative, where the unlabeled arrows are structure-erasing interpretations. For $t \in [0,1]$ let $\widehat{\xi}_t \stackrel{\text{def}}{=} \xi \otimes \psi_t$ be the independent coupling of ξ and the 2-coloring ψ_t of densities $(t, 1-t)$ (see Definition 4.1.8); note that the fact that (4.7) is commutative implies that $\widehat{\xi}_t^{\mathrm{id}_T \cup I'}$ is a coupling of ϕ and $(\psi \otimes \psi_t)^{I'}$.

We start by showing (4.6) in the case $p \leq q$. In this case, we take

$$I'(E)(x_1, \ldots, x_\ell) \stackrel{\text{def}}{=} E(x_1, \ldots, x_\ell) \wedge \bigwedge_{i \in [\ell]} \chi_1(x_i),$$

that is, I' keeps edges that are monochromatic in color 1. Let $t \stackrel{\text{def}}{=} (p/q)^{1/\ell}$ and note that for $n \geq \ell$ we have

$$(\psi \otimes \psi_t)^{I'}(K_n^{(\ell)}) = \psi(K_n^{(\ell)}) t^n = \psi(K_n^{(\ell)}) \left(\frac{p}{q}\right)^{n/\ell},$$

which in particular implies that $(\psi \otimes \psi_t)^{I'}(\rho_\ell) = p$. On the other hand, we also have $\widehat{\xi}_t^{\mathrm{id}_T \cup I'}(M) = \xi(M) t^m$, so unique p-inducibility of ϕ gives

$$\xi(M) t^m = \widehat{\xi}_t^{\mathrm{id}_T \cup I'}(M) = \phi(I(M))(\psi \otimes \psi_t)^{I'}(K_m^{(\ell)}) = \phi(I(M)) \psi(K_m^{(\ell)}) t^m,$$

from which (4.6) follows.

We now show (4.6) in the case $\ell \geq 2$ and $q < p$. In this case, we let

$$I'(E)(x_1, \ldots, x_\ell) \stackrel{\text{def}}{=} \left(E(x_1, \ldots, x_\ell) \land \bigwedge_{i \in [\ell]} \chi_1(x_i)\right) \lor \bigwedge_{i \in [\ell]} \chi_2(x_i)$$

that is, I' declares edges to be either old edges that are monochromatic in color 1 or any ℓ-set that is monochromatic in color 2. Let $f(x) \stackrel{\text{def}}{=} x^\ell q + (1-x)^\ell$ and note that $f(0) = 1$ and $f(1) = q$, so there exists $t \in (0,1)$ such that $f(t) = p$. Since $\ell \geq 2$, for $n \geq \ell$, we have

$$(\psi \otimes \psi_t)^{I'}(K_n^{(\ell)}) = \psi(K_n^{(\ell)})t^n + (1-t)^n,$$

which in particular implies that $(\psi \otimes \psi_t)^{I'}(\rho_\ell) = f(t) = p$. On the other hand, we also have $\widehat{\xi}_t^{\mathrm{id}_T \cup I'}(M) = \xi(M)t^m + \phi(I(M))(1-t)^m$, so unique p-inducibility of ϕ gives

$$\begin{aligned}\xi(M)t^m + \phi(I(M))(1-t)^m = \widehat{\xi}_t^{\mathrm{id}_T \cup I'}(M) &= \phi(I(M))(\psi \otimes \psi_t)^{I'}(K_m^{(\ell)}) \\ &= \phi(I(M))(\psi(K_m^{(\ell)})t^m + (1-t)^m),\end{aligned}$$

from which (4.6) follows.

The case $q < p$ and $\ell = 1$ is more complicated as the construction analogous to the above does not work: cliques in arity 1 need not be monochromatic.

Let us prove first the sub-case $q = p^2$. The idea, roughly speaking, is that when $\ell = 1$, unique p-inducibility says that any "subset of vertices" of relative size p in ϕ induces ϕ and since a "subset of vertices" of relative size p^2 can be seen as having relative size p in some "subset of vertices" that itself has relative size p in the whole space, it must also induce ϕ.

It is worth noting that this idea can be implemented almost literally in the geometric language. But that would require working with theons that have different ground sets in different coordinates so we prefer to present a syntactic argument instead, similar to the one above.

We work with the theory $T_{2\text{-Coloring}}$ instead of $T_{1\text{-Hypergraph}}$ (see Remark 8). Let ξ be a coupling of ϕ and $\psi \stackrel{\text{def}}{=} \psi_{p^2} \in \operatorname{Hom}^+(\mathcal{A}[T_{2\text{-Coloring}}], \mathbb{R})$; we want to show that for every $M \in \mathcal{M}[T \cup T_{2\text{-Coloring}}]$ with $R_{\chi_1}(M) = V(M)$, we have

$$\xi(M) = \phi(I(M))p^{2m},$$

where $m \stackrel{\text{def}}{=} |M|$ and $I\colon T \rightsquigarrow T \cup T_{2\text{-Coloring}}$ is the structure-erasing interpretation.

Let $I_1, I_2\colon T_{2\text{-Coloring}} \rightsquigarrow T_{3\text{-Coloring}}$ be the interpretations given by

$$I_1(\chi_1)(x) \stackrel{\text{def}}{=} \chi_1(x) \lor \chi_2(x); \qquad I_2(\chi_1)(x) \stackrel{\text{def}}{=} \chi_1(x);$$
$$I_1(\chi_2)(x) \stackrel{\text{def}}{=} \chi_3(x); \qquad I_2(\chi_2)(x) \stackrel{\text{def}}{=} \chi_2(x) \lor \chi_3(x).$$

Let $\widehat{\psi} \stackrel{\text{def}}{=} \psi_{(p^2, p-p^2, 1-p)} \in \operatorname{Hom}^+(\mathcal{A}[T_{3\text{-Coloring}}], \mathbb{R})$ and note that $\widehat{\psi}^{I_i} = \psi_{p^i}$ for $i \in [2]$.

Let $J\colon T_{2\text{-Coloring}} \rightsquigarrow T \cup T_{2\text{-Coloring}}$ and $\widehat{J}\colon T_{3\text{-Coloring}} \rightsquigarrow T \cup T_{3\text{-Coloring}}$ be the structure-erasing interpretations. Our definitions ensure that the following diagram is commutative.

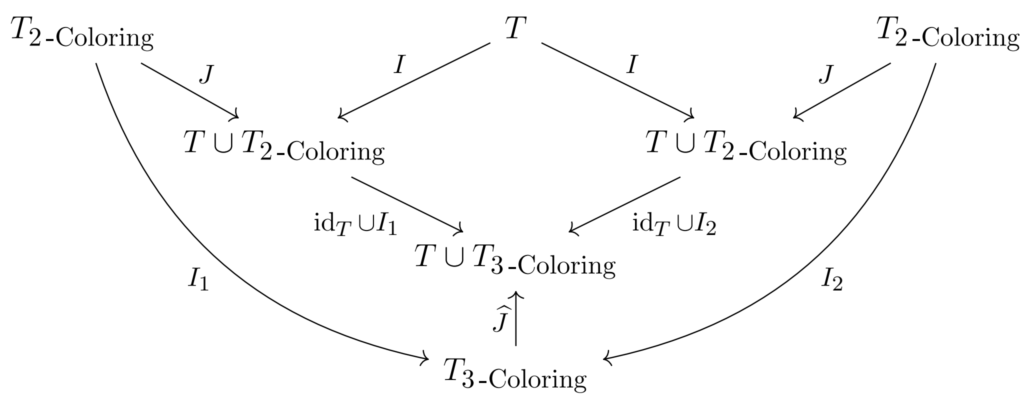

(4.8)

For every $n \in \mathbb{N}$, let $C_n \in \mathcal{M}_n[T_{2\text{-Coloring}}]$ be the unique model with all vertices satisfying χ_1.

Since $\widehat{\psi}^{I_2} = \psi$, by Proposition 3.2.6, there exists a coupling $\widehat{\xi}$ of ϕ and $\widehat{\psi}$ such that $\widehat{\xi}^{\mathrm{id}_T \cup I_2} = \xi$. We now make use of the operator $\pi^{(\neg\chi_3, \mathrm{id}_T \cup I_2)}\colon \mathcal{A}[T \cup T_{2\text{-Coloring}}] \to \mathcal{A}_u[T \cup T_{3\text{-Coloring}}]$ [59, Definition 4], where $u = \sum\{N \in \mathcal{M}_1[T \cup T_{3\text{-Coloring}}] \mid I_1(\widehat{J}(N)) \cong C_1\}$

and $\mathcal{A}_u[T \cup T_{3\text{-Coloring}}]$ is the localization by the multiplicative system $\{u, u^2, \ldots, u^n, \ldots\}$. Intuitively, it corresponds to applying the interpretation $\mathrm{id}_T \cup I_2$, followed by throwing away vertices of color 3. (All densities have to be re-normalized by a power of u, this is why we need to localize.) Since

$$\widehat{\xi}(u) = \widehat{\xi}^{\widehat{J} \circ I_1}(C_1) = \widehat{\psi}^{I_1}(C_1) = p > 0, \tag{4.9}$$

we can apply [59, Theorem 2.6] and form the element $\zeta \stackrel{\text{def}}{=} \widehat{\xi} \circ \pi^{(\neg\chi_3, \mathrm{id}_T \cup I_2)} \in \mathrm{Hom}^+(\mathcal{A}[T \cup T_{2\text{-Coloring}}], \mathbb{R})$. We claim that $\zeta^I = \phi$.

To see this, note that for $N \in \mathcal{M}[T]$, we have

$$\zeta^I(N) = \frac{\sum_{N'} \widehat{\xi}(N')}{\widehat{\xi}(u)^{|N|}},$$

where the sum is over all $N' \in \mathcal{M}_{|N|}[T \cup T_{3\text{-Coloring}}]$ such that $I((\mathrm{id}_T \cup I_2)(N')) \cong N$ and $J((\mathrm{id}_T \cup I_1)(N')) \cong C_{|N|}$. But since (4.8) is commutative, the condition $I((\mathrm{id}_T \cup I_2)(N')) \cong N$ is equivalent to $I((\mathrm{id}_T \cup I_1)(N')) \cong N$, which together with (4.9) gives

$$\zeta^I(N) = \frac{\widehat{\xi}^{\mathrm{id}_T \cup I_1}(\widehat{N})}{p^{|N|}},$$

where $\widehat{N} \in \mathcal{M}_{|N|}[T \cup T_{2\text{-Coloring}}]$ is the unique model such that $I(\widehat{N}) \cong N$ and $J(\widehat{N}) \cong C_{|N|}$.

Since $\widehat{\xi}^{(\mathrm{id}_T \cup I_1) \circ J}(C_1) = \widehat{\psi}^{I_1}(C_1) = p$ and $\widehat{\xi}^{(\mathrm{id}_T \cup I_1) \circ I} = \xi^I = \phi$, unique p-inducibility of ϕ implies that $\widehat{\xi}^{\mathrm{id}_T \cup I_1}(\widehat{N}) = p^{|N|}\phi(N)$ and thus $\zeta^I = \phi$.

Now we claim that $\zeta^J = \psi_p$. Indeed, note that

$$\begin{aligned}\zeta^J(C_1) &= \frac{\sum\{\widehat{\xi}(N) \mid N \in \mathcal{M}_1[T \cup T_{3\text{-Coloring}}] \wedge J((\mathrm{id}_T \cup I_2)(N)) \cong J((\mathrm{id}_T \cup I_1)(N)) \cong C_1\}}{\widehat{\xi}(u)} \\ &= \frac{\widehat{\xi}^{\widehat{J}}(\widehat{C}_1)}{p} = \frac{\widehat{\psi}(\widehat{C}_1)}{p} = p,\end{aligned}$$

where $\widehat{C}_1 \in \mathcal{M}_1[T_{3\text{-Coloring}}]$ is the model whose unique vertex satisfies χ_1, hence $\zeta^J = \psi_p$.

This means that ζ is a coupling of ϕ and ψ_p, so for our fixed $M \in \mathcal{M}_m[T \cup T_{2\text{-Coloring}}]$ with $R_{\chi_1}(M) = V(M)$, unique p-inducibility of ϕ gives

$$\begin{aligned}\xi(M) &= \widehat{\xi}^{\mathrm{id}_T \cup I_2}(M) = \widehat{\xi}(\pi^{(\neg\chi_3, \mathrm{id}_T \cup I_2)}(M)) \cdot \widehat{\xi}(u)^m \\ &= \zeta(M) \cdot p^m = \phi(I(M)) \cdot p^{2m},\end{aligned}$$

as desired.

From the case $\ell = 1$ and $q = p^2 < p$, with a simple induction, we can derive the case when $\ell = 1$ and $q = p^{2^k} < p$ for some $k \in \mathbb{N}_+$.

Finally, for the case $\ell = 1$ and arbitrary $q < p$, we let $k \in \mathbb{N}_+$ be large enough so that $p^{2^k} < q$ and putting together the previous cases gives that unique p-inducibility implies unique p^{2^k}-inducibility, which in turn implies unique q-inducibility. ■

The rest of this section is devoted to various relations between the unique inducibility and the clique discrepancy for hypergraphons; we will also use our findings to prove the last remaining equivalence (i)⇔(iii) in Theorem 4.2.11.

It was proved in [66, 1] that for $\ell < k$, `CliqueDisc`$[\ell]$ is equivalent to the non-induced labeled density of every ℓ-linear hypergraph H (i.e., hypergraphs whose edges have pairwise intersections of size at most ℓ) being $p^{e(H)}$. We restate below this result in the language of exchangeable arrays.

Theorem 4.4.2 ([66, 1])**.** *Let $\ell \in [k-1]$, let $\phi \in \mathrm{Hom}^+(\mathcal{A}[T_{k\text{-Hypergraph}}], \mathbb{R})$ and let $\boldsymbol{K}$ be the corresponding exchangeable array. Then $\phi \in$ `CliqueDisc`$[\ell]$ if and only if for every finite collection $(V_i)_{i \in I}$ of finite subsets of $\mathbb{N}_+$ of size k each and with pairwise intersections of size at most ℓ we have*

$$\mathbb{P}[\forall i \in I, \boldsymbol{K}|_{V_i} \cong \rho_k] = \prod_{i \in I} \mathbb{P}[\boldsymbol{K}|_{V_i} \cong \rho_k].$$

Even though this theorem only makes sense in the theory of hypergraphs, we can derive the implication (iii) $\implies$ (i) of Theorem 4.2.11 for general theories from it.

Lemma 4.4.3 (Theorem 4.2.11(iii) $\implies$ (i))**.** *If* $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ *is symmetrically* ℓ*-local, then* $\phi \in$ `UInduce`$[\ell]$.

Proof. Let $I\colon T \rightsquigarrow T \cup T_{\ell\text{-Hypergraph}}$ and $J\colon T_{\ell\text{-Hypergraph}} \rightsquigarrow T \cup T_{\ell\text{-Hypergraph}}$ be the structure-erasing interpretations.

Our objective is to show that for every $\psi \in \operatorname{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$, every coupling ξ of ϕ and ψ, every $m \in \mathbb{N}$ and every $M \in \mathcal{M}_m[T \cup T_{\ell\text{-Hypergraph}}]$ with $J(M) \cong K_m^{(\ell)}$, we have

$$\xi(M) = \phi(I(M))\psi(K_m^{(\ell)}). \tag{4.10}$$

Let us first consider the case $m \leq \ell$. In this case, note that for the exchangeable array $\boldsymbol{K}$ corresponding to ϕ, by letting $V_1 = V_2 = [m]$, symmetric ℓ-locality of ϕ gives

$$\phi(I(M)) = \mathbb{P}[\boldsymbol{K}|_{[m]} \cong I(M)] = \mathbb{P}[\boldsymbol{K}|_{[m]} \cong I(M)]^2 = \phi(I(M))^2,$$

so $\phi(I(M)) \in \{0,1\}$, hence (4.10) follows.

Suppose now that $m > \ell$ and let $I'\colon T_{m\text{-Hypergraph}} \rightsquigarrow T$ be the open interpretation that declares m-edges to be *isomorphic* copies of $I(M)$, that is, it is given by

$$I'(E)(x_1, \ldots, x_m) \stackrel{\text{def}}{=} \bigvee_{\sigma \in S_m} D_{\text{open}}(I(M))(x_{\sigma(1)}, \ldots, x_{\sigma(m)}).$$

Let us show that $\phi^{I'} \in \operatorname{Hom}^+(\mathcal{A}[T_{m\text{-Hypergraph}}], \mathbb{R})$ satisfies `CliqueDisc`$[\ell]$. Let $\boldsymbol{K}$ be the exchangeable array corresponding to ϕ so that $I'(\boldsymbol{K})$ is the exchangeable array corresponding to $\phi^{I'}$. Then if $(V_i)_{i \in [t]}$ is a finite collection of finite subsets of $\mathbb{N}_+$ of size m each and with pairwise intersections of size at most ℓ, then

$$\begin{aligned}\mathbb{P}[\forall i \in [t], I'(\boldsymbol{K})|_{V_i} \cong \rho_m] &= \mathbb{P}[\forall i \in [t], \boldsymbol{K}|_{V_i} \cong M] \\ &= \prod_{i\in[t]} \mathbb{P}[\boldsymbol{K}|_{V_i} \cong M] = \prod_{i\in[t]} \mathbb{P}[I'(\boldsymbol{K})|_{V_i} \cong \rho_m],\end{aligned}$$

where the second equality follows from the fact that ϕ is symmetrically ℓ-local. By Theorem 4.4.2, it follows that $\phi^{I'}$ satisfies `CliqueDisc`$[\ell]$.

Note now that the diagram

$$\begin{array}{ccccc} T_m\text{-Hypergraph} & \longrightarrow & T_m\text{-Hypergraph} \cup T_\ell\text{-Hypergraph} & \longleftarrow & T_\ell\text{-Hypergraph} \\ {\scriptstyle I'}\downarrow & & \downarrow{\scriptstyle I' \cup \mathrm{id}_{T_\ell\text{-Hypergraph}}} & \swarrow{\scriptstyle J} & \\ T & \xrightarrow[I]{} & T \cup T_\ell\text{-Hypergraph} & & \end{array}$$

is commutative, where the unlabeled arrows are structure-erasing interpretations. This implies that $\xi^{I' \cup \mathrm{id}_{T_\ell\text{-Hypergraph}}}$ is a coupling of $\phi^{I'}$ and ψ, so we get

$$\xi(M) = \xi^{I' \cup \mathrm{id}_{T_\ell\text{-Hypergraph}}}(K_m^{(m,\ell)}) = \phi^{I'}(\rho_m)\psi(K_m^{(\ell)}) = \phi(I(M))\psi(K_m^{(\ell)}),$$

where the second equality follows from $\phi^{I'} \in$ `CliqueDisc`$[\ell]$. ■

Let us now prove an important fact about `CliqueDisc`$[\ell]$ and ℓ-flattenings defined below.

Definition 4.4.4. For a peon $\mathcal{N}$ over $\Omega = (X, \mathcal{A}, \mu)$ and $\ell \in \mathbb{N}$, the *ℓ-flattening* of $\mathcal{N}$ is the function $W_{\mathcal{N}}^{\ell} \colon \mathcal{E}_{k,\ell}(\Omega) \to [0, 1]$ defined by

$$W_{\mathcal{N}}^{\ell}(x) \stackrel{\text{def}}{=} \mu(\{y \in X^{\binom{[k]}{>\ell}} \mid (x, y) \in \mathcal{N}\}),$$

and defined arbitrarily when the set above is not measurable.

Note that the constructions in (3.2) and (3.13) are precisely r-flattenings, and so is the construction of a graphon in the ordinary sense from T_{Graph}-on (cf. (2.5) and (3.1)).

Lemma 4.4.5. *Let $\mathcal{N}$ be a $T_{k\text{-Hypergraph}}$-on over $\Omega = (X, \mathcal{A}, \mu)$ such that $\phi_{\mathcal{N}}$ satisfies* `CliqueDisc`$[\ell]$. *Then $W^\ell_{\mathcal{N}} = \phi_{\mathcal{N}}(\rho_k)$ a.e.*

Proof. It is sufficient to prove that the two measures on $X^{r(k,\ell)}$ given by $Y \mapsto \int_Y W^\ell_{\mathcal{N}}\, d\mu$ and $\nu(Y) \stackrel{\text{def}}{=} \phi_{\mathcal{N}}(\rho_k)\mu(Y)$ coincide, and for that we only have to consider the basis of our σ-algebra, i.e., sets of the form

$$Y = \prod_{A \in r(k,\ell)} V_A.$$

In other words, for every collection $V_A \subseteq X$ $(A \in r(k,\ell))$ of measurable sets we have to prove that

$$\int_Y W^\ell_{\mathcal{N}}\, d\mu = \phi_{\mathcal{N}}(\rho_k) \cdot \mu(Y). \tag{4.11}$$

Recall from [66, 1] that `CliqueDisc`$[\ell]$ is equivalent to `Disc`$[\binom{[k]}{\ell}]$ (see Definition 4.1.4) and for the language $\mathcal{L}_{\binom{[k]}{\ell}}$ containing one predicate symbol P_A of arity ℓ for each $A \in \binom{[k]}{\ell}$, define the $T_{\mathcal{L}_{\binom{[k]}{\ell}}} \cup T_{k\text{-Hypergraph}}$-on $\mathcal{H}$ over Ω by

$$\mathcal{H}_E \stackrel{\text{def}}{=} \mathcal{N}_E; \qquad \mathcal{H}_{P_A} \stackrel{\text{def}}{=} \iota^*_A(Y) = \{x \in \mathcal{E}_\ell(\Omega) \mid \forall A' \in r(A), x_{\iota_A^{-1}(A')} \in V_{A'}\}.$$

Let then $\boldsymbol{K}$ be the exchangeable array corresponding to $\mathcal{H}$. Since $\phi_{\mathcal{N}}$ satisfies `CliqueDisc`$[\ell] =$ `Disc`$[\binom{[k]}{\ell}]$, we get

$$\begin{aligned}
\int_Y W^\ell_{\mathcal{N}}\, d\mu &= \mathbb{P}\left[(1,\ldots,k) \in R_E(\boldsymbol{K}) \land \forall A \in \binom{[k]}{\ell}, \iota_A \in R_{P_A}(\boldsymbol{K})\right] \\
&= \phi_{\mathcal{N}}(\rho_k) \cdot \mathbb{P}\left[\forall A \in \binom{[k]}{\ell}, \iota_A \in R_{P_A}(\boldsymbol{K})\right] \\
&= \phi_{\mathcal{N}}(\rho_k) \cdot \mu(Y),
\end{aligned}$$

as desired. ■

To prove the final implication (i) $\implies$ (iii) in Theorem 4.2.11, we will need a small generalization of the easier direction of Theorem 4.4.2 for disjoint unions of theories of hypergraphs.

Definition 4.4.6 ($\vec{k}$-hypergraphs)**.** Given $\vec{k} = (k_1, \ldots, k_t) \in \mathbb{N}_+^t$, we let $T_{\vec{k}\text{-Hypergraph}} \stackrel{\text{def}}{=} \bigcup_{i\in[t]} T_{k_i\text{-Hypergraph}}$ and in this theory, we denote the predicate symbol corresponding to the i-th hypergraph by E_i. Models of $T_{\vec{k}\text{-Hypergraph}}$ will be called *$\vec{k}$-hypergraphs* and for one such model M, we let $E_i(M) \stackrel{\text{def}}{=} \{\operatorname{im}(\alpha) \mid \alpha \in R_{E_i}(M)\}$ be its *i-th edge set*. We also denote by $I_i\colon T_{k_i\text{-Hypergraph}} \rightsquigarrow T_{\vec{k}\text{-Hypergraph}}$ the structure-erasing interpretation corresponding to the i-th edge set.

Proposition 4.4.7. *Let $\vec{k} = (k_1, \ldots, k_t)$, let $\ell \leq \min_{i\in[t]} k_i$, let $i_1, \ldots, i_s \in [t]$ and let $(V_j)_{j=1}^s$ be such that $V_j \in \binom{\mathbb{N}_+}{k_{i_j}}$ and $|V_j \cap V_{j'}| \leq \ell$, whenever $j \neq j'$.*

Let $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_{\vec{k}\text{-Hypergraph}}], \mathbb{R})$ be such that all ϕ^{I_i} ($i \in [t]$) satisfy `CliqueDisc`$[\ell]$ and let $\boldsymbol{K}$ be the corresponding exchangeable array. Then

$$\mathbb{P}[\forall j \in [s], V_j \in E_{i_j}(\boldsymbol{K})] = \prod_{j\in[s]} \mathbb{P}[V_j \in E_{i_j}(\boldsymbol{K})].$$

Proof. Let $\mathcal{N}$ be a $T_{k\text{-Hypergraph}}$-on such that $\phi_{\mathcal{N}} = \phi$ and note that

$$\begin{aligned}\mathbb{P}[\forall j \in [s], V_j \in E_{i_j}(\boldsymbol{K})] &= \lambda\left(\bigcap_{j\in[s]} (\alpha_j^*)^{-1}(\mathcal{N}_{E_{i_j}})\right)\\ &= \lambda(\{x \in \mathcal{E}_{\mathbb{N}_+} \mid \forall j \in [s], \alpha_j^*(x) \in \mathcal{N}_{E_{i_j}}\}),\end{aligned}$$

where $\alpha_j \in (\mathbb{N}_+)_{k_{i_j}}$ is such that $\operatorname{im}(\alpha_j) = V_j$. Since the sets V_j have pairwise intersections of size at most ℓ, in the set above, the coordinates x_A with $|A| > \ell$ are only constrained by at most one of the α_j^*, so Fubini's Theorem gives

$$\mathbb{P}[\forall j \in [s], V_j \in E_{i_j}(\boldsymbol{K})] = \int_{\mathcal{E}_{V,\ell}} \prod_{j\in[s]} W^{\ell}_{\mathcal{N}_{E_{i_j}}}(\alpha_j^*(x))\, d\lambda(x),$$

where $V \stackrel{\text{def}}{=} \bigcup_{j\in[s]} V_j$.

Since each ϕ^{I_i} satisfies `CliqueDisc`$[\ell]$, by Lemma 4.4.5, it follows that $W^{\ell}_{\mathcal{N}_{E_i}} = \phi^{I_i}(\rho_{k_i})$ a.e., so we get

$$\mathbb{P}[\forall j \in [s], V_j \in E_{i_j}(\boldsymbol{K})] = \prod_{j\in[s]} \phi^{I_{i_j}}(\rho_{k_{i_j}}) = \prod_{j\in[s]} \mathbb{P}[V_j \in E_{i_j}(\boldsymbol{K})],$$

as desired. ■

Proposition 4.4.7 (and Theorem 4.2.15) will be sufficient to handle the case in the definition of symmetric ℓ-locality when all sets have size at least ℓ. For smaller sets, we need the notion of categoricity of elements of $\operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ defined below.

Definition 4.4.8. Recall from Definition 3.6.1 that for $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$, $\operatorname{Th}(\phi)$ is the theory of all models that have positive density in ϕ. Recall also that in model theory a theory T is called *ℓ-categorical* if it has exactly one model of size ℓ up to isomorphism. We say that $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is *ℓ-categorical* if $\operatorname{Th}(\phi)$ is ℓ-categorical.

Remark 9. Since $\sum_{M\in\mathcal{M}_\ell[T]} \phi(M) = 1$, it follows that ϕ is ℓ-categorical if and only if $\phi(M) \in \{0, 1\}$ for every $M \in \mathcal{M}_\ell[T]$.

Lemma 4.4.9. *Let $I : T_1 \rightsquigarrow T_2$ be an open interpretation and let $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$ be ℓ-categorical. Then ϕ^I is ℓ-categorical.*

Proof. Since for $M \in \mathcal{M}_\ell[T_1]$, we have $\phi^I(M) = \sum\{\phi(N) \mid N \in \mathcal{M}_\ell[T_2] \land I(N) \cong M\}$, it follows that $\phi^I(M) > 0$ if and only if $M \cong I(N_0)$ for the unique model $N_0 \in \mathcal{M}_\ell[\operatorname{Th}(\phi)]$. ■

Lemma 4.4.10. *If $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_{k\text{-Hypergraph}}], \mathbb{R})$ is ℓ-categorical for $\ell \geq k$ then $\phi(\rho_k) \in \{0, 1\}$, that is, the hypergraphon ϕ is either empty or complete.*

Proof. Let M be the unique k-hypergraph on ℓ vertices such that $\phi(M) = 1$. Then $M \in \{K^{(k)}_\ell, \overline{K}^{(k)}_\ell\}$ as $\phi(K^{(k)}_\ell) = \phi(\overline{K}^{(k)}_\ell) = 0$ would have contradicted Ramsey's Theorem. The lemma follows. ■

Lemma 4.4.11. *If $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is ℓ-categorical and $0 \leq \ell' \leq \ell$, then ϕ is ℓ'-categorical.*

Proof. Let $M \in \mathcal{M}_{\ell'}[T]$ and consider the open interpretation $I\colon T_{\ell'\text{-Hypergraph}} \rightsquigarrow T$ that declares m-edges to be isomorphic copies of M. By Lemma 4.4.9, it follows that ϕ^I is ℓ-categorical, and it follows from Lemma 4.4.10 that ϕ^I is either the empty or the complete hypergraphon. Now, ϕ is ℓ'-categorical by Remark 9. ■

Lemma 4.4.12. *If $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ satisfies $\texttt{UInduce}[\ell]$, then ϕ is ℓ'-categorical for every $0 \leq \ell' \leq \ell$.*

Proof. By Lemma 4.4.11, it is enough to show the case $\ell' = \ell$. Let $I\colon T \rightsquigarrow T \cup T_{\ell\text{-Hypergraph}}$ and $J\colon T_{\ell\text{-Hypergraph}} \rightsquigarrow T \cup T_{\ell\text{-Hypergraph}}$ be the structure-erasing interpretations. Let $\mathcal{N}$ be a T-on such that $\phi_{\mathcal{N}} = \phi$ and for $M \in \mathcal{M}_\ell[T]$, let $\mathcal{H}$ be the $T \cup T_{\ell\text{-Hypergraph}}$-on given by

$$\mathcal{H}_P \stackrel{\text{def}}{=} \mathcal{N}_P; \qquad \mathcal{H}_E \stackrel{\text{def}}{=} \bigcup_{\substack{K \in \mathcal{K}_\ell[T] \\ K \cong M}} T_{\text{ind}}(K, \mathcal{N})$$

for every predicate symbol P in the language of T.

Let $\widehat{M} \in \mathcal{M}_\ell[T \cup T_{\ell\text{-Hypergraph}}]$ be such that $I(\widehat{M}) \cong M$ and $J(\widehat{M}) \cong \rho_\ell$. Then

$$\phi(M) = \phi_{\mathcal{H}}(\widehat{M}) = \phi(M)\phi_{\mathcal{H}}^J(\rho_\ell) = \phi(M)^2,$$

where the second equality follows since $\phi \in \texttt{UInduce}[\ell]$. Hence $\phi(M) \in \{0, 1\}$ for every $M \in \mathcal{M}_\ell[T]$, so ϕ is ℓ-categorical by Remark 9. ■

Remark 10. The converse to Lemma 4.4.12 is very far from being true. For example, every graphon is 1-categorical, and, slightly less trivially, every tournamon is 2-categorical. They are seldom uniquely 1-inducible (see Theorem 4.2.14).

We can finally prove the last implication of Theorem 4.2.11.

Lemma 4.4.13 (Theorem 4.2.11(i) $\implies$ (iii))**.** *If $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ satisfies* $\texttt{UInduce}[\ell]$*, then ϕ is symmetrically ℓ-local.*

Proof. Let $\boldsymbol{K}$ be the exchangeable array corresponding to ϕ. We need to show that for every finite collection $(V_i)_{i\in[t]}$ of finite subsets of $\mathbb{N}_+$ with pairwise intersections of size at most ℓ and every collection $(M_i)_{i\in[t]}$ of models of T, we have

$$\mathbb{P}[\forall i \in [t], \boldsymbol{K}|_{V_i} \cong M_i] = \prod_{i\in[t]} \mathbb{P}[\boldsymbol{K}|_{V_i} \cong M_i].$$

By Lemma 4.4.12, we know that ϕ is ℓ'-categorical for every $0 \leq \ell' \leq \ell$, which implies that if $|V| \leq \ell$, then $\mathbb{P}[\boldsymbol{K}|_V \cong M] = \phi(M) \in \{0,1\}$, i.e., the event $\boldsymbol{K}|_V \cong M$ is trivial. So we may assume that $|V_i| > \ell$ for every $i \in [t]$.

Let $\vec{k} = (k_1, \ldots, k_t)$ be given by $k_i \stackrel{\text{def}}{=} |V_i|$ $(i \in [t])$ and consider the interpretation $I\colon T_{\vec{k}\text{-Hypergraph}} \rightsquigarrow T$ that declares E_i-edges to be isomorphic copies of M_i. In other words, I is given by

$$I(E_i)(x_1, \ldots, x_{k_i}) \stackrel{\text{def}}{=} \bigvee_{\sigma\in S_{k_i}} D_{\text{open}}(M_i)(x_{\sigma(1)}, \ldots, x_{\sigma(k_i)}).$$

By Theorem 4.2.3, we know that for every $i \in [t]$ we have $\phi^{I\circ I_i} \in \texttt{UInduce}[\ell]$ and by Theorem 4.2.15, it follows that $\phi^{I\circ I_i} \in \texttt{CliqueDisc}[\ell]$. Then we have

$$\begin{aligned}\mathbb{P}[\forall i \in [t], \boldsymbol{K}|_{V_i} \cong M_i] &= \mathbb{P}[\forall i \in [t], V_i \in E_i(I(\boldsymbol{K}))]\\ &= \prod_{i\in[t]} \mathbb{P}[V_i \in E_i(I(\boldsymbol{K}))] = \prod_{i\in[t]} \mathbb{P}[\boldsymbol{K}|_{V_i} \cong M_i],\end{aligned}$$

where the second equality follows from Proposition 4.4.7. ■

We finish this section with the (now trivial) proof of Theorem 4.2.1.

Proof of Theorem 4.2.1. The implications $\texttt{Independence}[\ell] \implies \texttt{Independence}[\ell-1]$ and

$\texttt{UCouple}[\ell] \implies \texttt{UCouple}[\ell-1]$ follow easily from definitions. The fact that $\texttt{UInduce}[\ell] \implies \texttt{UInduce}[\ell-1]$ follows since symmetric ℓ-locality trivially implies symmetric $(\ell-1)$-locality and from Lemmas 4.4.3 and 4.4.13. ■

4.5 Unique coupleability

In this section we prove Theorem 4.2.10. We start with the equivalence (i)⇔(ii)⇔(iii). While implications (i) $\implies$ (iii) and (iii) $\implies$ (ii) are fairly straightforward, the proof of the implication (ii) $\implies$ (i) is more involved and naturally splits into five rather independent parts:

1. Show that unique coupleability of ϕ with the quasirandom ℓ'-hypergraphon $\psi_{\ell',p}$ for some $p \in (0,1)$ implies the same statement for *every* $p \in (0,1)$.

2. Show that unique coupleability of ϕ with the quasirandom ℓ'-hypergraphon $\psi_{\ell',p}$ for all $p \in (0,1)$ implies that ϕ is unique coupleable with the quasirandom c-colored ℓ'-hypergraphon $\psi_{\ell',q}$ for every $c \geq 2$ and every $q \in \Pi_c$.

3. Show that unique coupleability of ϕ with all quasirandom colored ℓ'-hypergraphons for $\ell' \in [\ell]$ implies that ϕ is uniquely coupleable with all independent couplings $\psi_{1,p_1} \otimes \cdots \otimes \psi_{\ell,p_\ell}$ of quasirandom colored ℓ'-hypergraphons for $\ell' \in [\ell]$.

4. Show that in an arbitrary theory T', the set of elements that are uniquely coupleable with $\phi \in \operatorname{Hom}^+(\mathcal{A}[T],\mathbb{R})$ is closed in $\operatorname{Hom}^+(\mathcal{A}[T'],\mathbb{R})$ in the L_1-*topology*.

5. Show that for any pure canonical theory $T_\mathcal{L}$, the set of all elements of the form $(\psi_{1,p} \otimes \cdots \otimes \psi_{\ell,p})^I$, where $I\colon T_\mathcal{L} \rightsquigarrow \bigcup_{k=1}^{\ell} T_{c,k}$ is an open interpretation, is dense in the set of $\psi \in \operatorname{Hom}^+(\mathcal{A}[T_\mathcal{L}],\mathbb{R})$ of rank at most ℓ (again in the L_1-topology) and apply Theorem 4.2.3.

Let us point out that items 1, 2 and 3 combined show a strengthened version of implication (ii) $\implies$ (iii) that allows for multiple colors and arbitrary densities. Furthermore, most likely items 4 and 5 in this program can be replaced with an ad hoc argument but we prefer this more structured approach.

We start with item 1.

Lemma 4.5.1. *Let* $\ell \in \mathbb{N}_+$ *and* $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$. *If there exists* $p \in (0,1)$ *such that* ϕ *is uniquely coupleable with the quasirandom* ℓ*-hypergraphon* $\psi_{\ell,p}$, *then* ϕ *is uniquely coupleable with* $\psi_{\ell,q}$ *for every* $q \in (0,1)$.

Proof. Let $\mathcal{C}_q$ be the set of all couplings of ϕ with $\psi_{\ell,q}$. Our objective is to show that $|\mathcal{C}_q| = 1$. Without loss of generality, let us suppose that $p < q$ (otherwise, we can use the complementation automorphism $C\colon T_{\ell\text{-Hypergraph}} \rightsquigarrow T_{\ell\text{-Hypergraph}}$ given by $C(E)(\vec{x}) \stackrel{\text{def}}{=} \bigwedge_{1\le i<j\le \ell} x_i \neq x_j \land \neg E(\vec{x})$ and Theorem 4.2.3). Intuitively, we are going to "dilute" $\psi_{\ell,q}$ by a factor $t = p/q$ so that it will turn into $\psi_{\ell,p}$. The simplest way to make this intuition precise is by introducing yet another quasirandom hypergraphon $\psi_{\ell,t}$ on the same ground set and then taking its intersection with $\psi_{\ell,q}$.

Formally, we consider the commutative diagram

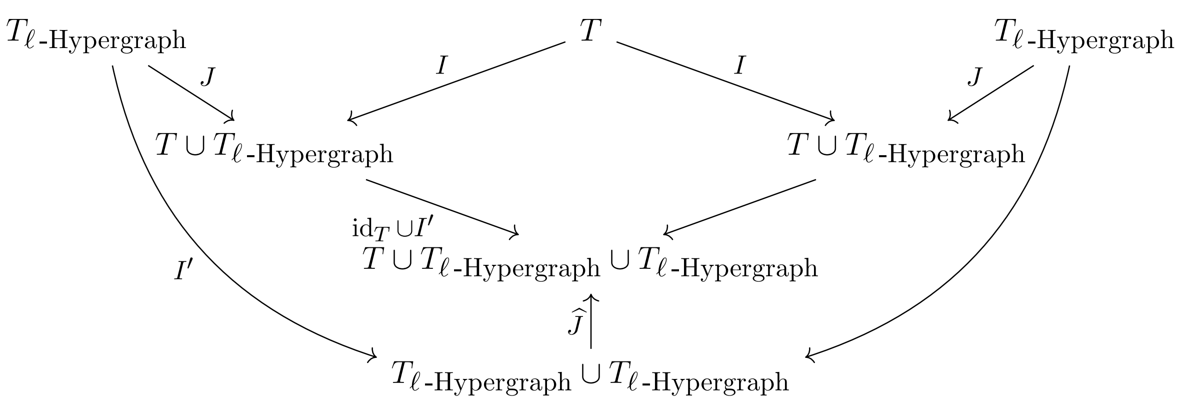

(4.12)

where I, J, $\widehat{J}$ and the unlabeled arrows are the structure-erasing interpretations, with the unlabeled arrows keeping the second copy of $T_{\ell\text{-Hypergraph}}$, and I' is given by

$$I'(E)(x_1, \ldots, x_\ell) = E(x_1, \ldots, x_\ell) \land E'(x_1, \ldots, x_\ell).$$

Here E corresponds to the first copy of $T_{\ell\text{-Hypergraph}}$ and E' corresponds to the second one.

We now define the dilution map $F\colon \mathcal{C}_q \to \mathcal{C}_p$ by

$$F(\xi) \stackrel{\text{def}}{=} (\xi \otimes \psi_{\ell,t})^{\mathrm{id}_T \cup I'},$$

where $t \stackrel{\text{def}}{=} p/q \in (0,1)$. The fact that $F(\xi)$ is indeed an element of $\mathcal{C}_p$ follows from

$$((\xi \otimes \psi_{\ell,t})^{\mathrm{id}_T \cup I'})^I = (\phi \otimes \psi_{\ell,t})^I = \phi;$$
$$((\xi \otimes \psi_{\ell,t})^{\mathrm{id}_T \cup I'})^J = (\psi_{\ell,q} \otimes \psi_{\ell,t})^{I'} = \psi_{\ell,p}.$$

For $M \in \mathcal{M}[T]$ and $U \subseteq \binom{V(M)}{\ell}$, let M_U be the model of $T \cup T_{\ell\text{-Hypergraph}}$ obtained from M by declaring the ℓ-hypergraph edge set to be U, that is, we have $I(M_U) = M$ and $E(J(M_U)) = U$. Then we have

$$F(\xi)(\langle M_U \rangle) = t^{|U|} \sum_{\substack{W \subseteq \binom{[m]}{\ell} \\ U \subseteq W}} (1-t)^{|W \setminus U|} \xi(\langle M_W \rangle).$$

By Möbius inversion, it follows that F is injective[3], hence $|\mathcal{C}_q| \leq |\mathcal{C}_p| = 1$ as claimed. ■

We now proceed to item 2 of our program.

Lemma 4.5.2. *Let $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ and $\ell \in \mathbb{N}_+$ and suppose that for every $p \in (0,1)$, ϕ is uniquely coupleable with the quasirandom ℓ-hypergraphon $\psi_{\ell,p}$. Then for every $c \geq 2$ and every $q \in \Pi_c$, ϕ is uniquely coupleable with the quasirandom c-colored ℓ-hypergraphon $\psi_{\ell,q}$.*

3. The left-inverse is given by

$$F^{-1}(\zeta)(\langle M_U \rangle) = t^{-|U|} \sum_{\substack{W \subseteq \binom{[m]}{\ell} \\ U \subseteq W}} \left(1 - \frac{1}{t}\right)^{|W \setminus U|} \zeta(\langle M_W \rangle).$$

Proof. For $i \in [c]$, consider the following commutative diagram

$$\begin{array}{ccccc}
T_{\ell\text{-Hypergraph}} & \xrightarrow{J} & T \cup T_{\ell\text{-Hypergraph}} & \xleftarrow{I} & T \\
{\scriptstyle I'_i}\downarrow & & {\scriptstyle \mathrm{id}_T \cup I'_i}\downarrow & \swarrow{\scriptstyle I_c} & \\
T_{c,\ell} & \xrightarrow[J_c]{} & T \cup T_{c,\ell} & &
\end{array}$$

where I, I_c, J and J_c are structure-erasing and I'_i is given by

$$I'_i(E)(x_1, \ldots, x_\ell) \stackrel{\text{def}}{=} E_i(x_1, \ldots, x_\ell).$$

The set $\mathcal{K}_m[T_{c,\ell}]$ of labeled models of size m can be naturally identified with functions $f\colon \binom{[m]}{\ell} \to [c]$: given $m \in \mathbb{N}$ and $f\colon \binom{[m]}{\ell} \to [c]$, $C_f \in \mathcal{K}_m[T_{c,\ell}]$ is given by

$$V(C_f) \stackrel{\text{def}}{=} [m]; \qquad\qquad R_{E_i}(C_f) \stackrel{\text{def}}{=} \{\alpha \in ([m])_\ell \mid f(\mathrm{im}(\alpha)) = i\} \quad (i \in [c]).$$

Let $F \stackrel{\text{def}}{=} C^{-1}$. Given further $K \in \mathcal{K}_m[T]$ and $f\colon \binom{[m]}{\ell} \to [c]$, let K_f be the alignment of K and C_f, that is, K_f is the unique model in $\mathcal{K}_m[T \cup T_{c,\ell}]$ such that $I_c(K_f) = K$ and $J_c(K_f) = C_f$. Similarly, given $U \subseteq \binom{[m]}{\ell}$, let $K_U \in \mathcal{K}_m[T \cup T_{\ell\text{-Hypergraph}}]$ be the unique model such that $I(K_U) = K$ and $R_E(K_U) = \{\alpha \in ([m])_\ell \mid \mathrm{im}(\alpha) \in U\}$.

Let $\psi \stackrel{\text{def}}{=} \psi_{\ell,q} \in \mathrm{Hom}^+(\mathcal{A}[T_{c,\ell}], \mathbb{R})$ and let ξ be a coupling of ϕ and ψ. Our goal is to show that

$$\xi(\langle K_f \rangle) = \psi(\langle C_f \rangle)\phi(\langle K \rangle) \tag{4.13}$$

for every $m \in \mathbb{N}$, every $K \in \mathcal{K}_m[T]$ and every $f\colon \binom{[m]}{\ell} \to [c]$. Note that to improve readability, here and in the forthcoming calculations, K and K_f are identified with their isomorphism classes $[K]$, $[K_f]$ in $\mathcal{M}_m$.

If $m < \ell$, then (4.13) holds trivially and if $\phi(\langle K \rangle) = 0$, then both sides of (4.13) are 0, so suppose $m \geq \ell$ and $\phi(\langle K \rangle) > 0$. Note that $\xi^{(\mathrm{id}_T \cup I'_i) \circ J} = \psi^{I'_i} = \psi_{\ell,q_i} \in$

$\mathrm{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$, hence $\xi^{\mathrm{id}_T \cup I'_i}$ is a coupling of ϕ and ψ_{ℓ,q_i}, so we must have $\xi^{\mathrm{id}_T \cup I'_i} = \phi \otimes \psi_{\ell,q_i}$. Note also that for $m \in \mathbb{N}$, $K \in \mathcal{K}_m[T]$ and $U \subseteq \binom{[m]}{\ell}$, we have

$$\pi^{\mathrm{id}_T \cup I'_i}(\langle K_U \rangle) = \sum_{\substack{f\colon \binom{[m]}{\ell} \to [c] \\ f^{-1}(i) = U}} \langle K_f \rangle. \tag{4.14}$$

Pick now $\boldsymbol{f}\colon \binom{[m]}{\ell} \to [c]$ at random according to the distribution

$$\mathbb{P}[\boldsymbol{f} = f] \stackrel{\mathrm{def}}{=} \frac{\xi(\langle K_f \rangle)}{\phi(\langle K \rangle)}.$$

The identity (4.14) allows us to compute, for $A \in \binom{[m]}{\ell}$ and $i \in [c]$, that

$$\begin{aligned} \mathbb{P}[\boldsymbol{f}(A) = i] &= \sum_{\substack{f\colon \binom{[m]}{\ell} \to [c] \\ f(A) = i}} \frac{\xi(\langle K_f \rangle)}{\phi(\langle K \rangle)} = \sum_{\substack{U \subseteq \binom{[m]}{\ell} \\ A \in U}} \frac{\xi^{\mathrm{id}_T \cup I'_i}(\langle K_U \rangle)}{\phi(\langle K \rangle)} \\ &= \sum_{\substack{U \subseteq \binom{[m]}{\ell} \\ A \in U}} q_i^{|U|}(1 - q_i)^{\binom{m}{\ell} - |U|} = q_i, \end{aligned}$$

where the the second equality follows from (4.14) and the third equality follows since $\xi^{\mathrm{id}_T \cup I'_i} = \phi \otimes \psi_{\ell,q_i}$. Since $\psi(\langle C_f \rangle) = \prod_{A \in \binom{[m]}{\ell}} q_{f(A)}$, to complete the proof of (4.13), it remains to show that the values $(\boldsymbol{f}(A) \mid A \in \binom{[m]}{\ell})$ of $\boldsymbol{f}$ are mutually independent.

For that purpose, it is in turn sufficient to prove that for every fixed $A_0 \in \binom{[m]}{\ell}$ and every fixed $i_0 \in [c]$, the event $\boldsymbol{f}(A_0) = i_0$ is independent from $\boldsymbol{f}|_W$, where $W \stackrel{\mathrm{def}}{=} \binom{[m]}{\ell} \setminus \{A_0\}$.

To do so, we will generate the distribution of $\boldsymbol{f}$ in a very specific way. Let $\mathcal{N}$ be a T-on such that $\phi = \phi_{\mathcal{N}}$ and note that $\psi_{\ell,q_{i_0}} = \phi_{\mathcal{N}'} \in \mathrm{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$ for the $(\ell - 1)$-independent $T_{\ell\text{-Hypergraph}}$-on $\mathcal{N}'$ given by

$$\mathcal{N}'_E \stackrel{\mathrm{def}}{=} \{x \in \mathcal{E}_\ell \mid x_{[\ell]} < q_{i_0}\}. \tag{4.15}$$

Since $\xi^{\mathrm{id}_T \cup I'_{i_0}} = \phi \otimes \psi_{\ell,q_{i_0}} = \phi_{\mathcal{N}\otimes\mathcal{N}'}$, by Proposition 3.2.1 applied to the interpretation $\mathrm{id}_T \cup I'_{i_0}$, there exists a $(T \cup T_{c,\ell})$-on $\mathcal{H}$ over $[0,1]^4$ such that $\phi_{\mathcal{H}} = \xi$ and

$$\begin{aligned} \mathcal{H}_P &= \mathcal{N}_P \times \mathcal{E}_{k(P)}([0,1]^3) \text{ a.e} \qquad (P \in \mathcal{L}); \\ \mathcal{H}_{E_{i_0}} &= \mathcal{E}_\ell \times \mathcal{N}'_E \times \mathcal{E}_\ell([0,1]^2) \text{ a.e.}, \end{aligned} \tag{4.16}$$

where $\mathcal{L}$ is the language of T.

Let now $(\boldsymbol{\theta^1}, \boldsymbol{\theta^2}, \boldsymbol{\theta^3}, \boldsymbol{\theta^4})$ be picked at random in $\mathcal{E}_{\mathbb{N}_+}([0,1]^4)$ according to λ and let $\boldsymbol{K}$ be the exchangeable array corresponding to $\mathcal{H}$ with respect to $(\boldsymbol{\theta^1}, \boldsymbol{\theta^2}, \boldsymbol{\theta^3}, \boldsymbol{\theta^4})$. Denote also $\boldsymbol{F} \stackrel{\text{def}}{=} F(J_c(\boldsymbol{K}|_{[m]}))$; $\boldsymbol{F} = (\boldsymbol{F}(A_0), \boldsymbol{F}|_W)$, and let E be the event $I_c(\boldsymbol{K}|_{[m]}) = K$. Then the function $\boldsymbol{f}$ is equidistributed with the function $\boldsymbol{F}$ conditioned by the event E. It remains to note that by (4.16), the event $\boldsymbol{F}(A_0) = i_0$ depends only on the coordinate $\boldsymbol{\theta}^{\boldsymbol{2}}_{A_0}$ (warning: we do *not* claim that the whole random variable $\boldsymbol{F}(A_0)$ depends only on $\boldsymbol{\theta}^{\boldsymbol{2}}_{A_0}$). On the other hand, both E and $\boldsymbol{F}|_W$ do not depend on it; more precisely, E depends only on $\boldsymbol{\theta}^1$ and $\boldsymbol{F}|_W$ depends on those $\boldsymbol{\theta}^{\boldsymbol{j}}_B$ with $j \in [4]$, $|B| \leq \ell$ and $B \neq A_0$. ■

We now address item 3 of our program (cf. the second remark made after the statement of Theorem 4.2.10).

Lemma 4.5.3. *Let $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ and $\psi_i \in \mathrm{Hom}^+(\mathcal{A}[T_i], \mathbb{R})$ for $i \in [t]$. Let also $\ell_1 \leq \cdots \leq \ell_t$ and suppose that the following hold.*

i. For every $i \in \{1, \ldots, t-1\}$, we have $\mathrm{rk}(\psi_i) \leq \ell_i$.

ii. For every $i \in \{2, \ldots, t\}$, we have $\psi_i \in$ `Independence`$[\ell_{i-1}]$.

iii. For every $i \in \{1, \ldots, t\}$, ϕ and ψ_i are uniquely coupleable.

Then $\phi, \psi_1, \ldots, \psi_t$ are uniquely coupleable.

Proof. The proof is by induction on t. For $t = 1$, the result is trivial. For $t = 2$, let $I_i\colon T \cup T_i \rightsquigarrow T \cup T_1 \cup T_2$, $J_i\colon T_i \rightsquigarrow T \cup T_1 \cup T_2$ and $J\colon T \rightsquigarrow T \cup T_1 \cup T_2$ be the structure-erasing interpretations. Let $\mathcal{L}$, $\mathcal{L}_1$ and $\mathcal{L}_2$ be the languages of T, T_1 and T_2, respectively. Let also $\mathcal{N}$ be a T-on with $\phi_{\mathcal{N}} = \phi$ and $\mathcal{H}^2$ be an ℓ_1-independent T_2-on with $\phi_{\mathcal{H}^2} = \psi_2$. Fix a coupling ξ of ϕ, ψ_1, ψ_2. Since ϕ and ψ_2 are uniquely coupleable, we know that $\xi^{I_2} = \phi \otimes \psi_2 = \phi_{\mathcal{N}\otimes\mathcal{H}^2}$. By Proposition 3.2.1, there exists a $(T \cup T_1 \cup T_2)$-on $\mathcal{G}$ over $[0,1]^4$ such that $\phi_{\mathcal{G}} = \xi$ and

$$\mathcal{G}_P = \begin{cases} \mathcal{N}_P \times \mathcal{E}_{k(P)}([0,1]^3), & \text{if } P \in \mathcal{L}; \\ \mathcal{E}_{k(P)} \times \mathcal{H}^2_P \times \mathcal{E}_{k(P)}([0,1]^2), & \text{if } P \in \mathcal{L}_2. \end{cases}$$

On the other hand, for the predicate symbols P in $\mathcal{L}_1$, by possibly changing zero-measure sets of the corresponding P-ons $\mathcal{G}_P$ using Proposition 3.1.2, we may suppose that $\operatorname{rk}(J_1(\mathcal{G})) \leq \operatorname{rk}(\psi_1) \leq \ell_1$.

Let us pick $\boldsymbol{\theta} \stackrel{\text{def}}{=} (\boldsymbol{\theta}^1, \boldsymbol{\theta}^2, \boldsymbol{\theta}^3, \boldsymbol{\theta}^4)$ at random in $\mathcal{E}_{\mathbb{N}_+}([0,1]^4)$ according to λ and let $\boldsymbol{K}$ be the exchangeable array corresponding to $\mathcal{G}$ with respect to $\boldsymbol{\theta}$. Then we know that $J(\boldsymbol{K})$ depends only on $\boldsymbol{\theta}^1$, $J_1(\boldsymbol{K})$ depends only on $((\boldsymbol{\theta}^1_A, \boldsymbol{\theta}^2_A, \boldsymbol{\theta}^3_A, \boldsymbol{\theta}^4_A) \mid |A| \leq \ell_1)$ and $J_2(\boldsymbol{K})$ depends only on $(\boldsymbol{\theta}^2_A \mid |A| > \ell_1)$ (as $\mathcal{H}^2$ is ℓ_1-independent), so $J_2(\boldsymbol{K})$ is independent from $(J(\boldsymbol{K}), J_1(\boldsymbol{K}))$. This means that for every $m \in \mathbb{N}$ and every $K \in \mathcal{K}_m[T \cup T_1 \cup T_2]$, we have

$$\begin{aligned}
\xi(\langle K \rangle) &= \mathbb{P}[\boldsymbol{K}|_{[m]} = K] \\
&= \mathbb{P}[J(\boldsymbol{K})|_{[m]} = J(K) \wedge J_1(\boldsymbol{K})|_{[m]} = J_1(K) \wedge J_2(\boldsymbol{K})|_{[m]} = J_2(K)] \\
&= \mathbb{P}[J(\boldsymbol{K})|_{[m]} = J(K) \wedge J_1(\boldsymbol{K})|_{[m]} = J_1(K)] \cdot \mathbb{P}[J_2(\boldsymbol{K})|_{[m]} = J_2(K)] \\
&= \mathbb{P}[I_1(\boldsymbol{K})|_{[m]} = I_1(K)] \cdot \mathbb{P}[J_2(\boldsymbol{K})|_{[m]} = J_2(K)] \\
&= \xi^{I_1}(\langle I_1(K) \rangle) \cdot \psi_2(\langle J_2(K) \rangle) \\
&= \phi(\langle J(K) \rangle) \cdot \psi_1(\langle J_1(K) \rangle) \cdot \psi_2(\langle J_2(K) \rangle),
\end{aligned}$$

where the last equality follows since ϕ is uniquely coupleable with ψ_1 and ξ^{I_1} is a coupling of ϕ and ψ_1. Therefore $\xi = \phi \otimes \psi_1 \otimes \psi_2$.

For the case $t \geq 3$, let $I\colon T \cup \bigcup_{i=2}^t T_i \rightsquigarrow T \cup \bigcup_{i=1}^t T_i$ be the structure-erasing interpretation and note that for a coupling ξ of $\phi, \psi_1, \ldots, \psi_t$, it follows that ξ^I is a coupling of $\phi, \psi_2, \ldots, \psi_t$. By inductive hypothesis, we must have $\xi^I = \phi \otimes \widehat{\psi}$, where $\widehat{\psi} \stackrel{\text{def}}{=} \bigotimes_{i=2}^t \psi_i$. In fact, since $\phi, \psi_2, \ldots, \psi_t$ are uniquely coupleable, it also follows that ϕ is uniquely coupleable with $\widehat{\psi}$ (as any coupling of ϕ with $\widehat{\psi}$ can be seen as a coupling of $\phi, \psi_2, \ldots, \psi_t$). But by Theorem 4.2.4, we know that $\widehat{\psi} \in \texttt{Independence}[\ell_1]$ and since ξ can also be seen as a coupling of $\phi, \psi_1, \widehat{\psi}$, we get $\xi = \phi \otimes \bigotimes_{i=1}^t \psi_i$ from the previous case. ∎

Lemma 4.5.4. *Let $c \geq 2$, $p \in \Pi_c$ and $k \in \mathbb{N}_+$. Then the quasirandom c-colored k-hypergraphon $\psi_{k,p}$ satisfies $\texttt{Independence}[k-1]$ and $\operatorname{rk}(\psi_{k,p}) = k$.*

Proof. Note that $\psi_{k,p}$ can be represented by the $T_{c,k}$-on $\mathcal{N}^{k,p}$ given by

$$\mathcal{N}_{E_i}^{k,p} \stackrel{\text{def}}{=} \left\{ x \in \mathcal{E}_k \;\middle|\; \sum_{j=1}^{i-1} p_j \leq x_{[k]} < \sum_{j=1}^{i} p_j \right\} \qquad (i \in [c]),$$

hence $\psi_{k,p} \in \texttt{Independence}[k-1]$ and $\operatorname{rk}(\psi_{k,p}) \leq k$. Since $c \geq 2$, it follows that $\operatorname{rk}(\psi_{k,p}) > 0$, so by Theorem 4.2.2 and Proposition 4.3.1, we must have $\operatorname{rk}(\psi_{k,p}) = k$. ∎

Before proceeding to item 4 in the program, let us remark why need the L_1-topology for it instead of the standard and much nicer density topology (i.e., the one induced by the inclusion $\operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R}) \subseteq [0,1]^{\mathcal{M}[T]}$ from the product topology). One simple explanation is that the set of all $\psi \in \operatorname{Hom}^+(\mathcal{A}[T'], \mathbb{R})$ that are uniquely coupleable with some $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is not closed in the latter.

Example 5. Let $\phi_p \in \operatorname{Hom}^+(\mathcal{A}[T_{\text{Graph}}], \mathbb{R})$ be the quasirandom graphon of density $p \in (0,1)$. If $(G_n)_{n\in\mathbb{N}}$ ($G_n \in \mathcal{M}_n[T_{\text{Graph}}]$) is a sequence of graphs converging to ϕ_p, then the associated step functions ψ_n converge to ϕ_p in the density topology. Since $\operatorname{rk}(\psi_n) = 1$ and $\phi_p \in \texttt{Independence}[1]$, it follows that ϕ_p and ψ_n are uniquely coupleable, but $\phi_p = \lim_{n\to\infty} \psi_n$ is obviously not uniquely coupleable with itself.

Lemma 4.5.5. *Let $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ and T' be an arbitrary theory. Then the set of $\psi \in \operatorname{Hom}^+(\mathcal{A}[T'], \mathbb{R})$ that are uniquely coupleable with ϕ is closed in the L_1-topology.*

Proof. Let $(\psi_n)_{n\in\mathbb{N}}$ be a sequence in $\mathrm{Hom}^+(\mathcal{A}[T'],\mathbb{R})$ converging to ψ in the L_1-topology and suppose every ψ_n is uniquely coupleable with ϕ. It is clear from the definition that $\delta_1(\phi\otimes\psi_n,\phi\otimes\psi) = \delta_1(\psi_n\otimes\psi)$, so $\phi\otimes\psi_n$ also converges to $\phi\otimes\psi$ in the L_1-topology. For each $n\in\mathbb{N}$, let ζ_n be a coupling of ψ and ψ_n attaining the L_1-distance in (3.12).

Let ξ be a coupling of ϕ and ψ; we have to show that $\xi = \phi\otimes\psi$. Let $I\colon T'\cup T' \rightsquigarrow T\cup T'\cup T'$ and $J_i\colon T' \rightsquigarrow T'\cup T'$ be the structure-erasing interpretations, where J_i keeps the i-th copy of T'. Since ξ is a coupling of ϕ and $\psi = \zeta_n^{J_1}$, by Proposition 3.2.6, there exists a coupling $\widehat{\xi}_n$ of ϕ and ζ_n such that $\widehat{\xi}_n^{\mathrm{id}_T\cup J_1} = \xi$. Note that $\widehat{\xi}_n$ can also be seen as a coupling of ϕ, ψ and ψ_n as $\widehat{\xi}_n^I = \zeta_n$.

Let now $\mathcal{N}^n$ be a $(T\cup T'\cup T')$-on such that $\widehat{\xi}_n = \phi_{\mathcal{N}^n}$. By considering the $(T\cup T')$-ons $(\mathrm{id}_T\cup J_1)(\mathcal{N}^n)$ and $(\mathrm{id}_T\cup J_2)(\mathcal{N}^n)$, since ψ_n is uniquely coupleable with ϕ, we conclude from (3.10) that

$$\delta_1(\xi,\phi\otimes\psi_n) \leq \sum_{P\in\mathcal{L}'} \lambda(J_1(I(\mathcal{N}^n))_P \mathbin{\triangle} J_2(I(\mathcal{N}^n))_P) = \zeta_n(d_{T'}) = \delta_1(\psi,\psi_n),$$

where $\mathcal{L}'$ is the language of T'. Since $\psi_n\to\psi$ and $\phi\otimes\psi_n\to\phi\otimes\psi$ in the L_1-topology, it follows that $\xi = \phi\otimes\psi$. ■

We proceed to the last item 5 in our program, which is to provide a way of approximating Euclidean structures with interpretations of independent couplings $\psi_{1,p}\otimes\cdots\otimes\psi_{\ell,p}$ of quasirandom colored hypergraphons in the L_1-topology.

Lemma 4.5.6. *Let $\mathcal{L}$ be a language, $\phi\in\mathrm{Hom}^+(\mathcal{A}[T_{\mathcal{L}}],\mathbb{R})$, $r \stackrel{\text{def}}{=} rk(\phi)$ and $\varepsilon > 0$. Then there exist $c\geq 2$, $p\in\Pi_c$ and an open interpretation $I\colon T_{\mathcal{L}} \rightsquigarrow \bigcup_{k=1}^r T_{c,k}$ such that $\delta_1(\phi,(\bigotimes_{k=1}^r \psi_{k,p})^I)\leq\varepsilon$.*

Proof. Let $\mathcal{N}$ be a $T_{\mathcal{L}}$-on such that $\phi_{\mathcal{N}} = \phi$ and $\mathrm{rk}(\mathcal{N}) = r$, that is, for each $P\in\mathcal{L}$, there exists $\mathcal{H}_P\subseteq\mathcal{E}_{k(P),r}$ such that $\mathcal{N}_P = \mathcal{H}_P\times[0,1]^{\binom{[k(P)]}{>r}}$. By standard measure theory arguments, for each $P\in\mathcal{L}$, there exists a finite family of pairwise disjoint closed cubes $(C_j^P)_{j=1}^{m_P}$ ($C_j^P\subseteq\mathcal{E}_{k(P),r}$) such that setting $\mathcal{H}'_P \stackrel{\text{def}}{=} \bigcup_{j=1}^{m_P} C_j^P$ gives $\lambda(\mathcal{H}_P\mathbin{\triangle}\mathcal{H}'_P)\leq\varepsilon/|\mathcal{L}|$.

Let X be the set of all coordinates of vertices of all cubes C_j^P for all $P \in \mathcal{L}$. The set X induces a partition of $[0,1]$ into intervals $J_1, \ldots, J_c$ of positive length (we can ensure $c \geq 2$ by including an extra point if necessary). Define then $p \in \Pi_c$ by letting $p_i \stackrel{\text{def}}{=} \lambda(J_i) > 0$ and define the $(\bigcup_{k=1}^r T_{c,k})$-on $\mathcal{G}$ by

$$\mathcal{G}_{E_i^k} \stackrel{\text{def}}{=} \{x \in \mathcal{E}_k \mid x_{[k]} \in J_i\} \qquad (i \in [c], k \in [r]),$$

where for each $k \in [r]$, the symbols $E_1^k, \ldots, E_c^k$ correspond to $T_{c,k}$.

Let $\psi \stackrel{\text{def}}{=} \phi_{\mathcal{G}}$ and note that ψ is a coupling of $\psi_{1,p}, \ldots, \psi_{r,p}$, so we must have $\psi = \bigotimes_{k=1}^r \psi_{k,p}$ by Lemmas 4.5.3 and 4.5.4.

Note now that from the definition of X, each cube $C_j^P \subseteq \mathcal{E}_{k(P),r}$ can be written as a finite union of the form $\bigcup_{u \in U_{P,j}} \prod_{A \in r(k(P),r)} J_{i_{P,u,A}}$. We then define $I\colon T_{\mathcal{L}} \rightsquigarrow \bigcup_{k=1}^r T_{c,k}$ by

$$I(P)(x_1, \ldots, x_{k(P)}) \stackrel{\text{def}}{=} \bigvee_{j=1}^{m_P} \bigvee_{u \in U_{P,j}} \bigwedge_{A \in r(k(P),r)} E^k_{i_{P,u,A}}(x_{\iota_A(1)}, \ldots, x_{\iota_A(|A|)}) \qquad (P \in \mathcal{L}).$$

Our definition ensures that

$$\begin{aligned} I(\mathcal{G})_P &= \bigcup_{j=1}^{m_P} \bigcup_{u \in U_{P,j}} \left(\prod_{A \in r(k(P),r)} J_{i_{P,u,A}} \times [0,1]^{\binom{[k(P)]}{>r}} \right) \\ &= \bigcup_{j=1}^{m_P} (C_j^P \times [0,1]^{\binom{[k(P)]}{>r}}) = \mathcal{H}'_P \times [0,1]^{\binom{[k(P)]}{>r}}. \end{aligned}$$

This implies that

$$\delta_1(\phi, \psi) \leq \sum_{P \in \mathcal{L}} \lambda(\mathcal{N}_P \mathbin{\triangle} (\mathcal{H}'_P \times [0,1]^{\binom{[k(P)]}{>r}})) = \sum_{P \in \mathcal{L}} \lambda(\mathcal{H}_P \mathbin{\triangle} \mathcal{H}'_P) \leq \varepsilon,$$

as desired. ∎

We now have all the ingredients to show the equivalence (i)⇔(ii)⇔(iii) of Theorem 4.2.10.

Lemma 4.5.7 (Theorem 4.2.10(i)⇔(ii)⇔(iii))**.** *Let $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ and $\ell \in \mathbb{N}_+$. Then the following are equivalent.*

i. $\phi \in$ ***UCouple***$[\ell]$.

ii. For every $\ell' \in [\ell]$, there exists $p \in (0,1)$ such that ϕ is uniquely coupleable with the quasirandom ℓ'-hypergraphon $\psi_{\ell',p}$.

iii. There exist $p_1, \ldots, p_\ell \in (0,1)$ such that ϕ is uniquely coupleable with the independent coupling $\psi_{1,p_1} \otimes \cdots \otimes \psi_{\ell,p_\ell}$ of quasirandom ℓ'-hypergraphons $\psi_{\ell',p_{\ell'}}$ for $\ell' \in [\ell]$.

Proof. Since ℓ'-hypergraphons have rank at most ℓ', by Proposition 3.1.2, we have $\operatorname{rk}(\psi_{1,p_1} \otimes \cdots \otimes \psi_{\ell,p_\ell}) \leq \ell$, so the implication (i) $\implies$ (iii) follows.

Implication (iii) $\implies$ (ii) follows from Theorem 4.2.3 by considering the structure-erasing interpretations $I_k\colon T_{k\text{-Hypergraph}} \rightsquigarrow \bigcup_{\ell'=1}^{\ell} T_{\ell'\text{-Hypergraph}}$.

For the non-trivial implication (ii) $\implies$ (i), we want to show that ϕ is uniquely coupleable with any $\psi \in \operatorname{Hom}^+(\mathcal{A}[T'], \mathbb{R})$ of rank at most ℓ. We can assume w.l.o.g. that $T' = T_{\mathcal{L}}$ for some language $\mathcal{L}$. Using Lemma 4.5.6, for each $n \in \mathbb{N}$, we can find $c_n \geq 2$, $p_n \in \Pi_{c_n}$ and $I_n\colon T_{\mathcal{L}} \rightsquigarrow \bigcup_{k=1}^{r} T_{c_n,k}$ such that $\delta_1(\phi, (\bigotimes_{k=1}^{r} \psi_{k,p_n})^{I_n}) \leq 1/n$.

By Lemmas 4.5.1, 4.5.2, 4.5.3 and 4.5.4, we know that ϕ is uniquely coupleable with $\bigotimes_{k=1}^{r} \psi_{k,p_n}$ and by Theorem 4.2.3, it follows that ϕ is also uniquely coupleable with $(\bigotimes_{k=1}^{r} \psi_{k,p_n})^{I_n}$.

Finally, since $((\bigotimes_{k=1}^{r} \psi_{k,p_n})^{I_n})_{n \in \mathbb{N}}$ converges to ψ in the L_1-topology, by Lemma 4.5.5, it follows that ϕ is uniquely coupleable with ψ. ■

We now proceed to add items (vi) and (vii) to the list of equivalent properties of Theorem 4.2.10 (recall that (i)⇔(iv)⇔(v) and (iv) $\implies$ (vi) were proved in Section 4.3).

Lemma 4.5.8 (Theorem 4.2.10(vi) $\implies$ (vii))**.** *If $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is ℓ-local, then $\phi \otimes \psi_{\text{lin}}$ satisfies* ***UInduce***$[\ell]$.

Proof. By Lemma 4.4.3, it is enough to show that $\phi \otimes \psi_{\text{lin}}$ is symmetrically ℓ-local. Let $\boldsymbol{K}$ be the exchangeable array corresponding to $\phi \otimes \psi_{\text{lin}}$, and fix a finite family of finite sets $(V_i)_{i\in[t]}$ ($V_i \subseteq \mathbb{N}_+$) with pairwise intersections of size at most ℓ. We let $\boldsymbol{K_i} \stackrel{\text{def}}{=} \boldsymbol{K}|_{V_i} \in \mathcal{K}_{V_i}[T \cup T_{\text{LinOrder}}]$ and let $\boldsymbol{M_i} \stackrel{\text{def}}{=} [\boldsymbol{K_i}] \in \mathcal{M}_{|V_i|}[T \cup T_{\text{LinOrder}}]$ be the isomorphism type of $\boldsymbol{K_i}$. We have to prove that $\boldsymbol{M_1}, \ldots, \boldsymbol{M_t}$ are mutually independent, and for that purpose we are going to apply Claim 4.3.5 again.

More specifically, let $I\colon T \rightsquigarrow T \cup T_{\text{LinOrder}}$ be the structure-erasing interpretation and $\boldsymbol{L_i} = I(\boldsymbol{K_i}) \in \mathcal{K}_{V_i}[T]$ be the results of erasing linear order. Likewise, let $J\colon T_{\text{LinOrder}} \rightsquigarrow T \cup T_{\text{LinOrder}}$, and let $\boldsymbol{\leq_i} = J(\boldsymbol{K_i})$ be the corresponding (random) linear order on V_i so that $\boldsymbol{K_i} = (\boldsymbol{L_i}, \boldsymbol{\leq_i})$. In Claim 4.3.5, we set $\boldsymbol{X} = (\boldsymbol{\leq_1}, \ldots, \boldsymbol{\leq_n})$, $\boldsymbol{Y_i} = \boldsymbol{L_i}$, and let $f_i(\leq_1, \ldots, \leq_n, L_i)$ be the function first computing K_i from L_i and $\leq_i$ and then taking its isomorphism type $M_i = [K_i]$.

We know that the tuple $(\boldsymbol{L_1}, \ldots, \boldsymbol{L_t})$ is independent from $\boldsymbol{X} = (\boldsymbol{\leq_1}, \ldots, \boldsymbol{\leq_t})$ (as the coupling of ϕ and ψ_{lin} is independent) and that $\boldsymbol{L_1}, \ldots, \boldsymbol{L_t}$ are mutually independent (as ϕ is ℓ-local). This gives us the first assumption in Claim 4.3.5: $\boldsymbol{X}, \boldsymbol{Y_1}, \ldots, \boldsymbol{Y_n}$ are mutually independent (note that we do *not* claim that $\boldsymbol{\leq_1}, \ldots, \boldsymbol{\leq_n}$ are mutually independent, this is in general not true). It remains to show that $(\boldsymbol{M_1}, \ldots, \boldsymbol{M_n})$ is independent from $(\boldsymbol{\leq_1}, \ldots, \boldsymbol{\leq_n})$, and it essentially follows from the observation that the function $f_i(X, Y_i)$ becomes invertible after fixing its first argument.

More specifically, we compute $\boldsymbol{L_i} = g_i(\boldsymbol{\leq_i}, \boldsymbol{M_i})$, where $g_i(\leq_i, M_i)$ is obtained by first aligning the internal order of $V(M_i)$ with the order $\leq_i$ on V_i, and then discarding it. The crucial property is that $L_i = g_i(\leq_i, M_i)$ if and only if $M_i = f_i((\leq_1, \ldots, \leq_n), L_i)$. Using this, fixing arbitrary models $M_i \in \mathcal{M}_{|V_i|}[T \cup T_{\text{LinOrder}}]$ and a particular tuple of values $(\leq_1, \ldots, \leq_t)$, we have the calculation

$$\begin{aligned}
&\mathbb{P}[\forall i \in [t], \boldsymbol{M_i} \cong M_i \mid \forall i \in [t], \boldsymbol{\leq_i} = \leq_i] \\
&\quad = \mathbb{P}[\forall i \in [t], \boldsymbol{L_i} = g_i(\leq_i, M_i) \mid \forall i \in [t], \boldsymbol{\leq_i} = \leq_i] \\
&\quad = \mathbb{P}[\forall i \in [t], \boldsymbol{L_i} = g_i(\leq_i, M_i)] \\
&\quad = \mathbb{P}[\forall i \in [t], \boldsymbol{M_i} \cong M_i].
\end{aligned}$$

This shows that $(\boldsymbol{M_1}, \ldots, \boldsymbol{M_t})$ is indeed independent from $(\boldsymbol{\leq_1}, \ldots, \boldsymbol{\leq_t})$. We are now in position to apply Claim 4.3.5 which completes the proof. ■

Lemma 4.5.9 (Theorem 4.2.10(vii) $\implies$ (ii))**.** *If the independent coupling $\phi \otimes \psi$ of $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$ with ψ_{lin} satisfies* `UInduce`$[\ell]$*, then for every $\ell' \in [\ell]$, ϕ is uniquely coupleable with the quasirandom ℓ'-hypergraphon $\psi_{\ell',1/2} \in \mathrm{Hom}^+(\mathcal{A}[T_{\ell'\text{-Hypergraph}}], \mathbb{R})$.*

Proof. Let $\mathcal{L}$ be the language of T and note that since `UInduce`$[\ell]$ implies `UInduce`$[\ell']$ (Theorem 4.2.1), it is sufficient to consider the case $\ell' = \ell$. Let us first assume $\ell \geq 2$.

Note that ψ_{lin} can be represented by the T_{LinOrder}-on $\mathcal{N}^<$ given by

$$\mathcal{N}^< \stackrel{\mathrm{def}}{=} \{x \in \mathcal{E}_2 \mid x_{\{1\}} < x_{\{2\}}\},$$

and that $\psi_{\ell,1/2}$ can be represented as

$$\mathcal{N}_E \stackrel{\mathrm{def}}{=} \{x \in \mathcal{E}_\ell \mid x_{[\ell]} \leq 1/2\}.$$

Let ξ be a coupling of ϕ and $\psi_{\ell,1/2}$ and let $\mathcal{N}$ be a $(T \cup T_{\ell\text{-Hypergraph}})$-on such that $\phi_{\mathcal{N}} = \xi$. As in the proof of Lemma 4.5.2, for every $m \in \mathbb{N}$ and every $U \subseteq \binom{[m]}{\ell}$, let $H_U \in \mathcal{K}_m[T_{\ell\text{-Hypergraph}}]$ be the hypergraph given by $V(H_U) \stackrel{\mathrm{def}}{=} [m]$ and $R_E(H_U) \stackrel{\mathrm{def}}{=} \{\alpha \in ([m])_\ell \mid \mathrm{im}(\alpha) \in U\}$. If we are further given $K \in \mathcal{K}_m[T]$, then we let $K_U \in \mathcal{K}_m[T \cup T_{\ell\text{-Hypergraph}}]$ be the alignment of K and H_U, that is, we have $R_P(K_U) \stackrel{\mathrm{def}}{=} R_P(K)$ $(P \in \mathcal{L})$

and $R_E(K_U) \stackrel{\text{def}}{=} R_E(H_U)$. Finally, we let $K_U^< \in \mathcal{K}_m[T \cup T_{\ell\text{-Hypergraph}} \cup T_{\text{LinOrder}}]$ be the model obtained from K_U by equipping it with the natural order of $[m]$. Note that while we do need labels in K to properly define the models K_U and $K_U^<$, in the computations below they are treated as unlabeled models $[K_U]$, $[K_U^<]$, i.e., labels are discarded.

To show that ξ is the independent coupling of ϕ and $\psi_{\ell,1/2}$, we need to show that for every $m \in \mathbb{N}$, every $K \in \mathcal{K}_m[T]$ and every $U \subseteq \binom{[m]}{\ell}$, we have

$$\xi(\langle K_U \rangle) = \phi(\langle K \rangle) \cdot \psi_{\ell,1/2}(\langle H_U \rangle) = \frac{\phi(\langle K \rangle)}{2^{\binom{m}{\ell}}}. \tag{4.17}$$

The assertion is trivial if $m < \ell$, so suppose $m \geq \ell$. Fix $U \subseteq \binom{[m]}{\ell}$ and for every $v \in [m]$, define

$$V_v \stackrel{\text{def}}{=} \begin{cases} \left[\frac{v-1}{m}, \frac{v}{m}\right), & \text{if } v < m; \\ \left[\frac{m-1}{m}, 1\right], & \text{if } v = m. \end{cases}$$

For $n \in \mathbb{N}$ and $y \in \mathcal{E}_n$, let $\alpha_y\colon [n] \to [m]$ be the unique function such that $y_{\{j\}} \in V_{\alpha_y(j)}$ for every $j \in [n]$. Finally, define the set

$$W_U \stackrel{\text{def}}{=} \left\{ (x, y) \in \mathcal{E}_\ell \times \mathcal{E}_\ell \;\middle|\; |\text{im}(\alpha_y)| = \ell \wedge \left(x_{[\ell]} \leq \frac{1}{2} \leftrightarrow \text{im}(\alpha_y) \in U \right) \right\};$$

clearly, W_U is S_ℓ-invariant. This means that we can define the $(T \cup T_{\ell\text{-Hypergraph}} \cup T_{\text{LinOrder}})$-on $\mathcal{H}^U$ over $[0,1]^2$ by

$$\mathcal{H}_P^U \stackrel{\text{def}}{=} \mathcal{N}_P \times \mathcal{E}_{k(P)} \qquad (P \in \mathcal{L}), \qquad\qquad \mathcal{H}_\prec^U \stackrel{\text{def}}{=} \mathcal{E}_2 \times \mathcal{N}^<, \qquad\qquad \mathcal{H}_E^U \stackrel{\text{def}}{=} W_U.$$

Obviously, if $(x, y) \in T_{\text{ind}}(K_m^{(\ell)}, W_U)$, then each $y_{\{j\}}$ must belong to a different V_v. Indeed, if there exist $j_1, j_2 \in [m]$ with $y_{\{j_1\}}, y_{\{j_2\}} \in V_v$ but $j_1 \neq j_2$, since $m \geq \ell \geq 2$, there exists $\beta \in ([m])_\ell$ with $j_1, j_2 \in \text{im}(\beta)$ and thus $(x, y) \notin (\beta^*)^{-1}(W_U)$, a contradiction.

Our claim and the definition of W_U then imply

$$
\begin{aligned}
&T_{\mathrm{ind}}(K_m^{(\ell)}, W_U)\\
&\quad = \left\{ (x,y) \in \mathcal{E}_m \times \mathcal{E}_m \;\middle|\; |\mathrm{im}(\alpha_y)| = m \wedge \bigwedge_{\beta \in ([m])_\ell} (\beta^*(x) \in \mathcal{N}_E \leftrightarrow \mathrm{im}(\alpha_{\beta^*(y)}) \in U) \right\}.
\end{aligned}
$$

Thus, denoting by $J_\ell \colon T_{\ell\text{-Hypergraph}} \rightsquigarrow T \cup T_{\ell\text{-Hypergraph}} \cup T_{\mathrm{LinOrder}}$ the structure-erasing interpretation, we get

$$
\phi_{J_\ell(\mathcal{H}^U)}(K_m^{(\ell)}) = \frac{m!}{m^m} \psi_{\ell,1/2}(H_U) = \frac{m!}{m^m \cdot 2^{\binom{m}{\ell}}}. \tag{4.18}
$$

Let now $J \colon T \rightsquigarrow T \cup T_{\ell\text{-Hypergraph}} \cup T_{\mathrm{LinOrder}}$ be another structure-erasing interpretation; we have

$$
\begin{aligned}
&T_{\mathrm{ind}}(K^{<}_{\binom{[m]}{\ell}}, \mathcal{H}^U)\\
&\quad = T_{\mathrm{ind}}(K, J(\mathcal{H}^U)) \cap T_{\mathrm{ind}}(K_m^{(\ell)}, J_\ell(\mathcal{H}^U)) \cap \{(x,y) \in \mathcal{E}_m \times \mathcal{E}_m \mid y_{\{1\}} < \cdots < y_{\{m\}}\}\\
&\quad = \{(x,y) \in \mathcal{E}_m \times \mathcal{E}_m \mid x \in T_{\mathrm{ind}}(K_U, \mathcal{N}) \wedge \forall v \in [m], y_{\{v\}} \in V_v\}.
\end{aligned}
$$

Since $\phi_{\mathcal{N}} = \xi$, we get

$$
\xi(\langle K_U \rangle) = m^m \cdot \phi_{\mathcal{H}^U}(\langle K^{<}_{\binom{[m]}{\ell}} \rangle) = \frac{m^m \cdot \phi(\langle K \rangle) \cdot \phi_{J_\ell(\mathcal{H}^U)}(K_m^{(\ell)})}{m!} = \frac{\phi(\langle K \rangle)}{2^{\binom{m}{\ell}}},
$$

where the second equality follows since $\phi_{\mathcal{H}^U} \in \mathrm{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}} T \cup T_{\mathrm{LinOrder}}], \mathbb{R})$ is a coupling of $\phi_{J_\ell(\mathcal{H}^U)} \in \mathrm{Hom}^+(\mathcal{A}[T_{\ell\text{-Hypergraph}}], \mathbb{R})$ and $\phi \otimes \psi_{\mathrm{lin}}$ (and the latter satisfies `UInduce`$[\ell]$), and the third equality follows from (4.18). Hence (4.17) holds.

Let us now show the case $\ell = 1$. In this case, since $T_{1\text{-Hypergraph}} \cong T_{2\text{-Coloring}}$, we will work with the latter theory. Let ξ be a coupling of ϕ and $\psi_{1/2} \in \mathrm{Hom}^+(\mathcal{A}[T_{2\text{-Coloring}}], \mathbb{R})$ and let $\mathcal{N}$ be a $(T \cup T_{2\text{-Coloring}})$-on such that $\phi_{\mathcal{N}} = \xi$.

For every $m \in \mathbb{N}$, every $K \in \mathcal{K}_m[T]$ and every $j \in \{0, \ldots, m\}$, let $K_j \in \mathcal{K}_m[T \cup T_{2\text{-Coloring}}]$ be the model obtained from K by coloring the first j vertices with color 1 and all others with color 2, that is, we have $R_P(K_j) \stackrel{\text{def}}{=} R_P(K)$ $(P \in \mathcal{L})$, $R_{\chi 1}(K_j) \stackrel{\text{def}}{=} [j]$ and $R_{\chi 2}(K_j) \stackrel{\text{def}}{=} \{j+1, \ldots, m\}$. Again, we let $K_j^{<} \in \mathcal{K}_m[T \cup T_{2\text{-Coloring}} \cup T_{\text{LinOrder}}]$ be the model obtained from K_j by equipping it with the natural order of $[m]$, and, again, in the computations below we view $K, K_j, K_j^{<}$ as unlabeled models.

Due to exchangeability, in order to show that ξ is the independent coupling of ϕ and $\psi_{1/2}$, it is sufficient to show that for every $m \in \mathbb{N}$, every $K \in \mathcal{K}_m[T]$ and every $j \in \{0, \ldots, m\}$, we have

$$\xi(\langle K_j \rangle) = \frac{\phi(\langle K \rangle)}{2^m}. \tag{4.19}$$

For every $t \in (0, 1)$, let

$$U_t \stackrel{\text{def}}{=} \{(x, y) \in \mathcal{E}_1 \times \mathcal{E}_1 \mid x \in \mathcal{N}_{\chi 1} \leftrightarrow y < t\}$$

(χ_1 corresponds to the first color) and note that $\lambda(U_t) = 1/2$. Define the $(T \cup T_{\text{LinOrder}} \cup T_{2\text{-Coloring}})$-on $\mathcal{H}^t$ over $[0, 1]^2$ by

$$\mathcal{H}_P^t \stackrel{\text{def}}{=} \mathcal{N}_P \times \mathcal{E}_{k(P)} \qquad (P \in \mathcal{L}), \qquad \mathcal{H}_\prec^t \stackrel{\text{def}}{=} \mathcal{E}_2 \times \mathcal{N}^{<},$$
$$\mathcal{H}_{\chi_1}^t \stackrel{\text{def}}{=} U_t, \qquad \mathcal{H}_{\chi_2}^t \stackrel{\text{def}}{=} (\mathcal{E}_1 \times \mathcal{E}_1) \setminus U_t.$$

Since $\phi_{\mathcal{H}^t}$ is a coupling of $\psi_{1/2}$ and $\phi \otimes \psi_{\text{lin}}$ and the latter satisfies `UInduce`[1], we get

$$\phi_{\mathcal{H}^t}(\langle K_m^{<} \rangle) = \frac{\phi(\langle K \rangle)}{m! \cdot 2^m}. \tag{4.20}$$

On the other hand, from the definition of $\mathcal{H}^t$, we have

$$\begin{aligned}\phi_{\mathcal{H}^t}(\langle K_m^{<}\rangle) &= \sum_{j=0}^{m} \frac{t^j(1-t)^{m-j}}{j!(m-t)!}\xi(\langle K_j\rangle)\\ &= \sum_{k=0}^{m}\left(\sum_{j=0}^{k}\frac{1}{j!(m-j)!}\binom{m-j}{k-j}(-1)^{k-j}\xi(\langle K_j\rangle)\right)t^k.\end{aligned}$$

Since this identity is true for any t, putting it together with (4.20) and comparing coefficients of the polynomials in t, we conclude that

$$\sum_{i=0}^{k}\frac{1}{i!(m-i)!}\binom{m-i}{k-i}(-1)^{k-i}\xi(\langle K_i\rangle) = \begin{cases}\dfrac{\phi(\langle K\rangle)}{m!\cdot 2^m}, & \text{if } k=0;\\ 0, & \text{if } k\in[m].\end{cases} \tag{4.21}$$

We can finally prove (4.19) by induction in $j \in \{0,\ldots,m\}$. For $j = 0$, the assertion follows from (4.21) for $k = 0$. Suppose then that $j \geq 1$ and by using the inductive hypothesis, note that (4.21) for $k = j$ gives

$$\begin{aligned}\xi(\langle K_j\rangle) &= -j!(m-j)!\sum_{i=0}^{j-1}\frac{1}{i!(m-i)!}\binom{m-i}{j-i}(-1)^{j-i}\frac{\phi(\langle K\rangle)}{2^m}\\ &= -\sum_{i=0}^{j-1}\binom{j}{i}(-1)^{j-i}\frac{\phi(\langle K\rangle)}{2^m} = \frac{\phi(\langle K\rangle)}{2^m}.\end{aligned}$$

Thus (4.19) holds. ■

4.6 Separations

In this section we prove all separation theorems.

Recall from Section 4.1.2 that for $x \in \mathcal{E}_n$, $\sigma_x \in S_n$ denotes the unique permutation such that $x_{\{\sigma_x^{-1}(1)\}} < \cdots < x_{\{\sigma_x^{-1}(n)\}}$ when the coordinates $(x_{\{i\}} \mid i \in [n])$ are distinct, and is defined arbitrarily otherwise.

Proof of Theorem 4.2.6. First note that the quasirandom $(\ell+1)$-tournamon $\psi_{\ell+1}$ can be represented by the $T_{(\ell+1)\text{-Tournament}}$-on

$$\mathcal{N} \stackrel{\text{def}}{=} \left\{ x \in \mathcal{E}_{\ell+1} \;\middle|\; x_{[\ell+1]} < \frac{1}{2} \leftrightarrow \operatorname{sgn}(\sigma_x) = 1 \right\}. \tag{4.22}$$

Let $\boldsymbol{K}$ be the exchangeable array corresponding to $\mathcal{N}$ with respect to $\boldsymbol{\theta}$ picked in $\mathcal{E}_{\mathbb{N}_+}$. By Theorem 4.2.10, to show that $\psi_{\ell+1} \in \texttt{UCouple}[\ell]$ it is sufficient to prove that $\psi_{\ell+1}$ is weakly ℓ-independent, that is for every $m \in \mathbb{N}$, the random variable $\boldsymbol{K}|_{[m]}$ is independent from $(\boldsymbol{\theta}_A \mid A \in r(m,\ell))$. Indeed, $\boldsymbol{K}|_{[m]}$ is completely determined by $\sigma_{\iota^*_{[m]}(\boldsymbol{\theta})}$ and $(\boldsymbol{\theta}_A \mid A \in \binom{[m]}{\ell+1})$, and any changes in the values of the signs $\operatorname{sgn}(\sigma_{\iota^*_A(\boldsymbol{\theta})})$ can be offset by flipping the corresponding variables $\boldsymbol{\theta}_A$ (cf. (4.22)) so that the *distribution* of $\boldsymbol{K}|_{[m]}$ does not change from fixing $\sigma_{\iota^*_{[m]}(\boldsymbol{\theta})}$.

Suppose now toward a contradiction that $\psi_{\ell+1} \in \texttt{Independence}[\ell]$, that is $\psi_{\ell+1} = \phi_{\mathcal{H}}$ for some $T_{(\ell+1)\text{-Tournament}}$-on $\mathcal{H}$ of the form $\mathcal{H} = \mathcal{E}_{\ell+1,\ell} \times \mathcal{G}$ for some $\mathcal{G} \subseteq [0,1]$. Note that for any $\sigma \in S_{\ell+1}$, we have $\mathcal{H} \cdot \sigma = \mathcal{H}$. But this is a contradiction as the axioms of $T_{k\text{-Tournament}}$ imply that $\lambda((\mathcal{H} \cdot \sigma) \cap \mathcal{H}) = 0$ whenever $\operatorname{sgn}(\sigma) = -1$. ∎

Proof of Theorem 4.2.7. Since ψ_{lin} is represented by the T_{LinOrder}-on $\mathcal{N} \stackrel{\text{def}}{=} \{x \in \mathcal{E}_2 \mid x_{\{1\}} < x_{\{2\}}\}$, we know $\operatorname{rk}(\psi_{\text{lin}}) = 1$, thus by Proposition 4.3.1, we have $\psi_{\text{lin}} \notin \texttt{UCouple}[1]$.

Since ψ_{lin} is n-categorical for every $n \in \mathbb{N}$, it is symmetrically ℓ-local for trivial reasons (namely, all events $\boldsymbol{K}|_{V_i} \cong M_i$ have probability 1), for any integer ℓ. Hence $\psi_{\text{lin}} \in \texttt{UInduce}[\ell]$ by Theorem 4.2.11. ∎

To prove Theorems 4.2.8 and 4.2.9, the alternating tournament defined below will play a key role.

Definition 4.6.1. Let $k \geq 1$. For $\alpha\colon [k] \rightarrowtail [k+1]$, denote by σ_α the unique extension of α to an element of S_{k+1}, and let $\operatorname{sgn}(\alpha) \stackrel{\text{def}}{=} \operatorname{sgn}(\sigma_\alpha)$. This definition behaves well with respect to the actions of S_k and S_{k+1}: for every $\eta \in S_k$ we have $\operatorname{sgn}(\alpha \circ \eta) = \operatorname{sgn}(\alpha)\operatorname{sgn}(\eta)$, and for every $\sigma \in S_{k+1}$ we have $\operatorname{sgn}(\sigma \circ \alpha) = \operatorname{sgn}(\sigma)\operatorname{sgn}(\alpha)$.

The *alternating k-tournament* is the model $A^{(k)}_{k+1} \in \mathcal{K}_{k+1}[T_{k\text{-Tournament}}]$ of $T_{k\text{-Tournament}}$ of size $k+1$ given by

$$V(A^{(k)}_{k+1}) \stackrel{\text{def}}{=} [k+1]; \qquad R_E(A^{(k)}_{k+1}) \stackrel{\text{def}}{=} \{\alpha \in ([k+1])_k \mid \operatorname{sgn}(\alpha) = 1\}.$$

For example, $A^{(2)}_3$ is the oriented cycle $\vec{C}_3$.

Proof of Theorem 4.2.8. For this proof, let us denote the predicate symbols of the theories $T_{(\ell+2)\text{-Hypergraph}}$ and $T_{(\ell+1)\text{-Tournament}}$ by E and P, respectively. Let also $\psi \stackrel{\text{def}}{=} \psi_{\ell+1} \in \operatorname{Hom}^+(\mathcal{A}[T_{(\ell+1)\text{-Tournament}}], \mathbb{R})$ be the quasirandom $(\ell+1)$-tournamon and let $I\colon T_{(\ell+2)\text{-Hypergraph}} \rightsquigarrow T_{(\ell+1)\text{-Tournament}}$ be given by

$$I(E)(x_1, \ldots, x_{\ell+2}) \stackrel{\text{def}}{=} \bigvee_{1 \le i_1 < \cdots < i_\ell \le \ell+2} (P(x_{i_1}, \ldots, x_{i_\ell}, x_{j_1}) \leftrightarrow P(x_{i_1}, \ldots, x_{i_\ell}, x_{j_2})),$$

where $j_1, j_2 \in [\ell+2]$ are such that $\{i_1, \ldots, i_\ell, j_1, j_2\} = [\ell+2]$. By Theorems 4.2.3 and 4.2.6, we know that $\phi \stackrel{\text{def}}{=} \psi^I \in \operatorname{Hom}^+(\mathcal{A}[T_{(\ell+2)\text{-Hypergraph}}], \mathbb{R})$ satisfies $\texttt{UCouple}[\ell]$.

To show that $\phi \notin \texttt{Independence}[\ell]$, we will make use of the theory T (isomorphic to $T_{(\ell+1)\text{-Tournament}}$) that is obtained from $T_{(\ell+2)\text{-Hypergraph}} \cup T_{(\ell+1)\text{-Tournament}}$ by adding the axiom

$$\forall \vec{x}, E(x_1, \ldots, x_{\ell+2}) \leftrightarrow I(E)(x_1, \ldots, x_{\ell+2}) \tag{4.23}$$

and the commutative diagram

$$\begin{array}{ccc} T_{(\ell+2)\text{-Hypergraph}} & \xrightarrow{\quad I \quad} & T_{(\ell+1)\text{-Tournament}} \\ {\scriptstyle S}\big\downarrow & & \big\downarrow{\scriptstyle J} \\ T_{(\ell+2)\text{-Hypergraph}} \cup T_{(\ell+1)\text{-Tournament}} & \xrightarrow[\quad A \quad]{} & T \end{array}$$

where S is the structure-erasing interpretation, A is the axiom-adding interpretation and J is the isomorphism mentioned above that acts identically on P (the inverse J^{-1} acts identically on P and acts as I on E). Let $\xi = \psi^{J^{-1}}$ so that $\psi = \xi^J$ and $\phi = \xi^{A\circ S}$.

Suppose toward a contradiction that $\phi \in \texttt{Independence}[\ell]$ and let $\mathcal{N}$ be an ℓ-independent $T_{(\ell+2)\text{-Hypergraph}}$-on over Ω such that $\phi_{\mathcal{N}} = \phi = \psi^I$. By Proposition 3.2.1, there exists a T-on $\mathcal{N}'$ over $\Omega \times \Omega$ such that $\phi_{\mathcal{N}'} = \xi$ and $S(A(\mathcal{N}'))_E = \mathcal{N}'_E = \mathcal{N}_E \times \mathcal{E}_{\ell+2}$ a.e. Note that $\operatorname{rk}(\phi) \leq \operatorname{rk}(\psi) \leq \ell + 1$, so by possibly changing zero-measure sets using Proposition 3.1.2, we may also suppose that $\operatorname{rk}(\mathcal{N}') \leq \ell + 1$. By applying a measure-isomorphism between $\Omega \times \Omega$ and $[0,1]$, we conclude that there exists a T-on $\mathcal{H}$ (over $[0,1]$) such that $\phi_{\mathcal{H}} = \xi$, $\operatorname{rk}(\mathcal{H}) \leq \ell + 1$ and the peon $\mathcal{H}_E$ is ℓ-independent.

Since $\mathcal{H}_E$ has rank at most $\ell + 1$ and is ℓ-independent, we can write it as $\mathcal{H}_E = \mathcal{E}_{\ell+2,\ell} \times \mathcal{G} \times [0,1]^{\{\ell+2\}}$ for some measurable $\mathcal{G} \subseteq [0,1]^{\binom{[\ell+2]}{\ell+1}}$. Using the symmetry axiom (2.2) of $T_{(\ell+2)\text{-Hypergraph}}$ and making a zero-measure change in $\mathcal{G}$, we may assume that it is $S_{\ell+2}$-invariant.

For every $t \in [\ell + 2]$, define the sets

$$V_t^{\ell+1} \stackrel{\text{def}}{=} \left\{ A \in \binom{[\ell+2]}{t} \;\middle|\; \ell+1 \in A \land \ell+2 \notin A \right\};$$
$$V_t^{\ell+2} \stackrel{\text{def}}{=} \left\{ A \in \binom{[\ell+2]}{t} \;\middle|\; \ell+1 \notin A \land \ell+2 \in A \right\};$$
$$V_t^{\ell+1,\ell+2} \stackrel{\text{def}}{=} \left\{ A \in \binom{[\ell+2]}{t} \;\middle|\; \ell+1, \ell+2 \in A \right\}.$$

Define also the sets

$$W_t^{\ell+1} \stackrel{\text{def}}{=} [0,1]^{V_t^{\ell+1}}; \qquad W_t^{\ell+2} \stackrel{\text{def}}{=} [0,1]^{V_t^{\ell+2}}; \qquad W_t^{\ell+1,\ell+2} \stackrel{\text{def}}{=} [0,1]^{V_t^{\ell+1,\ell+2}};$$
$$Y^{\ell+1} \stackrel{\text{def}}{=} \prod_{t=1}^{\ell} W_t^{\ell+1}; \qquad Y^{\ell+2} \stackrel{\text{def}}{=} \prod_{t=1}^{\ell} W_t^{\ell+2}; \qquad Z \stackrel{\text{def}}{=} \prod_{t=1}^{\ell+2} W_t^{\ell+1,\ell+2}.$$

Note that

$$\mathcal{E}_{\ell+1} = \mathcal{E}_\ell \times Y^{\ell+1} \times W^{\ell+1}_{\ell+1};$$
$$\mathcal{E}_{\ell+2} = \mathcal{E}_\ell \times Y^{\ell+1} \times W^{\ell+1}_{\ell+1} \times Y^{\ell+2} \times W^{\ell+2}_{\ell+1} \times Z.$$

Let $\iota\colon [\ell] \cup \{\ell+2\} \to [\ell+1]$ be the function that maps $\ell+2$ to $\ell+1$ and fixes all other points and note that ι induces maps $\iota^*\colon Y^{\ell+1} \to Y^{\ell+2}$ and $\iota^*_{\ell+1}\colon W^{\ell+1}_{\ell+1} \to W^{\ell+2}_{\ell+1}$ (given by $\iota^*(y)_A \overset{\text{def}}{=} y_{\iota(A)}$ and $\iota^*_{\ell+1}(w)_A \overset{\text{def}}{=} w_{\iota(A)}$).

For every $x \in \mathcal{E}_\ell$ and every $w \in W^{\ell+1}_{\ell+1}$, define the sections

$$\mathcal{H}^\alpha_P(x,w) \overset{\text{def}}{=} \{y \in Y^{\ell+1} \mid (x,y,w) \in \mathcal{H}_P\};$$
$$\mathcal{H}^\beta_P(x,w) \overset{\text{def}}{=} \{y \in Y^{\ell+1} \mid (x,y,w) \notin \mathcal{H}_P\};$$

and for every $x \in \mathcal{E}_\ell$, define

$$\mathcal{H}^\alpha_P(x) \overset{\text{def}}{=} \{w \in W^{\ell+1}_{\ell+1} \mid \lambda(\mathcal{H}^\alpha_P(x,w)) > 0\};$$
$$\mathcal{H}^\beta_P(x) \overset{\text{def}}{=} \{w \in W^{\ell+1}_{\ell+1} \mid \lambda(\mathcal{H}^\beta_P(x,w)) > 0\}.$$

It is clear that

$$\mathcal{H}^\alpha_P(x) \cup \mathcal{H}^\beta_P(x) = W^{\ell+1}_{\ell+1} \tag{4.24}$$

for every $x \in \mathcal{E}_\ell$.

Note that the axiom (4.23) of T and an application of Fubini's Theorem imply that for a.e. $x \in \mathcal{E}_\ell$, a.e. $w, \widehat{w} \in W^{\ell+1}_{\ell+1}$, a.e. $y \in \mathcal{H}^\alpha_P(x,w)$, a.e. $\widehat{y} \in \mathcal{H}^\alpha_P(x,\widehat{w})$ and a.e. $z \in Z$, we have

$$(x, y, w, \iota^*(\widehat{y}), \iota^*_{\ell+1}(\widehat{w}), z) \in \mathcal{H}_E. \tag{4.25}$$

Since the definition of $I(P)$ is invariant under negating P, the same assertion also holds with β in place of α.

Recalling that $\mathcal{H}_E = \mathcal{E}_{\ell+2,\ell} \times \mathcal{G} \times [0,1]^{\{\ell+2\}}$, (4.25) implies that for a.e. $x \in \mathcal{E}_\ell$, a.e. $w, \widehat{w} \in \mathcal{H}_P^\alpha(x)$ and a.e. $z \in W_{\ell+1}^{\ell+1,\ell+2}$, we have

$$(w, \iota_{\ell+1}^*(\widehat{w}), z) \in \mathcal{G}. \tag{4.26}$$

Again, the analogous statement with β in place of α also holds.

From (4.24) and (4.26), it follows that there exists $x_0 \in \mathcal{E}_\ell$ such that the following hold for $W^\alpha \stackrel{\text{def}}{=} \mathcal{H}_P^\alpha(x_0)$ and $W^\beta \stackrel{\text{def}}{=} \mathcal{H}_P^\beta(x_0)$.

i. We have $W^\alpha \cup W^\beta = W_{\ell+1}^{\ell+1}$.

ii. For a.e. $w, \widehat{w} \in W^\alpha$ and a.e. $z \in W_{\ell+1}^{\ell+1,\ell+2}$, we have $(w, \iota_{\ell+1}^*(\widehat{w}), z) \in \mathcal{G}$.

iii. For a.e. $w, \widehat{w} \in W^\beta$ and a.e. $z \in W_{\ell+1}^{\ell+1,\ell+2}$, we have $(w, \iota_{\ell+1}^*(\widehat{w}), z) \in \mathcal{G}$.

Since $|V_{\ell+1}^{\ell+1}| = 1$, let us for simplicity identify $W_{\ell+1}^{\ell+1}$ with $[0,1]$ and let $h \stackrel{\text{def}}{=} \mathbb{1}_{W^\alpha}$ be the indicator function of $W^\alpha \subseteq [0,1]$. For every $A \in \binom{[\ell+2]}{\ell+1}$, let $\pi_A \colon [0,1]^{\binom{[\ell+2]}{\ell+1}} \to [0,1]$ be the projection on the A-th coordinate and note that the properties above imply that for a.e. $u \in [0,1]^{\binom{[\ell+2]}{\ell+1}}$, if $h(\pi_{[\ell+1]}(u)) = h(\pi_{[\ell]\cup\{\ell+2\}})$, then $u \in \mathcal{G}$. Since $\mathcal{G}$ is $S_{\ell+2}$-invariant, this in turn implies that for a.e. $u \in [0,1]^{\binom{[\ell+2]}{\ell+1}}$, if there exist $j_1, j_2 \in [\ell+2]$ distinct such that $h(\pi_{[\ell+2]\setminus\{j_1\}}(u)) = h(\pi_{[\ell+2]\setminus\{j_2\}})$, then $u \in \mathcal{G}$. But since at least two of the values $h(\pi_{[\ell+1]}(u))$, $h(\pi_{[\ell+2]\setminus\{\ell+1\}}(u))$ and $h(\pi_{[\ell+2]\setminus\{\ell\}}(u))$ must be equal, it follows that $\lambda(\mathcal{G}) = 1$. So we must have

$$\phi(\rho_{\ell+2}) = \lambda(\mathcal{H}_E) = \lambda(\mathcal{G}) = 1,$$

which implies $\phi(\overline{K}_{\ell+2}^{(\ell+2)}) = 0$.

However, note that for the alternating $(\ell+1)$-tournament $A_{\ell+2}^{(\ell+1)}$, we have $I(A_{\ell+2}^{(\ell+1)}) \cong$

$\overline{K}^{(\ell+2)}_{\ell+2}$, hence

$$\phi(\overline{K}^{(\ell+2)}_{\ell+2}) \geq \psi(A^{(\ell+1)}_{\ell+2}) = \frac{(\ell+2)!}{2^{\ell+2}|\mathrm{Aut}(A^{(\ell+1)}_{\ell+2})|} = \frac{1}{2^{\ell+1}},$$

a contradiction. ■

The following is needed for the proof of Theorem 4.2.9.

Lemma 4.6.2. *If $M \in \mathcal{M}_{k+2}[T_{k\text{-Tournament}}]$ is a k-tournament on $k+2$ vertices, then M has at most two (unlabeled) copies of the alternating k-tournament $A^{(k)}_{k+1}$.*

Proof. Suppose toward a contradiction that $M \in \mathcal{M}_{k+2}[T_{k\text{-Tournament}}]$ contains three copies of $A^{(k)}_{k+1}$ and without loss of generality, let us suppose that these three copies are induced by $V_1 \stackrel{\text{def}}{=} [k+1]$, $V_2 \stackrel{\text{def}}{=} [k]\cup\{k+2\}$ and $V_3 \stackrel{\text{def}}{=} [k-1]\cup\{k+1,k+2\}$. Let $\alpha_{12}, \alpha_{13}, \alpha_{23} \in ([k+2])_k$ be given by

$$\alpha_{12}(v) \stackrel{\text{def}}{=} v; \qquad \alpha_{13}(v) \stackrel{\text{def}}{=} \begin{cases} v, & \text{if } v < k; \\ k+1 & \text{if } v = k; \end{cases} \qquad \alpha_{23}(v) \stackrel{\text{def}}{=} \begin{cases} v, & \text{if } v < k; \\ k+2 & \text{if } v = k; \end{cases}$$

and note that $\mathrm{im}(\alpha_{ij}) = V_i \cap V_j$.

But then $M|_{V_1} \cong A^{(k)}_{k+1}$, $M|_{V_2} \cong A^{(k)}_{k+1}$ and $M|_{V_3} \cong A^{(k)}_{k+1}$ imply respectively that

$$\alpha_{12} \in R_E(M) \iff \alpha_{13} \notin R_E(M),$$
$$\alpha_{12} \in R_E(M) \iff \alpha_{23} \notin R_E(M),$$
$$\alpha_{13} \in R_E(M) \iff \alpha_{23} \notin R_E(M).$$

This is a contradiction as all three equivalences above cannot be true at the same time. ■

Proof of Theorem 4.2.9. For this proof, let us again denote the predicate symbols of the theories $T_{(\ell+2)\text{-Hypergraph}}$ and $T_{(\ell+1)\text{-Tournament}}$ by E and P, respectively. For $p \in [0,1]$, let

$\mathcal{N}^p$ be the $T_{(\ell+1)\text{-Tournament}}$-on given by

$$\mathcal{N}^p_E \stackrel{\text{def}}{=} \left\{x \in \mathcal{E}_{\ell+1} \;\middle|\; x_{[\ell+1]} < p \leftrightarrow \operatorname{sgn}(\sigma_x) = 1\right\}$$

(note that for $p = 1/2$ this is precisely the theon (4.22) representing the quasirandom $(\ell+1)$-tournamon).

Let $I\colon T_{(\ell+2)\text{-Hypergraph}} \rightsquigarrow T_{(\ell+1)\text{-Tournament}}$ be the interpretation that declares $(\ell+2)$-edges to be isomorphic copies of $A^{(\ell+1)}_{\ell+2}$, and let $\phi_p \stackrel{\text{def}}{=} \phi^I_{\mathcal{N}^p} \in \operatorname{Hom}^+(\mathcal{A}[T_{(\ell+2)\text{-Hypergraph}}], \mathbb{R})$. We will show that ϕ_p satisfies $\texttt{UInduce}[\ell]$ for every $p \in [0,1]$, but does not satisfy $\texttt{UCouple}[1]$ unless $p \in \{0, 1/2, 1\}$.

To show the first property, recall that the quasirandom $(\ell+1)$-hypergraphon $\psi_{\ell+1,p} \in \operatorname{Hom}^+(\mathcal{A}[T_{(\ell+1)\text{-Hypergraph}}], \mathbb{R})$ satisfies $\texttt{Independence}[\ell]$ (cf. Lemma 4.5.4) and hence also satisfies $\texttt{UCouple}[\ell]$ (by Theorem 4.2.2). Note also that $\phi_{\mathcal{N}^p} = (\psi_{\ell+1,p} \otimes \psi_{\text{lin}})^{I'}$ and $I'\colon T_{(\ell+1)\text{-Tournament}} \rightsquigarrow T_{(\ell+1)\text{-Hypergraph}} \cup T_{\text{LinOrder}}$ is given by[4]

$$\begin{aligned} I'(P)(x_1,\ldots,x_{\ell+1}) \stackrel{\text{def}}{=} & \left(\bigwedge_{1\le i<j\le \ell+1} x_i \neq x_j\right) \\ & \wedge \left(E(x_1,\ldots,x_{\ell+1}) \leftrightarrow \bigvee_{\substack{\sigma\in S_{\ell+1}\\ \operatorname{sgn}(\sigma)=1}} \bigwedge_{1\le i<j\le \ell+1} x_{\sigma(i)} \prec x_{\sigma(j)}\right). \end{aligned}$$

By Theorem 4.2.10(i) $\Longrightarrow$ (vii), we know that $\psi_{\ell+1,p}\otimes\psi_{\text{lin}} \in \texttt{UInduce}[\ell]$ and by Theorem 4.2.3, we get that $\phi_p = (\psi_{\ell+1,p} \otimes \psi_{\text{lin}})^{I'\circ I}$ satisfies $\texttt{UInduce}[\ell]$.

Let us now show that for every $p \in (0,1) \setminus \{1/2\}$, ϕ_p does not satisfy $\texttt{UCouple}[1]$. Since ψ_{lin} has rank 1, it is enough to show that ϕ_p is not uniquely coupleable with ψ_{lin}. Consider

4. This is a generalization of the "arc-orientation" interpretation used implicitly in the implications $P_{10} \Longrightarrow P_{11} \Longrightarrow P_1(s)$ of [15].

the $(T_{(\ell+1)\text{-Tournament}} \cup T_{\text{LinOrder}})$-on $\mathcal{N}^{p,<}$ given by

$$\mathcal{N}^{p,<}_P \stackrel{\text{def}}{=} \mathcal{N}^p_P; \qquad\qquad \mathcal{N}^{p,<}_{\prec} \stackrel{\text{def}}{=} \{x \in \mathcal{E}_2 \mid x_{\{1\}} < x_{\{2\}}\}$$

and note that $\phi_{\mathcal{N}^{p,<}}$ is a coupling of $\phi_{\mathcal{N}^p}$ and ψ_{lin}, hence $\xi \stackrel{\text{def}}{=} \phi^{I\cup\text{id}_{T_{\text{LinOrder}}}}_{\mathcal{N}^{p,<}}$ is a coupling of ϕ_p and ψ_{lin}. We will show that $\xi \neq \phi_p \otimes \psi_{\text{lin}}$ by a direct computation exhibiting an $(\ell+2)$-hypergraph H and two different orders on it such that $\xi(H_1) \neq \xi(H_2)$ for the corresponding models of the theory $T_{(\ell+2)\text{-Hypergraph}} \cup T_{\text{LinOrder}}$. That will suffice since, clearly, $(\phi_p \otimes \psi_{\text{lin}})(H_1) = (\phi_p \otimes \psi_{\text{lin}})(H_2)$.

Let $H \in \mathcal{K}_{\ell+3}[T_{(\ell+2)\text{-Hypergraph}}]$ be given by

$$V(H) \stackrel{\text{def}}{=} [\ell+3]; \qquad\qquad E(H) \stackrel{\text{def}}{=} \{[k+1], [k] \cup \{k+2\}\};$$

and let $H_1, H_2 \in \mathcal{K}_{\ell+3}[T_{(\ell+2)\text{-Hypergraph}} \cup T_{\text{LinOrder}}]$ be obtained from H by equipping it with the orders $\prec_1$ and $\prec_2$, respectively, where $\prec_1$ is the natural order of $[\ell+3]$ and $\prec_2$ is obtained from $\prec_1$ by swapping the order position of $\ell+1$ and $\ell+3$, that is, we have

$$1 \prec_2 2 \prec_2 \cdots \prec_2 \ell \prec_2 \ell+3 \prec_2 \ell+2 \prec_2 \ell+1.$$

Let $\boldsymbol{\theta}$ be picked at random in $\mathcal{E}_{\mathbb{N}_+}$ according to λ and let $\boldsymbol{K}$ be the exchangeable array corresponding to $\mathcal{N}^{p,<}$ with respect to $\boldsymbol{\theta}$ (so that $(I \cup \text{id}_{T_{\text{LinOrder}}})(\boldsymbol{K})$ corresponds to $(I \cup \text{id}_{T_{\text{LinOrder}}})(\mathcal{N}^{p,<})$). Let $\boldsymbol{\sigma} \stackrel{\text{def}}{=} \sigma_{\iota^*_{[\ell+3]}(\boldsymbol{\theta})}$. Then we have

$$\xi(\langle H_1 \rangle) = \mathbb{P}[I(J(\boldsymbol{K}|_{[\ell+3]})) = H \wedge \boldsymbol{\sigma} = \text{id}_{\ell+3}];$$
$$\xi(\langle H_2 \rangle) = \mathbb{P}[I(J(\boldsymbol{K}|_{[\ell+3]})) = H \wedge \boldsymbol{\sigma} = \tau];$$

where $J\colon T_{(\ell+1)\text{-Tournament}} \rightsquigarrow T_{(\ell+1)\text{-Tournament}} \cup T_{\text{LinOrder}}$ is the structure-erasing interpretation and τ is the transposition that swaps $\ell+1$ and $\ell+3$. Then by Lemma 4.6.2,

$I(J(\boldsymbol{K}|_{[\ell+3]})) = H$ is equivalent to

$$J(\boldsymbol{K}|_{[\ell+2]}) \cong J(\boldsymbol{K}|_{[\ell+1]\cup\{\ell+3\}}) \cong A_{\ell+2}^{(\ell+1)}. \tag{4.27}$$

Since $\operatorname{Aut}(A_{\ell+2}^{(\ell+1)})$ is the alternating group on $[\ell+2]$, on any fixed set of $\ell+2$ vertices, there are exactly two models M_1 and M_2 that are isomorphic to $A_{\ell+2}^{(\ell+1)}$ and they satisfy $R_P(M_1) \cap R_P(M_2) = \varnothing$. This means that on the event (4.27), out of the a priori four presentations of $A_{\ell+2}^{(\ell+1)}$ induced on $[\ell+2]$ and $[\ell+1]\cup\{\ell+3\}$, only two are actually possible. Since ℓ is odd, a straightforward calculation gives

$$\xi(\langle H_1\rangle) = p^{(\ell+2)}(1-p)^{\ell+1} + p^{\ell+1}(1-p)^{\ell+2} = p^{\ell+1}(1-p)^{\ell+1};$$
$$\xi(\langle H_2\rangle) = p^{\ell}(1-p)^{\ell+3} + p^{\ell+3}(1-p)^{\ell} = p^{\ell}(1-p)^{\ell}(3p^2-3p+1).$$

Thus we get

$$\begin{aligned}\xi(\langle H_2\rangle) - \xi(\langle H_1\rangle) &= p^{\ell}(1-p)^{\ell}(4p^2-4p+1)\\ &= p^{\ell}(1-p)^{\ell}\,(2p-1)^2\,,\end{aligned}$$

which is non-zero as long as $p \in (0,1)\setminus\{1/2\}$. ■

Proof of Theorem 4.2.16. For $p \in (0,1)$, let $\mathcal{N}$ be the $T_{k\text{-Hypergraph}}$-on given by

$$\begin{aligned}\mathcal{N} \stackrel{\text{def}}{=} \Bigg\{x \in \mathcal{E}_k \;\Bigg|\; &(\min\{x_{\{v\}} \mid v\in[k]\} < 1/2 \wedge x_{[k]} < p)\\ &\vee \left(\min\{x_{\{v\}} \mid v \in [k]\} \geq 1/2 \wedge \sum_{v\in[k]} x_{[k]\setminus\{v\}} \bmod 1 < p\right)\Bigg\}.\end{aligned}$$

Let us show that $\phi \stackrel{\text{def}}{=} \phi_{\mathcal{N}}$ satisfies $\texttt{Dev}[k-1]$; recall that $\texttt{Dev}[k-1] = \texttt{Disc}[\mathcal{A}_{k-1}]$, where $\mathcal{A}_{k-1} \stackrel{\text{def}}{=} \{A \in \binom{[k]}{k-1} \mid \{1\} \subseteq A\} = \binom{[k]}{k-1} \setminus \{[k]\setminus\{1\}\}$ (see Definition 4.1.4) and for $\psi \in \operatorname{Hom}^+(\mathcal{A}[T_{\mathcal{L}_{\mathcal{A}_{k-1}}}], \mathbb{R})$, let ξ be a coupling of ϕ and ψ. By Proposition 3.2.1, there exists a $(T \cup T_{\mathcal{L}_{\mathcal{A}_{k-1}}})$-on $\mathcal{H}$ over $[0,1]^2$ such that $\phi_{\mathcal{H}} = \xi$ and $\mathcal{H}_E = \mathcal{N} \times \mathcal{E}_k$.

Let $(\boldsymbol{\theta^1}, \boldsymbol{\theta^2})$ be picked in $\mathcal{E}_{\mathbb{N}_+}([0,1]^2)$ according to λ and let $\boldsymbol{K}$ be the exchangeable array corresponding to $\mathcal{H}$ with respect to $(\boldsymbol{\theta^1}, \boldsymbol{\theta^2})$. Our objective is to show that the events $(1, 2, \ldots, k) \in R_E(\boldsymbol{K})$ and $\forall A \in \mathcal{A}_{k-1}, \iota_A \in R_{P_A}(\boldsymbol{K})$ are independent.

Since the event $\iota_A \in R_{P_A}(\boldsymbol{K})$ is completely determined by $((\boldsymbol{\theta^1_B}, \boldsymbol{\theta^2_B}) \mid B \subseteq A)$, it is sufficient to show that the event $(1, \ldots, k) \in R_E(\boldsymbol{K})$ is independent from $((\boldsymbol{\theta^1_B}, \boldsymbol{\theta^2_B}) \mid B \in r(k, k-1) \wedge B \neq [k] \setminus \{1\})$. But the event $(1, \ldots, k) \in R_E(\boldsymbol{K})$ is equivalent to $(\boldsymbol{\theta^1_B})_{B \in r(k)} \in \mathcal{N}$, and it is easy to see that the conditional probability of $(1, \ldots, k) \in R_E(\boldsymbol{K})$ given $((\boldsymbol{\theta^1_B}, \boldsymbol{\theta^2_B}) \mid B \in r(k, k-1) \wedge B \neq [k] \setminus \{1\})$ is p a.e. Hence ϕ satisfies $\texttt{Dev}[k-1]$.

Let us now show that ϕ does not satisfy $\texttt{UInduce}[1]$. To do so, for each $i \in [2]$, we consider the $(T_{k\text{-Hypergraph}} \cup T_{2\text{-Coloring}})$-on $\mathcal{H}^i$ (see Remark 8) given by

$$\begin{aligned} \mathcal{H}_E &= \mathcal{N}; \\ \mathcal{H}_{\chi_i} &= \{x \in \mathcal{E}_1 \mid x_{\{1\}} < 1/2\}; \\ \mathcal{H}_{\chi_{3-i}} &= \{x \in \mathcal{E}_1 \mid x_{\{1\}} \geq 1/2\}. \end{aligned}$$

Then by a straightforward calculation, for every $H \in \mathcal{M}[T_{k\text{-Hypergraph}} \cup T_{2\text{-Coloring}}]$ with $R_{\chi_1}(H) = V(H)$, we have

$$\phi_{\mathcal{H}^1}(H) = \frac{\psi_{k,p}(I(H))}{2^{|H|}}; \qquad \phi_{\mathcal{H}^2}(H) = \frac{\phi_{\mathcal{N}'}(I(H))}{2^{|H|}};$$

where $I \colon T_{k\text{-Hypergraph}} \rightsquigarrow T_{k\text{-Hypergraph}} \cup T_{2\text{-Coloring}}$ is the structure-erasing interpretation, $\psi_{k,p}$ is the quasirandom k-hypergraphon (see Definition 4.1.7) and $\mathcal{N}'$ is the $T_{k\text{-Hypergraph}}$-on given by

$$\mathcal{N}' = \left\{ x \in \mathcal{E}_k \;\middle|\; \sum_{v \in [k]} x_{[k] \setminus \{v\}} \bmod 1 < p \right\}.$$

Since $\phi_{\mathcal{N}'} \neq \psi_{k,p}$ (since $\operatorname{rk}(\psi_{k,p}) = k > k-1 \geq \operatorname{rk}(\psi_{\mathcal{N}'})$), it follows that $\phi_{\mathcal{H}^1}(H) \neq \phi_{\mathcal{H}^2}(H)$ for some $H \in \mathcal{M}[T_{k\text{-Hypergraph}} \cup T_{2\text{-Coloring}}]$ with $R_{\chi_1}(H) = V(H)$, hence ϕ does not satisfy $\texttt{UInduce}[1]$. ■

Proof of Theorem 4.2.17. For $p \in (0,1)$, let $\mathcal{N}$ be the $T_{k\text{-Hypergraph}}$-on given by

$$\mathcal{N} \stackrel{\text{def}}{=} \left\{ x \in \mathcal{E}_k \;\middle|\; \max\left\{ x_A \;\middle|\; A \in \binom{[k]}{\ell+1} \right\} < p \right\}.$$

It is clear that $\phi \stackrel{\text{def}}{=} \phi_{\mathcal{N}}$ satisfies $\texttt{Independence}[\ell]$. Consider now the $T_{\mathcal{L}_{\{[\ell+1]\}}}$-on $\mathcal{H}$ given by

$$\mathcal{H}_E \stackrel{\text{def}}{=} \mathcal{N}; \qquad\qquad \mathcal{H}_{P_{[\ell+1]}} \stackrel{\text{def}}{=} \{x \in \mathcal{E}_{\ell+1} \mid x_{[\ell+1]} \geq p\}$$

and note that if $\boldsymbol{K}$ is the exchangeable array corresponding to $\mathcal{H}$, then

$$\begin{aligned}\mathbb{P}[(1,\ldots,k) \in R_E(\boldsymbol{K}) \wedge (1,\ldots,\ell+1) \in R_{P_{[\ell+1]}}(\boldsymbol{K})] = 0 \\ \neq p^{\binom{k}{\ell+1}} \cdot (1-p) = \phi(\rho_k) \cdot \mathbb{P}[(1,\ldots,\ell+1) \in R_{P_{[\ell+1]}}(\boldsymbol{K})],\end{aligned}$$

so ϕ does not satisfy $\texttt{Disc}[\{[\ell+1]\}]$. ■

Proof of Theorem 4.2.5. Follows from Theorems 4.2.15 ($\texttt{UInduce}[\ell+1] \implies \texttt{CliqueDisc}[\ell+1]$) and 4.2.17 ($\texttt{Independence}[\ell] \not\Rightarrow \texttt{Disc}[\{[\ell+1]\}]$), and the fact that $\texttt{CliqueDisc}[\ell+1] \implies \texttt{Disc}[\{[\ell+1]\}]$ (see [66, 1]). ■

4.7 Top level quasirandomness

In this section we prove Theorems 4.2.12 and 4.2.13, which completely characterize the properties $\texttt{Independence}[k-1]$ and $\texttt{UCouple}[k-1]$, respectively when all arities are at most k. These can be seen as analogues of full quasirandomness for arbitrary universal theories (just as $\texttt{Dev}[k] = \texttt{CliqueDisc}[k-1] = \texttt{Disc}[\binom{[k]}{k-1}]$ gives full quasirandomness in $T_{k\text{-Hypergraph}}$).

We also show the weaker Theorem 4.2.14, which is an analogue of the above for `UInduce` when $k = 2$.

Proof of Theorem 4.2.12. By Lemma 4.5.4, we know that $\psi_{k,p} \in \mathrm{Hom}^+(\mathcal{A}[T_{c,k}], \mathbb{R})$ satisfies `Independence`$[k-1]$, so the backward direction follows from Theorem 4.2.3.

For the forward direction, first we claim that it is enough to show the case when $T = T_{\mathcal{L}}$. (This is not completely immediate as $I\colon T \rightsquigarrow T_{c,k}$ is required to satisfy $T_{c,k} \vdash \forall \vec{x}, I(F)(\vec{x})$ for every axiom $\forall \vec{x}, F(\vec{x})$ of T.) Let $A\colon T_{\mathcal{L}} \rightsquigarrow T$ be the axiom-adding interpretation and suppose ϕ^A (which satisfies `UCouple`$[k-1]$ by Theorem 4.2.3) can be written as $\phi^A = \psi_{k,p}^J$ for some $c \geq 2$, some $p \in \Pi_c$ and some $J\colon T_{\mathcal{L}} \rightsquigarrow T_{c,k}$, then we define $I\colon T \rightsquigarrow T_{c,k}$ to act as J and we have to show that it is indeed an interpretation, i.e., that $T_{c,k} \vdash \forall \vec{x}, I(F)(\vec{x})$ for every axiom $\forall \vec{x}, F(\vec{x})$ of T ($\psi_{k,p}^I = \phi$ will then follow trivially). Equivalently, we have to show that if $M \in \mathcal{M}[T_{c,k}]$, then $J(M) \in \mathcal{M}[T]$. But since all p_i are positive, we have $\psi_{k,p}(M) > 0$, so $\phi^A(J(M)) > 0$, hence trivially $J(M) \in \mathcal{M}[T]$.

Let us now prove the case $T = T_{\mathcal{L}}$. Let $\mathcal{N}$ be a $(k-1)$-independent $T_{\mathcal{L}}$-on such that $\phi_{\mathcal{N}} = \phi$. Note that if $P \in \mathcal{L}$ is such that $k(P) \leq k-1$, then $\mathcal{N}_P$ must be either $\varnothing$ or $\mathcal{E}_{k(P)}$, so we can write $\mathcal{L} = \mathcal{L}' \cup \mathcal{L}_0 \cup \mathcal{L}_1$, where

$$\mathcal{L}' \stackrel{\text{def}}{=} \{P \in \mathcal{L} \mid k(P) = k\};$$
$$\mathcal{L}_0 \stackrel{\text{def}}{=} \{P \in \mathcal{L} \mid k(P) \leq k-1 \land \mathcal{N}_P = \varnothing\};$$
$$\mathcal{L}_1 \stackrel{\text{def}}{=} \{P \in \mathcal{L} \mid k(P) \leq k-1 \land \mathcal{N}_P = \mathcal{E}_{k(P)}\}.$$

Recall from Definition 4.4.8 that $\mathcal{K}_k[\mathrm{Th}(\phi)] = \{K \in \mathcal{K}_k[T_{\mathcal{L}}] \mid \phi(K) > 0\}$ and enumerate its elements as $K_1, \ldots, K_c$. Note that since $\mathcal{N}$ is $(k-1)$-independent, it follows that every peon $\mathcal{N}_P$ with $P \in \mathcal{L}'$ is S_k-invariant, hence we must have $\mathrm{Aut}(K_i) = S_k$ for every $i \in [c]$. Suppose first that $c \geq 2$ and define $p \in \Pi_c$ by $p_i = \phi(K_i) > 0$ and let $I\colon T_{\mathcal{L}} \rightsquigarrow T_{c,k}$ be given by

$$I(P)(x_1,\ldots,x_{k(P)}) \stackrel{\text{def}}{=} x_1 \neq x_1 \qquad (P \in \mathcal{L}_0);$$
$$I(P)(x_1,\ldots,x_{k(P)}) \stackrel{\text{def}}{=} \bigwedge_{1\leq i<j\leq k(P)} x_i \neq x_j \qquad (P \in \mathcal{L}_1);$$
$$I(P)(x_1,\ldots,x_k) \stackrel{\text{def}}{=} \bigvee_{\substack{i\in[c]\\ \mathrm{id}_k\in R_P(K_i)}} E_i(x_1,\ldots,x_k). \qquad (P \in \mathcal{L}'). \tag{4.28}$$

Since $\mathcal{N}$ is $(k-1)$-independent, it follows that each $T_{\mathrm{ind}}(K_i,\mathcal{N})$ is $(k-1)$-independent and has measure p_i, which implies that the $T_{c,k}$-on $\mathcal{H}$ defined by $\mathcal{H}_{E_i} \stackrel{\text{def}}{=} T_{\mathrm{ind}}(K_i,\mathcal{N})$ $(i \in [c])$ satisfies $\phi_{\mathcal{H}} = \psi_{k,p}$ and since clearly $I(\mathcal{H}) = \mathcal{N}$, it follows that $\psi^I_{k,p} = \phi$.

If $c = 1$, then we can define I by replacing (4.28) with

$$I(P)(x_1,\ldots,x_k) \stackrel{\text{def}}{=} \bigwedge_{1\leq i<j\leq k(P)} x_i \neq x_j \qquad (P \in \mathcal{L}', \mathrm{id}_k \in R_P(K_1));$$
$$I(P)(x_1,\ldots,x_k) \stackrel{\text{def}}{=} x_1 \neq x_1 \qquad (P \in \mathcal{L}', \mathrm{id}_k \notin R_P(K_1))$$

instead and we trivially get $\phi = \psi^I_{k,p}$ for any $p \in \Pi_{c'}$ with $c' \geq 2$ as we must have $T_{\mathrm{ind}}(K_1,\mathcal{N}) = \mathcal{E}_k$ a.e. ■

Before we show Theorem 4.2.13, let us first see that the (Θ, p)-quasirandom homomorphisms $\psi_{\Theta,p} \in \mathrm{Hom}^+(\mathcal{A}[T_\Theta],\mathbb{R})$ from Definition 4.1.5 are well-defined (i.e., their definition as $\psi_{\Theta,p} \stackrel{\text{def}}{=} \phi_{\mathcal{N}^Z}$ is independent of the choice of Z) and satisfy `UCouple`$[k-1]$ when p is Θ-invariant.

Proposition 4.7.1. *With the notation and conditions of Definition 4.1.5, $\mathcal{N}^Z$ is a T_Θ-on.*

Furthermore, if p is Θ-invariant, then

$$\phi_{\mathcal{N}^Z}(\langle M\rangle) = \prod_{P\in\mathcal{L}} p_P^{|R_P(M)|/k!} \tag{4.29}$$

for every $M \in \mathcal{M}[T_\Theta]$ *and* $\psi_{\Theta,p} \stackrel{\text{def}}{=} \phi_{\mathcal{N}^Z}$ *satisfies* $\texttt{UCouple}[k-1]$.

Proof. First, let us show that $\mathcal{N}^Z$ is indeed a T_Θ-on.

Note first that T_Θ trivially proves that

$$\neg P(x, y, \ldots, t) \qquad (P \in \mathcal{L}, \text{the tuple } (x, y, \ldots, t) \text{ contains repeated variables}) \tag{4.30}$$

and if we add (4.30) to the axioms of T_Θ, then it becomes substitutionally closed, then by Theorem 2.4.1, to show that $\mathcal{N}^Z$ is a T_Θ-on, it is enough to show that $\mathcal{N}^Z$ satisfies the axioms of T_Θ and (4.30) a.e. It is trivial that $\mathcal{N}^Z$ satisfies (4.30) a.e.

Note that the fact that Z is a partition implies that there exists a unique $P_x \in \mathcal{L}$ such that $x_{[k]} \in Z_{P_x}$, thus there exists a unique $Q_x \in \mathcal{L}$ such that $x \in \mathcal{N}^Z_{Q_x}$, namely $Q_x = \sigma_x^{-1} \cdot P_x$ (where σ_x is as in Definition 4.1.5). This implies that $\mathcal{N}^Z$ satisfies axioms (4.1) and (4.3) a.e.

Note now that if $\tau \in S_k$, then we have $\sigma_{x\cdot\tau} = \sigma_x \circ \tau$, hence

$$x \cdot \tau \in \mathcal{N}^Z_P \iff x_{[k]} \in Z_{\sigma_{x\cdot\tau}\cdot P} \iff x_{[k]} \in Z_{\sigma_x \cdot (\tau \cdot P)} \iff x \in \mathcal{N}^Z_{\tau\cdot P},$$

so $\mathcal{N}^Z$ also satisfies axiom (4.2) a.e., hence $\mathcal{N}^Z$ is a T_Θ-on.

Suppose now that p is Θ-invariant and let us prove (4.29). Let $\boldsymbol{K}$ be the exchangeable array corresponding to $\mathcal{N}^Z$ with respect to $\boldsymbol{\theta}$ picked in $\mathcal{E}_{\mathbb{N}_+}$ according to λ. Since for $m \in \mathbb{N}$ and $K \in \mathcal{K}_m[T_\Theta]$, we have $\phi_{\mathcal{N}^Z}(\langle K \rangle) = \mathbb{P}[\boldsymbol{K}|_{[m]} = K]$, if we show that for every measurable $U \subseteq \mathcal{E}_{m,k-1}$ with $\lambda(U) > 0$, we have

$$\mathbb{P}[\boldsymbol{K}|_{[m]} = K \mid E] = \prod_{P \in \mathcal{L}} p_P^{|R_P(K)|/k!}, \tag{4.31}$$

where E is the event $(\boldsymbol{\theta}_B \mid B \in r(m, k-1)) \in U$, then both (4.29) and $\psi_{\Theta,p} \in \texttt{UCouple}[k-1]$ will follow (the former follows by taking $U = \mathcal{E}_{m,k-1}$ and the latter implies weak $(k-1)$-independence of $\mathcal{N}^Z$, which is equivalent to $\phi_{\mathcal{N}^Z} \in \texttt{UCouple}[k-1]$ by Theorem 4.2.10).

If $m < k$, (4.31) trivially holds, so suppose $m \geq k$ and note that the axioms of T_Θ imply that for each $\alpha\colon [k] \rightarrowtail [m]$, there exists a unique $P_\alpha \in \mathcal{L}$ such that $\alpha \in R_{P_\alpha}(K)$ and we must further have $P_\alpha = \tau \cdot P_{\alpha\circ\tau}$ for every $\tau \in S_k$. Note that for any choice of $(\alpha_A)_{A\in\binom{[m]}{k}}$ with $\alpha_A\colon [k] \rightarrowtail [m]$ and $\operatorname{im}(\alpha_A) = A$, we have

$$\begin{aligned}\mathbb{P}[\boldsymbol{K}|_{[m]} = K \mid E] &= \mathbb{P}\left[\forall\alpha \in ([m])_k, \alpha \in R_{P_\alpha}(\boldsymbol{K}) \;\middle|\; E\right] \\ &= \mathbb{P}\left[\forall A \in \binom{[m]}{k}, \alpha_A \in R_{P_{\alpha_A}}(\boldsymbol{K}) \;\middle|\; E\right].\end{aligned}$$

Now, the event $\alpha_A \in R_{P_{\alpha_A}}(\boldsymbol{K})$ depends only on the relative order of $(\boldsymbol{\theta}_{\{i\}} \mid i \in A)$ and on the variable $\boldsymbol{\theta}_A$ and, since p is Θ-invariant, we have $\lambda(Z_{\sigma\cdot P_\alpha}) = p_{P_\alpha}$ for every $\sigma \in S_k$ and every $\alpha\colon [k] \rightarrowtail [m]$. This means that if $\leq$ is an ordering of A and $E_\leq$ is the event that says that the relative order of $(\boldsymbol{\theta}_{\{i\}} \mid i \in A)$ is $\leq$, then $\mathbb{P}[\alpha \in R_{P_\alpha}(\boldsymbol{K}) \mid E \wedge E_\leq] = p_{P_\alpha}$ and thus

$$\mathbb{P}[\boldsymbol{K}|_{[m]} = K \mid E] = \prod_{A\in\binom{[m]}{k}} p_{P_{\alpha_A}}.$$

Since this holds for any choice of $(\alpha_A)_{A\in\binom{[m]}{k}}$ with $\operatorname{im}(\alpha_A) = A$, by considering all possible $k!^{\binom{m}{k}}$ such choices we get

$$\mathbb{P}[\boldsymbol{K}|_{[m]} = K \mid E]^{k!^{\binom{m}{k}}} = \prod_{P\in\mathcal{L}} p_P^{k!^{\binom{m}{k}-1}\cdot|R_P(K)|},$$

from which (4.31) follows. ■

Definition 4.7.2. Given a T-on $\mathcal{N}$ over $\Omega = (X, \mathcal{A}, \mu)$ and $K \in \mathcal{K}_V[T]$, we define the function $W^K_{\mathcal{N}}\colon \mathcal{E}_{V,|V|-1}(\Omega) \to [0, 1]$ by

$$W^K_{\mathcal{N}}(x) \stackrel{\text{def}}{=} \mu(\{y \in X \mid (x, y) \in T_{\text{ind}}(K, \mathcal{N})\}).$$

Note that $W^K_{\mathcal{N}}$ is essentially a $(|V| - 1)$-flattening of the peon $T_{\text{ind}}(K, \mathcal{N}) \subseteq \mathcal{E}_V(\Omega)$ (see

Definition 4.4.4).

The next two simple lemmas are fundamental in the proof of Theorem 4.2.13.

Lemma 4.7.3. *Let $k \in \mathbb{N}_+$ and suppose that $k(P) \leq k$ for all $P \in \mathcal{L}$. Let T be a theory over $\mathcal{L}$ and $\mathcal{N}$ be a T-on over $\Omega = (X, \mathcal{A}, \mu)$. Then for every $m \in \mathbb{N}$ and every $K \in \mathcal{K}_m[T]$, we have*

$$\phi_{\mathcal{N}}(\langle K \rangle) = \int_{X^{r(m,k-1)}} \prod_{A \in \binom{[m]}{k}} W_{\mathcal{N}}^{K|_A}(\pi_A(x))\ d\mu(x),$$

where $\pi_A \colon \mathcal{E}_{m,k-1}(\Omega) \to \mathcal{E}_{A,k-1}(\Omega)$ is the projection on the coordinates indexed by $r(A, k-1)$.

Proof. Follows by considering the exchangeable array corresponding to $\mathcal{N}$ with respect to $\boldsymbol{\theta}$ picked in $\mathcal{E}_{\mathbb{N}_+}(\Omega)$ according to μ, noting that $\boldsymbol{K}|_{[m]} = K$ is equivalent to $\forall A \in \binom{[m]}{k}, \boldsymbol{K}|_A = K|_A$ (since $k(P) \leq k$ for every $P \in \mathcal{L}$) and integrating out the top variables $(\boldsymbol{\theta}_A \mid A \in \binom{[m]}{k})$. ■

Lemma 4.7.4. *If a T-on $\mathcal{N}$ over Ω is such that $\phi_{\mathcal{N}}$ satisfies `UCouple`$[\ell]$ and $K \in \mathcal{K}_V[T]$ with $|V| \leq \ell + 1$, then $W_{\mathcal{N}}^K$ is a.e. constant.*

Proof. Without loss of generality, we may suppose that $V = [m]$. Write $\Omega = (X, \mathcal{A}, \mu)$. Then it is sufficient to show that for every measurable $U \in \mathcal{E}_{m,\ell}(\Omega)$, we have $\int_U W_{\mathcal{N}}^K\ d\mu = \mu(U)\phi_{\mathcal{N}}(\langle K \rangle)$. But for the exchangeable array $\boldsymbol{K}$ corresponding to $\mathcal{N}$ with respect to $\boldsymbol{\theta}$ picked in $\mathcal{E}_{\mathbb{N}_+}(\Omega)$ according to μ, it follows that

$$\begin{aligned}\int_U W_{\mathcal{N}}^K\ d\mu &= \mathbb{P}[\boldsymbol{K}|_{[m]} = K \wedge (\boldsymbol{\theta}_A \mid A \in r(m, k-1)) \in U] \\ &= \mathbb{P}[\boldsymbol{K}|_{[m]} = K] \cdot \mathbb{P}[(\boldsymbol{\theta}_A \mid A \in r(m, k-1)) \in U] = \mu(U)\phi_{\mathcal{N}}(\langle K \rangle),\end{aligned}$$

where the second equality follows since $\mathcal{N}$ is weakly ℓ-independent by Theorem 4.2.10. ■

Proof of Theorem 4.2.13. The backward direction follows from Proposition 4.7.1 and Theorem 4.2.3.

For the forward direction, we will show that in fact we can take $p = (p_P)_{P\in\mathcal{L}}$ satisfying $p_P > 0$ for every $P \in \mathcal{L}$. Note that when $p_P > 0$ for every $P \in \mathcal{L}$, we have $\psi_{\Theta,p}(M) > 0$ for every $M \in \mathcal{M}[T_\Theta]$, so by an argument analogous to that of the proof of Theorem 4.2.12, it is enough to consider the case when $T = T_{\mathcal{L}}$.

Suppose then that $T = T_{\mathcal{L}}$ and let $\mathcal{N}$ be a T-on such that $\phi_{\mathcal{N}} = \phi$. Note that if $P \in \mathcal{L}$ is such that $k(P) \leq k - 1$, then $\operatorname{rk}(\mathcal{N}_P) \leq k - 1$, so by Theorem 4.2.3 and Proposition 4.3.1, it follows that $\operatorname{rk}(\mathcal{N}_P) = 0$, that is, $\lambda(\mathcal{N}_P) \in \{0, 1\}$. This means that we can write $\mathcal{L} = \widehat{\mathcal{L}} \cup \mathcal{L}_0 \cup \mathcal{L}_1$, where

$$\widehat{\mathcal{L}} \stackrel{\text{def}}{=} \{P \in \mathcal{L} \mid k(P) = k\};$$
$$\mathcal{L}_i \stackrel{\text{def}}{=} \{P \in \mathcal{L} \mid k(P) \leq k - 1 \land \lambda(\mathcal{N}_P) = i\} \qquad (i \in \{0, 1\}).$$

Consider the (left) action of S_k on $\mathcal{K}_k[\operatorname{Th}(\phi)]$ given by letting $\sigma \cdot K \in \mathcal{K}_k[\operatorname{Th}(\phi)]$ ($\sigma \in S_k$, $K \in \mathcal{K}_k[\operatorname{Th}(\phi)]$) be the model obtained from K by permuting its vertices by σ, that is, we have

$$R_P(\sigma \cdot K) \stackrel{\text{def}}{=} \{\sigma \circ \alpha \mid \alpha \in R_P(K)\} \qquad (P \in \widehat{\mathcal{L}});$$
$$R_P(\sigma \cdot K) \stackrel{\text{def}}{=} \varnothing \qquad (P \in \mathcal{L}_0);$$
$$R_P(\sigma \cdot K) \stackrel{\text{def}}{=} ([k])_{k(P)} \qquad (P \in \mathcal{L}_1).$$

Note that this definition ensures that for a.e. $x \in \mathcal{E}_k$ and every $\sigma \in S_k$, we have

$$x \cdot \sigma \in T_{\text{ind}}(K, \mathcal{N}) \iff x \in T_{\text{ind}}(\sigma \cdot K, \mathcal{N}). \tag{4.32}$$

It is also clear that for a.e. $x \in \mathcal{E}_k$, there exists exactly one $K \in \mathcal{K}_k[\operatorname{Th}(\phi)]$ such that $x \in T_{\text{ind}}(K, \mathcal{N})$.

Let then $\mathcal{L}'$ be a language containing one predicate symbol P_K of arity k for each $K \in \mathcal{K}_k[\operatorname{Th}(\phi)]$ and let $\Theta\colon S_k \times \mathcal{L}' \to \mathcal{L}'$ be the induced action $\sigma \cdot P_K \stackrel{\text{def}}{=} P_{\sigma\cdot K}$ ($\sigma \in S_k$,

$K \in \mathcal{K}_k[\mathrm{Th}(\phi)]$). Define then $\mathcal{H}$ by

$$\mathcal{H}_{P_K} \stackrel{\text{def}}{=} T_{\text{ind}}(K, \mathcal{N})$$

and note that (4.32) and the remark below it ensure that $\mathcal{H}$ is a T_Θ-on.

Define $I\colon T \rightsquigarrow T_\Theta$ by

$$I(P)(x_1, \ldots, x_{k(P)}) \stackrel{\text{def}}{=} \begin{cases} \displaystyle\bigvee_{\substack{K \in \mathcal{K}_k[\mathrm{Th}(\phi)] \\ \mathrm{id}_k \in R_P(K)}} P_K(x_1, \ldots, x_{k(P)}), & \text{if } P \in \widehat{\mathcal{L}}; \\ x_1 \neq x_1, & \text{if } P \in \mathcal{L}_0; \\ \displaystyle\bigwedge_{1 \leq i < j \leq k(P)} x_i \neq x_j, & \text{if } P \in \mathcal{L}_1. \end{cases}$$

and note that we trivially have $I(\mathcal{H})_P = \mathcal{N}_P$ a.e. for every $P \in \mathcal{L}$, hence $\phi^I_{\mathcal{H}} = \phi$.

For every $K \in \mathcal{K}_k[\mathrm{Th}(\phi)]$, let $p_{P_K} \stackrel{\text{def}}{=} \lambda(\mathcal{H}_{P_K}) = \phi(\langle K \rangle) > 0$ and note that the definition of Θ implies that p is Θ-invariant and $\sum_{K \in \mathcal{K}_k[\mathrm{Th}(\phi)]} p_{P_K} = 1$. To conclude the proof, we will show that $\phi_{\mathcal{H}} = \psi_{\Theta,p}$. To do so, for every $K \in \mathcal{K}_k[\mathrm{Th}(\phi)]$, let $K_K \in \mathcal{K}_k[T_\Theta]$ be the unique model such that $\mathrm{id}_k \in R_{P_K}(K_K)$ and note that the axioms of T_Θ imply that $W^{K_K}_{\mathcal{H}}$ is a.e. equal to the $(k-1)$-flattening $W^{k-1}_{\mathcal{H}_{P_K}}$ of the peon $\mathcal{H}_{P_K}$, which in turn is a.e. equal to $W^K_{\mathcal{N}}$. But then from Lemma 4.7.4, it follows that $W^{K_K}_{\mathcal{H}} = \phi(\langle K \rangle) = p_{P_K}$ a.e. Since the T_Θ-on $\mathcal{N}^Z$ of Definition 4.1.5 and Proposition 4.7.1 also clearly satisfies $W^{K_K}_{\mathcal{N}^Z} = W^{k-1}_{\mathcal{N}^Z_{P_K}} = p_{P_K}$ a.e., from Lemma 4.7.3, it follows that $\phi_{\mathcal{H}} = \phi_{\mathcal{N}^Z} = \psi_{\Theta,p}$. ∎

The rest of this section is devoted to showing that a limit object satisfying `UInduce`[1] on a language in which all predicates have arity at most 2 must essentially be a (not necessarily unbiased) (Θ, p)-quasirandom homomorphism of Definition 4.1.5. We start with the following simple proposition that says that (possibly biased) (Θ, p)-quasirandom homomorphisms are essentially independent couplings of a quasirandom colored hypergraphon with a linear order.

Proposition 4.7.5. *Let $\mathcal{L}$ be a language containing only predicate symbols of arity exactly*

$k \in \mathbb{N}_+$, let $\Theta\colon S_k \times \mathcal{L} \to \mathcal{L}$ be a (left) action of S_k on $\mathcal{L}$ and let $p = (p_P)_{P\in\mathcal{L}} \in [0,1]^{\mathcal{L}}$ with $\sum_{P\in\mathcal{L}} p_P = 1$.

Then there exists $c \in \mathbb{N}_+$, $q \in \Pi_c$ and an open interpretation $I\colon T_\Theta \rightsquigarrow T_{c,k} \cup T_{\text{LinOrder}}$ such that $\psi_{\Theta,p} = (\psi_{k,q} \otimes \psi_{\text{lin}})^I$, for the quasirandom c-colored k-hypergraphon $\psi_{k,q} \in \operatorname{Hom}^+(\mathcal{A}[T_{c,k}], \mathbb{R})$.

In particular, $\psi_{\Theta,p}$ satisfies `UInduce`$[k-1]$.

Proof. The case $k = 1$ is trivial as $T_\Theta \cong T_{|\mathcal{L}|\text{-Coloring}} \cong T_{|\mathcal{L}|,1}$, so suppose $k \geq 2$.

Let $\mathcal{L}' \stackrel{\text{def}}{=} \{P \in \mathcal{L} \mid p_P > 0\}$. If $|\mathcal{L}'| = 1$, then $\operatorname{rk}(\psi_{\Theta,p}) = 0$ and the desired $I\colon T_\Theta \rightsquigarrow T_{c,k} \cup T_{\text{LinOrder}}$ is defined trivially, by declaring the unique $P \in \mathcal{L}'$ to be true everywhere and all other $P \in \mathcal{L} \setminus \mathcal{L}'$ to be false everywhere (and any $c \in \mathbb{N}_+$ and $q \in \Pi_c$ works).

Suppose then that $c \stackrel{\text{def}}{=} |\mathcal{L}'| \geq 2$. Let $P_1, \ldots, P_t$ be a transversal of the orbits of Θ and for each $i \in [t]$, let $G_i \stackrel{\text{def}}{=} (S_k)_{P_i}$ be the stabilizer of P_i. Let also

$$U \stackrel{\text{def}}{=} \{(P_i, \sigma G_i) \mid i \in [t], \sigma \in S_k\}$$

be the set of pairs (P_i, C), where $i \in [t]$ and C is a (left) coset of G_i. Given $P \in \mathcal{L}'$, we let $\widehat{P}$ be the unique P_i that is in the same orbit of P and let $G_P \stackrel{\text{def}}{=} (S_k)_P$ be the stabilizer of P (so $G_i = G_{P_i}$).

Then there is a natural bijection $f\colon U \to \mathcal{L}'$ given by

$$f(P_i, \sigma G_i) \stackrel{\text{def}}{=} \sigma P_i.$$

By using this bijection along with an enumeration of $\mathcal{L}'$, we can index more conveniently the coordinates of q and the predicates E_i of $T_{c,k}$ by U instead of $[c] = [|\mathcal{L}'|]$. Define then $q \in \Pi_c$ by

$$q_{(P_i,\sigma G_i)} \stackrel{\text{def}}{=} p_{\sigma \cdot P_i} \qquad (i \in [t], \sigma \in S_k)$$

and define the translation I from the language of T_Θ to the language of $T_{c,k} \cup T_{\text{LinOrder}}$ by

$$I(\sigma \cdot P_i)(x_1, \ldots, x_k) \stackrel{\text{def}}{=} \bigvee_{\tau \in S_k} (F_\tau(x_1, \ldots, x_k) \wedge E_{(P_i, \tau \circ \sigma G_i)}(x_1, \ldots, x_k)) \qquad (i \in [t], \sigma \in S_k);$$

$$I(P)(x_1, \ldots, x_k) \stackrel{\text{def}}{=} x \neq x \qquad (P \in \mathcal{L} \setminus \mathcal{L}');$$

where

$$F_\tau(x_1, \ldots, x_k) \stackrel{\text{def}}{=} \bigwedge_{1 \leq i < j \leq k} x_{\tau^{-1}(i)} \prec x_{\tau^{-1}(j)}.$$

Let us show that I is an open interpretation from T_Θ to $T_{c,k} \cup T_{\text{LinOrder}}$. The fact that I satisfies the axioms (4.1) and (4.3) follows trivially from the axioms of $T_{c,k} \cup T_{\text{LinOrder}}$ and our representation in terms of the transversal of the orbits and cosets of stabilizers.

For the axiom (4.2), first note that for $x_1, \ldots, x_k$ distinct, the axioms of T_{LinOrder} imply that there is a unique $\tau_x \in S_k$ such that $F_{\tau_x}(x_1, \ldots, x_k)$ and it further satisfies $\tau_{x \cdot \sigma} = \tau_x \circ \sigma$, where $x \cdot \sigma \stackrel{\text{def}}{=} (x_{\sigma(1)}, \ldots, x_{\sigma(k)})$ for $\sigma \in S_k$. Thus, the theory $T_{c,k} \cup T_{\text{LinOrder}}$ can prove the following chain of equivalences

$$\begin{aligned} I(P_i)(x_{\sigma(1)}, \ldots, x_{\sigma(k)}) &\iff E_{(P_i, \tau_{x \cdot \sigma} G_i)}(x_{\sigma(1)}, \ldots, x_{\sigma(k)}) \\ &\iff E_{(P_i, \tau_x \circ \sigma G_i)}(x_1, \ldots, x_k) \\ &\iff I(\sigma \cdot P_i)(x_1, \ldots, x_k) \end{aligned}$$

for every $i \in [t]$, from which (4.2) follows for every $P \in \mathcal{L}'$. For $P \in \mathcal{L} \setminus \mathcal{L}'$, the axiom (4.2) follows trivially.

Finally, let us show that $(\psi_{k,q} \otimes \psi_{\text{lin}})^I = \psi_{\Theta,p}$.

Let $Z = (Z_{(P_i, \sigma G_i)})_{(P_i, \sigma G_i) \in U}$ be a measurable partition of $[0, 1]$ such that $\lambda(Z_{(P_i, \sigma G_i)}) =$

$q_{(P_i,\sigma G_i)}$ and define the $T_{c,k}\cup T_{\mathrm{LinOrder}}$-on $\mathcal{H}$ by

$$\mathcal{H}_{E_{(P_i,\sigma G_i)}} \stackrel{\mathrm{def}}{=} \{x \in \mathcal{E}_k \mid x_{[k]} \in Z_{(P_i,\sigma G_i)}\} \qquad ((P_i,\sigma G_i)\in U);$$
$$\mathcal{H}_{\prec} \stackrel{\mathrm{def}}{=} \{x \in \mathcal{E}_2 \mid x_{\{1\}} < x_{\{2\}}\}.$$

Then it is straightforward to check that $I(\mathcal{H}) = \mathcal{N}^Z$ for $\mathcal{N}^Z$ as in (4.4). It remains to note that $\phi_{\mathcal{H}}$ is a coupling of $\psi_{k,q}$ and ψ_{lin} and since $\psi_{k,q} \in$ `Independence`$[k-1]$ and $k-1 \geq 1 = \mathrm{rk}(\psi_{\mathrm{lin}})$, by Theorem 4.2.2, we have $\phi_{\mathcal{H}} = \psi_{k,q}\otimes\psi_{\mathrm{lin}}$. Therefore $\psi_{\Theta,p} = (\psi_{k,q}\otimes\psi_{\mathrm{lin}})^I$.

Finally, since $\psi_{k,q}\in\mathrm{Hom}^+(\mathcal{A}[T_{c,k}],\mathbb{R})$ satisfies `Independence`$[k-1]$ (see Lemma 4.5.4), by Theorems 4.2.2, 4.2.3 and 4.2.10(i) $\implies$ (vii), it follows that $\psi_{\Theta,p}$ satisfies `UInduce`$[k-1]$. ■

Our next objective is to show that in T_Θ on arity $k=2$, the (Θ,p)-quasirandom homomorphisms are the only elements of $\mathrm{Hom}^+(\mathcal{A}[T_\Theta],\mathbb{R})$ that satisfy `UInduce`$[1]$. To do so, we will use the following bound for L_1-distance of T_Θ-ons in terms of the usual L_1-distance of the functions $W^K_{\mathcal{N}}$ of Definition 4.7.2 (which we can still prove in general arities).

Lemma 4.7.6. *Let $\mathcal{L}$ be a language containing only predicate symbols of arity exactly $k\in\mathbb{N}_+$ and let $\Theta\colon S_k\times\mathcal{L}\to\mathcal{L}$ be a (left) action of S_k on $\mathcal{L}$. For each $P\in\mathcal{L}$, let $K_P\in\mathcal{K}_k[T_\Theta]$ be the unique model of T_Θ such that $\mathrm{id}_k\in R_P(K_P)$.*

Suppose $\mathcal{N}$ and $\mathcal{H}$ are T_Θ-ons on the same space $\Omega=(X,\mathcal{A},\mu)$. Then

$$\delta_1(\phi_{\mathcal{N}},\phi_{\mathcal{H}}) \leq \sum_{P\in\mathcal{L}}\int_{\mathcal{E}_{k,k-1}(\Omega)} |W^{K_P}_{\mathcal{N}}(x) - W^{K_P}_{\mathcal{H}}(x)|\, d\mu(x).$$

Proof. First, we claim that it is enough to show the case when $\Omega=[0,1]$. Indeed, if $F\colon[0,1]\to\Omega$ is a measure-isomorphism and we let

$$\mathcal{N}'_P \stackrel{\mathrm{def}}{=} \{x\in\mathcal{E}_k \mid (F(x_A))_{A\in r(k)}\in\mathcal{N}_P\} \qquad (P\in\mathcal{L});$$
$$\mathcal{H}'_P \stackrel{\mathrm{def}}{=} \{x\in\mathcal{E}_k \mid (F(x_A))_{A\in r(k)}\in\mathcal{H}_P\} \qquad (P\in\mathcal{L});$$

then $\phi_{\mathcal{N}'} = \phi_{\mathcal{N}}$ and $\phi_{\mathcal{H}'} = \phi_{\mathcal{H}}$, so if the result holds when $\Omega = [0,1]$, we get

$$\begin{aligned}\delta_1(\phi_{\mathcal{N}}, \phi_{\mathcal{H}}) &\leq \sum_{P \in \mathcal{L}} \int_{\mathcal{E}_{k,k-1}} |W_{\mathcal{N}'}^{K_P}(x) - W_{\mathcal{H}'}^{K_P}(x)| d\lambda(x) \\ &= \sum_{P \in \mathcal{L}} \int_{\mathcal{E}_{k,k-1}(\Omega)} |W_{\mathcal{N}}^{K_P}(x) - W_{\mathcal{H}}^{K_P}(x)| d\mu(x).\end{aligned}$$

Let us prove the case $\Omega = [0,1]$. Let us construct for each $x \in \mathcal{E}_{k,k-1}$ two measurable partitions $Z^{\mathcal{N},x} = (Z_P^{\mathcal{N},x})_{P \in \mathcal{L}}$ and $Z^{\mathcal{H},x} = (Z_P^{\mathcal{H},x})_{P \in \mathcal{L}}$ of $[0,1]$ such that

i. $\lambda(Z_P^{\mathcal{N},x}) = W_{\mathcal{N}}^{K_P}(x)$ and $\lambda(Z_P^{\mathcal{H},x}) = W_{\mathcal{H}}^{K_P}(x)$ for a.e. $x \in \mathcal{E}_{k,k-1}$ and every $P \in \mathcal{L}$.

ii. $\lambda(Z_P^{\mathcal{N},x} \cap Z_P^{\mathcal{H},x}) = \min\{W_{\mathcal{N}}^{K_P}(x), W_{\mathcal{H}}^{K_P}(x)\}$ for a.e. $x \in \mathcal{E}_{k,k-1}$ and every $P \in \mathcal{L}$.

iii. For each $P \in \mathcal{L}$, the functions $f_P, g_P \colon \mathcal{E}_k \to \{0,1\}$ given by

$$\begin{aligned}f_P(x,y) &\stackrel{\text{def}}{=} \mathbb{1}[y \in Z^{\mathcal{N},x}] \qquad (x \in \mathcal{E}_{k,k-1}, y \in [0,1]); \\ g_P(x,y) &\stackrel{\text{def}}{=} \mathbb{1}[y \in Z^{\mathcal{H},x}] \qquad (x \in \mathcal{E}_{k,k-1}, y \in [0,1])\end{aligned}$$

are measurable.

To do this, enumerate the predicate symbols of $\mathcal{L}$ as $P_1, \ldots, P_t$ and for each $x \in \mathcal{E}_{k,k-1}$, let

$$u(x) \stackrel{\text{def}}{=} \sum_{P \in \mathcal{L}} \min\{W_{\mathcal{N}}^{K_P}(x), W_{\mathcal{H}}^{K_P}(x)\}.$$

Define also the points $a_0(x) \leq a_1(x) \leq \cdots \leq a_t(x)$, $b_0^{\mathcal{N}}(x) \leq b_1^{\mathcal{N}}(x) \leq \cdots \leq b_t^{\mathcal{N}}(x)$ and

$b_0^{\mathcal{H}}(x) \leq b_1^{\mathcal{H}}(x) \leq \cdots \leq b_t^{\mathcal{H}}(x)$ inductively as follows

$$\begin{aligned}
a_0(x) &\stackrel{\text{def}}{=} 0;\\
a_{i+1}(x) &\stackrel{\text{def}}{=} a_i(x) + \min\{W_{\mathcal{N}}^{K_{P_{i+1}}}(x), W_{\mathcal{H}}^{K_{P_{i+1}}}(x)\};\\
b_0^{\mathcal{N}} \stackrel{\text{def}}{=} b_0^{\mathcal{H}} &\stackrel{\text{def}}{=} u(x);\\
b_{i+1}^{\mathcal{N}}(x) &\stackrel{\text{def}}{=} b_i^{\mathcal{N}} + \max\{W_{\mathcal{N}}^{K_{P_{i+1}}}(x) - W_{\mathcal{H}}^{K_{P_{i+1}}}(x), 0\};\\
b_{i+1}^{\mathcal{H}}(x) &\stackrel{\text{def}}{=} b_i^{\mathcal{N}} + \max\{W_{\mathcal{H}}^{K_{P_{i+1}}}(x) - W_{\mathcal{N}}^{K_{P_{i+1}}}(x), 0\}.
\end{aligned}$$

Note that this definition ensures that $a_t(x) = u(x)$ and $b_t^{\mathcal{N}}(x) = b_t^{\mathcal{H}}(x) = 1$. Furthermore, the functions functions $u, a_i, b_i^{\mathcal{N}}, b_i^{\mathcal{H}}$ are measurable.

We then define the partitions $Z^{\mathcal{N},x}$ and $Z^{\mathcal{H},x}$ by

$$\begin{aligned}
Z_{P_i}^{\mathcal{N},x} &\stackrel{\text{def}}{=} [a_{i-1}(x), a_i(x)) \cup [b_{i-1}^{\mathcal{N}}(x), b_i^{\mathcal{N}}(x)) \qquad (i \in [t-1]);\\
Z_{P_t}^{\mathcal{N},x} &\stackrel{\text{def}}{=} [a_{t-1}(x), a_t(x)) \cup [b_{t-1}^{\mathcal{N}}(x), b_t^{\mathcal{N}}(x)];\\
Z_{P_i}^{\mathcal{H},x} &\stackrel{\text{def}}{=} [a_{i-1}(x), a_i(x)) \cup [b_{i-1}^{\mathcal{H}}(x), b_i^{\mathcal{H}}(x)) \qquad (i \in [t-1]);\\
Z_{P_t}^{\mathcal{H},x} &\stackrel{\text{def}}{=} [a_{t-1}(x), a_t(x)) \cup [b_{t-1}^{\mathcal{H}}(x), b_t^{\mathcal{H}}(x)].
\end{aligned}$$

Now we define the Euclidean structures $\mathcal{N}'$ and $\mathcal{H}'$ on $\mathcal{L}$ by

$$\begin{aligned}
\mathcal{N}'_P &\stackrel{\text{def}}{=} \{x \in \mathcal{E}_k \mid x_{[k]} \in Z_{\sigma_x \cdot P}^{\mathcal{N}, x\cdot\sigma_x^{-1}}\};\\
\mathcal{H}'_P &\stackrel{\text{def}}{=} \{x \in \mathcal{E}_k \mid x_{[k]} \in Z_{\sigma_x \cdot P}^{\mathcal{H}, x\cdot\sigma_x^{-1}}\}.
\end{aligned}$$

Note that by Proposition 4.7.1, both $\mathcal{N}'$ and $\mathcal{H}'$ are T_Θ-ons.

Finally, it is easy to see from construction that $W_{\mathcal{N}'}^{K_P} = W_{\mathcal{N}}^{K_P}$ a.e. and $W_{\mathcal{H}'}^{K_P} = W_{\mathcal{H}}^{K_P}$

a.e., so by Lemma 4.7.3, we have $\phi_{\mathcal{N}'} = \phi_{\mathcal{N}}$ and $\phi_{\mathcal{H}'} = \phi_{\mathcal{H}}$, thus

$$\begin{aligned}\delta_1(\phi_{\mathcal{N}}, \phi_{\mathcal{H}}) &\leq \sum_{P\in\mathcal{L}} \lambda(\mathcal{N}'_P \mathbin{\triangle} \mathcal{H}'_P)\\ &= \sum_{i=1}^{t} \int_{\mathcal{E}_{k,k-1}} (b_i^{\mathcal{N}}(x) - b_{i-1}^{\mathcal{N}}(x) + b_i^{\mathcal{H}}(x) - b_{i-1}^{\mathcal{H}}(x)) d\lambda(x)\\ &= \sum_{P\in\mathcal{L}} \int_{\mathcal{E}_{k,k-1}} |W_{\mathcal{N}}^{K_{P_i}}(x) - W_{\mathcal{H}}^{K_{P_i}}(x)| d\lambda(x),\end{aligned}$$

as desired. ■

Recall that a Lebesgue point of a function $f\colon \mathbb{R}^k \to \mathbb{R}$ is a point $z \in \mathbb{R}^k$ such that

$$\lim_{\varepsilon\to 0} \frac{1}{\lambda(B(z,\varepsilon))} \int_{B(z,\varepsilon)} |f(x) - f(z)|\ dx = 0,$$

where $B(z,\varepsilon)$ denotes the ℓ_∞-ball[5] of radius ε around z. The Lebesgue Differentiation Theorem [39, Theorem 2.9.7] says that if f is integrable, then almost every point $z \in \mathbb{R}^k$ is a Lebesgue point of f.

Proposition 4.7.7. *Let $\mathcal{L}$ be a language containing only predicate symbols of arity exactly 2, let $\Theta\colon S_2 \times \mathcal{L} \to \mathcal{L}$ be a (left) action of S_2 on $\mathcal{L}$ and let $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_\Theta], \mathbb{R})$. Then $\phi \in$ `UInduce`$[1]$ if and only if there exists $p = (p_P)_{P\in\mathcal{L}}$ with $\sum_{P\in\mathcal{L}} p_P = 1$ such that $\phi = \psi_{\Theta,p}$.*

Proof. The backward direction follows from Proposition 4.7.5.

For the forward direction, let $\mathcal{N}$ be a T_Θ-on such that $\phi_{\mathcal{N}} = \phi$ and for each $K \in \mathcal{K}_2[T_\Theta]$, let $P_K \in \mathcal{L}$ be the unique predicate symbol such that $(1,2) \in R_P(K)$.

By the Lebesgue Differentiation Theorem, there exists $z \in \mathcal{E}_2$ that is a Lebesgue point of all functions $W_{\mathcal{N}}^{K_P}\colon \mathcal{E}_2 \to [0,1]$ for $P \in \mathcal{L}$. Without loss of generality, we assume that $0 < z_{\{1\}} < z_{\{2\}} < 1$ and that $W_{\mathcal{N}}^{K_P}(z_{\{2\}}, z_{\{1\}}) = W_{\mathcal{N}}^{K_{(2,1)\cdot P}}(z)$ for every $P \in \mathcal{L}$.

5. Again, one can use other norms to define Lebesgue points and get an a.e. equivalent definition, but for us it will be slightly more convenient to use the ℓ_∞-norm.

Define then $p = (p_P)_{P\in\mathcal{L}}$ by $p_P \stackrel{\text{def}}{=} W_{\mathcal{N}}^{K_P}(z)$. We claim that $\sum_{P\in\mathcal{L}} p_P = 1$. Indeed, the axioms of T_Θ ensure that for $\varepsilon > 0$ we have

$$\left|1 - \sum_{P\in\mathcal{L}} p_P\right| \leq \left|\sum_{P\in\mathcal{L}} \frac{1}{\lambda(B(z,\varepsilon))} \int_{B(z,\varepsilon)} (W_{\mathcal{N}}^{K_P}(x) - W_{\mathcal{N}}^{K_P}(z))\, dx\right|,$$

so letting $\varepsilon \to 0$ gives $\sum_{P\in\mathcal{L}} p_P = 1$.

We will show that $\phi = \psi_{\Theta,p}$ by showing that their L_1-distance is 0 with aid of Lemma 4.7.6. To this purpose, fix $\varepsilon > 0$ and let $\varepsilon' > 0$ be small enough so that

$$\frac{1}{\lambda(B(z,\varepsilon'))} \int_{B(z,\varepsilon')} |W_{\mathcal{N}}^{K_P}(x) - W_{\mathcal{N}}^{K_P}(z)|\, dx \leq \frac{\varepsilon}{|\mathcal{L}|} \tag{4.33}$$

for all $P \in \mathcal{L}$. Without loss of generality, we also assume that $B(z,\varepsilon') \subseteq \mathcal{E}_2$ and that $z_{\{1\}} + \varepsilon' < z_{\{2\}} - \varepsilon'$.

Given an open set $U \subseteq [0,1]$, let us consider the space Ω_U obtained by equipping U with the measure $\mu_U \stackrel{\text{def}}{=} \lambda/\lambda(U)$ and let $F^U\colon \Omega_U \to [0,1]$ be a measure-isomorphism. Define also the T_Θ-on $\mathcal{N}^U$ over Ω_U by

$$\mathcal{N}_P^U \stackrel{\text{def}}{=} \{x \in \mathcal{E}_2(\Omega_U) \mid (x_{\{1\}}, x_{\{2\}}, F^U(x_{\{1,2\}})) \in \mathcal{N}_P\} \qquad (P \in \mathcal{L}).$$

Intuitively, $\mathcal{N}^U$ corresponds to the restriction of $\mathcal{N}$ to the vertices in U.

Note that if $\widehat{\mathcal{N}}^U$ is the $(T_\Theta \cup T_{2\text{-Coloring}})$-on obtained from $\mathcal{N}$ by declaring $\widehat{\mathcal{N}}_{\chi_1}^U \stackrel{\text{def}}{=} U$ and $\widehat{\mathcal{N}}_{\chi_2}^U \stackrel{\text{def}}{=} [0,1] \setminus U$, then $\phi_{\mathcal{N}^U} = \phi_{\widehat{\mathcal{N}}^U} \circ \pi^{(\chi_1, \mathrm{id}_{T_\Theta})}$. Since $\phi \in \mathtt{UInduce}[1]$, for every $M \in \mathcal{M}[T_\Theta]$ we have

$$\phi_{\mathcal{N}^U}(M) = (\phi_{\widehat{\mathcal{N}}^U} \circ \pi^{(\chi_1, \mathrm{id}_{T_\Theta})})(M) = \frac{\phi_{\widehat{\mathcal{N}}^U}(\widehat{M})}{\lambda(U)^{|M|}} = \phi(M),$$

where $\widehat{M} \in \mathcal{M}[T_\Theta \cup T_{2\text{-Coloring}}]$ is obtained from M by coloring all of its vertices with color 1. Thus $\phi_{\mathcal{N}^U} = \phi$.

For $i \in [2]$, let $U_i \stackrel{\text{def}}{=} (z_{\{i\}} - \varepsilon', z_{\{i\}} + \varepsilon')$ and note that $U_1 \times U_2 = B(z, \varepsilon')$. Also, by Remark 2, there exists a T_Θ-on $\mathcal{H}^i$ over $\Omega_{U_i} \times \Omega_{U_i}$ such that $\phi_{\mathcal{H}^i} = \psi_{\Theta,p}$ and

$$\delta_1(\phi, \psi_{\Theta,p}) = \sum_{P \in \mathcal{L}} \mu_{U_i}((\mathcal{N}_P^{U_i} \times \mathcal{E}_2(\Omega_{U_i})) \mathbin{\triangle} \mathcal{H}^i). \tag{4.34}$$

Let now $U_0 \stackrel{\text{def}}{=} U_1 \cup U_2$ and define the T_Θ-on $\mathcal{H}$ over $\Omega_{U_0} \times \Omega_{U_0}$ by

$$\begin{aligned}
\mathcal{H}_P \stackrel{\text{def}}{=} \bigl\{ & (x, y) \in \mathcal{E}_2(\Omega_{U_0}) \times \mathcal{E}_2(\Omega_{U_0}) \mid \\
& \bigl(x_{\{1\}} \in U_1 \land x_{\{2\}} \in U_1 \\
& \quad \land (x_{\{1\}}, x_{\{2\}}, G^{U_1}(x_{\{1,2\}}), G^{U_1}(y_{\{1\}}), G^{U_1}(y_{\{2\}}), G^{U_1}(y_{\{1,2\}})) \in \mathcal{H}^1\bigr) \\
& \lor \bigl(x_{\{1\}} \in U_2 \land x_{\{2\}} \in U_2 \\
& \quad \land (x_{\{1\}}, x_{\{2\}}, G^{U_2}(x_{\{1,2\}}), G^{U_2}(y_{\{1\}}), G^{U_2}(y_{\{2\}}), G^{U_2}(y_{\{1,2\}})) \in \mathcal{H}^2\bigr) \\
& \lor \bigl(x_{\{1\}} \in U_1 \land x_{\{2\}} \in U_2 \land F^{U_0}(x_{\{1,2\}}) \in Z_P\bigr) \\
& \lor \bigl(x_{\{1\}} \in U_1 \land x_{\{2\}} \in U_2 \land F^{U_0}(x_{\{1,2\}}) \in Z_{(2,1)\cdot P}\bigr)\bigr\},
\end{aligned}$$

where $G^{U_i} \stackrel{\text{def}}{=} (F^{U_i})^{-1} \circ F^{U_0}$ and $Z = (Z_P)_{P \in \mathcal{L}}$ is a measurable partition of $[0,1]$ with $\lambda(Z_P) = p_P$. The intuition is that $\mathcal{H}$ mimics $\mathcal{H}^i$ in $U_i \times U_i$ and mimics a version over $[0,1]^2$ of the standard T_Θ-on $\mathcal{N}^Z$ representing $\psi_{\Theta,p}$ in $(U_1 \times U_2) \cup (U_2 \times U_1)$ given by (4.4).

Consider now the T_Θ-on $\mathcal{N}'$ over $\Omega_{U_0} \times \Omega_{U_0}$ obtained from $\mathcal{N}^{U_0}$ by adding a dummy variable ($\mathcal{N}'_P \stackrel{\text{def}}{=} \mathcal{N}_P^{U_0} \times \mathcal{E}_2(\Omega_{U_0})$) and note that Lemma 4.7.6 gives

$$\delta_1(\phi, \phi_{\mathcal{H}}) \leq \sum_{P \in \mathcal{L}} \int_{\mathcal{E}_{2,1}(\Omega_{U_0} \times \Omega_{U_0})} |W_{\mathcal{N}'}^{K_P}(x) - W_{\mathcal{H}}^{K_P}| \, d(\mu_{U_0} \otimes \mu_{U_0})(x).$$

We will bound the integral above by splitting its domain into four parts:

$$\mathcal{E}_{2,1}(\Omega_{U_0} \times \Omega_{U_0}) = (U_1 \times U_1) \cup (U_2 \times U_2) \cup (U_1 \times U_2) \cup (U_2 \times U_1)$$

We can deal with the first two parts at the same time. Note that

$$\begin{aligned}
&\sum_{P\in\mathcal{L}} \int_{U_i\times U_i} |W_{\mathcal{N}'}^{K_P}(x) - W_{\mathcal{H}}^{K_P}(x)|\, d(\mu_{U_0}\otimes\mu_{U_0})(x) \\
&\quad \le \sum_{P\in\mathcal{L}} \frac{1}{4}\mu_{U_i}((\mathcal{N}_P^{U_i}\times\mathcal{E}_2(\Omega_{U_i}))\bigtriangleup\mathcal{H}^i) \\
&\quad = \frac{1}{4}\delta_1(\phi,\psi_{\Theta,p}),
\end{aligned}$$

where the last equality follows from (4.34).

For the third part, note that

$$\begin{aligned}
&\sum_{P\in\mathcal{L}} \int_{U_1\times U_2} |W_{\mathcal{N}'}^{K_P}(x) - W_{\mathcal{H}}^{K_P}(x)|\, d(\mu_{U_0}\otimes\mu_{U_0})(x) \\
&\quad = \frac{1}{4\lambda(B(z,\varepsilon'))}\int_{B(z,\varepsilon')} |W_{\mathcal{N}}^{K_P}(x) - p_P|\, d\lambda(x) \\
&\quad \le \frac{\varepsilon}{4},
\end{aligned}$$

where the last inequality follows from (4.33).

Finally, for the fourth part, by using the axiom (4.2) of T_Θ and the fact that $p_{(2,1)\cdot P} = W_{\mathcal{N}}^{K_{(2,1)\cdot P}}(z) = W_{\mathcal{N}}^{K_P}(z_{\{2\}}, z_{\{1\}})$, note that

$$\begin{aligned}
&\sum_{P\in\mathcal{L}} \int_{U_2\times U_1} |W_{\mathcal{N}'}^{K_P}(x) - W_{\mathcal{H}}^{K_P}(x)|\, d(\mu_{U_0}\otimes\mu_{U_0})(x) \\
&\quad = \sum_{P\in\mathcal{L}} \int_{U_2\times U_1} |W_{\mathcal{N}'}^{K_{(2,1)\cdot P}}(x_{\{2\}}, x_{\{1\}}) - p_{(2,1)\cdot P}|\, d(\mu_{U_0}\otimes\mu_{U_0})(x) \\
&\quad = \sum_{P\in\mathcal{L}} \int_{U_1\times U_2} |W_{\mathcal{N}'}^{K_P}(x) - W_{\mathcal{H}}^{K_P}(x)|\, d(\mu_{U_0}\otimes\mu_{U_0})(x) \le \frac{\varepsilon}{4},
\end{aligned}$$

where the last inequality follows from the bound for the third part.

Putting these together, we conclude that

$$\delta_1(\phi,\phi_{\mathcal{H}}) \le \frac{\delta_1(\phi,\psi_{\Theta,p})}{2} + \frac{\varepsilon}{2}.$$

This means that if we show that $\phi_{\mathcal{H}} = \psi_{\Theta,p}$, then by letting $\varepsilon \to 0$ in the above, we conclude that $\delta_1(\phi, \psi_{\Theta,p}) = 0$, so $\phi = \psi_{\Theta,p}$ as desired.

Let us then show that $\phi_{\mathcal{H}} = \psi_{\Theta,p}$. For this, since $\phi_{\mathcal{H}^i} = \psi_{\Theta,p}$, by theon uniqueness, Theorem 2.6.1, there exists a family $h^i = (h^i_1, h^i_2)$ of functions ($h^i_d\colon \mathcal{E}_d(\Omega_{U_i}) \times \mathcal{E}_d(\Omega_{U_i}) \times \mathcal{E}_d(\Omega_{U_i}) \times \mathcal{E}_d(\Omega_{U_i}) \to [0,1]$) symmetric and measure preserving on h.o.a. such that

$$(x,y) \in \mathcal{H}^i_P \iff \widehat{h}^i_2(x,y,u,v) \in \mathcal{N}^Z_P$$

for every $P \in \mathcal{L}$ and a.e. $(x,y,u,v) \in \mathcal{E}_2(\Omega_{U_i}) \times \mathcal{E}_2(\Omega_{U_i}) \times \mathcal{E}_2(\Omega_{U_i}) \times \mathcal{E}_2(\Omega_{U_i})$, where $\mathcal{N}^Z$ is the standard T_Θ-on representing $\psi_{\Theta,p}$ given by (4.4).

Then consider the family $h = (h_1, h_2)$ ($h_d\colon \mathcal{E}_d(\Omega_{U_0}) \times \mathcal{E}_d(\Omega_{U_0}) \times \mathcal{E}_d(\Omega_{U_0}) \times \mathcal{E}_d(\Omega_{U_0}) \to [0,1]$) defined by

$$h_1(x,y,u,v) = \begin{cases} h^1_1(x, G^{U_1}(y), G^{U_1}(u), G^{U_1}(v)), & \text{if } x \in U_1;\\ h^2_1(x, G^{U_2}(y), G^{U_2}(u), G^{U_2}(v)), & \text{if } x \in U_2; \end{cases}$$

$$h_2(x,y,u,v) = \begin{cases} h^1_1(x, H^{U_1}(y), H^{U_1}(u), H^{U_1}(v)), & \text{if } x_{\{1\}}, x_{\{2\}} \in U_1;\\ h^2_1(x, H^{U_2}(y), H^{U_2}(u), H^{U_2}(v)), & \text{if } x_{\{1\}}, x_{\{2\}} \in U_2;\\ F^{U_0}(x_{\{1,2\}}), & \text{otherwise}; \end{cases}$$

where $H^{U_i}(w_{\{1\}}, w_{\{2\}}, w_{\{1,2\}}) \stackrel{\text{def}}{=} (G^{U_i}(w_{\{1\}}), G^{U_i}(w_{\{2\}}), G^{U_i}(w_{\{1,2\}}))$ (recall that $G^{U_i} \stackrel{\text{def}}{=} (F^{U_i})^{-1} \circ F^{U_0}$). Then it is straightforward to check that h_1 and h_2 are symmetric and measure-preserving on h.o.a. and we have

$$(x,y) \in \mathcal{H}_P \iff \widehat{h}_2(x,y,u,v) \in \mathcal{N}^Z_P$$

for every $P \in \mathcal{L}$ and a.e. $(x,y,u,v) \in \mathcal{E}_2(\Omega_{U_0}) \times \mathcal{E}_2(\Omega_{U_0}) \times \mathcal{E}_2(\Omega_{U_0}) \times \mathcal{E}_2(\Omega_{U_0})$ such that $x_{\{1\}} \in U_i$ and $x_{\{2\}} \in U_{3-i}$ for some $i \in [2]$.

This implies that

$$W_{\mathcal{H}}^{K_P}((x_{\{1\}},y_{\{1\}}),(x_{\{2\}},y_{\{2\}})) = W_{\mathcal{N}^Z}(\widehat{h}_1(x_{\{1\}},y_{\{1\}},u_{\{1\}},v_{\{1\}}),\widehat{h}_1(x_{\{2\}},y_{\{2\}},u_{\{2\}},v_{\{2\}}))$$

for a.e. $(x,y,u,v) \in \mathcal{E}_{2,1}(\Omega_{U_0})$. Therefore, since $\widehat{h}_1$ is measure-preserving, by Lemma 4.7.3, we have $\phi_{\mathcal{H}} = \phi_{\mathcal{N}^Z} = \psi_{\Theta,p}$. ■

Before proving Theorem 4.2.14, by comparing it to Theorem 4.2.13, note that in item (ii) we can only produce a translation rather than an open interpretation from T to T_Θ; only in item (iii), we can actually produce an open interpretation from T to $T_{c,2}$. The following simple example shows why we cannot get an open interpretation in the former item.

Example 6. Consider the case $T = T_{\text{LinOrder}}$. By Theorem 4.2.7, the linear order $\psi \in \text{Hom}^+(\mathcal{A}[T_{\text{LinOrder}}],\mathbb{R})$ satisfies `UInduce`[1]. On the other hand, we claim that there is no open interpretation $I\colon T_{\text{LinOrder}} \rightsquigarrow T_\Theta$ for any action $\Theta\colon S_2 \times \mathcal{L} \to \mathcal{L}$ over some language $\mathcal{L}$ with predicate symbols of arity 2.

Suppose for a contradiction that one such I exists. Let $P_0 \in \mathcal{L}$. If P_0 is a fixed point of the action Θ, then for the $K \in \mathcal{K}_2[T_\Theta]$ given by $R_{P_0}(K) \stackrel{\text{def}}{=} ([2])_2$ (and $R_P(K) \stackrel{\text{def}}{=} \varnothing$ for $P \neq P_0$), $I(K)$ violates anti-symmetry axiom of T_{LinOrder}. If P_0 is not a fixed point of the action Θ, then for the $K \in \mathcal{K}_3[T_\Theta]$ given by $R_{P_0}(K) \stackrel{\text{def}}{=} \{(1,2),(2,3),(3,1)\}$ (and $R_P(K) \stackrel{\text{def}}{=} \varnothing$ for $P \neq P_0$), $I(K)$ violates the transitivity axiom of T_{LinOrder}.

Proof of Theorem 4.2.14. We start with the implication (ii) $\implies$ (iii).

By Proposition 4.7.5, there exist $c \in \mathbb{N}_+$, $q \in \Pi_c$ and an open interpretation $J\colon T_\Theta \rightsquigarrow T_{c,2} \cup T_{\text{LinOrder}}$ such that $\psi_{\Theta,p} = (\psi_{k,q} \otimes \psi_{\text{lin}})^J$. Then for the translation $J \circ I$ from $\mathcal{L}$ to the language of $T_{c,2}$, we have $\phi^A = (\psi_{k,q} \otimes \psi_{\text{lin}})^{J \circ I}$.

Since $A\colon T_{\mathcal{L}} \rightsquigarrow T$ is the axiom-adding interpretation, the result will follow if we show that $J \circ I$ is an open interpretation from T to $T_{c,2} \cup T_{\text{LinOrder}}$ (even though I might not be an open interpretation from T to T_Θ). But indeed, since $\text{Th}(\psi_{k,q} \otimes \psi_{\text{lin}}) = T_{c,2} \cup T_{\text{LinOrder}}$, it follows that for every $M \in \mathcal{M}[T_{c,2} \cup T_{\text{LinOrder}}]$, we have $(\psi_{k,q} \otimes \psi_{\text{lin}})(M) > 0$, so $(\psi_{k,q} \otimes \psi_{\text{lin}})^{J\circ I}((J \circ I)(M)) > 0$, hence $(J \circ I)(M) \in \mathcal{M}[T]$ as it has positive density in ϕ.

Implication (iii) $\implies$ (i) follows from Theorems 4.2.2, 4.2.3 and 4.2.10(i) $\implies$ (vii) and the fact that $\psi_{k,p} \in \texttt{Independence}[k-1]$ (see Lemma 4.5.4).

Finally, let us show (i) $\implies$ (ii). Without loss of generality, we may suppose that $T = \text{Th}(\phi)$. Let $\mathcal{N}$ be a T-on such that $\phi_{\mathcal{N}} = \phi$. We claim that if $P \in \mathcal{L}$ has $k(P) = 1$, then $\text{rk}(\mathcal{N}_P) = 0$. Indeed, this follows since ϕ is 1-categorical by Lemma 4.4.12. This means that we can write $\mathcal{L} = \widehat{\mathcal{L}} \cup \mathcal{L}_0 \cup \mathcal{L}_1$, where

$$\begin{aligned}
\widehat{\mathcal{L}} &\stackrel{\text{def}}{=} \{P \in \mathcal{L} \mid k(P) = 2\};\\
\mathcal{L}_i &\stackrel{\text{def}}{=} \{P \in \mathcal{L} \mid k(P) = 1 \land \lambda(\mathcal{N}_P) = i\} \qquad (i \in \{0,1\}).
\end{aligned}$$

Let $\mathcal{L}'$ have one predicate symbol P_K of arity 2 for each $K \in \mathcal{K}_2[T]$ and let us define the action $\Theta\colon S_2 \times \mathcal{L}' \to \mathcal{L}'$ based on the natural action of S_2 on $\mathcal{K}_2[T]$, that is, we let $\sigma \cdot P_K \stackrel{\text{def}}{=} P_{\sigma\cdot K}$ $(\sigma \in S_2, K \in \mathcal{K}_2[T])$, where $\sigma \cdot K \in \mathcal{K}_2[T]$ is given by

$$\begin{aligned}
R_P(\sigma \cdot K) &\stackrel{\text{def}}{=} \{\sigma \circ \alpha \mid \alpha \in R_P(K)\} && (P \in \widehat{\mathcal{L}});\\
R_P(\sigma \cdot K) &\stackrel{\text{def}}{=} \varnothing && (P \in \mathcal{L}_0);\\
R_P(\sigma \cdot K) &\stackrel{\text{def}}{=} \{(1)\} && (P \in \mathcal{L}_1).
\end{aligned}$$

Then we have the natural translation $I\colon T_{\mathcal{L}} \rightsquigarrow T_{\mathcal{L}'}$ from $\mathcal{L}$ to $\mathcal{L}'$ given by

$$\begin{aligned}
I(P)(x_1, x_2) &\stackrel{\text{def}}{=} \bigvee_{\substack{K\in\mathcal{K}_2[T]\\(1,2)\in R_P(K)}} P_K(x_1, x_2) && (P \in \widehat{\mathcal{L}});\\
I(P)(x) &\stackrel{\text{def}}{=} x \neq x && (P \in \mathcal{L}_0);\\
I(P)(x) &\stackrel{\text{def}}{=} x = x && (P \in \mathcal{L}_1).
\end{aligned}$$

Note that I has a left-inverse, namely, the translation $J\colon T_{\mathcal{L}'} \rightsquigarrow T_{\mathcal{L}}$ from $\mathcal{L}'$ to $\mathcal{L}$ given by

$$J(P_K)(x_1, x_2) \stackrel{\text{def}}{=} D_{\text{open}}(K)(x_1, x_2).$$

Note now that our definition of Θ ensures that J is an open interpretation from T_Θ to T. More formally, we have an open interpretation $\widehat{J}\colon T_\Theta \rightsquigarrow T$ that acts as J on $\mathcal{L}'$. Since J is a left-inverse of I, it follows that $\widehat{J} \circ A' \circ I = A$, where $A'\colon T_{\mathcal{L}'} \rightsquigarrow T_\Theta$ and $A\colon T_{\mathcal{L}} \rightsquigarrow T$ are the axiom-adding interpretations. But by Theorem 4.2.3, we know that $\phi^{\widehat{J}} \in \operatorname{Hom}^+(\mathcal{A}[T_\Theta], \mathbb{R})$ satisfies `UInduce`[1], so by Proposition 4.7.7, we have $\phi^{\widehat{J}} = \psi_{\Theta,p}$ for some $p = (p_Q)_{Q\in\mathcal{L}'}$ with $\sum_{Q\in\mathcal{L}'} p_Q = 1$ and thus $\phi^A = \phi^{\widehat{J}\circ A'\circ I} = \psi_{\Theta,p}^{A'\circ I}$. ■

4.8 Compatibility

In this section, we explore the following generalizations of the notions rank, `Independence`, weak independence and `UCouple`.

Definition 4.8.1. For $B \subseteq \mathbb{N}_+$, we say that a peon $\mathcal{N}$ over $\Omega = (X, \mathcal{A}, \mu)$ is *B-compatible* if it only depends on coordinates that are indexed by sets A with $|A| \in B$, that is, it can be written as $\mathcal{N} = \mathcal{G} \times X^{\bigcup_{b\in[k(P)]\setminus B}\binom{[k(P)]}{b}}$ for some $\mathcal{G} \subseteq X^{\bigcup_{b\in B}\binom{[k(P)]}{b}}$. We say that an Euclidean structure is *B-compatible* if all its peons are so and we say that $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is *B-compatible* if there exists a B-compatible T-on $\mathcal{N}$ with $\phi = \phi_{\mathcal{N}}$.

We say that an Euclidean structure $\mathcal{N}$ over Ω is *weakly B-independent* if the exchangeable array $\boldsymbol{K}$ corresponding to $\mathcal{N}$ with respect to $\boldsymbol{\theta}$ picked in $\mathcal{E}_{\mathbb{N}_+}(\Omega)$ according to μ is independent from $(\boldsymbol{\theta}_A \mid A \in \binom{\mathbb{N}_+}{b}, b \in B)$ as a random variable. We say that $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$ is *weakly B-independent* if there exists a weakly B-independent T-on $\mathcal{N}$ such that $\phi_{\mathcal{N}} = \phi$ and we say that ϕ is *completely weakly B-independent* if *every* T-on $\mathcal{N}$ such that $\phi_{\mathcal{N}} = \phi$ is weakly B-independent.

Finally, let say that ϕ is *uniquely B-coupleable* if it is uniquely coupleable with any B-compatible $\psi \in \operatorname{Hom}^+(\mathcal{A}[T'], \mathbb{R})$ (for any T'). We will use the abbreviation `UCouple`$[B]$

for this property.

Note that $\mathrm{rk}(\phi) \leq \ell$ is the same as $[\ell]$-compatibility, $\texttt{Independence}[\ell]$ is the same as $(\mathbb{N}_+ \setminus [\ell])$-compatibility and $\texttt{UCouple}[\ell]$ is the same as $\texttt{UCouple}[[\ell]]$. Furthermore, note that weak ℓ-independence is the same as weak $[\ell]$-independence and by Theorem 4.2.10(iv)$\Leftrightarrow$(v) is the same as complete weak $[\ell]$-independence.

In this section, we sketch how the results of this chapter can be used to prove the following modest generalizations of Theorems 4.2.2, 4.2.3, 4.2.10, and 4.2.6. Except for Theorem 4.8.2, all others have proofs that are either trivial (at this point) or are obtained from their analogues mutatis mutandi.

Theorem 4.8.2. *For $B \subseteq \mathbb{N}_+$, B-compatibility implies $\texttt{UCouple}[\mathbb{N}_+ \setminus B]$.*

Let us also note that for $B' \subseteq B \subseteq \mathbb{N}_+$, it is obvious that B'-compatibility implies B-compatibility and that $\texttt{UCouple}[B]$ implies $\texttt{UCouple}[B']$.

Theorem 4.8.3. *Let $I\colon T_1 \rightsquigarrow T_2$ be an open interpretation and let $B \subseteq \mathbb{N}_+$. The following hold for any $\phi \in \mathrm{Hom}^+(\mathcal{A}[T_2], \mathbb{R})$.*

i. If ϕ is B-compatible, then ϕ^I is B-compatible.

ii. If $\phi \in \texttt{UCouple}[B]$, then $\phi^I \in \texttt{UCouple}[B]$.

Theorem 4.8.4. *Let $B \subseteq \mathbb{N}_+$. The following are equivalent for $\phi \in \mathrm{Hom}^+(\mathcal{A}[T], \mathbb{R})$.*

i. $\phi \in \texttt{UCouple}[B]$.

ii. For every $\ell \in B$, there exists $p \in (0,1)$ such that ϕ is uniquely coupleable with the quasirandom ℓ-hypergraphon $\psi_{\ell,p}$.

iii. There exist $(p_\ell)_{\ell \in B} \in (0,1)^B$ such that ϕ is uniquely coupleable with the independent coupling $\bigotimes_{\ell \in B} \psi_{\ell,p_\ell}$ of the quasirandom ℓ-hypergraphons ψ_{ℓ,p_ℓ} for $\ell \in B$.

iv. ϕ is completely weakly B-independent.

Note that the theorem above does not have analogues of Theorem 4.2.10 concerning (plain) weak independence, locality or unique inducibility of the independent coupling with the linear order.

Theorem 4.8.5. *Let $B \subseteq \mathbb{N}_+$ be non-empty and let $\ell = \min B$. If $\ell \geq 2$, then the quasirandom ℓ-tournamon ψ_ℓ satisfies $\texttt{UCouple}[\mathbb{N}_+ \setminus \{\ell\}]$ but is not $(B \setminus \{\ell\})$-compatible.*

We start with Theorem 4.8.3, whose proof is trivial at this point.

Proof of Theorem 4.8.3. Item (i) follows trivially since I preserves B-compatibility of Euclidean structures and item (ii) follows trivially from Proposition 3.2.9. ■

Before we proceed, note that the proof of Lemma 4.5.6 can be used to show the following lemma mutatis mutandis.

Lemma 4.8.6. *Let $\mathcal{L}$ be a language, $\phi \in \operatorname{Hom}^+(\mathcal{A}[T_\mathcal{L}], \mathbb{R})$ be B-compatible and $\varepsilon > 0$. Then there exist $c \geq 2$, $p \in \Pi_c$ and an open interpretation $I\colon T_\mathcal{L} \rightsquigarrow \bigcup_{\ell \in B} T_{c,\ell}$ such that $\delta_1(\phi, (\bigotimes_{\ell \in B} \psi_{\ell,p})^I) \leq \varepsilon$.*

Proof of Theorem 4.8.4 (sketch). The proof of equivalence between (i), (iii) and (ii) is analogous to that of Lemma 4.5.7 for Theorem 4.2.10 but replacing Theorem 4.2.3 and Lemma 4.5.6 with Theorem 4.8.3 and Lemma 4.8.6, respectively (and noting that $\bigotimes_{\ell \in B} \psi_{\ell,p_\ell}$ is trivially B-compatible).

The proof of equivalence between (i) and (iv) is analogous to that of Lemma 4.3.2 for Theorem 4.2.10, except that instead of using Proposition 3.1.2 to argue that any representation of a theon $\mathcal{N}$ can be changed in a zero-measure set to have the correct rank, we use *complete* weak B-independence (as opposed to plain weak B-independence). ■

Proof of Theorem 4.8.2. By Theorem 4.8.4, it is sufficient to prove that for a theory T on a language $\mathcal{L}$ and for a B-compatible $\phi \in \operatorname{Hom}^+(\mathcal{A}[T], \mathbb{R})$, we must have that ϕ is completely weakly $(\mathbb{N}_+ \setminus B)$-independent. To this purpose, let $\mathcal{N}$ be a T-on over some space $\Omega = (X, \mathcal{A}, \mu)$

such that $\phi = \phi_{\mathcal{N}}$. Since ϕ is B-compatible, there exists a B-compatible T-on $\mathcal{N}'$ over some space $\Omega' = (X', \mathcal{A}', \mu')$ such that $\phi = \phi_{\mathcal{N}'}$.

By theon uniqueness, Theorem 2.6.1, there exists a family $h = (h_1, \ldots, h_k)$ of symmetric functions measure-preserving on h.o.a. ($h_d \colon \mathcal{E}_d(\Omega \times \Omega) \to \Omega'$) such that

$$\widehat{h}_{k(P)}(x, y) \in \mathcal{N}'_P \iff x \in \mathcal{N}_P \tag{4.35}$$

for every $P \in \mathcal{L}$ and a.e. $(x, y) \in \mathcal{E}_{k(P)}(\Omega) \times \mathcal{E}_{k(P)}(\Omega)$.

Pick $(\boldsymbol{\theta}, \boldsymbol{\eta})$ in $\mathcal{E}_{\mathbb{N}_+}(\Omega) \times \mathcal{E}_{\mathbb{N}_+}(\Omega)$ according to $\mu \otimes \mu$, let $\boldsymbol{K}$ be the exchangeable array corresponding to $\mathcal{N}$ with respect to $\boldsymbol{\theta}$ and let us show that $\boldsymbol{K}$ is independent from $(\boldsymbol{\theta}_A \mid A \in \binom{\mathbb{N}_+}{b}, b \in B)$. It is sufficient to show that for every $m \in \mathbb{N}$ every $K \in \mathcal{K}_m[T]$ and every measurable set $U \subseteq X^{\bigcup_{b \in [m] \setminus B} \binom{[m]}{b}}$ with $\mu(U) > 0$, the events $E_1 \stackrel{\text{def}}{=} [\boldsymbol{K}|_{[m]} = K]$ and $E_2 \stackrel{\text{def}}{=} [(\boldsymbol{\theta} \mid A \in \binom{[m]}{b}, b \in [m] \setminus B) \in U]$ are independent. But note that (4.35) implies that E_1 is a.e. equivalent to

$$\forall P \in \mathcal{L}, \forall \alpha \in ([m])_{k(P)}, \alpha \in R_P(K) \leftrightarrow \widehat{h}_{k(P)}(\alpha^*(\boldsymbol{\theta}), \alpha^*(\boldsymbol{\eta})) \in \mathcal{N}'_P.$$

Since $\mathcal{N}'_P$ is B-compatible, the property above does not depend on the coordinates of $\widehat{h}_{k(P)}(\alpha^*(\boldsymbol{\theta}), \alpha^*(\boldsymbol{\eta}))$ indexed by sets A with $|A| \notin B$. Even though the property above may depend on $(\boldsymbol{\theta}_A \mid |A| \notin B)$, since the h_t are measure preserving on h.o.a., it follows that the conditional distribution given E_2 of $(\widehat{h}_{k(P)}(\alpha^*(\boldsymbol{\theta}), \alpha^*(\boldsymbol{\eta}))_A \mid A \in \binom{[m]}{b}, b \in B)$ is μ' and thus by letting $\boldsymbol{\theta}'$ be picked in $\mathcal{E}_{\mathbb{N}_+}(\Omega')$ according to μ', we have

$$\begin{aligned} \mathbb{P}[E_1 \mid E_2] &= \mathbb{P}[\forall P \in \mathcal{L}, \forall \alpha \in ([m])_{k(P)}, \alpha \in R_P(K) \leftrightarrow \alpha^*(\boldsymbol{\theta}') \in \mathcal{N}'_P] \\ &= \phi_{\mathcal{N}'}(\langle K \rangle) = \phi_{\mathcal{N}}(\langle K \rangle) = \mathbb{P}[E_1], \end{aligned}$$

as desired. ■

Proof of Theorem 4.8.5. By Theorem 4.2.6, we know that ψ_ℓ satisfies $\texttt{UCouple}[\ell - 1]$ but does not satisfy `Independence`$[\ell-1]$. Since $\mathrm{rk}(\psi_\ell) \leq \ell = \min B$, it follows that ψ_ℓ is not $(B \setminus \{\ell\})$-compatible. By Theorem 4.8.4, we know that ψ_ℓ is completely weakly $[\ell-1]$-independent and since $\mathrm{rk}(\psi_\ell) \leq \ell$, it follows that ψ_ℓ is completely weakly $(\mathbb{N}_+ \setminus \{\ell\})$-independent, which gives $\psi_\ell \in$ `UCouple`$[\mathbb{N}_+ \setminus \{\ell\}]$ again by Theorem 4.8.4. ■

4.9 Concluding remarks and open problems

In this chapter we have attempted to build a general theory of quasirandomness that is uniformly applicable to arbitrary combinatorial structures and is invariant under their "natural transformations". While our basic definitions deliberately avoided mentioning specific densities, it turned out, in the vein of the previous research in the area, that our quasirandom properties can be characterized in several equivalent ways, *including* such densities. We have shown how to arrange these properties into a hierarchy and, with one or two notable exceptions, have been able to prove that this hierarchy is proper. Finally, we have compared our quasirandom properties to what has been studied before for hypergraphs (with the focus on specific densities) and have found that these two frameworks are essentially incomparable.

One topic that we touched tangentially in the proof of Theorem 4.2.10, more specifically with Example 5 and Lemma 4.5.5, is the closedness of our properties with respect to both the density topology and L_1-topology (Definition 3.3.1). The aforementioned example and lemma show that in general unique coupleability with a particular collection of limit objects is closed in L_1-topology but not necessarily closed in the density topology. On the other hand, alternative syntactic descriptions of `UCouple`$[\ell]$ and `UInduce`$[\ell]$ (as ℓ-locality and symmetric ℓ-locality, respectively) imply that these classes are closed even in the density topology. So in a sense we have a satisfactory overall picture for the classes based on the "extrinsic" notion of coupleability.

As we briefly mentioned in Section 3.8, we do not know how the class `Independence`$[\ell]$ interacts with the different topologies, or even if it has a very clean and natural "intrinsic" definition. We reiterate the question of that section here: is $\texttt{Independence}[\ell]$ closed in the density, or at least L_1-topology? One sensible approach to this question might consist in developing an alternative, and perhaps more concrete, characterization of this class that might be interesting in its own right.

If ϕ_1 and ϕ_2 are uniquely coupleable with *all* theons of rank $\leq \ell$, then the same is true for $\phi_1 \otimes \phi_2$ (Theorem 4.2.4 (ii)). We do not know if the same remains true after replacing this class of tests with individual tests, and when we needed this in one of our proofs, we had to take a considerable detour (see item 3 in our program at the beginning of Section 4.5). Thus comes our second open question: assume that ϕ_1 and ψ, as well as ϕ_2 and ψ are uniquely coupleable. Does it imply that $\phi_1 \otimes \phi_2$ is also uniquely coupleable with ψ?

Under the additional assumption that ϕ_1, ϕ_2 are themselves uniquely coupleable, the question takes a particularly nice and symmetric form: assume that ϕ_1, ϕ_2 and $\phi_3 (= \psi)$ are pairwise uniquely coupleable. Does it imply that ϕ_1, ϕ_2, ϕ_3 are (mutually) uniquely coupleable? While the analogy with independence for random variables is now visible, it is not immediately clear how useful it might turn out here.

Another interesting question is whether unique coupleability establishes a Galois correspondence between $\texttt{UCouple}[\ell]$ and limit objects of rank at most ℓ. In other words, is it true that if ϕ is uniquely coupleable with every $\psi \in \texttt{UCouple}[\ell]$, then $\text{rk}(\phi) \leq \ell$?

As we mentioned before, the results of Theorems 4.2.1, 4.2.2, 4.2.5, 4.2.6 and 4.2.7 almost complete the Hasse diagram of implications between the families `Independence`, `UCouple` and `UInduce`. After personal communication with Henry Towsner, we obtained an argument for $\texttt{UCouple}[\ell] \implies \texttt{Independence}[\ell - 1]$, which will appear in a future joint work. Along with the aforementioned theorems, this completes the Hasse diagram of implications between the hierarchies of properties `Independence`, `UCouple` and `UInduce`, with the first two hierarchies with the first two intercalated.

Recall that Theorem 4.2.10(i)$\Leftrightarrow$(vii) says that $\phi \in \texttt{UCouple}[\ell]$ is equivalent to $\phi \otimes \psi_{\text{lin}} \in$

$\texttt{UInduce}[\ell]$. Let us now draw attention to three interesting open problems that can be extracted from this equivalence.

The first is whether a "converse" of this is true in the spirit of Theorems 4.2.12 and 4.2.13: can every $\phi \in \texttt{UInduce}[\ell]$ be written as $\phi = (\widehat{\phi} \otimes \psi_{\text{lin}})^I$ for some $\widehat{\phi} \in \texttt{UCouple}[\ell]$ and some open interpretation $I\colon T \rightsquigarrow T' \cup T_{\text{LinOrder}}$?

The second problem is an analogue of Theorems 4.2.12 and 4.2.13 themselves in the context of unique inducibility. We conjecture that Theorem 4.2.14 can be generalized to characterize $\texttt{UInduce}[k-1]$ when all arities are at most k (of course, this would follow from a positive answer to the previous problem).

The third question is more open-ended. In the three scenarios discussed in Section 4.2.1 (permutations, words and Latin squares), the quasirandom object is "straightforward" but does not satisfy even the weakest of our properties $\texttt{UInduce}[1]$. Hence we might reasonably ask if the theory of "natural" (understood as in the introduction) quasirandomness properties can be extended beyond $\texttt{UInduce}[1]$. One possibility would be to consider the closure of $\texttt{UInduce}[\ell]$ under independent couplings and open interpretations. Both the quasirandom permuton $\psi_{\text{lin}} \otimes \psi_{\text{lin}}$ and the quasirandom Latin square $\psi_{\text{lin}} \otimes \psi_{\text{lin}} \otimes \psi_{\text{lin}}$ belong to this class (for every ℓ). This definition, however, is of the same distinctly ad hoc nature we have been trying to avoid in this paper. Are there any "reasonable" descriptions of this class, be them extrinsic or intrinsic? The only thing we can prove (and even that is non-trivial) is that this class is proper, i.e., there are theons that do not belong to it, for an arbitrary ℓ. If the conjectures from the previous two paragraphs are true, this would also form another interesting hierarchy: starting from $\texttt{UCouple}[\ell]$, we can get progressively weaker families of natural quasirandomness properties by taking independent coupling with the linear order ψ_{lin}.

Another possible approach would be to start with quasirandom permutations that is by far the most widely studied class, and from their known properties [22, 23, 49, 10]. However, in comparison to their (hyper)graph and tournament counterparts, the theory of permutation quasirandomness provides a much smaller variety of quasirandomness formulations as candidates for natural generalizations, essentially boiling down to only three types: explicit density notions, discrepancy notions based on intervals and spectral notions. Let us also note that there is still a whole host of properties [32, 12] that random permutations satisfy and that have not yet been fully explored in the quasirandom setting. In fact, some of these properties are so fine-grained that it is not even clear if they can be encoded by subpermutation densities.

In Section 4.8, we saw that the notions of rank and `Independence` are naturally generalized by the notion of compatibility and several of the results of Section 4.2 carry over. The main difference is how only some of the items of Theorem 4.2.10 are generalized to Theorem 4.8.4. Namely, in Theorem 4.8.4, we only have an analogue (item (iv)) of *complete* weak independence (item (v) of Theorem 4.2.10) but not weak independence (item (iv)). We also do not have any analogue of locality (item (vi)), which in particular would give a syntactic description of `Independence`$[\ell]$ (i.e., $\mathbb{N}_+ \setminus [\ell]$-compatibility). Finally, an analogue of item (vii) would require also providing an analogue of `UInduce` in the setting of compatibility. Filling any of these omissions (or showing that the analogous items are not equivalent to `UCouple`$[B]$) in Theorem 4.8.4 is an interesting problem.

CHAPTER 5

ABSTRACT CHROMATIC NUMBER

As we mentioned in Chapter 1, the celebrated Erdős–Stone–Simonovits Theorem and its later generalization by Alon–Shikhelman stated below characterize the maximum asymptotic density of t-cliques K_t in graphs without non-induced copies of graphs in a family $\mathcal{F}$ in terms of the chromatic number $\chi(\mathcal{F})$.

Theorem 5.0.1 (Erdős–Stone–Simonovits [38, 37], Alon–Shikhelman [2])**.** *Let $t \in \mathbb{N}$ and let $\mathcal{F}$ be a non-empty family of finite non-empty graphs. The maximum number of copies of t-cliques K_t in a graph G with n vertices and without any non-induced copies of elements of $\mathcal{F}$ is*

$$\prod_{j=1}^{t-1}\left(1 - \frac{j}{\chi(\mathcal{F}) - 1}\right)\binom{n}{t} + o(n^t),$$

where $\chi(\mathcal{F}) \overset{\text{def}}{=} \min\{\chi(F) \mid F \in \mathcal{F}\}$ is the minimum chromatic number of a graph in $\mathcal{F}$.

In this chapter we provide a generalization that answers the following question: given an open interpretation $I\colon T_{\text{Graph}} \rightsquigarrow T$, what is the maximum asymptotic density of t-cliques K_t in graphs of the form $I(M)$ for $M \in \mathcal{M}[T]$? We will see that an analogue of Theorem 5.0.1 above holds by replacing $\chi(\mathcal{F})$ with an *abstract chromatic number* $\chi(I)$ and Theorem 5.0.1 can then be retrieved by simply letting T be the theory of graphs without non-induced copies of graphs in $\mathcal{F}$ and I be the axiom-adding interpretation. We will also show how to retrieve analogues of Theorem 5.0.1 from the literature of ordered graphs [57], cyclically ordered graphs [9] and edge-ordered graphs [41].

The case $t = 2$ of such generalization was first shown in [24, Examples 25 and 31]. However, the formula for $\chi(I)$ presented in [24, Equation (16)] (see (5.2) in Section 5.1 below) is considerably abstract and it was left open if $\chi(I)$ is (algorithmically) computable even when T is assumed to be finitely axiomatizable. In this chapter, we will also prove an alternative, more concrete formula for $\chi(I)$ (Theorem 5.2.2). Such formula allows us to deduce that when T is *finitely* axiomatizable, then $\chi(I)$ is (algorithmically) computable from a list of the axioms of T and a description of I (Theorem 5.2.3). Our alternative formula is based on a partite version of Ramsey's Theorem (Theorem 5.1.8) for universal theories that informally says that given $\ell, m \in \mathbb{N}$, there exists $n \in \mathbb{N}$ such that for every model M and every partition of M into ℓ parts all of size at least n must have a "uniform" submodel on the same partition with all parts of size m (this version of Ramsey's Theorem for disjoint unions of theories of hypergraphs follows from [42, Section 5] and the non-partite version, when $\ell = 1$, for general theories follows from the general Ramsey theory for systems of [55]; see Section 5.1.2 for more details). By using these different formulas for $\chi(I)$, we can retrieve the results of [57, 9, 41] on ordered graphs, cyclically ordered graphs and edge-ordered graphs, respectively from the general theory (see Section 5.7).

5.1 Preliminaries

5.1.1 The general Turán density and the abstract chromatic number

In the theory of graphs T_{Graph}, we denote the *complete graph* on n vertices by $K_n \in \mathcal{M}_n[T_{\text{Graph}}]$, that is, we have $R_E(K_n) \stackrel{\text{def}}{=} (V(K_n))_2$; we denote the *empty graph* on n vertices by $\overline{K}_n$, that is, we have $R_E(\overline{K}_n) \stackrel{\text{def}}{=} \varnothing$; and we denote the *$\ell$-partite Turán graph* of size n by $T_{n,\ell} \in \mathcal{M}_n[T_{\text{Graph}}]$, that is, $T_{n,\ell}$ is the complete ℓ-partite graph with parts of sizes either $\lfloor n/\ell \rfloor$ or $\lceil n/\ell \rceil$, or in a formula, we have $R_E(T_{n,\ell}) \stackrel{\text{def}}{=} \{\alpha \in ([n])_2 \mid \alpha_1 \not\equiv \alpha_2 \pmod{\ell}\}$. For graphs G and H, we write $G \subseteq H$ if H has a *non-induced copy* of G, that is, if there is a positive embedding of G in H (i.e., if there exists $f\colon V(G) \rightarrowtail V(H)$ that maps edges of G to edges of H, or in formulas, for every $\alpha \in R_E(G)$, we have $f \circ \alpha \in R_E(H)$). Recall that a proper coloring of a graph G is a function $f\colon V(G) \to [\ell]$ such that $\forall \alpha \in R_E(G), f(\alpha_1) \neq f(\alpha_2)$ and the *chromatic number* of G is the minimum $\ell \in \mathbb{N}$ such that there exists a proper coloring of G of the form $f\colon V(G) \to [\ell]$.

Definition 5.1.1 (Abstract Turán density)**.** For an open interpretation $I\colon T_{\text{Graph}} \rightsquigarrow T$ and $t \in \mathbb{N}$, the *t-Turán density* of I is defined as

$$\pi_I^t \stackrel{\text{def}}{=} \lim_{n\to\infty} \sup_{N\in\mathcal{M}_n[T]} p(K_t, I(N)). \tag{5.1}$$

The existence of the limit in (5.1) follows from the fact that the sequence is non-increasing (for $n \geq t$). This can be proved by the standard averaging argument of extremal combinatorics: if T is degenerate, then the sequence is eventually constant equal to $-\infty$; otherwise, if $N_0 \in \mathcal{M}_{n+1}[T]$ ($n \geq t$) maximizes $p(K_t, I(N_0))$, then picking uniformly at random a subset $\boldsymbol{U}$ of $V(N_0)$ of size n, we conclude that

$$\sup_{N\in\mathcal{M}_n[T]} p(K_t, I(N)) \geq \mathbb{E}[p(K_t, I(N_0|_{\boldsymbol{U}}))] = p(K_t, I(N_0)) = \sup_{N\in\mathcal{M}_{n+1}[T]} p(K_t, I(N)).$$

Note also that since π_I^t is stated in terms of densities, when we count copies of K_t instead, we incur an $o(n^t)$ error.

Definition 5.1.2 (Abstract chromatic number [24, Equation (16)])**.** For an open interpretation $I\colon T_{\text{Graph}} \rightsquigarrow T$, the *abstract chromatic number of* I is defined as[1]

$$\chi(I) \stackrel{\text{def}}{=} \sup\{\ell \in \mathbb{N}_+ \mid \forall n \in \mathbb{N}, \exists N \in \mathcal{M}_n[T], T_{n,\ell} \subseteq I(N)\} \cup \{0\} + 1. \tag{5.2}$$

Note that $\chi(I) \in \mathbb{N}_+ \cup \{\infty\}$ because the set in (5.2) always contains 0. Furthermore, note that if T is degenerate, then $\chi(I) = 1$ as the set in (5.2) is $\{0\}$.

The usual Turán density studied in Theorem 5.0.1 is $\pi_{I_{\mathcal{F}}}^t$ for the axiom-adding interpretation $I_{\mathcal{F}}\colon T_{\text{Graph}} \rightsquigarrow \operatorname{Forb}^+_{T_{\text{Graph}}}(\mathcal{F})$, where $\operatorname{Forb}^+_{T_{\text{Graph}}}(\mathcal{F})$ is the theory obtained from T_{Graph}

1. The formula in (5.2) is actually a slight modification of [24, Equation (16)], forcing 0 to belong to the set. This is done so that we can also cover degenerate theories T.

by adding for each $F \in \mathcal{F}$ the axiom

$$\forall x_1 \cdots \forall x_m, \neg \left(\bigwedge_{1 \leq i < j \leq m} x_i \neq x_j \wedge \bigwedge_{\alpha \in R_E(F)} E(x_{\alpha_1}, x_{\alpha_2}) \right),$$

where we rename the vertices of F so that $V(F) = [m]$. We will see in Proposition 5.7.1 that in this case $\chi(I_{\mathcal{F}})$ is equal to the usual chromatic number $\chi(\mathcal{F}) \stackrel{\text{def}}{=} \inf\{\chi(F) \mid F \in \mathcal{F}\}$ except for when $\mathcal{F}$ is empty or contains an empty graph; more precisely, we have $\chi(I_{\mathcal{F}}) = \max\{\chi(\mathcal{F}), 1\}$.

5.1.2 Partite Ramsey numbers

As we mentioned in the beginning of the chapter, our alternative formula for the abstract chromatic number is based on a partite version of Ramsey's Theorem for universal theories. The first step to this version is identifying what are the "uniform" structures that are unavoidable in a large structure. Let us start with the easier case in which all predicate symbols are symmetric: this is captured by the theories of $\vec{k}$-hypergraphs of Definition 4.4.6.

Any ordered partition $(V_1, \ldots, V_\ell)$ of a set V can be described alternatively by the function $f\colon V \to [\ell]$ such that $v \in V_{f(v)}$ for every $v \in V$. We can then classify the subsets $e \subseteq V$ according to how many points e contains in each of the parts V_i. The notions of Ramsey patterns and uniform $\vec{k}$-hypergraphs defined below explore this classification .

Definition 5.1.3 ($\vec{k}$-hypergraph Ramsey patterns and uniform $\vec{k}$-hypergraphs)**.** Recall that for $\ell, k \in \mathbb{N}_+$, a *weak composition* of k of length ℓ is an ℓ-tuple $q = (q_j)_{j=1}^{\ell} \in \mathbb{N}^{\ell}$ such that $\sum_{j=1}^{\ell} q_j = k$. We denote the set of weak compositions of k of length ℓ by $\mathcal{C}_{\ell,k}$.

For $\vec{k} = (k_1, \ldots, k_t) \in \mathbb{N}_+^t$ and $\ell \in \mathbb{N}_+$, a *$\vec{k}$-hypergraph ℓ-Ramsey pattern* is a t-tuple $Q = (Q_i)_{i \in [t]}$ such that $Q_i \subseteq \mathcal{C}_{\ell,k_i}$ for every $i \in [t]$. We let $\mathcal{P}_{\ell,\vec{k}}$ be the set of all $\vec{k}$-hypergraph ℓ-Ramsey patterns.

Given a $\vec{k}$-hypergraph ℓ-Ramsey pattern $Q \in \mathcal{P}_{\ell,\vec{k}}$, a $\vec{k}$-hypergraph H and a function $f\colon V(H) \to [\ell]$, we say that H is *Q-uniform with respect to f* if for every $i \in [t]$, the E_i-edges of H are precisely those e such that there exists some $q \in Q_i$ such that e contains exactly q_j

points in $f^{-1}(j)$, or in formulas we have

$$E_i(H) = \left\{ e \in \binom{V(H)}{k_i} \;\middle|\; (|e \cap f^{-1}(j)|)_{j\in[t]} \in Q_i \right\},$$

which is in turn equivalent to

$$R_{E_i}(H) = \{\alpha \in (V(H))_{k_i} \mid (|(f \circ \alpha)^{-1}(j)|)_{j\in[t]} \in Q_i\}.$$

The partite version of Ramsey's Theorem for $\vec{k}$-hypergraphs (Theorem 5.1.5 below) says that uniform $\vec{k}$-hypergraphs cannot all be avoided as long as the parts of the partition are sufficiently large.

Definition 5.1.4 (Thickness and $\vec{k}$-hypergraph Ramsey numbers)**.** The *thickness* of a function $f\colon V \to [\ell]$ is $\mathrm{th}(f) \stackrel{\text{def}}{=} \min\{|f^{-1}(i)| \mid i \in [\ell]\}$.

Given $\ell \in \mathbb{N}_+$ and $m \in \mathbb{N}$, the $(\ell, \vec{k}, m)$-*Ramsey number* $R_{\ell,\vec{k}}(m)$ is defined as the least $n \in \mathbb{N}$ such that for every $\vec{k}$-hypergraph H and every $f\colon V(H) \to [\ell]$ with $\mathrm{th}(f) \geq n$, there exists $Q \in \mathcal{P}_{\ell,\vec{k}}$ and a set $W \subseteq V(H)$ such that $\mathrm{th}(f|_W) \geq m$ and $H|_W$ is Q-uniform with respect to $f|_W$.

Theorem 5.1.5. *For every $\ell \in \mathbb{N}_+$, every $m \in \mathbb{N}$ and every $\vec{k} \in \mathbb{N}_+^t$, the $(\ell, \vec{k}, m)$-Ramsey number $R_{\ell,\vec{k}}(m)$ is finite.*

Theorem 5.1.5 above can be obtained e.g. by repeatedly applying [42, Theorem 5 of Section 5], but we provide a proof via a reduction to Ramsey's original theorem for hypergraphs in Section 5.4.

For the case of general universal theories, we have an extra technicality: predicate symbols are not necessarily symmetric. The correct way of addressing this issue is illustrated by the case of the theory of tournaments $T_{\text{Tournament}}$. The unavoidable "uniform" models here are the transitive tournaments Tr_n (with $R_E(\mathrm{Tr}_n) \stackrel{\text{def}}{=} \{\alpha \in ([n])_2 \mid \alpha_1 < \alpha_2\}$): for every $k \in \mathbb{N}$, every sufficiently large tournament M must contain a transitive tournament of size k as a subtournament [63, 36]. Another way of seeing a transitive tournament is that there is an underlying order $\leq$ of its vertices such that we can decide whether $\alpha \in ([n])_2$ is in $R_E(\mathrm{Tr}_n)$ based only on the relative order of α_1 and α_2 with respect to $\leq$. In the ℓ-partite case, the role of the order $\leq$ is played by the ℓ-split orders defined below, which are tuples $(f, \preceq)$ such that $f\colon V \to [\ell]$ encodes an ℓ-partition and $\preceq$ orders each of the parts of this partition.

Definition 5.1.6 (Split orders)**.** For $\ell \in \mathbb{N}_+$ and a set V, an *ℓ-split order* over V is a pair $(f, \preceq)$, where $f\colon V \to [\ell]$ and $\preceq$ is a partial order on V such that

$$\forall v, w \in V, (f(v) = f(w) \leftrightarrow v \preceq w \vee w \preceq v),$$

that is, two elements of V are comparable under $\preceq$ if and only if they have the same image under f. We let $\mathcal{S}_{\ell,V}$ be the set of all ℓ-split orders over V and for $k \in \mathbb{N}$, we use the shorthand $\mathcal{S}_{\ell,k} \stackrel{\text{def}}{=} \mathcal{S}_{\ell,[k]}$.

When $\ell = 1$, we will typically omit f from the notation as it must be the constant function; with this abuse, we will think of $\mathcal{S}_{1,V}$ as the set of all total orders on V.

For a partial order $\preceq$ on a set V and an injective function $g\colon W \rightarrowtail V$, we let $\preceq_g$ be the partial order on W defined by

$$w_1 \preceq_g w_2 \iff g(w_1) \preceq g(w_2).$$

If $W \subseteq V$, then we let $\preceq_W \stackrel{\text{def}}{=} \preceq_{\iota_W}$, where $\iota_W\colon W \rightarrowtail V$ is the canonical injection, that is, $\preceq_W$ is just the restriction $\preceq \cap (W \times W)$ of $\preceq$ to W.

Note that for $g\colon W \to V$ and $h\colon U \to W$ and for a partial order $\preceq$ on V, we have $(\preceq_g)_h = \preceq_{g \circ h}$. Furthermore, if $(f, \preceq) \in \mathcal{S}_{\ell,V}$, then $(f \circ g, \preceq_g) \in \mathcal{S}_{\ell,W}$. Finally, note that there are finitely many ℓ-split orders over $[k]$.

Given an ℓ-split order $(f, \preceq) \in \mathcal{S}_{\ell,V}$ over V, we can classify the tuples $\alpha \in (V)_k$ according to $(f \circ \alpha, \preceq_\alpha)$, that is, $f \circ \alpha$ captures the values of f on the image of α and $\preceq_\alpha$ captures the partial order induced by $\preceq$ on the image of α. Just as in the case of $\vec{k}$-hypergraphs, the notions of Ramsey patterns, uniform structures and Ramsey numbers defined below explore this classification.

Definition 5.1.7 (Ramsey patterns, uniform structures and Ramsey number)**.** Fix $\ell \in \mathbb{N}_+$ and a language $\mathcal{L}$. An *ℓ-Ramsey pattern on $\mathcal{L}$* is a function Q that maps each predicate symbol $P \in \mathcal{L}$ to a collection $Q_P \subseteq \mathcal{S}_{\ell,k(P)}$ of ℓ-split orders on $[k(P)]$. We let $\mathcal{P}_{\ell,\mathcal{L}}$ be the set of all ℓ-Ramsey patterns on $\mathcal{L}$.

Given an ℓ-Ramsey pattern $Q \in \mathcal{P}_{\ell,\mathcal{L}}$ on $\mathcal{L}$, a canonical structure M on $\mathcal{L}$ and an ℓ-split order $(f, \preceq) \in \mathcal{S}_{\ell,V(M)}$ on $V(M)$, we say that M is *Q-uniform with respect to* $(f, \preceq)$ if for every $P \in \mathcal{L}$, we have

$$R_P(M) = \{\alpha \in (V(M))_{k(P)} \mid (f \circ \alpha, \preceq_\alpha) \in Q_P\}.$$

For a canonical structure M on $\mathcal{L}$, the *ℓ-Ramsey uniformity set* of M is the set $\mathcal{U}_\ell(M)$ of all ℓ-Ramsey patterns $Q \in \mathcal{P}_{\ell,\mathcal{L}}$ such that M is Q-uniform with respect to some $(f, \preceq) \in \mathcal{S}_{\ell,V(M)}$. We extend this definition to a family $\mathcal{F}$ of canonical structures as $\mathcal{U}_\ell(\mathcal{F}) \stackrel{\text{def}}{=} \bigcup_{M\in\mathcal{F}} \mathcal{U}_\ell(M)$.

Given a canonical theory T over $\mathcal{L}$ and $m \in \mathbb{N}$, the *(ℓ, T, m)-Ramsey number* $R_{\ell,T}(m)$ is defined as the least $n \in \mathbb{N}$ such that for every model M of T and every ℓ-split order $(f, \preceq) \in \mathcal{S}_{\ell,V(M)}$ on $V(M)$ with $\text{th}(f) \geq n$, there exists an ℓ-Ramsey pattern $Q \in \mathcal{P}_{\ell,\mathcal{L}}$ over $\mathcal{L}$ and a set $W \subseteq V(M)$ such that $\text{th}(f|_W) \geq m$ and $M|_W$ is Q-uniform with respect to $(f|_W, \preceq_W)$.

Note that since $\mathcal{L}$ is finite, there are only finitely many ℓ-Ramsey patterns on $\mathcal{L}$. Note also that the definition of $R_{\ell,T}(m)$ is strong in the sense that *every* ℓ-split order of $V(M)$ is required to yield a uniform submodel. This is slightly stronger than our motivating example of tournaments: our definition for $T_{\text{Tournament}}$ with $\ell = 1$ requires that every ordering $\leq$ of the vertices of M yields a tournament of size m whose edges either all match the order $\leq$ or all disagree with $\leq$.

Example 7. In the language $\mathcal{L}$ containing a single predicate symbol E of arity $k(E) = 2$, for every $n \geq 2$, there are exactly three (up to isomorphism) canonical structures M of size n that are Q-uniform for some 1-Ramsey pattern $Q \in \mathcal{P}_{1,\mathcal{L}}$ with respect to some $(f, \preceq) \in \mathcal{S}_{1,V(M)}$: the complete graph K_n, the empty graph $\overline{K}_n$ and the transitive tournament Tr_n. Note also that for $n \geq 2$, both $\mathcal{U}_1(K_n)$ and $\mathcal{U}_1(\overline{K}_n)$ have a single element but $\mathcal{U}_1(\mathrm{Tr}_n)$ has two elements.

In the same language, canonical structures M that are Q-uniform for some ℓ-Ramsey pattern Q with respect to some $(f, \preceq)$ are precisely those in which each level set $f^{-1}(i)$ of f induces either a complete graph $K_{|f^{-1}(i)|}$, an empty graph $\overline{K}_{|f^{-1}(i)|}$ or a transitive tournament $\mathrm{Tr}_{|f^{-1}(i)|}$ and (directed) edges between $v, w \in V(M)$ in different level sets of f are completely determined by $f(v)$ and $f(w)$. See Figure 5.1.

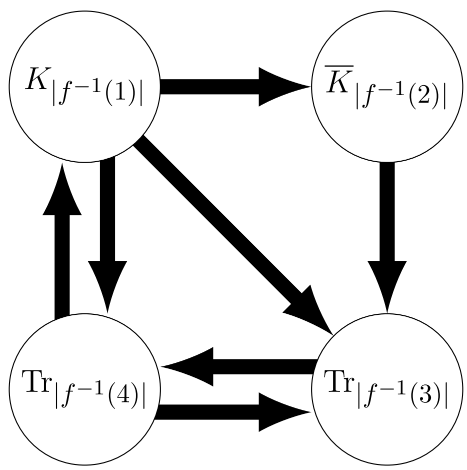

Figure 5.1: Pictorial view of a Q-uniform model for the Ramsey pattern $Q \in \mathcal{P}_{4,\{E\}}$ $(k(E) = 2)$ given by

$$\begin{aligned} Q_E \stackrel{\text{def}}{=} \{&((1,1), \leq), ((1,1), \geq), ((3,3), \leq), ((4,4), \geq), \\ &((1,2), \preceq_0), ((1,3), \preceq_0), ((1,4), \preceq_0), \\ &((2,3), \preceq_0), ((3,4), \preceq_0), ((4,3), \preceq_0), ((4,1), \preceq_0)\}, \end{aligned}$$

where $\leq$ is the usual order on $[2]$, $\geq$ is its reverse and $\preceq_0$ is the trivial partial order on $[2]$, and the functions $f\colon [2] \to [4]$ are represented as $(f(1), f(2))$. An arrow from a part A to a part B in the figure means that $(a, b) \in R_E(M)$ for every $a \in A$ and every $b \in B$.

Theorem 5.1.8. *For every $\ell \in \mathbb{N}_+$, every $m \in \mathbb{N}$ and every canonical theory T, the (ℓ, T, m)-Ramsey number $R_{\ell,T}(m)$ is finite.*

We provide a proof of Theorem 5.1.8 via a reduction to Theorem 5.1.5 in Section 5.4. Let us also note that the case $\ell = 1$ of Theorem 5.1.8 follows from the very general Ramsey Theory for systems of [55].

We will typically be working in theories of the form $T_{\text{Graph}} \cup T$ and two types of Ramsey patterns will play an important role in the alternative formula for the abstract chromatic number.

Definition 5.1.9 (Complete patterns and Turán patterns)**.** Fix $\ell \in \mathbb{N}_+$ and a language $\mathcal{L}$ and let $E \in \mathcal{L}$ be a binary predicate symbol.

A 1-Ramsey pattern $Q \in \mathcal{P}_{1,\mathcal{L}}$ on $\mathcal{L}$ is called *E-complete* if $Q_E = \mathcal{S}_{1,2}$. We let $\mathcal{C}^E_{\mathcal{L}}$ be the set of all E-complete 1-Ramsey patterns on $\mathcal{L}$.

An ℓ-Ramsey pattern $Q \in \mathcal{P}_{\ell,\mathcal{L}}$ on $\mathcal{L}$ is called *E-Turán* if

$$Q_E = \{(g, \preceq) \in \mathcal{S}_{\ell,2} \mid g \text{ is injective}\}.$$

We let $\mathcal{T}^E_{\ell,\mathcal{L}}$ be the set of all E-Turán ℓ-Ramsey patterns on $\mathcal{L}$.

Note that if $I\colon T_{\{E\}} \rightsquigarrow T_{\mathcal{L}}$ is the structure-erasing interpretation, then M is Q-uniform with respect to some $(f, \preceq) \in \mathcal{S}_{1,V(M)}$ for some E-complete $Q \in \mathcal{C}^E_{\mathcal{L}}$ if and only if $I(M) \cong K_{|M|}$. Analogously, M is Q-uniform with respect to some $(f, \preceq) \in \mathcal{S}_{\ell,V(M)}$ for some E-Turán $Q \in \mathcal{T}^E_{\ell,\mathcal{L}}$ if and only if $I(M)$ is a complete ℓ-partite graph with respect to the partition given by the level sets of f.

5.1.3 Non-induced setting

As we mentioned before, the abstract chromatic number works in the general setting of induced submodels. For the non-induced setting, we will be able to provide a slightly simpler formula for the abstract chromatic number in terms of proper split orderings defined below.

Definition 5.1.10 (E-upward closures and proper split orderings)**.** Let $\mathcal{L}$ be a language and let E be the predicate symbol corresponding to T_{Graph} in the language $\mathcal{L} \cup \{E\}$ of

$T_{\text{Graph}} \cup T_{\mathcal{L}}$.

Given a family $\mathcal{F}$ of models of $T_{\text{Graph}} \cup T_{\mathcal{L}}$, the *$E$-upward closure* of $\mathcal{F}$ is the family $\mathcal{F}\!\uparrow^E$ of all F' that can be obtained from some $F \in \mathcal{F}$ by possibly adding edges, that is, all models F' of $T_{\text{Graph}} \cup T_{\mathcal{L}}$ such that there exists $F \in \mathcal{F}$ with

$$V(F') = V(F); \qquad R_E(F') \supseteq R_E(F); \qquad R_P(F') = R_P(F) \quad (P \in \mathcal{L}).$$

Let $I\colon T_{\text{Graph}} \rightsquigarrow T_{\text{Graph}} \cup T_{\mathcal{L}}$ and $J\colon T_{\mathcal{L}} \rightsquigarrow T_{\text{Graph}} \cup T_{\mathcal{L}}$ be the structure-erasing interpretations. Given $\ell \in \mathbb{N}_+$, an ℓ-Ramsey pattern $Q \in \mathcal{P}_{\ell,\mathcal{L}}$ on $\mathcal{L}$ and a model M of $T_{\text{Graph}} \cup T_{\mathcal{L}}$, an *$E$-proper Q-split ordering* of M is a split order $(f, \preceq) \in \mathcal{S}_{\ell,V(M)}$ such that $J(M)$ is Q-uniform with respect to $(f, \preceq)$ and f is a proper coloring of the graph $I(M)$. The *E-proper ℓ-split ordering set* of M is the set $\chi_\ell^E(M)$ of all ℓ-Ramsey patterns $Q \in \mathcal{P}_{\ell,\mathcal{L}}$ such that M has an E-proper Q-split ordering. We extend this definition to a family $\mathcal{F}$ of canonical structures as $\chi_\ell^E(\mathcal{F}) \stackrel{\text{def}}{=} \bigcup_{M \in \mathcal{F}} \chi_\ell^E(M)$.

Note that in the definition of E-proper Q-split orderings, the predicate symbol E is excluded from the uniformity condition. Note also that if the language $\mathcal{L}$ is empty, then $\mathcal{P}_{\ell,\mathcal{L}}$ has a unique element Q and an E-proper Q-split ordering of M consists of any ℓ-split order $(f, \preceq)$ in which f is a proper coloring of the graph $I(M)$.

5.2 Main results on abstract chromatic number

In this section we formalize the main results. We start with the generalization of Theorem 5.0.1 to the setting of open interpretations. The case when $t = 2$ and T is non-degenerate was done in [24, Example 31].

Theorem 5.2.1. *Let $t \in \mathbb{N}_+$ and let $I\colon T_{\text{Graph}} \rightsquigarrow T$ be an open interpretation. Then*

$$\pi_I^t = \begin{cases} \prod_{i=1}^{t-1}\left(1 - \frac{j}{\chi(I)-1}\right), & \text{if } \chi(I) \geq 2;\\ -\infty, & \text{if } \chi(I) \leq 1. \end{cases} \tag{5.3}$$

The next theorem gives an alternative formula for the abstract chromatic number based on the Ramsey uniformity sets of the forbidden models.

Theorem 5.2.2. *Let $I\colon T_{\text{Graph}} \rightsquigarrow T$ be an open interpretation and let T' be the theory obtained from $T_{\text{Graph}} \cup T$ by adding the axiom*

$$\forall x \forall y, E(x,y) \leftrightarrow I(E)(x,y).$$

Let $\mathcal{L}$ be the language of T' and let $\mathcal{F}$ be such that $T' = \mathrm{Forb}_{T_{\mathcal{L}}}(\mathcal{F})$. Then

$$\chi(I) = \begin{cases} \infty, & \text{if } \mathcal{C}_{\mathcal{L}}^E \not\subseteq \mathcal{U}_1(\mathcal{F});\\ \min\{\ell \in \mathbb{N}_+ \mid \mathcal{T}_{\ell,\mathcal{L}}^E \subseteq \mathcal{U}_\ell(\mathcal{F})\}, & \text{otherwise.} \end{cases} \tag{5.4}$$

Furthermore, if T is itself obtained from $T_{\text{Graph}} \cup T''$ by adding axioms and I acts identically on E, then the same result holds by taking $T' = T$ instead.

Remark 11. In fact, we show that the set in (5.4) is either empty or an infinite interval of $\mathbb{N}_+$ (with the empty case only happening when $\chi(I) = \infty$), and thus we also have

$$\chi(I) = \begin{cases} \infty, & \text{if } \mathcal{C}_{\mathcal{L}}^E \not\subseteq \mathcal{U}_1(\mathcal{F});\\ \max\{\ell \in \mathbb{N}_+ \mid \mathcal{T}_{\ell,\mathcal{L}}^E \not\subseteq \mathcal{U}_\ell(\mathcal{F})\} \cup \{0\} + 1, & \text{otherwise.} \end{cases} \tag{5.5}$$

The alternative formula provided by the theorem above can be used to algorithmically compute $\chi(I)$ when T is finitely axiomatizable.

Theorem 5.2.3. *There exists an algorithm that computes $(\chi(I), \pi_I^t)$ for $I\colon T_{\text{Graph}} \rightsquigarrow T$ for* *a finitely axiomatizable T from a list of the axioms of T, a description of I and $t \in \mathbb{N}$.*

For the case when the theory is the theory of graphs with extra structure with some forbidden submodels that are *non-induced* in the graph part, we can provide slightly simpler formulas for $\chi(I)$. The first theorem provides a formula based on the usual chromatic number, but as abstract as (5.2) and the second provides formulas in terms of proper split orderings.

Theorem 5.2.4. *Let $\mathcal{L}$ be a language, let E be the predicate symbol corresponding to T_{Graph} in the language $\mathcal{L} \cup \{E\}$ of $T_{\text{Graph}} \cup T_{\mathcal{L}}$. Let $\mathcal{F}$ be a family of models of $T_{\text{Graph}} \cup T_{\mathcal{L}}$ and let $I\colon T_{\text{Graph}} \rightsquigarrow \text{Forb}_{T_{\text{Graph}} \cup T_{\mathcal{L}}}(\mathcal{F}{\uparrow}^E)$ act identically on E.*

Then we have

$$\chi(I) = \inf\{\chi(G) \mid G \in \mathcal{M}[T_{\text{Graph}}] \land \forall M \in \mathcal{M}[\text{Forb}_{T_{\text{Graph}} \cup T_{\mathcal{L}}}(\mathcal{F}{\uparrow}^E)], I(M) \ncong G\}. \quad (5.6)$$

Theorem 5.2.5. *Let $\mathcal{L}$ be a language, let E be the predicate symbol corresponding to T_{Graph} in the language $\mathcal{L} \cup \{E\}$ of $T_{\text{Graph}} \cup T_{\mathcal{L}}$ and let J interpret$T_{\mathcal{L}}T_{\text{Graph}} \cup T_{\mathcal{L}}$ be the structure-erasing interpretation. Let $\mathcal{F}$ be a family of models of $T_{\text{Graph}} \cup T_{\mathcal{L}}$ and let $I\colon T_{\text{Graph}} \rightsquigarrow \text{Forb}_{T_{\text{Graph}} \cup T_{\mathcal{L}}}(\mathcal{F}{\uparrow}^E)$ act identically on E.*

Then we have

$$\chi(I) = \inf\{\ell \in \mathbb{N}_+ \mid \mathcal{P}_{\ell,\mathcal{L}} \subseteq \chi_\ell^E(\mathcal{F})\}. \quad (5.7)$$

Furthermore, we have $\chi(I) < \infty$ if and only if $\mathcal{P}_{1,\mathcal{L}} \subseteq \mathcal{U}_1(J(\mathcal{F}))$, where $J(\mathcal{F}) \stackrel{\text{def}}{=} \{J(F) \mid F \in \mathcal{F}\}$.

Remark 12. Just as in the case of Theorem 5.2.2, the set in (5.7) is either empty or an infinite interval of $\mathbb{N}_+$, and thus we also have

$$\chi(I) = \sup\{\ell \in \mathbb{N}_+ \mid \mathcal{P}_{\ell,\mathcal{L}} \nsubseteq \chi_\ell^E(\mathcal{F})\} \cup \{0\} + 1. \quad (5.8)$$

5.3 Abstract Turán densities from abstract chromatic number

The objective of this section is to prove Theorem 5.2.1. Before we do so, we show that the set in the definition of $\chi(I)$ in (5.2) is a non-empty initial interval of $\mathbb{N}$.

Lemma 5.3.1. *Given an open interpretation $I\colon T_{\text{Graph}} \rightsquigarrow T$, the set*

$$\{\ell \in \mathbb{N}_+ \mid \forall n \in \mathbb{N}, \exists N \in \mathcal{M}_n[T], I(N) \supseteq T_{n,\ell}\} \cup \{0\} \tag{5.9}$$

is a non-empty initial interval of $\mathbb{N}$.

In particular, we have

$$\chi(I) = \inf\{\ell \in \mathbb{N}_+ \mid \exists n \in \mathbb{N}, \forall N \in \mathcal{M}_n[T], I(N) \not\supseteq T_{n,\ell}\}. \tag{5.10}$$

Proof. Let X be the set in (5.9). It is clear that $0 \in X$. On the other hand, if $\ell \in X \cap \mathbb{N}_+$, then for every $n \in \mathbb{N}$, there exists $N \in \mathcal{M}_n[T]$ such that $I(N) \supseteq T_{n,\ell}$. So if $\ell' \in [\ell]$ and $n \in \mathbb{N}$, then since $T_{n,\ell'} \subseteq T_{\ell\cdot\lceil n/\ell'\rceil,\ell}$, it follows that there exists $N' \in \mathcal{M}_n[T]$ such that $I(N') \supseteq T_{n,\ell'}$, hence $\ell' \in X$.

Since $\chi(I) = \sup X + 1$ by (5.2) and X is a non-empty initial interval of $\mathbb{N}$, we get $\chi(I) = \inf \mathbb{N} \setminus X$, so (5.10) follows. ■

Proof of Theorem 5.2.1. If $\chi(I) = \infty$, then for every $n \in \mathbb{N}$, there exists $N_n \in \mathcal{M}_n[T]$ such that $I(N_n) \supseteq T_{n,n} = K_n$, so $\pi_I^t = 1$, hence (5.3) holds.

On the other hand, if $\chi(I) = 1$, then by Lemma 5.3.1, there exists $n \in \mathbb{N}$ such that for every $N \in \mathcal{M}_n[T]$, we have $I(N) \not\supseteq T_{n,1} = \overline{K}_n$. But since every graph on n vertices contains a non-induced copy of $\overline{K}_n$, we must have $\mathcal{M}_n[T] = \varnothing$. This means that T is degenerate, hence $\pi_I^t = -\infty$, so (5.3) holds.

Suppose then that $2 \leq \chi(I) < \infty$. For every $n \in \mathbb{N}$, let $N_n \in \mathcal{M}_n[T]$ be such that

$I(N_n) \supseteq T_{n,\chi(I)-1}$. Then we get

$$\pi_I^t \geq \liminf_{n\to\infty} p(K_t, I(N_n)) \geq \liminf_{n\to\infty} p(K_t, T_{n,\chi(I)-1}) = \prod_{j=0}^{t-1}\left(1 - \frac{j}{\chi(I)-1}\right).$$

Suppose now toward a contradiction that $(N_m)_{m\in\mathbb{N}}$ is a sequence of models of T with $|N_m| < |N_{m+1}|$ such that $\lim_{m\to\infty} p(K_t, I(N_m)) > \prod_{j=0}^{t-1}(1 - j/(\chi(I)-1))$. Fix $n \in \mathbb{N}$ and note that Theorem 5.0.1 for $\mathcal{F} \stackrel{\text{def}}{=} \{T_{n,\chi(I)}\}$ implies that there exists $m_n \in \mathbb{N}$ such that $I(N_{m_n}) \supseteq T_{n,\chi(I)}$. By restricting N_{m_n} to a set V of size n such that $I(N_{m_n})|_V \supseteq T_{n,\chi(I)}$, we conclude that there exists $N'_n \in \mathcal{M}_n[T]$ such that $I(N'_n) \supseteq T_{n,\chi(I)}$ so $\chi(I) \geq \chi(I) + 1$, a contradiction (as $\chi(I) < \infty$). ■

5.4 Partite Ramsey numbers

The objective of this section is to prove Theorems 5.1.5 and 5.1.8.

Proof of Theorem 5.1.5. The proof is by induction in the length t of the tuple $\vec{k} = (k_1, \ldots, k_t)$.

For the case $t = 1$, let us denote k_1 simply by k and let us identify $\mathcal{P}_{\ell,\vec{k}}$ with $2^{\mathcal{C}_{\ell,k}}$. Let $c \stackrel{\text{def}}{=} |\mathcal{P}_{\ell,\vec{k}}| < \infty$ and let $n \stackrel{\text{def}}{=} R(k, c, \ell m) < \infty$ be the usual Ramsey number corresponding to finding monochromatic cliques of size ℓm in colorings of k-uniform complete hypergraphs with c colors. We will show that $R_{\ell,\vec{k}}(m) \leq n$.

Suppose H is a $\vec{k}$-hypergraph and $f\colon V(H) \to [\ell]$ has $\operatorname{th}(f) \geq n$. For every $j \in [\ell]$, let $v(1, j), \ldots, v(n, j)$ be distinct vertices in $f^{-1}(j)$ and let $V \stackrel{\text{def}}{=} \{v(i, j) \mid i \in [n] \wedge j \in [\ell]\}$.

Recall that for a set $A \in \binom{[n]}{k}$, $\iota_A\colon [k] \rightarrowtail [n]$ denotes the injective function that enumerates A in increasing order and if we are further given a weak composition $q = (q_j)_{j=1}^{\ell} \in \mathcal{C}_{\ell,k}$, let $A_q \subseteq V$ be defined by

$$A_q \stackrel{\text{def}}{=} \left\{ v(\iota_A(i), j) \;\middle|\; i \in [k] \wedge j \in [\ell] \wedge \sum_{r=1}^{j-1} q_r < i \leq \sum_{r=1}^{j} q_r \right\}. \tag{5.11}$$

Note that $|A_q| = k$ and $|f^{-1}(j) \cap A_q| = q_j$ for every $j \in [\ell]$. Furthermore, if $q \neq q'$, then $A_q \neq A_{q'}$.

Define the coloring $g\colon \binom{[n]}{k} \to \mathcal{P}_{\ell,\vec{k}}$ by letting

$$g(A) \stackrel{\text{def}}{=} \{q \in \mathcal{C}_{\ell,k} \mid A_q \in E(H)\},$$

where $E(H)$ is the edge set of H. By the definition of $n = R(k, c, \ell m)$, there exists $U \subseteq [n]$ of size $|U| = \ell m$ such that $g|_{\binom{U}{k}}$ is monochromatic, say, of color $Q \in \mathcal{P}_{\ell,\vec{k}}$.

Let us enumerate the elements of U in increasing order $u_1 < \cdots < u_{\ell m}$ and let

$$W \stackrel{\text{def}}{=} \{v(u_{(j-1)m+r}, j) \mid j \in [\ell] \wedge r \in [m]\}.$$

Clearly, for every $j \in [\ell]$, we have $W \cap f^{-1}(j) = \{v(u_{(j-1)m+r}, j) \mid r \in [m]\}$, which has size m, so $\operatorname{th}(f|_W) = m$.

We claim that $H|_W$ is Q-uniform with respect to $f|_W$. To show this, we need to show that for every $B \in \binom{W}{k}$, we have

$$B \in E(H) \iff q^B \in Q, \tag{5.12}$$

where $q^B \in \mathcal{C}_{\ell,k}$ is given by $q^B_j \stackrel{\text{def}}{=} |f^{-1}(j) \cap B|$.

Note that the definition of W implies that there exists an increasing function $\eta_B\colon [k] \rightarrowtail [n]$ with $\operatorname{im}(\eta_B) \subseteq U$ such that

$$B = \left\{ v(\eta_B(i), j) \;\middle|\; i \in [k] \wedge j \in [\ell] \wedge \sum_{r=1}^{j-1} q^B_r < i \leq \sum_{r=1}^{j} q^B_r \right\}.$$

Since $\iota_{\operatorname{im}(\eta_B)} = \eta_B$, from (5.11) we get $\operatorname{im}(\eta_B)_{q^B} = B$ and for every $q \in \mathcal{C}_{\ell,k} \setminus \{q^B\}$, we have

$\operatorname{im}(\eta_B)_q \neq B$. Since $g|_{\binom{U}{k}}$ is monochromatic of color Q, we have

$$Q = g(\operatorname{im}(\eta_B)) = \{q \in \mathcal{C}_{\ell,k} \mid \operatorname{im}(\eta_B)_q \in E(H)\},$$

so (5.12) follows, concluding the proof of case $t = 1$.

Suppose now that $t \geq 2$ and, by inductive hypothesis, suppose $m' \stackrel{\text{def}}{=} R_{\ell,(k_1,\ldots,k_{t-1})}(m)$ is finite. Let also $n \stackrel{\text{def}}{=} R_{\ell,(k_t)}(m')$, which by the case above is also finite. We will show that $R_{\ell,\vec{k}}(m) \leq n$.

Suppose H is a $\vec{k}$-hypergraph and $f\colon V(H) \to [\ell]$ has $\operatorname{th}(f) \geq n$. By the definition of $n = R_{\ell,(k_t)}(m')$, there exists $Q' \in \mathcal{P}_{\ell,(k_t)}$ and $W' \subseteq V(H)$ such that $\operatorname{th}(f|_{W'}) \geq m'$ and the k_t-hypergraph part of $H|_{W'}$ is Q'-uniform with respect to $f|_{W'}$. In turn, by the definition of $m' = R_{\ell,(k_1,\ldots,k_{t-1})}(m)$, there exists $Q'' \in \mathcal{P}_{\ell,(k_1,\ldots,k_{t-1})}$ and $W \subseteq W'$ such that $\operatorname{th}(f|_W) \geq m$ and the $(k_1, \ldots, k_{t-1})$-hypergraph part of $H|_W$ is Q''-uniform with respect to $f|_W$. By letting $Q \in \mathcal{P}_{\ell,\vec{k}}$ be given by

$$Q_j \stackrel{\text{def}}{=} \begin{cases} Q''_j, & \text{if } j \in [t-1]; \\ Q', & \text{if } j = t; \end{cases}$$

it follows that $H|_W$ is Q-uniform with respect to $f|_W$. ■

Before we can finally prove Theorem 5.1.8, we need one more definition.

Definition 5.4.1. If $\leq$ is a total order on a set V and $f\colon V \to [\ell]$, we let $\leq\downarrow_f \stackrel{\text{def}}{=} \leq \cap \bigcup_{i\in[\ell]} f^{-1}(i) \times f^{-1}(i)$ be the restriction of $\leq$ to the level sets of f, that is, it is the unique partial order such that $(f, \leq\downarrow_f)$ is an ℓ-split order and $\leq$ is an extension of it.

Proof of Theorem 5.1.8. Consider the set

$$K \stackrel{\text{def}}{=} \{(P, \leq) \mid P \in \mathcal{L} \wedge \leq \text{ is a total order on } [k(P)]\},$$

enumerate the elements of K as $(P_1, \leq^1), \ldots, (P_t, \leq^t)$ and define $\vec{k} = (k_1, \ldots, k_t)$ by letting $k_i \stackrel{\text{def}}{=} k(P_i)$.

Let $n \stackrel{\text{def}}{=} R_{\ell,\vec{k}}(m)$, which is finite by Theorem 5.1.5. We claim that $R_{\ell,T}(m) \leq n$. Suppose M is a model of T and $(f, \preceq) \in \mathcal{S}_{\ell,V(M)}$ is an ℓ-split order on $V(M)$ with $\operatorname{th}(f) \geq n$. Define the relation $\leq$ on $V(M)$ by

$$v \leq w \iff f(v) < f(w) \lor v \preceq w.$$

Since $(f, \preceq)$ is a split order, it follows that $\leq$ is a total order extending $\preceq$. Note that f becomes non-decreasing with respect to $\leq$ on $V(M)$ and the usual order on $[\ell]$, that is, we have

$$v \leq w \to f(v) \leq f(w) \tag{5.13}$$

for every $v, w \in V(M)$.

Define now the $\vec{k}$-hypergraph H with vertex set $V(H) \stackrel{\text{def}}{=} V(M)$ by letting the i-th edge set be

$$E_i(H) \stackrel{\text{def}}{=} \left\{ A \in \binom{V(H)}{k_i} \;\middle|\; \iota^i_A \in R_{P_i}(M) \right\},$$

where $\iota^i_A \colon [k(P_i)] \rightarrowtail V(M)$ is the unique function with $\operatorname{im}(\iota^i_A) = A$ that is increasing with respect to the order $\leq^i$ on $[k(P_i)]$ and the order $\leq$ on $V(M)$ (the latter condition is equivalent to $\leq_{\iota^i_A} = \leq^i$). For every $P \in \mathcal{L}$, let $I_P \stackrel{\text{def}}{=} \{i \in [t] \mid P_i = P\}$ and note that

$$R_P(M) = \{\alpha \in (V(M))_{k(P)} \mid i \in I_P \land \operatorname{im}(\alpha) \in E_i(H) \land \leq_\alpha = \leq^i\}. \tag{5.14}$$

By the definition of $n = R_{\ell,\vec{k}}(m)$, there exists $Q' \in \mathcal{P}_{\ell,\vec{k}}$ and a set $W \subseteq V(H)$ such that $\operatorname{th}(f|_W) \geq m$ and $H|_W$ is Q'-uniform with respect to $f|_W$. Define then the ℓ-Ramsey pattern

$Q \in \mathcal{P}_{\ell,\mathcal{L}}$ on $\mathcal{L}$ by

$$Q_P \stackrel{\text{def}}{=} \{(g, \leq^i\downarrow_g) \mid g\colon [k(P)] \to [\ell] \wedge q^g \in Q'_i \wedge i \in I_P^g\}, \tag{5.15}$$

where $q^g \in \mathcal{C}_{\ell,k(P)}$ is the weak composition given by $q_j^g \stackrel{\text{def}}{=} |g^{-1}(j)|$ and

$$I_P^g \stackrel{\text{def}}{=} \{i \in I_P \mid \forall j_1, j_2 \in [k(P)], (j_1 \leq^i j_2 \to g(j_1) \leq g(j_2))\}.$$

We claim that $M|_W$ is Q-uniform with respect to $(f|_W, \preceq_W)$. To show this, we have to show that

$$R_P(M|_W) = \{\alpha \in (W)_{k(P)} \mid (f \circ \alpha, \preceq_\alpha) \in Q_P\}.$$

Let $\alpha \in R_P(M|_W)$ and let us show that $(f \circ \alpha, \preceq_\alpha) \in Q_P$. By (5.14), there exists $i \in I_P$ such that $\operatorname{im}(\alpha) \in E_i(H)$ and $\leq_\alpha = \leq^i$. Note that if $j_1, j_2 \in [k(P)]$ are such that $j_1 \leq^i j_2$, then we must have $\alpha(j_1) \leq \alpha(j_2)$, hence (5.13) implies $f(\alpha(j_1)) \leq f(\alpha(j_2))$, so $i \in I_P^{f\circ\alpha}$. On the other hand, since $\leq$ extends $\preceq$ and $(f, \preceq)$ is a split order, it follows that $\preceq_\alpha = \leq^i\downarrow_{f\circ\alpha}$. Note also that since $H|_W$ is Q'-uniform with respect to $f|_W$ and $\operatorname{im}(\alpha) \in E_i(H)$, we must have $q^{f\circ\alpha} \in Q'_i$. Putting everything together, we have that there exists $i \in I_P^{f\circ\alpha}$ such that $q^{f\circ\alpha} \in Q'_i$ and $\preceq_\alpha = \leq^i\downarrow_{f\circ\alpha}$, so (5.15) gives $(f \circ \alpha, \preceq_\alpha) \in Q_P$.

Suppose now that $\alpha \in (W)_{k(P)}$ is such that $(f \circ \alpha, \preceq_\alpha) \in Q_P$ and let us show that $\alpha \in R_P(M|_W)$. From (5.15), we know that there exists $i \in I_P^{f\circ\alpha}$ such that $q^{f\circ\alpha} \in Q'_i$ and $\preceq_\alpha = \leq^i\downarrow_{f\circ\alpha}$. The fact that $H|_W$ is Q'-uniform with respect to $f|_W$ then implies that $\operatorname{im}(\alpha) \in E_i(H)$ and the fact that $i \in I_P^{f\circ\alpha}$ along with (5.13) implies $\leq_\alpha = \leq^i$. Putting everything together, since $I_P^{f\circ\alpha} \subseteq I_P$, we have that there exists $i \in I_P$ such that $\operatorname{im}(\alpha) \in E_i(H)$ and $\leq_\alpha = \leq^i$, so by (5.14), we get $\alpha \in R_P(M|_W)$.

Therefore $M|_W$ is Q-uniform with respect to $(f|_W, \preceq_W)$. ■

5.5 Ramsey-based formula for the abstract chromatic number

In this section we prove Theorems 5.2.2 and 5.2.3.

Proof of Theorem 5.2.2. Recall from [24, Remark 2] that we can write $I = J \circ A \circ S$, where $S\colon T_{\text{Graph}} \rightsquigarrow T_{\text{Graph}} \cup T$ is the structure-erasing interpretation, $A\colon T_{\text{Graph}} \cup T \rightsquigarrow T'$ is the axiom-adding interpretation and $J\colon T' \rightsquigarrow T$ is the isomorphism that acts identically on predicate symbols of T and acts as I on E (the inverse $J^{-1}\colon T \rightsquigarrow T'$ acts identically on the predicate symbols of T).

We start by characterizing when $\chi(I)$ is finite. Suppose first that $\mathcal{C}_{\mathcal{L}}^{E} \not\subseteq \mathcal{U}_1(\mathcal{F})$ and let us show that $\chi(I) = \infty$. Let $Q \in \mathcal{C}_{\mathcal{L}}^{E} \setminus \mathcal{U}_1(\mathcal{F})$ and for every $n \in \mathbb{N}$, let N_n be the unique structure on $\mathcal{L}$ with vertex set $[n]$ that is Q-uniform with respect to the usual order $\leq$ on $[n]$, that is, we have

$$R_P(N_n) \stackrel{\text{def}}{=} \{\alpha \in ([n])_{k(P)} \mid \leq_\alpha \in Q_P\}.$$

Our choice of Q ensures that N_n is a model of $T' = \operatorname{Forb}_{T_{\mathcal{L}}}(\mathcal{F})$. Since $Q \in \mathcal{C}_{\mathcal{L}}^{E}$, it follows that $S(A(N_n))$ is the complete graph K_n, so $I(J^{-1}(N_n)) \supseteq T_{n,\ell}$ for every $\ell \in \mathbb{N}_+$, so $\chi(I) = \infty$ by (5.2).

Suppose now that $\mathcal{C}_{\mathcal{L}}^{E} \subseteq \mathcal{U}_1(\mathcal{F})$ and let us show that $\chi(I) < \infty$ and that the minimum in (5.4) is attained (i.e., that the set in (5.4) is non-empty). For every $Q \in \mathcal{C}_{\mathcal{L}}^{E}$, let $F_Q \in \mathcal{F}$ and $\preceq^Q \in \mathcal{S}_{1,V(F_Q)}$ be such that F_Q is Q-uniform with respect to $\preceq^Q$. Let $m \stackrel{\text{def}}{=} \max\{|F_Q| \mid Q \in \mathcal{C}_{\mathcal{L}}^{E}\} \cup \{2\}$ and let $n \stackrel{\text{def}}{=} R_{1,T_{\mathcal{L}}}(m)$ (which is finite by Theorem 5.1.8).

We will show that $\chi(I) \leq n$. By (5.10) of Lemma 5.3.1, it is enough to show that every $N \in \mathcal{M}_n[T]$ satisfies $I(N) \not\supseteq T_{n,n}$. Suppose not and for a violating N let $M = J(N)$ be the associated model of T'. The definition of $n = R_{1,T_{\mathcal{L}}}(m)$ implies that there exists $W \subseteq V(M)$ such that $|W| \geq m$ and $M|_W$ is Q-uniform with respect to $\leq_W$ where $\leq$ is the usual order over $[n]$. Since $I(N) \supseteq T_{n,n} = K_n$ and $m \geq 2$, it follows that $Q \in \mathcal{C}_{\mathcal{L}}^{E}$. But this is a contradiction as $M|_W$ must then contain an induced copy of $F_Q \in \mathcal{F}$ (as $|W| \geq m \geq |F_Q|$), hence $\chi(I) < \infty$.

To show that the minimum in (5.4) is attained, it is enough to show that for $\ell \geq n$, we have $\mathcal{T}^E_{\ell,\mathcal{L}} \subseteq \mathcal{U}_\ell(\mathcal{F})$ (as this implies that the set in (5.4) is non-empty). Fix $Q \in \mathcal{T}^E_{\ell,\mathcal{L}}$ and let N_Q be the unique structure on $\mathcal{L}$ with vertex set $[\ell]$ that is Q-uniform with respect to the unique element of $\mathcal{S}_{\ell,\ell}$ of the form $(\mathrm{id}_\ell, \preceq_0)$, where $\mathrm{id}_\ell(i) \stackrel{\text{def}}{=} i$ for every $i \in [\ell]$ and $\preceq_0$ is the trivial partial order, that is, we have

$$R_P(N_Q) \stackrel{\text{def}}{=} \{\alpha \in ([\ell])_{k(P)} \mid (\alpha, \preceq_0) \in Q_P\}.$$

Since $\ell \geq n = R_{1,T_\mathcal{L}}(m)$, we know that there exists $W \subseteq [\ell]$ with $|W| = m$ and some $Q' \in \mathcal{P}_{1,\mathcal{L}}$ such that $N_Q|_W$ is Q'-uniform with respect to $\leq_W$, where $\leq$ is the usual order on $[\ell]$. Since N_Q is Q-uniform with respect to $(\mathrm{id}_\ell, \preceq_0)$, Q is an E-Turán pattern and $m \geq 2$, it follows that Q' must be E-complete. But then since $|F_{Q'}| \leq m = |W|$, there exists $U \subseteq W$ such that $N_Q|_U \cong F_{Q'}$. As $F_{Q'}$ is an induced submodel of N_Q and $Q \in \mathcal{U}_\ell(N_Q)$, we get $Q \in \mathcal{U}_\ell(F_{Q'})$, hence $\mathcal{T}^E_{\ell,\mathcal{L}} \subseteq \mathcal{U}_\ell(\mathcal{F})$, so the minimum in (5.4) is attained.

To finish the proof of (5.4), it remains to show that if $\chi(I) < \infty$ and $\ell_0 < \infty$ is the minimum in (5.4), then $\chi(I) = \ell_0$. We start by showing $\chi(I) \leq \ell_0$.

Since $\mathcal{C}^E_\mathcal{L}$ is finite and $\mathcal{C}^E_\mathcal{L} \subseteq \mathcal{U}_1(\mathcal{F})$, we know there exists a finite $\mathcal{F}' \subseteq \mathcal{F}$ such that $\mathcal{C}^E_\mathcal{L} \subseteq \mathcal{U}_1(\mathcal{F}')$. Since $\ell_0 < \infty$, we have $\mathcal{T}^E_{\ell_0,\mathcal{L}} \subseteq \mathcal{U}_{\ell_0}(\mathcal{F})$, that is, for every $Q \in \mathcal{T}^E_{\ell_0,\mathcal{L}}$, there exists $F_Q \in \mathcal{F}$ and $(f_Q, \preceq^Q) \in \mathcal{S}_{\ell_0,V(F_Q)}$ such that F_Q is Q-uniform with respect to $(f_Q, \preceq^Q)$.

Let

$$m \stackrel{\text{def}}{=} \max\{|F_Q| \mid Q \in \mathcal{T}^E_{\ell_0}\} \cup \{|F| \mid F \in \mathcal{F}'\} \cup \{2\}$$

and let $n \stackrel{\text{def}}{=} \ell_0 \cdot R_{\ell_0,T_\mathcal{L}}(m)$ (which is finite by Theorem 5.1.8). By (5.10) of Lemma 5.3.1, to show $\chi(I) \leq \ell_0$, it is enough to show that every $N \in \mathcal{M}_n[T]$ satisfies $I(N) \not\supseteq T_{n,\ell_0}$. Suppose not and for a violating $N \in \mathcal{M}_n[T]$, let $f_N\colon V(N) \to [\ell_0]$ be a function whose level sets are the parts of the natural partition of T_{n,ℓ_0} so that $\mathrm{th}(f_N) = n/\ell_0 = R_{\ell_0,T_{\mathcal{L}}}(m)$.

Let $M \stackrel{\text{def}}{=} J(N)$ be the associated model of T' and let $\preceq^N$ be any partial order such that $(f_N, \preceq^N)$ is an ℓ_0-split order. Since $\mathrm{th}(f_N) = R_{\ell_0,T_{\mathcal{L}}}(m)$, there exists an ℓ_0-Ramsey pattern $Q \in \mathcal{P}_{\ell_0,\mathcal{L}}$ and some $W \subseteq [n]$ such that $\mathrm{th}(f_N|_W) \geq m$ and $M|_W$ is Q-uniform with respect to $(f_N|_W, \preceq^N_W)$.

We claim that Q is an E-Turán pattern. Suppose not. Since the definition of f_N ensures that Q_E contains all $(g, \preceq) \in \mathcal{S}_{\ell,2}$ with g injective, for Q to not be an E-Turán pattern there must exist $i \in [\ell]$ such that Q_E contains at least one of $(g_i, \leq), (g_i, \geq) \in \mathcal{S}_{\ell,2}$, where $g_i(1) = g_i(2) = i$ and $\leq$ is the usual order on $[2]$ and $\geq$ is its reverse. From the symmetry of E and the fact that $\mathrm{th}(f_N|_W) \geq m \geq 2$, it follows that Q_E must in fact contain both $(g_i, \leq)$ and $(g_i, \geq)$. Let $Q' \in \mathcal{C}^E_{\mathcal{L}}$ be given by

$$Q'_E \stackrel{\text{def}}{=} \mathcal{S}_{1,2}; \qquad\qquad Q'_P \stackrel{\text{def}}{=} \{\preceq \mid (f, \preceq) \in Q_P \wedge \mathrm{im}(f) = \{i\}\};$$

for every $P \in \mathcal{L} \setminus \{E\}$. Let $U \stackrel{\text{def}}{=} f_N^{-1}(i) \cap W$ and note that $M|_U$ is Q'-uniform with respect to $\preceq^N_U$. Since $|U| \geq \mathrm{th}(f_N|_W) \geq m \geq \max\{|F| \mid F \in \mathcal{F}'\}$ and since $Q' \in \mathcal{C}^E_{\mathcal{L}}$, there exists $F \in \mathcal{F}'$ such that $M|_U$ contains a copy of F, so M is not a model of $T' = \mathrm{Forb}_{T_{\mathcal{L}}}(\mathcal{F})$, a contradiction. Thus Q must be an E-Turán pattern.

Since $Q \in \mathcal{T}^E_{\ell,\mathcal{L}}$, it follows that $M|_W$ must contain an induced copy of $F_Q \in \mathcal{F}$, namely, such copy can be produced by taking exactly $|f_Q^{-1}(i)|$ vertices in $f_N^{-1}(i) \cap W$ for each $i \in [\ell_0]$ (this is possible since $|f_Q^{-1}(i)| \leq |F_Q| \leq m \leq \mathrm{th}(f_N|_W)$). This contradicts the fact that M is a model of $T' = \mathrm{Forb}_{T_{\mathcal{L}}}(\mathcal{F})$, hence $\chi(I) \leq \ell_0$.

Let us now show that $\chi(I) \geq \ell_0$. If $\ell_0 = 1$, then the inequality trivially holds, so suppose $\ell_0 \geq 2$. From the definition of ℓ_0, there exists $Q \in \mathcal{T}^E_{\ell_0-1,\mathcal{L}} \setminus \mathcal{U}_{\ell_0-1}(\mathcal{F})$. For every $n \in \mathbb{N}$, let $f_n\colon [n] \to [\ell_0 - 1]$ be any function with $\mathrm{th}(f_n) = \lfloor n/(\ell_0 - 1)\rfloor$ and let N_n be the unique structure on $\mathcal{L}$ with vertex set $[n]$ that is Q-uniform with respect to $(f_n, \leq\downarrow_{f_n})$, where $\leq$ is the usual order on $[n]$, that is, we have

$$R_P(N_n) \stackrel{\text{def}}{=} \{\alpha \in ([n])_{k(P)} \mid (f_n \circ \alpha, (\leq\downarrow_{f_n})_\alpha) \in Q_P\}.$$

Our choice of Q ensures that N_n is a model of $T' = \mathrm{Forb}_{T_{\mathcal{L}}}(\mathcal{F})$.

Since $\mathrm{th}(f_n) = \lfloor n/(\ell_0 - 1) \rfloor$ and $Q \in \mathcal{T}^E_{\ell_0-1,\mathcal{L}}$, it follows that $S(A(N_n))$ is isomorphic to the Turán graph T_{n,ℓ_0-1}, which implies that $I(J^{-1}(N_n)) \cong T_{n,\ell_0-1}$, so by (5.2), we have $\chi(I) \geq \ell_0$.

This concludes the proof of (5.4).

Finally, let us consider the case when T is itself obtained from $T_{\mathrm{Graph}} \cup T''$ by adding axioms and I acts identically on the predicate symbol E of T_{Graph}. To apply the previous case of the theorem, note that to form $T_{\mathrm{Graph}} \cup T$, we add a new predicate symbol E' corresponding to the new copy of T_{Graph} and the theory T' is defined from $T_{\mathrm{Graph}} \cup T$ by adding the axiom

$$\forall x \forall y, E'(x, y) \leftrightarrow E(x, y).$$

But then the isomorphism $J\colon T' \rightsquigarrow T$ simply copies E to E', which means that we can replace T' with T and use E from T in place of the newly added E' from T'. ■

Remark 13. One of the consequences of Theorem 5.2.2 is that to compute $\chi(I)$, models $F \in \mathcal{F}$ such that the graph part $I(J^{-1}(F))$ contains an induced copy of $\overline{P}_3$ (the graph on 3 vertices with exactly 1 edge) are completely irrelevant as such models are never uniform for complete patterns nor for Turán patterns.

Proof of Remark 11. We want to show that the set

$$X \stackrel{\text{def}}{=} \{\ell \in \mathbb{N}_+ \mid \mathcal{T}^E_{\ell,\mathcal{L}} \subseteq \mathcal{U}_\ell(\mathcal{F})\}$$

in (5.4) is either empty or an infinite interval of $\mathbb{N}_+$. To show this, it is enough to show that if $\ell \in \mathbb{N}_+ \setminus X$ and $\ell' \in [\ell]$, then $\ell' \notin X$. But if $\ell \in \mathbb{N}_+ \setminus X$ then there exists $Q \in \mathcal{T}^E_{\ell,\mathcal{L}} \setminus \mathcal{U}_\ell(\mathcal{F})$. Let then $Q' \in \mathcal{T}^E_{\ell',\mathcal{L}}$ be given by

$$Q'_P \stackrel{\text{def}}{=} \{(f, \preceq) \in Q_P \mid \text{im}(f) \subseteq [\ell']\} \qquad (P \in \mathcal{L}),$$

where we reinterpret functions $f\colon [k(P)] \to [\ell]$ with $\text{im}(f) \subseteq [\ell']$ as $f\colon [k(P)] \to [\ell']$. We claim that $Q' \notin \mathcal{U}_{\ell'}(\mathcal{F})$. Indeed, if $F \in \mathcal{F}$ was Q'-uniform with respect to some $(f, \preceq) \in \mathcal{S}_{\ell',V(F)}$, then it would also be Q-uniform with respect to $(\widehat{f}, \preceq)$, where $\widehat{f}$ is obtained from f by simply extending the codomain to $[\ell]$. Hence $\ell' \notin X$.

Since X is either empty or an infinite interval of $\mathbb{N}_+$, it follows that $\inf X = \sup \mathbb{N} \setminus X + 1$. If we further assume that $\chi(I) < \infty$, then X is non-empty so $\min X = \max \mathbb{N} \setminus X + 1$, hence (5.4) and (5.5) are equal. ∎

Before showing Theorem 5.2.3, let us first address a small technicality on axiomatization of universal theories.

Lemma 5.5.1. *If T be a universal theory that is finitely axiomatizable, then it has a finite axiomatization in which all of its axioms are universal. Furthermore, such finite axiomatization with universal axioms can be algorithmically computed from any finite axiomatization of T.*

Proof. Let A be a finite list of axioms of T. Since T is universal, the set S of all universal formulas that are theorems of T is an axiomatization of T, hence $S \vdash \bigwedge_{\phi \in A} \phi$, which implies that there must exist a finite set $S' \subseteq S$ such that $S' \vdash \bigwedge_{\phi \in A} \phi$, so S' is a finite axiomatization of T by universal formulas.

To algorithmically compute S' as above, we can enumerate all universal formulas ϕ that are theorems of T in parallel (by also enumerating possible proofs of ϕ from A in parallel) and also check in parallel whether finite subsets S' of the S enumerated so far satisfy $S' \vdash \bigwedge_{\phi \in A} \phi$ (by also enumerating possible proofs in parallel). The reasoning above shows that such algorithm must eventually find a satisfying S'. ∎

Proof of Theorem 5.2.3. Using the notation of Theorem 5.2.2, note that the fact that T is finitely axiomatizable implies that T' is also finitely axiomatizable and the list of axioms of T' can trivially be computed from the list of axioms of T and a description of I. By Lemma 5.5.1, we may compute an axiomatization A of T' in which every axiom is a universal formula.

Let k be the maximum number of variables appearing in an axiom in A and let $\mathcal{F}$ be the (finite) set of all canonical structures M on $\mathcal{L}$ with vertex set $[t]$ for some $t \leq k$ that are not models of T'. Our choice of k ensures that $T' = \mathrm{Forb}_{T_{\mathcal{L}}}(\mathcal{F})$. We then check if $\mathcal{C}_{\mathcal{L}}^{E} \subseteq \mathcal{U}_1(\mathcal{F})$. If this is false, then Theorem 5.2.2 guarantees that $\chi(I) = \infty$. Otherwise, we know that $\chi(I) < \infty$ and is given by (5.4), which means that we can compute it by finding the smallest $\ell \in \mathbb{N}_+$ such that $\mathcal{T}_{\ell,\mathcal{L}}^{E} \subseteq \mathcal{U}_\ell(\mathcal{F})$; Theorem 5.2.2 ensures that such ℓ exists and is precisely $\chi(I)$.

Finally, we can compute π_I^t from $\chi(I)$ and t using formula (5.3) in Theorem 5.2.1. Note that this is a valid algorithm as all sets and searches above are finite. ■

5.6 The non-induced case

In this section, we prove Theorems 5.2.4 and 5.2.5, which provide simpler formulas for the abstract chromatic number in the setting of graphs with extra structure with some forbidden submodels that are non-induced in the graph part.

For this section, let us fix a language $\mathcal{L}$, let E be the predicate symbol of T_{Graph} in the language $\mathcal{L} \cup \{E\}$ of $T_{\mathrm{Graph}} \cup T_{\mathcal{L}}$, let $J\colon T_{\mathcal{L}} \rightsquigarrow T_{\mathrm{Graph}} \cup T_{\mathcal{L}}$ be the structure-erasing interpretation and let $\mathcal{F}$ be a family of models of $T_{\mathrm{Graph}} \cup T_{\mathcal{L}}$.

Proof of Theorem 5.2.4. Let ℓ_0 be the right-hand side of (5.6).

Suppose $G \in \mathcal{M}[T_{\mathrm{Graph}}]$ is such that for every $M \in \mathcal{M}[\mathrm{Forb}_{T_{\mathrm{Graph}} \cup T_{\mathcal{L}}}(\mathcal{F}{\uparrow}^{E})]$, we have $I(M) \not\cong G$. Since for $n \stackrel{\mathrm{def}}{=} |G|\chi(G)$, we have $T_{n,\chi(G)} \supseteq G$, from the definition of $\mathcal{F}{\uparrow}^{E}$, it follows that for every $M \in \mathcal{M}[\mathrm{Forb}_{T_{\mathrm{Graph}} \cup T_{\mathcal{L}}}]$, we have $I(M) \not\supseteq T_{n,\chi(G)}$, so by (5.10) of

Lemma 5.3.1, we have $\chi(I) \leq \ell_0$.

On the other hand, if $\ell \in \mathbb{N}_+$ is such that there exists $n \in \mathbb{N}_+$ such that for all $N \in \mathcal{M}_n[\mathrm{Forb}_{T_{\mathrm{Graph}} \cup T_{\mathcal{L}}}(\mathcal{F}\!\uparrow^E)]$, we have $T_{n,\ell} \not\subseteq I(N)$, then we must also have that $I(N) \not\cong T_{n,\ell}$ for every $N \in \mathcal{M}[\mathrm{Forb}_{T_{\mathrm{Graph}} \cup T_{\mathcal{L}}}(\mathcal{F}\!\uparrow^E)]$, hence from (5.10) of Lemma 5.3.1 we also get $\ell_0 \leq \chi(I)$. ■

To prove Theorem 5.2.5, we first need to relate uniformity of over $\mathcal{L} \cup \{E\}$ with uniformity and E-proper split orders over $\mathcal{L}$.

Claim 5.6.1. *For $Q \in \mathcal{C}^E_{\mathcal{L} \cup \{E\}}$, we have $Q \in \mathcal{U}_1(\mathcal{F}\!\uparrow^E)$ if and only if $Q|_{\mathcal{L}} \in \mathcal{U}_1(J(\mathcal{F}))$, where $Q|_{\mathcal{L}} \in \mathcal{P}_{1,\mathcal{L}}$ is the restriction of Q to $\mathcal{L}$ and $J(\mathcal{F}) \stackrel{\text{def}}{=} \{J(F) \mid F \in \mathcal{F}\}$.*

Proof. Suppose $Q \in \mathcal{U}_1(\mathcal{F}\!\uparrow^E)$, that is, there exists some $F \in \mathcal{F}\!\uparrow^E$ and some $\preceq \in \mathcal{S}_{1,V(F)}$ such that F is Q-uniform with respect to $\preceq$. From the definition of $\mathcal{F}\!\uparrow^E$, there exists $F' \in \mathcal{F}$ such that $V(F') = V(F)$, $R_E(F') \subseteq R_E(F)$ and $R_P(F') = R_P(F)$ for every $P \in \mathcal{L}$. Since F is Q-uniform with respect to $\preceq$, it follows that $J(F) = J(F')$ is $Q|_{\mathcal{L}}$-uniform with respect to $\preceq$, so $Q|_{\mathcal{L}} \in \mathcal{U}_1(J(F'))$.

Suppose now that $Q|_{\mathcal{L}} \in \mathcal{U}_1(J(\mathcal{F}))$, that is, there exists some $F \in \mathcal{F}$ and some $\preceq \in \mathcal{S}_{1,V(F)}$ such that $J(F)$ is $Q|_{\mathcal{L}}$-uniform with respect to $\preceq$. Let F' be defined by $V(F') \stackrel{\text{def}}{=} V(F)$, $R_P(F') \stackrel{\text{def}}{=} R_P(F)$ for every $P \in \mathcal{L}$ and $R_E(F') \stackrel{\text{def}}{=} (V(F'))_2$. Note that $F' \in \mathcal{F}\!\uparrow^E$ and F' is Q-uniform with respect to $\preceq$, so $Q \in \mathcal{U}_1(F)$. ■

Claim 5.6.2. *For $Q \in \mathcal{T}^E_{\ell,\mathcal{L} \cup \{E\}}$, we have $Q \in \mathcal{U}_\ell(\mathcal{F}\!\uparrow^E)$ if and only if $Q|_{\mathcal{L}} \in \chi^E_\ell(\mathcal{F})$, where $Q|_{\mathcal{L}} \in \mathcal{P}_{\ell,\mathcal{L}}$ is the restriction of Q to $\mathcal{L}$.*

Proof. Let $I\colon T_{\mathrm{Graph}} \rightsquigarrow T_{\mathrm{Graph}} \cup T_{\mathcal{L}}$ be the structure-erasing interpretation.

Suppose $Q \in \mathcal{U}_\ell(\mathcal{F}\!\uparrow^E)$, that is, there exists some $F \in \mathcal{F}\!\uparrow^E$ and some $(f, \preceq) \in \mathcal{S}_{\ell,V(F)}$ such that F is Q-uniform with respect to $(f, \preceq)$. From the definition of $\mathcal{F}\!\uparrow^E$, there exists $F' \in \mathcal{F}$ such that $V(F') = V(F)$, $R_E(F') \subseteq R_E(F)$ and $R_P(F') = R_P(F)$ for every $P \in \mathcal{L}$. Since F is Q-uniform with respect to $(f, \preceq)$, it follows that $J(F) = J(F')$ is $Q|_{\mathcal{L}}$-uniform with respect to $(f, \preceq)$. Since Q is an E-Turán pattern, we also get that f is a proper coloring of $I(F)$, hence also of $I(F')$, so $Q|_{\mathcal{L}} \in \chi_\ell^E(F')$.

Suppose now that $Q|_{\mathcal{L}} \in \chi_\ell^E(\mathcal{F})$, that is, there exists some $F \in \mathcal{F}$ and some E-proper $Q|_{\mathcal{L}}$-split ordering $(f, \preceq) \in \mathcal{S}_{\ell,V(F)}$ of F. Define F' by letting $V(F') \stackrel{\text{def}}{=} V(F)$, $R_P(F') \stackrel{\text{def}}{=} R_P(F)$ for every $P \in \mathcal{L}$ and

$$R_E(F') \stackrel{\text{def}}{=} \{\alpha \in (V(F'))_2 \mid f(\alpha(1)) \neq f(\alpha(2))\}.$$

Note that since f is a proper coloring of $I(F)$, it follows that $R_E(F') \supseteq R_E(F)$, so $F' \in \mathcal{F}{\uparrow}^E$. Note also that F' is Q-uniform with respect to $(f, \preceq)$ as $J(F)$ is $Q|_{\mathcal{L}}$-uniform with respect to $(f, \preceq)$, so $Q \in \mathcal{U}_\ell(F')$. ■

Proof of Theorem 5.2.5. Note that first that the restriction function $\mathcal{C}^E_{\mathcal{L}\cup\{E\}} \to \mathcal{P}_{1,\mathcal{L}}$ given by $Q \mapsto Q|_{\mathcal{L}}$ is bijective, so Claim 5.6.1 implies that $\mathcal{C}^E_{\mathcal{L}\cup\{E\}} \subseteq \mathcal{U}_1(\mathcal{F})$ is equivalent to $\mathcal{P}_{1,\mathcal{L}} \subseteq \mathcal{U}_1(J(\mathcal{F}))$, so the characterization of $\chi(I) < \infty$ of Theorem 5.2.5 follows from the characterization of $\chi(I) < \infty$ of Theorem 5.2.2.

On the other hand, the restriction function $\mathcal{T}^E_{\ell,\mathcal{L}\cup\{E\}} \to \mathcal{P}_{\ell,\mathcal{L}}$ given by $Q \mapsto Q|_{\mathcal{L}}$ is also a bijection. This along with Claim 5.6.2 implies that $\mathcal{T}^E_{\ell,\mathcal{L}\cup\{E\}} \not\subseteq \mathcal{U}_\ell(\mathcal{F}{\uparrow}^E)$ is equivalent to $\mathcal{P}^E_{\ell,\mathcal{L}} \not\subseteq \chi_\ell^E(\mathcal{F})$, so from (5.4) of Theorem 5.2.2, we get that if $\chi(I) < \infty$, then (5.7) holds.

It remains to prove that (5.7) also holds when $\chi(I) = \infty$, that is, we need to show that if $\mathcal{P}_{1,\mathcal{L}} \not\subseteq \mathcal{U}_1(J(\mathcal{F}))$, then $\mathcal{P}_{\ell,\mathcal{L}} \not\subseteq \chi_\ell^E(\mathcal{F})$ for every $\ell \in \mathbb{N}_+$.

Let $Q \in \mathcal{P}_{1,\mathcal{L}} \setminus \mathcal{U}_1(J(\mathcal{F}))$ and fix $\ell \in \mathbb{N}_+$. Given $(f, \preceq) \in \mathcal{S}_{\ell,V}$, let $\preceq^f \in \mathcal{S}_{1,V}$ be the total order on V given by

$$v \preceq^f w \iff f(v) < f(w) \lor v \preceq w.$$

Clearly $\preceq^f{\downarrow}_f = \preceq$.

Let $Q' \in \mathcal{P}_{\ell,\mathcal{L}}$ be given by

$$Q'_P \stackrel{\text{def}}{=} \{(g, \preceq) \in \mathcal{S}_{\ell,k} \mid \preceq^g \in Q_P\}.$$

We claim that $Q' \notin \chi_\ell^E(\mathcal{F})$. Suppose not, that is, suppose there exists $F \in \mathcal{F}$ and an E-proper Q'-split ordering $(f, \preceq) \in \mathcal{S}_{\ell,V(F)}$ of F. Note that for every $P \in \mathcal{L}$, we have

$$\begin{aligned} R_P(F) &= \{\alpha \in (V(F))_{k(P)} \mid (f \circ \alpha, \preceq_\alpha) \in Q'_P\} \\ &= \{\alpha \in (V(F))_{k(P)} \mid (\preceq_\alpha)^{f\circ\alpha} \in Q_P\} \\ &= \{\alpha \in (V(F))_{k(P)} \mid (\preceq^f)_\alpha \in Q_P\}, \end{aligned}$$

hence $J(F)$ is Q-uniform with respect to $\preceq^f$, contradicting the fact that $Q \notin \mathcal{U}_1(J(F))$. Hence $Q' \notin \chi_\ell^E(\mathcal{F})$ as desired. ■

Proof of Remark 12. In the proof above, we determined that $\mathcal{T}^E_{\ell,\mathcal{L}\cup\{E\}} \nsubseteq \mathcal{U}_\ell(\mathcal{F}{\uparrow}^E)$ is equivalent to $\mathcal{P}^E_{\ell,\mathcal{L}} \nsubseteq \chi_\ell^E(\mathcal{F})$, so from Remark 11 it follows that the set X in (5.7) is either empty or an infinite interval of $\mathbb{N}_+$ and thus $\inf X = \sup \mathbb{N}_+ \setminus X + 1$. ■

5.7 Applications to concrete theories

In this section we illustrate how to use the general theory to obtain easier formulas for the abstract chromatic number for some specific theories. We start with the easy example of recovering the original setting of Theorem 5.0.1: graphs with forbidden non-induced subgraphs.

Proposition 5.7.1. *Let $\mathcal{F}$ be a family of graphs and $\mathrm{Forb}^+_{T_{\mathrm{Graph}}}(\mathcal{F})$ be the theory of all graphs that do not have any non-induced copy of graphs in $\mathcal{F}$. Then for the axiom-adding* *interpretation* $I^+_{\mathcal{F}}\colon T_{\text{Graph}} \rightsquigarrow \text{Forb}^+_{T_{\text{Graph}}}(\mathcal{F})$*, we have*

$$\chi(I^+_{\mathcal{F}}) = \max\{\chi(\mathcal{F}), 1\},$$

where $\chi(\mathcal{F}) \stackrel{\text{def}}{=} \inf\{\chi(F) \mid F \in \mathcal{F}\}$ *is the infimum of the chromatic numbers of elements of* $\mathcal{F}$*.*

Proof. Let $\mathcal{L} \stackrel{\text{def}}{=} \varnothing$ be the empty language and note that in the notation of Theorem 5.2.5 we have $\text{Forb}^+_{T_{\text{Graph}}}(\mathcal{F}) = \text{Forb}_{T_{\text{Graph}} \cup T_{\mathcal{L}}}(\mathcal{F}{\uparrow}^E)$, so we get

$$\chi(I^+_{\mathcal{F}}) = \sup\{\ell \in \mathbb{N}_+ \mid \mathcal{P}_{\ell,\mathcal{L}} \not\subseteq \chi^E_\ell(\mathcal{F})\} \cup \{0\} + 1.$$

But since $\mathcal{L}$ is empty, each $\mathcal{P}_{\ell,\mathcal{L}}$ has a unique element (namely, the empty pattern) and this unique element is in $\chi^E_\ell(F)$ if and only if there exists a proper coloring of F with ℓ colors, hence

$$\chi(I^+_{\mathcal{F}}) = \sup\{\ell \in \mathbb{N}_+ \mid \forall F \in \mathcal{F}, \ell < \chi(F)\} \cup \{0\} + 1 = \max\{\chi(\mathcal{F}), 1\},$$

as desired. ■

We now show how the picture changes when the forbidden subgraphs are induced instead of non-induced.

Proposition 5.7.2. *Let* $\mathcal{F}$ *be a family of graphs and let* $I_{\mathcal{F}}\colon T_{\text{Graph}} \rightsquigarrow \text{Forb}_{T_{\text{Graph}}}(\mathcal{F})$ *be the axiom-adding interpretation. If* $\mathcal{F}$ *contains a complete graph, then*

$$\begin{aligned}\chi(I_{\mathcal{F}}) &= \max\{\ell \in \mathbb{N}_+ \mid \mathcal{F} \textit{ does not contain a complete } \ell\textit{-partite graph}\} \cup \{0\} + 1\\ &= \min\{\ell \in \mathbb{N}_+ \mid \mathcal{F} \textit{ contains a complete } \ell\textit{-partite graph}\};\end{aligned}$$

otherwise, we have $\chi(I_{\mathcal{F}}) = \infty$*.*

Proof. In the notation of Theorem 5.2.2, we can view $\text{Forb}_{T_{\text{Graph}}}(\mathcal{F})$ as obtained from the theory $T_{\text{Graph}} \cup T_0$ by adding axioms, where T_0 is the trivial theory over the empty language. Then taking $T' = T$ (so $\mathcal{L} = \{E\}$), note that $\mathcal{C}^E_{\mathcal{L}}$ contains a single element Q_0 and we have $Q_0 \in \mathcal{U}_1(F)$ if and only if F is complete, so Theorem 5.2.2 gives $\chi(I_{\mathcal{F}}) < \infty$ if and only if $\mathcal{F}$ has a complete graph.

Suppose then that $\mathcal{F}$ contains a complete graph (so $\chi(I_{\mathcal{F}}) < \infty$) and note that for every $\ell \in \mathbb{N}_+$, $\mathcal{T}^E_{\ell,\mathcal{L}}$ also contains a single element Q_ℓ and we have $Q_\ell \in \mathcal{U}_\ell(F)$ if and only if F is a complete ℓ-partite graph, hence from (5.5) and (5.4), we get

$$\begin{aligned}\chi(I_{\mathcal{F}}) &= \max\{\ell \in \mathbb{N}_+ \mid \mathcal{F} \text{ does not contain a complete } \ell\text{-partite graph}\} \cup \{0\} + 1\\ &= \min\{\ell \in \mathbb{N}_+ \mid \mathcal{F} \text{ contains a complete } \ell\text{-partite graph}\},\end{aligned}$$

as desired. ■

For our next example, we will recover the interval chromatic number used for ordered graphs in [57] from our result.

Definition 5.7.3 (Interval chromatic number [57])**.** An *ordered graph* is a model of the theory $T_{\text{Graph}} \cup T_{\text{LinOrder}}$. A *proper interval coloring* of an ordered graph G is a proper coloring of the graph part of G such that each color class is an interval of the order part of G. Formally, a proper interval coloring of G is a function $f\colon V(G) \to [\ell]$ such that

$$\begin{gathered}\forall v, w \in V(G), (v, w) \in R_E(G) \implies f(v) \neq f(w);\\ \forall u, v, w \in V(G), (u, v) \in R_<(G) \land (v, w) \in R_<(G) \land f(u) = f(w) \implies f(u) = f(v).\end{gathered}$$

The *interval chromatic number* $\chi_<(G)$ of an ordered graph G is the minimum ℓ such that there exists a proper interval coloring of G of the form $f\colon V(G) \to [\ell]$.

Proposition 5.7.4. *Let $\mathcal{F}$ be a family of ordered graphs and $\operatorname{Forb}^+_{T_{\text{Graph}} \cup T_{\text{LinOrder}}}(\mathcal{F})$ be the theory of all ordered graphs that do not have any non-induced copy of ordered graphs in $\mathcal{F}$. Then for the interpretation $I^{\leq}_{\mathcal{F}}\colon T_{\text{Graph}} \rightsquigarrow \operatorname{Forb}^+_{T_{\text{Graph}} \cup T_{\text{LinOrder}}}(\mathcal{F})$ that acts identically on*

E, we have

$$\chi(I_{\mathcal{F}}^{\leq}) = \max\{\chi_<(\mathcal{F}), 1\},$$

where $\chi_<(\mathcal{F}) \stackrel{\text{def}}{=} \inf\{\chi_<(F) \mid F \in \mathcal{F}\}$ is the infimum of the interval chromatic numbers of elements of $\mathcal{F}$.

Proof. Let $\mathcal{L} \stackrel{\text{def}}{=} \{<\}$ with $k(<) \stackrel{\text{def}}{=} 2$ and let further F_1, F_2, F_3 be the structures on $\{E\} \cup \mathcal{L}$ defined by

$$\begin{aligned} V(F_1) &\stackrel{\text{def}}{=} [2]; & V(F_2) &\stackrel{\text{def}}{=} [2]; & V(F_3) &\stackrel{\text{def}}{=} [3]; \\ R_E(F_1) &\stackrel{\text{def}}{=} \varnothing; & R_E(F_2) &\stackrel{\text{def}}{=} \varnothing; & R_E(F_3) &\stackrel{\text{def}}{=} \varnothing; \\ R_<(F_1) &\stackrel{\text{def}}{=} \varnothing; & R_<(F_2) &\stackrel{\text{def}}{=} ([2])_2; & R_<(F_3) &\stackrel{\text{def}}{=} \{(1,2),(2,3),(3,1)\}; \end{aligned}$$

Define also $\widehat{\mathcal{F}} \stackrel{\text{def}}{=} \mathcal{F} \cup \{F_1, F_2, F_3\}$ and note that in the notation of Theorem 5.2.5, we have $\mathrm{Forb}^+_{T_{\mathrm{Graph}} \cup T_{\mathrm{LinOrder}}}(\mathcal{F}) = \mathrm{Forb}_{T_{\mathrm{Graph}} \cup T_{\mathcal{L}}}(\widehat{\mathcal{F}}\uparrow^E)$, so we get

$$\chi(I_{\mathcal{F}}^{\leq}) = \sup\{\ell \in \mathbb{N}_+ \mid \mathcal{P}_{\ell,\mathcal{L}} \not\subseteq \chi_\ell^E(\widehat{\mathcal{F}})\} \cup \{0\} + 1.$$

For $i, j \in [\ell]$, let $\mathcal{S}_{\ell,i,j} \stackrel{\text{def}}{=} \{(f, \preceq) \in \mathcal{S}_{\ell,2} \mid \mathrm{im}(f) = \{i, j\}\}$. Note that $\mathcal{S}_{\ell,i,j} = \mathcal{S}_{\ell,j,i}$ and, regardless of whether $i \neq j$, we have $|\mathcal{S}_{\ell,i,j}| = 2$.

Fix an ℓ-Ramsey pattern $Q \in \mathcal{P}_{\ell,\mathcal{L}}$ on $\mathcal{L}$. Let us call a pair $(i, j) \in [\ell]^2$ *empty in Q* if $\mathcal{S}_{\ell,i,j} \cap Q_< = \varnothing$ and let us call (i, j) *full in Q* if $\mathcal{S}_{\ell,i,j} \subseteq Q_<$.

Note that if (i, j) is empty in $Q \in \mathcal{P}_{\ell,\mathcal{L}}$, then any $(f, \preceq) \in \mathcal{S}_{\ell,2}$ with $\mathrm{im}(f) = \{i, j\}$ is an E-proper Q-split ordering of F_1. Conversely, note that if $(f, \preceq)$ is an E-proper Q-split ordering of F_1 then $(f(1), f(2))$ is empty in Q. Hence $Q \in \mathcal{P}_{\ell,\mathcal{L}}$ has an empty pair if and only if $Q \in \chi_\ell^E(F_1)$. With an analogous argument, we can show that $Q \in \mathcal{P}_{\ell,\mathcal{L}}$ has a full pair if and only if $Q \in \chi_\ell^E(F_2)$.

Let then $\mathcal{P}'_{\ell,\mathcal{L}}$ be the set of all ℓ-Ramsey patterns that do not have any empty pairs nor any full pairs. To each $Q \in \mathcal{P}'_{\ell,\mathcal{L}}$, let us associate a tournament T_Q given by $V(T_Q) \stackrel{\text{def}}{=} [\ell]$ and

$$E(T_Q) \stackrel{\text{def}}{=} \{(v, w) \in [\ell]^2 \mid v \neq w \land \exists (f, \preceq) \in Q_<, f(1) = v \land f(2) = w\}.$$

Note that the fact that Q does not have any empty or full pairs ensures that T_Q is indeed a tournament.

We claim that for $Q \in \mathcal{P}_{\ell,\mathcal{L}}$, the tournament T_Q has a cycle if and only if $Q \in \chi_\ell^E(F_3)$.

For the forward direction, since T_Q has a cycle, it must have a 3-cycle, say $(u, v, w) \in [\ell]^3$ with $(u, v), (v, w), (w, u) \in E(T_Q)$. Then any $(f, \preceq) \in \mathcal{S}_{\ell,3}$ with $f(1) = u$, $f(2) = v$ and $f(3) = w$ is an E-proper Q-split ordering of F_3. For the backward direction, if $(f, \preceq) \in \mathcal{S}_{\ell,3}$ is an E-proper Q-split ordering of F_3, then $(f(1), f(2), f(3))$ is a 3-cycle in T_Q.

Let then $\mathcal{P}''_{\ell,\mathcal{L}} \stackrel{\text{def}}{=} \{Q \in \mathcal{P}'_{\ell,\mathcal{L}} \mid T_Q \text{ is transitive}\}$ and note that our claims above show that

$$\chi(I_{\mathcal{F}}^{\leq}) = \sup\{\ell \in \mathbb{N}_+ \mid \mathcal{P}''_{\ell,\mathcal{L}} \not\subseteq \chi_\ell^E(\mathcal{F})\} \cup \{0\} + 1.$$

We now claim that for $Q \in \mathcal{P}''_{\ell,\mathcal{L}}$ and $F \in \mathcal{F}$, we have $Q \in \chi_\ell^E(F)$ if and only if $\ell \geq \chi_<(F)$.

For the forward direction, we claim that if $(f, \preceq) \in \mathcal{S}_{\ell,V(F)}$ is an E-proper Q-split ordering of F, then $f \colon V(F) \to [\ell]$ is a proper interval coloring of F. Since f is a proper coloring of the graph part of F, we need to show that its color classes are intervals of the order part of F. Suppose not, that is, suppose there exist $u, v, w \in V(F)$ such that $(u, v), (v, w) \in R_<(F)$ and $f(u) = f(w) \neq f(v)$. But then $(u, v) \in R_<(F)$ implies $(f(u), f(v)) \in E(T_Q)$ and $(v, w) \in R_<(F)$ implies $(f(v), f(w)) \in E(T_Q)$, contradicting the fact that T_Q does not have anti-parallel edges.

For the backward direction, suppose $f \colon V(F) \to [\ell]$ is a proper interval coloring of F. Since T_Q is transitive, by possibly permuting the colors of f, we may suppose that the color classes of f are in the same order in F as the colors are in T_Q, that is, we may suppose that

$$\forall v, w \in V(F), (f(v) \neq f(w) \to ((v, w) \in R_<(F) \leftrightarrow (f(v), f(w)) \in E(T_Q))). \tag{5.16}$$

For $i \in [\ell]$, let $(g_i, \leq) \in \mathcal{S}_{\ell,2}$ be the ℓ-split order over $[2]$ given by $g_i(1) = g_i(2) = i$ and $1 \leq 2$. Define the partial order $\preceq$ over $V(F)$ as

$$v \preceq w \iff f(v) = f(w) \wedge ((v, w) \in R_<(F) \leftrightarrow (g_{f(v)}, \leq) \in Q_<).$$

It is clear that $(f, \preceq)$ is an ℓ-split order over $V(F)$.

We claim that $(f, \preceq)$ is an E-proper Q-split order of F. We know that f is a proper coloring of the graph part of F, so we need to show that the order part of F is Q-uniform with respect to $(f, \preceq)$. But this follows from the definition of $\preceq$ and (5.16); this concludes the proof of our claim.

From our claim, it follows that

$$\begin{aligned} \chi(I_{\mathcal{F}}^{\leq}) &= \sup\{\ell \in \mathbb{N}_+ \mid \forall F \in \mathcal{F}, \ell < \chi_<(F)\} \cup \{0\} + 1 \\ &= \max\{\chi_<(\mathcal{F}), 1\}, \end{aligned}$$

as desired. ■

Let us note that the result of [9] that proves an analogue of Theorem 5.0.1 in terms of the cyclic interval chromatic number (which has the same definition as the interval chromatic number, but intervals are considered in the cyclic order) can also be retrieved from Theorem 5.2.5 with a similar proof to that of Proposition 5.7.4.

Finally, the result of [41] that proves an analogue of Theorem 5.0.1 in terms of the edge-order chromatic number follows trivially from Theorem 5.2.4.

5.8 Concluding remarks and open problems

In this chapter, we have shown how the Erdős–Stone–Simonovits and Alon–Shikhelman Theorems (Theorem 5.0.1) generalize to the setting of open interpretations $I\colon T_{\text{Graph}} \rightsquigarrow T$ via the abstract chromatic number $\chi(I)$ and we have shown how an alternative formula for $\chi(I)$ based on Ramsey Theory can be used to algorithmically compute $\chi(I)$ when T is finitely axiomatizable. We have also shown how to retrieve the particular chromatic numbers of [9, 57, 41] from the abstract chromatic number.

One property that the usual chromatic number satisfies is principality: in the setting of Proposition 5.7.1 (i.e., the setting of the original Theorem 5.0.1), we have $\chi(I^+_{\mathcal{F}}) = \max\{\chi(\mathcal{F}), 1\} = \min\{\chi(I^+_{\{F\}}) \mid F \in \mathcal{F}\} \cup \{1\}$, that is, the chromatic number corresponding to a non-empty family of graphs is simply the minimum of the chromatic numbers corresponding to its elements.

In the more general setting of Theorem 5.2.5, let $\mathcal{L}$ be a language and let $\mathcal{F}_0$ be a family of structures on $\mathcal{L} \cup \{E\}$. For another family $\mathcal{F}$ of structures on $\mathcal{L} \cup \{E\}$, we let $I_{\mathcal{F}}\colon T_{\text{Graph}} \rightsquigarrow \operatorname{Forb}_{T_{\text{Graph}} \cup T_{\mathcal{L}}}((\mathcal{F}_0 \cup \mathcal{F})\!\uparrow^E)$ act identically on E. We say that $T \stackrel{\text{def}}{=} \operatorname{Forb}_{T_{\text{Graph}} \cup T_{\mathcal{L}}}(\mathcal{F}_0\!\uparrow^E)$ satisfies the *principality property* if

$$\chi(I_{\mathcal{F}}) = \min\{\chi(I_{\{F\}}) \mid F \in \mathcal{F}\}$$

for every non-empty $\mathcal{F}$.

The setting of Proposition 5.7.1 shows that T_{Graph} satisfies the principality property. Proposition 5.7.4 shows that $T_{\text{Graph}} \cup T_{\text{LinOrder}}$ satisfies the principality property as well. Since an analogous result to Proposition 5.7.4 holds for cyclically ordered graphs (see [9]) in terms of the cyclic interval chromatic number, it follows that the theory of cyclically ordered graphs $T_{\text{Graph}} \cup T_{\text{CycOrder}}$ also satisfies the principality property. However, it was observed in [41] that the theory of edge-ordered graphs does not satisfy the principality property. A natural question then is what theories satisfy the principality property?

Let us also note that just as Theorem 5.0.1, Theorem 5.2.1 also fails to completely characterize the asymptotic behavior of the maximum number of copies of K_t in $I(M)$ for $M \in \mathcal{M}[T]$ when $\chi(I) \leq t$. Even for the case $t = 2$, the study of this problem when $\chi(I) \leq 2$ has been done in a case by case manner and we refer the interested reader again to [57, 9, 41, 64] for some of these results for graphs with extra structure.

In Section 5.4 we proved the finiteness of the partite Ramsey numbers, but we made no attempt at optimizing the upper bounds that can be derived from its proof. Just as with the classical Ramsey numbers, providing good upper bounds is a very interesting problem in its own right and some work has been done in the non-partite case for some specific theories [55, 26, 19, 27, 6].

Let us also point out that the partite Ramsey numbers that we studied can be viewed as the diagonal case. The non-diagonal case can be defined as follows: given a function $h\colon \mathcal{P}_{\ell,\mathcal{L}} \to \mathbb{N}^\ell$ and $\vec{n} \stackrel{\text{def}}{=} (n_1, \ldots, n_\ell) \in \mathbb{N}^\ell$, we write $\vec{n} \xrightarrow{T} h$ if for every model M of T and every ℓ-split order $(f, \preceq) \in \mathcal{S}_{\ell,V(M)}$ on $V(M)$ with $|f^{-1}(i)| \geq n_i$ for all $i \in [\ell]$, there exists an ℓ-Ramsey pattern $Q \in \mathcal{P}_{\ell,\mathcal{L}}$ over $\mathcal{L}$ and a set $W \subseteq V(M)$ such that $|f^{-1}(i) \cap W| \geq h(Q)_i$ for all $i \in [\ell]$ and $M|_W$ is Q-uniform with respect to $(f|_W, \preceq_W)$. It follows that for $m \stackrel{\text{def}}{=} \max\{h(Q)_i \mid Q \in \mathcal{P}_{\ell,\mathcal{L}} \wedge i \in [\ell]\}$, if $\min\{n_i \mid i \in [\ell]\} \geq R_{\ell,T}(m)$, then $\vec{n} \xrightarrow{T} h$. Just as in the classical Ramsey theory, studying the off-diagonal case is an interesting problem as well.

References

[1] Elad Aigner-Horev, David Conlon, Hiệp Hàn, Yury Person, and Mathias Schacht. Quasirandomness in hypergraphs. *Electron. J. Combin.*, 25(3):Paper 3.34, 22, 2018.

[2] Noga Alon and Clara Shikhelman. Many T copies in H-free graphs. *J. Combin. Theory Ser. B*, 121:146–172, 2016.

[3] Noga Alon and Joel H. Spencer. *The probabilistic method*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons, Inc., Hoboken, NJ, third edition, 2008. With an appendix on the life and work of Paul Erdős.

[4] Ashwini Aroskar and James Cummings. Limits, regularity and removal for finite structures. Technical Report arXiv:1412.8084 [math.LO], arXiv e-print, 2014.

[5] T. Austin and T. Tao. Testability and repair of hereditary hypergraph properties. *Random Structures Algorithms*, 36(4):373–463, 2010.

[6] Martin Balko and Máté Vizer. Edge-ordered Ramsey numbers. *European J. Combin.*, 87:103100, 11, 2020.

[7] J. Barwise. *Handbook of Mathematical Logic. Studies in Logic and the Foundations of Mathematics, vol. 90.* North-Holland, 1977.

[8] V. I. Bogachev. *Measure theory. Vol. I, II.* Springer-Verlag, Berlin, 2007.

[9] Peter Brass, Gyula Károlyi, and Pavel Valtr. A Turán-type extremal theory of convex geometric graphs. In *Discrete and computational geometry*, volume 25 of *Algorithms Combin.*, pages 275–300. Springer, Berlin, 2003.

[10] Timothy F. N. Chan, Daniel Král', Jonathan A. Noel, Yanitsa Pehova, Maryam Sharifzadeh, and Jan Volec. Characterization of quasirandom permutations by a pattern sum. *Random Structures Algorithms*, 57(4):920–939, 2020.

[11] C. C. Chang and H. J. Keisler. *Model Theory. Studies in Logic and the Foundations of Mathematics, vol. 73.* North-Holland, 1973.

[12] Sourav Chatterjee and Persi Diaconis. A central limit theorem for a new statistic on permutations. *Indian J. Pure Appl. Math.*, 48(4):561–573, 2017.

[13] Artem Chernikov and Henry Towsner. Hypergraph regularity and higher arity vc-dimension. Technical Report arXiv:2010.00726 [math.CO], arXiv e-print, 2020.

[14] F. R. K. Chung and R. L. Graham. Quasi-random hypergraphs. *Random Structures Algorithms*, 1(1):105–124, 1990.

[15] F. R. K. Chung and R. L. Graham. Quasi-random tournaments. *J. Graph Theory*, 15(2):173–198, 1991.

[16] F. R. K. Chung and R. L. Graham. Maximum cuts and quasirandom graphs. In *Random graphs, Vol. 2 (Poznań, 1989)*, Wiley-Intersci. Publ., pages 23–33. Wiley, New York, 1992.

[17] F. R. K. Chung, R. L. Graham, and R. M. Wilson. Quasi-random graphs. *Combinatorica*, 9(4):345–362, 1989.

[18] Fan R. K. Chung. Quasi-random classes of hypergraphs. *Random Structures Algorithms*, 1(4):363–382, 1990.

[19] David Conlon, Jacob Fox, Choongbum Lee, and Benny Sudakov. Ordered Ramsey numbers. *J. Combin. Theory Ser. B*, 122:353–383, 2017.

[20] David Conlon, Jacob Fox, and Benny Sudakov. Hereditary quasirandomness without regularity. *Math. Proc. Cambridge Philos. Soc.*, 164(3):385–399, 2018.

[21] Jacob W. Cooper, Daniel Král', Ander Lamaison, and Samuel Mohr. Quasirandom Latin squares. Technical Report arXiv:2011.07572 [math.CO], arXiv e-print, 2020.

[22] Joshua N. Cooper. Quasirandom permutations. *J. Combin. Theory Ser. A*, 106(1):123–143, 2004.

[23] Joshua N. Cooper. Quasirandom arithmetic permutations. *J. Number Theory*, 114(1):153–169, 2005.

[24] L. N. Coregliano and A. A. Razborov. Semantic limits of dense combinatorial objects. *Uspekhi Mat. Nauk*, 75(4(454)):45–152, 2020.

[25] Leonardo N. Coregliano and Alexander A. Razborov. On the density of transitive tournaments. *J. Graph Theory*, 85(1):12–21, 2017.

[26] Christopher Cox and Derrick Stolee. Ordered Ramsey numbers of loose paths and matchings. *Discrete Math.*, 339(2):499–505, 2016.

[27] Christopher Cox and Derrick Stolee. Ramsey numbers for partially-ordered sets. *Order*, 35(3):557–579, 2018.

[28] Domingos Dellamonica, Jr. and Vojtěch Rödl. Hereditary quasirandom properties of hypergraphs. *Combinatorica*, 31(2):165–182, 2011.

[29] József Dénes and Donald Keedwell. *Latin squares and their applications.* New York-London: Academic Press, 1974.

[30] P. Diaconis, S. Holmes, and S. Janson. Interval graph limits. *Ann. Comb.*, 17(1):27–52, 2013.

[31] P. Diaconis and S. Janson. Graph limits and exchangeable random graphs. *Rendiconti di Matematica, Serie VII*, 28:33–61, 2008.

[32] Persi Diaconis, Steven N. Evans, and Ron Graham. Unseparated pairs and fixed points in random permutations. *Adv. in Appl. Math.*, 61:102–124, 2014.

[33] Persi Diaconis and David Freedman. On the statistics of vision: the Julesz conjecture. *J. Math. Psych.*, 24(2):112–138, 1981.

[34] Gábor Elek and Balázs Szegedy. A measure-theoretic approach to the theory of dense hypergraphs. *Adv. Math.*, 231(3-4):1731–1772, 2012.

[35] P. Erdős. On the number of complete subgraphs contained in certain graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 7:459–464, 1962.

[36] P. Erdős and L. Moser. On the representation of directed graphs as unions of orderings. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 9:125–132, 1964.

[37] P. Erdős and M. Simonovits. A limit theorem in graph theory. *Studia Sci. Math. Hungar.*, 1:51–57, 1966.

[38] P. Erdös and A. H. Stone. On the structure of linear graphs. *Bull. Amer. Math. Soc.*, 52:1087–1091, 1946.

[39] Herbert Federer. *Geometric measure theory*. Die Grundlehren der mathematischen Wissenschaften, Band 153. Springer-Verlag New York Inc., New York, 1969.

[40] Frederik Garbe, Robert Hancock, Jan Hladký, and Maryam Sharifzadeh. Limits of Latin squares. Technical Report arXiv:2010.07854 [math.CO], arXiv e-print, 2020.

[41] D. Gerbner, A. Methuku, D. T. Nagy, D. Pálvölgyi, G. Tardos, and M. Vizer. Edge ordered Turán problems. *Acta Math. Univ. Comenian. (N.S.)*, 88(3):717–722, 2019.

[42] Ronald L. Graham, Bruce L. Rothschild, and Joel H. Spencer. *Ramsey theory*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons, Inc., New York, second edition, 1990. A Wiley-Interscience Publication.

[43] Hiep Hán, Marcos Kiwi, and Matías Pavez-Signé. Quasi-random words and limits of word sequences. Technical Report arxiv:2003.03664 [math.CO], arXiv e-print, 2020.

[44] C. Hoppen, Y. Kohayakawa, C. G. Moreira, B. Ráth, and R. M. Sampaio. Limits of permutation sequences. *Journal of Combinatorial Theory, ser. B*, 103:93–113, 2013.

[45] Svante Janson. Quasi-random graphs and graph limits. *European J. Combin.*, 32(7):1054–1083, 2011.

[46] Subrahmanyam Kalyanasundaram and Asaf Shapira. A note on even cycles and quasirandom tournaments. *J. Graph Theory*, 73(3):260–266, 2013.

[47] Yoshiharu Kohayakawa, Brendan Nagle, Vojtěch Rödl, and Mathias Schacht. Weak hypergraph regularity and linear hypergraphs. *J. Combin. Theory Ser. B*, 100(2):151–160, 2010.

[48] Yoshiharu Kohayakawa, Vojtěch Rödl, and Jozef Skokan. Hypergraphs, quasi-randomness, and conditions for regularity. *J. Combin. Theory Ser. A*, 97(2):307–352, 2002.

[49] Daniel Král' and Oleg Pikhurko. Quasirandom permutations are characterized by 4-point densities. *Geom. Funct. Anal.*, 23(2):570–579, 2013.

[50] John Lenz and Dhruv Mubayi. Eigenvalues and linear quasirandom hypergraphs. *Forum Math. Sigma*, 3:e2, 26, 2015.

[51] John Lenz and Dhruv Mubayi. The poset of hypergraph quasirandomness. *Random Structures Algorithms*, 46(4):762–800, 2015.

[52] John Lenz and Dhruv Mubayi. Eigenvalues of non-regular linear quasirandom hypergraphs. *Discrete Math.*, 340(2):145–153, 2017.

[53] László Lovász. *Large networks and graph limits*, volume 60 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2012.

[54] László Lovász and Balázs Szegedy. Limits of dense graph sequences. *J. Combin. Theory Ser. B*, 96(6):933–957, 2006.

[55] Jaroslav Nešetřil and Vojtěch Rödl. The partite construction and Ramsey set systems. volume 75, pages 327–334. 1989. Graph theory and combinatorics (Cambridge, 1988).

[56] John C. Oxtoby. *Measure and category*, volume 2 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, second edition, 1980. A survey of the analogies between topological and measure spaces.

[57] János Pach and Gábor Tardos. Forbidden paths and cycles in ordered graphs and matrices. *Israel J. Math.*, 155:359–380, 2006.

[58] F. P. Ramsey. On a Problem of Formal Logic. *Proc. London Math. Soc. (2)*, 30(4):264–286, 1929.

[59] Alexander A. Razborov. Flag algebras. *J. Symbolic Logic*, 72(4):1239–1282, 2007.

[60] Asaf Shapira. Quasi-randomness and the distribution of copies of a fixed graph. *Combinatorica*, 28(6):735–745, 2008.

[61] Miklós Simonovits and Vera T. Sós. Hereditarily extended properties, quasi-random graphs and not necessarily induced subgraphs. *Combinatorica*, 17(4):577–596, 1997.

[62] Miklós Simonovits and Vera T. Sós. Hereditary extended properties, quasi-random graphs and induced subgraphs. volume 12, pages 319–344. 2003. Combinatorics, probability and computing (Oberwolfach, 2001).

[63] Richard Stearns. The voting problem. *Amer. Math. Monthly*, 66:761–763, 1959.

[64] Gábor Tardos. Extremal theory of vertex or edge ordered graphs. In *Surveys in combinatorics 2019*, volume 456 of *London Math. Soc. Lecture Note Ser.*, pages 221–236. Cambridge Univ. Press, Cambridge, 2019.

[65] Andrew Thomason. Pseudorandom graphs. In *Random graphs '85 (Poznań, 1985)*, volume 144 of *North-Holland Math. Stud.*, pages 307–331. North-Holland, Amsterdam, 1987.

[66] Henry Towsner. σ-algebras for quasirandom hypergraphs. *Random Structures Algorithms*, 50(1):114–139, 2017.

[67] Paul Turán. Eine Extremalaufgabe aus der Graphentheorie. *Mat. Fiz. Lapok*, 48:436–452, 1941.

[68] V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity*, pages 11–30. Springer, Cham, 2015. Reprint of Theor. Probability Appl. **16** (1971), 264–280.

[69] V. N. Vapnik and A. Ja. Červonenkis. The uniform convergence of frequencies of the appearance of events to their probabilities. *Dokl. Akad. Nauk SSSR*, 181:781–783, 1968.

[70] V. N. Vapnik and A. Ja. Červonenkis. The uniform convergence of frequencies of the appearance of events to their probabilities. *Teor. Verojatnost. i Primenen.*, 16:264–279, 1971.

[71] Raphael Yuster. Quasi-randomness is determined by the distribution of copies of a fixed graph in equicardinal large sets. *Combinatorica*, 30(2):239–246, 2010.